Automated Exchange Economies*

Bryan R. Routledge
Tepper School of Business
Carnegie Mellon University

Yikang Shen
Tepper School of Business
Carnegie Mellon University

Ariel Zetlin-Jones
Tepper School of Business
Carnegie Mellon University

February 29, 2024

Abstract:
The canonical mechanism for financial asset exchange is the limit-order book. In decentralized blockchain ledgers (DeFi), costs and delays in appending new blocks to the ledger render a limit-order book impractical. Instead, a "pricing curve" is specified (e.g., the "constant product pricing function") and implemented using smart contracts deployed to the ledger. We develop a framework to study the equilibrium properties of such markets. Our framework provides new insights into how informational frictions distort liquidity provision in DeFi markets.

*We thank Bruno Biais, Cyril Monnet, and participants at conferences and seminar audiences at the CMU Secure Blockchain Summit, Ripple UBRI Partners Conference, HEC Paris Blockchain and Digital Assets Conference, and the Richmond Fed for useful comments and discussions. We are especially grateful to Julien Prat for an insightful discussion. We are also thankful for the financial support of the Ripple Foundation. Contact authors at: azj@andrew.cmu.edu

1 Introduction

Decentralized finance, or DeFi, summarizes classical financial intermediation activities that take place on blockchains and in other decentralized marketplaces. DeFi represents a growing and sizable portion of trade in blockchain-based assets. One of the core activities of DeFi is automated market-making. Automated markets are smart-contract-based exchanges that facilitate trading or swaps of tokens that are native to a decentralized distributed ledger. Due to the costly nature of blockchain-based communication, running a limit order book on a blockchain is prohibitively expensive. As a consequence, we have observed a rapid pace of innovation and development in efforts to provide intermediation services, known as automated market makers (AMMs), on blockchains.

While AMMs economize on transaction costs, by design, they are passive and do not react to the information available in the market. As a result, they are susceptible to losses. This is called impermanent losses due to adverse selection (see Glosten and Milgrom (1985)). In this paper, we undertake a systematic analysis of the design of AMMs.

One might expect these automated markets to resemble a central limit order book that runs as code—smart contracts—that happens to be executed on the blockchain by the decentralized network of ledger validators or "miners." A challenge to establishing a central limit order book is that this method of intermediation requires an incredibly high volume of messages to be recorded. Each bid–even each update to a bid such as a modified price or a cancellation—is a message that changes the state of the blockchain ledger. Since each change of the blockchain state requires a (non-negligible) transaction cost, replicating a central limit order book on a blockchain is prohibitively expensive.

To facilitate trading on blockchains then, AMMs, such as Uniswap or Curve, have instead deployed ad-hoc pricing functions to the ledger that define terms of trade between liquidity providers (depositors) and liquidity takers. Specifically, suppose liquidity providers have contributed quantities Q_A and Q_B of tokens A and B and a liquidity taker would like to swap q_A units of token A for token B. Then the AMM pricing function G dictates that the liquidity taker may withdraw $q_B = G(q_A; Q_A, Q_B)$ units of token B (for some arbitrarily specified function G.[1] A portion of the contributed tokens q_A augment the quantity of tokens in the pool Q_A and a portion is paid out to liquidity providers as fees. A commonly used functional form is the geometric mean function wherein the geometric mean of the pre-trade positions of the AMM in the two tokens equals to that of the post-trade positions—save for the fees charged by the AMM.

While we observe a great deal of trial and error of different pricing functions, fee structures, and code development for AMMs, there is little systematic analysis of the underlying market micro-structure of AMMs. What are the gains to trade between liquidity providers and liquidity takers? What are the potential losses from the inflexibility of the functional form in the face of informed traders? How does the design of the AMM pricing function impact AMM volume and the division of surplus between liquidity providers and liquidity takers? And finally, what is the optimal design of AMM pricing, which is robust to a variety of beliefs about the potential returns to holding tokens? There is a nascent literature studying the outcomes of AMMs and, specifically, the opportunity cost of providing liquidity—sometimes referred to as impermanent loss. See Milionis et al. (2022) for a recent example. In this paper, we develop the first comprehensive framework to examine these questions.

Our proposed framework begins by specifying potential gains to trade between liquidity providers and liquidity takers—something essentially absent from the emerging literature on AMMs. These gains arise in our model due to heterogeneous beliefs in the (relative) value of a pair of tokens as in Harrison and Kreps (1978). Implicitly, we think of liquidity providers as "slow" traders who are less able to obtain transaction priority on the blockchain and, therefore, less able to take advantage of high-frequency arbitrage opportunities.[2] Such agents are the natural liquidity providers in our environment.

Liquidity takers, on the other hand, we model as "fast" and able to attain priority for blockchain execution. In our framework, liquidity takers may be "informed" or "unin-

[1] In theory, one could represent a central limit order book in this fashion where the function G depends on the entire set of messages—bids and asks—relayed to the exchange. However, here, consistent with what we observe at AMMs, we focus on G that depend solely on quantities

[2] Existing work on impermanent loss implicitly assumes the opportunity cost of liquidity provision is the ability to profit from such high-frequency opportunities Milionis et al. (2022).

formed,", giving rise to a classic form of adverse selection in asset markets—see Glosten and Milgrom (1985). While we use the language of informed and uninformed trading, our preferred interpretation is rather that uninformed traders trade for reasons that are orthogonal to liquidity providers' beliefs about the value of the tokens. Instead, informed traders trade for reasons that are correlated with (changes) in liquidity providers' beliefs about the value of the tokens.

To the extent that AMM pricing cannot flexibly react to the news or information available in the market, and to the extent that informed traders are able to trade at the AMM *before* the liquidity providers may withdraw their deposits—again, liquidity providers are slow traders relative to liquidity takers—informed traders create losses for liquidity providers.

We explore how this form of adverse selection distorts the amount of liquidity contributed by providers who must balance profits they earn from uninformed liquidity takers (noise traders) with the losses that arise from trading with informed liquidity takers. Our results provide modern analogs to those in Glosten and Milgrom (1985) in a smart contract setting and offer a new interpretation of impermanent loss—committing to trade with informed liquidity takers at "stale" prices—stemming from a traditional notion of adverse selection. While in Glosten and Milgrom (1985) liquidity providers distort *prices* to protect themselves from informed trading losses, such distortions may only manifest in the quantities of deposits liquidity providers post in the AMM.

The conventional wisdom shared as guidance on major AMMs is that liquidity providers should deposit liquidity in equal US dollar value amounts. For example, Uniswap explains that liquidity providers "are incentivized to deposit an equal value of both tokens into the pool. To see why, consider the case where the first liquidity provider deposits tokens at a ratio different from the current market rate. This immediately creates a profitable arbitrage opportunity, which is likely to be taken by an external party." [3] This simple logic, while correct, ignores the concept that some fast traders may trade for reasons orthogonal to current market prices should they require liquidity. Effectively, conventional

[3] See Uniswap-V2 (2023) https://docs.uniswap.org/contracts/v2/concepts/core-concepts/pools

wisdom assumes there is a single agreed-upon "market price".

Instead, we argue that (fast) liquidity takers may have heterogeneous beliefs or heterogeneous reasons for trade and therefore, depositing at equal *value* may not be optimal. Indeed, we show that in any equilibrium where liquidity providers in sum earn strictly positive profits, they prefer to distort their deposit ratio away from equal value. Such changes allow them to earn higher profits per trade should the first trader be uninformed. We find that in if the fraction of informed traders is such that the profits from uninformed traders exactly balance the losses from informed traders, then it is optimal for liquidity providers to deposit tokens in equal value (according to the liquidity providers' expected valuation of the tokens).

We go on to explore how the shape of the pricing function G impacts gains to trade and liquidity provider's profits. Analogous to results in Milionis, Moallemi and Roughgarden (2023*b*), we find that in the presence of only uninformed traders, convex prices impede ex post trading volumes and reduce ex-ante profits of liquidity providers. Hence, in such a case, linear pricing is optimal. However, the presence of informed traders complicates this analysis because convex prices also limit the losses liquidity providers realize from informed trades. Nonetheless, we show that reducing the (local) convexity of the pricing function improves liquidity provider's profits as long as liquidity provision is profitable. Specifically, we construct a perturbation of the pricing function that decreases its convexity around the liquidity provider's deposit point and scales the gains from uninformed trades at the same rate as losses from adverse selection. If the original CPMM function induces positive ex-ante gains for the liquidity traders, then less locally convex prices increase ex-ante gains for both liquidity providers and liquidity traders, thus improving efficiency.

The early, existing literature on AMMs has focused on examining how AMMs perform along-side the presence of deep, liquid, centralized exchanges. One of the earliest examples is Angeris and Chitra (2020) who when a class of AMM mechanisms reflect "true" prices—those observed on an infinitely deep centralized limit order book. Angeris et al. (2021) presents a more specific analysis of the leading AMM Uniswap and show that the exchange rate on Uniswap matches the exogenous prices up to the interval of fee level. Aoyagi (2022) extends these frameworks to consider the effect of information asymmetry in AMMs shows that the equilibrium liquidity supply is stable under the assumptions that liquidity provision is perfectly competitive and one token in the pool is stable (its value has zero volatility). Also under the assumption of a known, true price of tokens, Fabi and Prat (2023) demonstrates how to use consumer choice theory to study how liquidity providers and liquidity takers exert externalities on each other. They use their framework to examine how the shape of constant function market makers impacts adverse selection costs faced by liquidity providers and execution costs faced by liquidity takers. More recently, Lehar and Parlour (2023) show how AMM fees can balance losses imposed by liquidity traders who conduct such an arbitrage. They argue that pool sizes should decrease with the severity of this arbitrage risk and find empirical support for this observation. Since these papers assume the opportunity to conduct a perfect arbitrage between the AMM and the centralized exchange, liquidity providers necessarily deposit tokens in equal value according to the centralized exchange prices. One of our contributions is to relax the assumption of perfect arbitrage and examine optimal liquidity provision when the notion of equal values is not clear because perfect price discovery is not possible.

A related literature has emerged studying the costs imposed by traders who arbitrage between centralized exchange prices and AMM prices. For example, Capponi and Jia (2021) studies competition for priority among traders who would like to conduct such an arbitrage and characterizes the joint determination of gas fees and liquidity pool sizes. Hasbrouck, Rivera and Saleh (2023) study the impact of trading fees on trading volume and show how an increase in the fees, by attracting more liquidity provision and thus reducing traders' execution costs may lead to increased trading volumes. Milionis et al. (2022) use a continuous-time Black-Scholes analysis to estimate these arbitrage losses for liquidity providers using a stablecoin pool and decomposes the losses into risky and predictable components. Milionis, Moallemi and Roughgarden (2023*a*) extend the model to involve trading fees and provide results on the arbitrager's behavior and profits accordingly. They also conduct a cost-benefit analysis on the LP's side with the new features.

In our model in the absence of a true price, the AMM generates gains to trade and so liquidity provision may be sustained even in the absence of direct fees.

In terms of the design and efficiency of the price function, Park (2023) demonstrates that constant function market makers may causes economically meaningless and costly trading, such as front running. Bergault et al. (2023) shows that the return of LP is always smaller than holding by duality theorem and a constant product formula with a proportional fee is not efficient from the mean-variance perspective. Goyal et al. (2023) focus on the design of convex pricing functions that maximize the fraction of trades that with only uninformed trades. Milionis, Moallemi and Roughgarden (2023*b*) uses the optimal auction framework to show that a linear price curve maximizes the expected return of the liquidity provider when one token is a stablecoin. Our results on the optimal shape of the design function are similar to those in Milionis, Moallemi and Roughgarden (2023*b*) but hold under a wider set of assumptions on traders' beliefs about the token valuations.

2 Model

An AMM is a blockchain-based automated market that uses smart contracts to permit individuals to exchange cryptocurrencies. The smart contract is a computer script stored on the blockchain. Liquidity providers and liquidity takers post transactions that are then executed by a decentralized network of validators (or "miners"). Functionally, the AMM smart contract is immutable, so the trade process is fixed. The typical AMM smart contract is specific to two coins; we will call these coins A and B. The contract defines a Liquidity Provider (LP) as one who deposits both coins A and B. Later, the LP may withdraw both coins A and B. In contrast, the contract defines a Liquidity Taker (LT) by the transaction of depositing one coin (e.g., A) and withdrawing the other (e.g., B). The smart contract also specifies the function that maps the number of coins the LP deposits to the exchange rate between the two coins received by the LT.

To model the costs and benefits of the AMM setting, we model traders' relative value of the coins as a mix of common and private values. The common value component is public information. However, the trade arrival is sequential and so some traders will be "informed" in that they have arrived at the same time as new information. The private value component motivates gains to trade. The public component creates the potential for an "adverse selection" cost. This cost is sometimes called "impermanent loss" in the AMM setting.

Information. Time is discrete, $t \in \{0, 1, ..., T\}$, and coins have a terminal value at date T given by $\exp(d_{i,T})$. We interpret the terminal value $\exp(d_{i,T})$ as either the future "price" of token i or possibly the service flow attainable by holding 1 unit of token i. For example, 1 unit of the Ethereum cryptocurrency may be "spent" on the execution of smart contracts on the Ethereum blockchain or 1 unit of the stablecoin USDC may be redeemed for 1 US dollar by trading with the company Circle who issues USDC (Cicle (2023)). We assume the "dividend" or terminal common payoff of coin i evolves according to

$$d_{i,T} = \sum_{t=0}^{T-1} y_{i,t} + \epsilon_i$$

with the public information at each date $y_{i,t}$ and the residual uncertainty, ϵ_i realized in period T independently and satisfying $E[\exp(\epsilon_i)] = 1$.

In particular, assume public information $\{y_{i,t}\}$ arrives independently across date. With probability $\hat{\pi}$, $y_{A,t} = y_{B,t} = 0$. With probability $1 - \hat{\pi}$, $y_{i,t} \in \{-\Delta_l, +\Delta_h\}$ where each is equally likely. We assume Δ_l, Δ_h are positive and $\frac{1}{2}e^{-\Delta_l} + \frac{1}{2}e^{\Delta_h} = 1$ such that the expected value is the same as the current realized value. At the beginning of each period, both LPs and LTs have beliefs about the terminal, common value component of each token given by $\mu_{i,t}$ where

$$\mu_{i,t} = E[\exp(d_{i,T})|y_0, \ldots, y_{t-1}] \equiv E_t[\exp(d_{i,T})].$$

Timing. In each period, before public information is realized, the LP decides how much of each token to deposit in the AMM smart contract. Once the LP deposits tokens, public information is realized. After public information is realized, LTs value the tokens according to

$$\hat{\mu}_{i,t} = E_{t+1}[\exp(d_{i,T})]\exp(\eta_i)$$

where η_i reflects a private value component of owning token i realized by the LT that trades in period t. Our timing implies that the LT is a "fast trader" who may modify their liquidity position before the LP may re-balance their supplied liquidity. Blockchain transaction ordering depends on the decisions of the miners who maintain the blockchain ledger. Because space for transactions in these ledgers is scarce relative to the demand to place transactions on the ledger, there is a market for priority. Our model captures the idea that the "natural" liquidity providers are those who are unlikely know how to obtain priority (or are unwilling to pay for priority) while those who are efficient at obtaining priority are likely to be liquidity takers.

Once the LT trades, a new period begins and the LP may re-balance the liquidity supplied to the AMM.

2.1 One-Period Model

We now specialize this information setting to a static, one-period model. While we model these as individual agents, we think of the LP and LT as representatives of a pool of potential liquidity providers and takers, respectively. The representative liquidity provider is initially endowed with a portfolio of tokens (E_A, E_B). We assume the LT has deep pockets and cares only about her net trading profits. At the beginning of the period, before any new public information is realized, both the liquidity providers and liquidity takers have the same beliefs, μ_i where we drop the time subscripts for ease of notation. (One may think of this as $\mu_{i,T} = E_T[\exp(d_{i,T})]$, the beliefs of LPs and LTs upon entering the last period of the dynamic game described above before any public information in period T is realized.) Once public information arrives, the LT who trades has valuation $\hat{\mu}_i$ distinct from μ_i. (Using notation from the last period of the game, we have

$\hat{\mu}_{i,T} = E_{T+1}[\exp(d_{i,T})] \exp(\eta_i)$.)

Here, we make one more specific assumption to leverage the insights from this setting. Recall that with probability $1 - \hat{\pi}$ trader j arrives where $y_{i,t} \in \{-\Delta_l, \Delta_h\}$. For that trader, we set $\eta_i = 0$. Under this specification, our model features two types of information events as in Glosten and Milgrom (1985). The first type of information event—analogous to uninformed trading in Glosten and Milgrom (1985)—occurs when $y_{i,t} = 0$ and represents a case where the LT's new beliefs of the tokens' values, $\hat{\mu}_i$ are uncorrelated with the LP's beliefs. That is, the LP believes the value of each token i will yield terminal value according to $E_{t+1}[\exp(d_{i,T})] = E_t[\exp(d_{i,T})]$ while the LT believes the value of each token i is distributed according to $\hat{\mu}_i = E_t[\exp(d_{i,T})] exp(\eta_i)$. When $\eta_i \neq 0$ under such an event, there are gains to trade between the LP and the LT. Following the literature, we interpret such an event as a "pure noise" trade where trade occurs for reasons orthogonal to the LP's beliefs about the potential returns to her tokens. We let $\pi \in [0, 1]$ denote the probability of this first type of information event which we describe as a *trade for tastes* or *uninformed trade*.

Instead, the second type of information event—analogous to informed trading in Glosten and Milgrom (1985)—occurs when $y_{i,t} \in \{-\Delta_l, \Delta_h\}$ (for some token i) and represents a case where the LT's new beliefs are correlated with the LP's new beliefs. In such a case both the LP and the LT now believe the value of each token has mean $\hat{\mu}_i = E_{t+1}[\exp(d_{i,T})]$ and hence there can be no gains to trade between the LT and the LP. Following the literature, we interpret such an event as pure information event that we describe as an *informed trade*. This imposed correlation between the information arrival and private values of the LT lets us isolate the idea that liquidity takers are either trading for "information" or are trading for "tastes."

Suppose the LP has deposited a portfolio (e_A, e_B) with the smart contract of the AMM. We let $G(\cdot)$ the embedded pricing function. That is, if the LT wishes to deposit (withdraw) q_A units of token A then the function specifies an amount q_B units of token B that the LT may withdraw (deposit) where

$$q_B = G(q_A | e_A, e_B).$$

We use the convention that if $q_A > 0$—the LT deposits token A—then $q_B < 0$—the LT may withdraw token B—and vice versa. Most AMM price functions also have the property that q_B is decreasing in q_A so that the LT must pay more of token A per unit of token B she wishes to withdraw. The price function G also requires $q_i \leqslant e_i$, at least implicitly, as capacity constraints.

The most common implementation of automated markets imposes the constant product market maker (CPMM):

$$(e_A + q_A)(e_B - q_B) = e_A e_B. \tag{1}$$

This particular function was originally proposed by Bergault et al. (2023) and was then adopted in Uniswap-V2 (2023).

To summarize the static model, at the beginning of the period, the LP deposits a portfolio (e_A, e_B) with the AMM given a pricing function $G(\cdot)$ given her beliefs (μ_a, μ_b). Next, the type of information event is realized according to π and the LT realizes a shock to her beliefs specified by $(\hat{\mu}_A, \hat{\mu}_B)$. With probability π the LT is uninformed and the LP's beliefs remain (μ_A, μ_B). With probability $1 - \pi$ the LT is informed and the LP's beliefs also shift to $(\hat{\mu}_A, \hat{\mu}_B)$. In either case, once information is realized the LT then chooses an amount to trade with the AMM. Finally, values and payoffs are realized according to the terminal portfolios of the LP and LT.

Next, we define the problem of the liquidity taker and the liquidity provider working backwards from the LT's problem. We maintain the Constant Product Market Making rule specified in Equation (1) through Section 2.2, 2.3, 3 below.

2.2 The Liquidity Taker's Problem

The LT—whether in an uninformed or informed trading event—observes liquidity on deposit at the AMM as well as her realization of $\hat{\mu}_i$. From her own perspective, the LT perceives an arbitrage opportunity as prices at the AMM do not reflect her realized beliefs of the token values.

The LT maximizes the expected value of her tokens:

$$\max_{q_A, q_B} - \hat{\mu}_A q_A + \hat{\mu}_B q_B \tag{2}$$
$$\text{s.t. } (e_A + q_A)(e_B - q_B) = e_A e_B$$

When $q_A > 0$, the LT's problem given in (2) represents a case where the LT "buys" token B from the AMM by depositing token A. She may wish to set $q_A < 0$ in which case she buys token A from the exchange by depositing some amount of token B. The constraint represents the effective price she faces on any trade. Under the Constant Product rule, the LT would have to deposit infinitely much of one token to withdraw all of the other (i.e. setting $q_B = e_B$, requires $q_A \to -\infty$) and hence the implicit capacity constraints are slack under such a rule.

The solution to the LT's problem is straightforward and satisfies

$$e_A + q_A = \sqrt{\frac{\hat{\mu}_B}{\hat{\mu}_A} e_A e_B}, \quad e_B - q_B = \sqrt{\frac{\hat{\mu}_A}{\hat{\mu}_B} e_A e_B}. \tag{3}$$

More importantly, for any beliefs $\hat{\mu}_i$, she will trade up until the relative price at the AMM equals her relative valuation of the tokens or

$$\frac{\hat{\mu}_B}{\hat{\mu}_A} = \frac{e_A + q_A}{e_B - q_B}. \tag{4}$$

If we let $x_A = e_A + q_A$ and $x_B = e_B - q_B$ denote the LP's post-trade portfolio, then (1) and

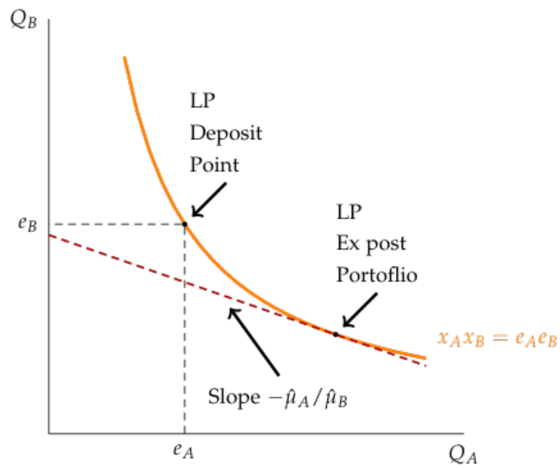

Figure 1: Liquidity Taker's Optimal Trade

(4) imply that LP's post-trade portfolio satisfies

$$x_A x_B = e_A e_B \tag{5}$$

$$\hat{\mu}_A x_A = \hat{\mu}_B x_B. \tag{6}$$

The liquidity provider internalizes that for any realization of beliefs of the LT, $\hat{\mu}_i$, her ex-post portfolio will satisfy (5)–(6). We may represent this behavior graphically as in Figure 1.

The convex curve represents the constant product market-making rule, and the point (e_A, e_B) represents the liquidity deposited by the LP. Any trade by the LT will move the LP's ex-post portfolio along the convex curve. Once the LT realizes her beliefs $\hat{\mu}_i$, she will trade up until the relative price at the AMM equals her relative valuation of the tokens (represented by the dashed line with slope $-\hat{\mu}_A/\hat{\mu}_B$.

2.3 The Liquidity Provider's Problem

Anticipating the behavior of the liquidity taker, the LP chooses her liquidity deposit to solve the following program.

$$\begin{aligned}\max_{e_A, e_B}\ & \pi(\mu_A \mathbb{E}[x_A - e_A] + \mu_B \mathbb{E}[x_B - e_B]) + \\ & \quad (1-\pi)(\mathbb{E}\hat{\mu}_A[x_A - e_A] + \mu_B \mathbb{E}\hat{\mu}_B[x_B - e_B]) \\ \text{s.t. }\ & (5)\text{–}(6), \\ & 0 \leqslant e_i \leqslant E_i, \quad \forall i\end{aligned} \tag{7}$$

where π is the probability of an uninformed trading event. Notice, regardless of whether the LP experiences an uninformed or informed trading event, the beliefs of the liquidity taker will result in an ex-post portfolio of the LP according to (5)–(6). These events differ, however, in how the LP perceives the value of these ex-post portfolios. When the LT represents an uninformed trade, the LP continues to value her ex-post portfolio according to her prior beliefs, μ_i. Instead, when the LT represents an informed trade, the LP values her ex-post portfolio according to the realized beliefs of the LT, $\hat{\mu}_i$. As we show below, the LP will trade off profits she earns on uninformed trades with losses on informed trades. Unlike in standard models of exchange subject to adverse selection where market makers post prices that reflect the extent of adverse selection, blockchain market makers must distort their quantity choices for liquidity provision to protect themselves from possible adverse selection.

3 Equilibrium AMM Liquidity Provision

In this section, we examine equilibrium liquidity provision by liquidity providers in our model. Our notion of equilibrium is standard subgame perfect equilibrium. We examine the usefulness of the conventional wisdom from existing automated marketplaces—that liquidity providers *should* deposit liquidity in equal (dollar) values—and find that such behavior is optimal for the representative liquidity provider only under special circumstances. We demonstrate how adverse selection distorts the quantities of liquidity deposited by providers on automated exchanges. To ease the analysis, we first consider two special cases of our model—when all trade is uninformed and when all trade is informed—before turning to the general case.

3.1 Liquidity Provision with Uninformed Trade Only

Suppose first that $\pi = 1$ so that there are only uninformed trades. Using straightforward algebra, the LP's problem (7) simplifies to

$$\max_{e_A, e_B} \mu_A \left(\mathbb{E}\sqrt{\frac{\hat{\mu}_B}{\hat{\mu}_A} e_A e_B} - e_A \right) + \mu_B \left(\mathbb{E}\sqrt{\frac{\hat{\mu}_A}{\hat{\mu}_B} e_A e_B} - e_B \right)$$
$$\text{s.t. } 0 \leqslant e_i \leqslant E_i, \quad \forall i.$$

Since the LP's deposit quantities, e_i, are not random, her objective may be written as

$$\left(\mathbb{E}\omega + \mathbb{E}\frac{1}{\omega} - 2 \right) \sqrt{\mu_A e_A}\sqrt{\mu_B e_B} - \left(\sqrt{\mu_A e_A} - \sqrt{\mu_B e_B}\right)^2 \tag{8}$$

where $\omega = \sqrt{\frac{\hat{\mu}_A/\mu_A}{\hat{\mu}_B/\mu_B}}$. Equation (8) shows how an LP facing only uninformed trade chooses the optimal liquidity to provide. By changing the quantities of tokens A and B she deposits, she adjusts the position of the pricing curve the LT will face ex-post.

To better understand (8), consider one possible (suboptimal) deposit choice for the LP: an equal value deposit, or e_A and e_B that satisfy $\mu_A e_A = \mu_B e_B$. Notice that all possible ex-post portfolios for the LP lie on the constant product price function that runs through the point (e_A, e_B). Moreover, at (e_A, e_B), the constant product price function has slope $-\mu_A/\mu_B$. Since the constant product price function is convex, any trade by the LT will appear to happen at favorable prices from the perspective of the LP—that is, terms of trade are better than $-\mu_A/\mu_B$ for the LP regardless of whether the LT is buying token A or token B. As a result, for such a deposit choice, the LP only stands to gain and suffers no losses.

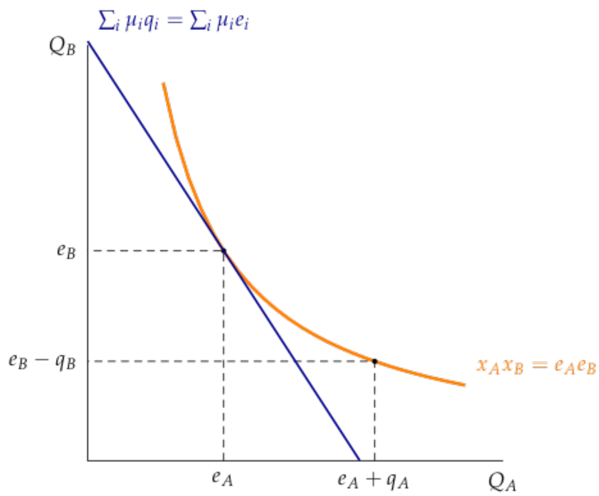

Figure 2: Liquidity Provider's "No-Loss" Deposit Choice

Figure 2 illustrates this result graphically. Given the LP's beliefs are fixed, facing only uninformed trades, the straight (blue) line with slope $-\mu_A/\mu_B$ reflects the LP's indifference curve. Since all terminal portfolios lie on the constant product price function, and this function lies above the LP's preferences, such a deposit choice by the LP ensures the LP only stands to gain from trade.

Should the LP provide liquidity different from an equal value deposit, then for small differences in beliefs from her own, the constant produce price function will provide prices that appear unfavorable from the perspective of the LP and yield second-order losses. For this reason, the LP faces a loss function—the second term in (8)—that depends on how her portfolio differs from an equal value ($\mu_A e_A = \mu_B e_B$) portfolio.

To the extent $\hat{\mu}_i$ differs from μ_i, there are gains to trade. The value of these gains depend on the term $\mathbb{E}\omega + \mathbb{E}\frac{1}{\omega} - 2 \geqslant 0$. (The inequality follows directly from Jensen's inequality.) As a result, from any equal value deposit, a small perturbation that raises e_A or e_B on the margin will induce second-order losses but incur first-order gains. As a result, equal-value deposits are generically not optimal for the LP. In general, the LP desires to provide as much liquidity as possible to facilitate gains to trade, and thus, her budget constraint must bind (either $e_A = E_A$ or $e_B = E_B$). We then have the following proposition.

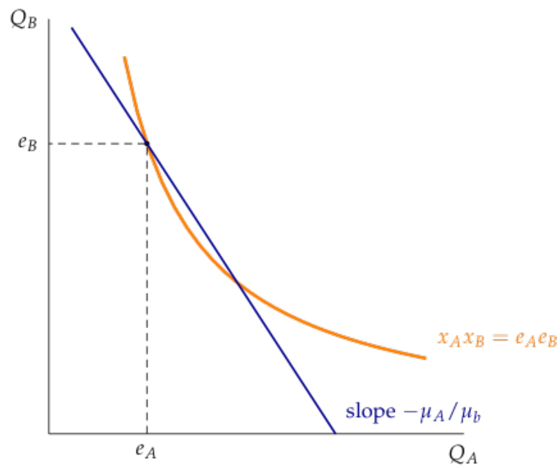

Figure 3: Liquidity Provider's Optimal Deposit Choice

Proposition 1: Optimal Liquidity with only Uninformed Trade. With only uninformed trade, the optimal liquidity deposit satisfies:

$$\begin{cases} e_A^* = E_A,\ e_B^* = \min\left\{\left(\frac{\mathbb{E}\omega + \mathbb{E}\frac{1}{\omega}}{2}\right)^2 \frac{\mu_A}{\mu_B} E_A, E_B\right\}, & \text{if } \mu_A E_A \leqslant \mu_B E_B \\ e_A^* = \min\left\{\left(\frac{\mathbb{E}\omega + \mathbb{E}\frac{1}{\omega}}{2}\right)^2 \frac{\mu_B}{\mu_A} E_B, E_A\right\},\ e_B^* = E_B, & \text{if } \mu_A E_A > \mu_B E_B. \end{cases}$$

Generically, then, the LP will prefer a deposit choice different from the equal value portfolio to maximize intermediation profits with uninformed traders. Such a choice is illustrated in Figure 3 where, according to Proposition 1 typically, we expect either $e_A = E_A$ or $e_B = E_B$.

3.2 Liquidity Provision with Informed Trade Only

Suppose next that $\pi = 0$ so that there are only informed trades. The LP's problem (7) simplifies to

$$\max_{e_A, e_B} \mathbb{E}\hat{\mu}_A \left(\sqrt{\frac{\hat{\mu}_B}{\hat{\mu}_A} e_A e_B} - e_A\right) + \mathbb{E}\hat{\mu}_B \left(\sqrt{\frac{\hat{\mu}_A}{\hat{\mu}_B} e_A e_B} - e_B\right) \tag{9}$$

$$\text{s.t. } 0 \leqslant e_i \leqslant E_i, \quad \forall i \tag{10}$$

If we impose a mild assumption that $\hat{\mu}_i$ is a mean preserving spread of μ_i, i.e. $\mathbb{E}\frac{\hat{\mu}_i}{\mu_i} = 1$, the LP's objective in this case may be written as

$$(2\mathbb{E}\psi - 2)\sqrt{\mu_A e_A}\sqrt{\mu_B e_B} - (\sqrt{\mu_A e_A} - \sqrt{\mu_B e_B})^2 \tag{11}$$

where $\psi = \sqrt{\frac{\hat{\mu}_A}{\mu_A}\frac{\hat{\mu}_B}{\mu_B}}$. Equation (11) shows how an LP facing only informed trade chooses the optimal liquidity to provide.

Since the LP and the LT hold the same ex-post belief, any gains of the LT must reflect losses borne by the LP. Moreover, since the LT only trades when it is beneficial for herself, all trades hurt the LP. As a result, the case of only informed trading reflects a case of pure adverse selection and induced losses for the LP relative to what the value of her wealth would have been had she simply held her portfolio rather than providing liquidity.[4]

Mathematically, the Cauchy-Schwarz inequality implies $\mathbb{E}\psi \leqslant \sqrt{\mathbb{E}\frac{\hat{\mu}_A}{\mu_A}\mathbb{E}\frac{\hat{\mu}_B}{\mu_B}}$ and holds with equality only when $\hat{\mu}_A$ and $\hat{\mu}_B$ are perfectly correlated. Since we impose $\mathbb{E}\hat{\mu}_i/\mu_i = 1$, the above inequality implies $\mathbb{E}\psi \leqslant 1$. Therefore, the LP's objective function is necessarily non-positive for any deposit amount, yielding our next proposition.

Proposition 2: No Liquidity Provision with Only Informed Trade. The optimal liquidity deposit satisfies:

$$e_A^* = e_B^* = 0.$$

3.3 Liquidity Provision with Uninformed and Informed Trading

We now use these results to understand better the general problem (7) with arbitrary π. We once again simplify the LP's objective function as

$$\left[\pi\left(\mathbb{E}\omega + \mathbb{E}\frac{1}{\omega}\right) + (1-\pi)2\mathbb{E}\psi - 2\right]\sqrt{\mu_A e_A}\sqrt{\mu_B e_B} - (\sqrt{\mu_A e_A} - \sqrt{\mu_B e_B})^2 \tag{12}$$

[4] Since we implicitly assume LPs are "slow" traders, we do not consider the opportunity cost of trading at an AMM herself. See Milionis et al. (2022) for such an analysis.

As before, we may write the LP's objective as the sum of a revenue function less losses that depend on how the LP's deposit portfolio differs from an equal value portfolio. The revenue function now reflects the probability of realizing an informed versus an uninformed trade. Similar to the previous cases, when uninformed trades occur the LP realizes profits and when informed trades occur, the LP realizes losses. If the gains from uninformed trades are larger than the loss from informed trades, i.e. $\pi\left(\mathbb{E}\omega + \mathbb{E}\frac{1}{\omega}\right) + (1-\pi)2\mathbb{E}\psi \geqslant 2$, then the LP will be willing to provide as much liquidity as possible—up to their ex-ante resource constraint. Otherwise, the LP will optimally choose to provide no liquidity. We summarize this result in the next proposition.

Proposition 3: Optimal Liquidity. The optimal liquidity deposit with π proportion of uninformed trade and $1-\pi$ proportion of informed trade satisfies

$$\begin{cases} e_A^* = E_A,\ e_B^* = \min\left\{\left(\pi\left(\frac{\mathbb{E}_U\omega + \mathbb{E}_U\frac{1}{\omega}}{2}\right) + (1-\pi)\mathbb{E}_I\psi\right)^2 \frac{\mu_A}{\mu_B}E_A, E_B\right\}, & \text{if } \mu_A E_A \leqslant \mu_B E_B \\ e_A^* = \min\left\{\left(\pi\left(\frac{\mathbb{E}_U\omega + \mathbb{E}_U\frac{1}{\omega}}{2}\right) + (1-\pi)\mathbb{E}_I\psi\right)^2 \frac{\mu_B}{\mu_A}E_B, E_A\right\},\ e_B^* = E_B, & \text{if } \mu_A E_A > \mu_B E_B \end{cases}$$

if $\pi\left(\mathbb{E}\omega + \mathbb{E}\frac{1}{\omega}\right) + (1-\pi)2\mathbb{E}\psi \geqslant 2$ and

$$e_A^* = e_B^* = 0$$

otherwise.

Given optimal liquidity provision, we next explore the optimality of the conventional wisdom that liquidity providers should deposit portfolios with equal values.

We write $\Pi = \pi\left(\frac{\mathbb{E}_U\omega + \mathbb{E}_U\frac{1}{\omega}}{2}\right) + (1-\pi)\mathbb{E}_I\psi$ to represent the LP's expected profit margin from liquidity provision. According to Proposition 3, if $\Pi > 1$, then the optimal value ratio $\mu_A e_A^* / \mu_B e_B^*$ satisfies

$$\frac{\mu_A e_A^*}{\mu_B e_B^*} = \begin{cases} \frac{1}{\Pi^2} & \text{if } E_A \leqslant \frac{1}{\Pi^2}\frac{\mu_B}{\mu_A} E_B \\ \frac{\mu_A E_A}{\mu_B E_B} & \text{if } \frac{1}{\Pi^2}\frac{\mu_B}{\mu_A} E_B < E_A < \Pi^2 \frac{\mu_B}{\mu_A} E_B \\ \Pi^2 & \text{if } \Pi^2 \frac{\mu_B}{\mu_A} E_B \leqslant E_A \end{cases} . \tag{13}$$

For $\Pi > 1$, unless $\mu_A E_A = \mu_B E_B$ then the optimal deposit ratio is always different from 1. However, Proposition 3 also reveals that as $\Pi \to 1$ then $\mu_A e_A^* \to \mu_B e_B^*$ for all values of E_A, E_B. In other words, only when the gains from uninformed trades exactly offset the losses from informed trades, then it is optimal for the LP to deposit a portfolio with equal values.

We note that the LP's expected profit margin Π is increasing in the probability that trades are uninformed, π. Hence, there is a minimal value π such that $\Pi = 1$. We then have the following Corollary.

Corollary 1: Optimal Value Share. Let $\underline{\pi}$ be such that $\Pi = 1$ and assume $\mu_A E_A \neq \mu_B E_B$.[5] The equal value deposit $\mu_A e_A = \mu_B e_B$ is optimal only when $\pi = \underline{\pi}$.

3.4 Break Even Proportion of Uninformed Trading

The threshold $\underline{\pi}$ also sheds light on the extent to which liquidity provision is profitable. The value of π such that $\Pi = 1$ depends critically on the distribution of the LT's beliefs specified by H_i. Since the term $\omega + \frac{1}{\omega}$ is not globally convex in $\hat{\mu}_i$, a mean preserving spread of the LT's beliefs $\hat{\mu}_i$ could increase or decrease the threshold $\underline{\pi}$. We instead explore how the profitability of liquidity provision varies with the distribution of the LT's beliefs via a numerical example.

[5] If the LP happens to be endowed with an equal value portfolio and profits from liquidity provision are increasing, then she may deposit in equal value simply because she is constrained. We rule out this uninteresting case with this assumption.

To simplify the numerical analysis, consider a special case where one token is a stablecoin whose value (purportedly) does not fluctuate over time such as USDC or Tether.[6] We let token B represent the stable coin and set $\hat{\mu}_B = \mu_B = 1$ and $h_B(\hat{\mu}_B) = 1$ if $\hat{\mu}_B = 1$. Then we have $\omega = \psi = \sqrt{\frac{\hat{\mu}_A}{\mu_A}}$. We assume $\frac{\hat{\mu}_A}{\mu_A}$ is a log-normally distributed random variable with $\mathbb{E}[\hat{\mu}_A/\mu_A] = 1$ and $Var[\hat{\mu}_A/\mu_A] = \sigma_A^2$. As a benchmark, we impose $\sigma_A^2 = 0.8$ consistent with variation in the daily price of ETH–the native cryptocurrency of the Ethereum blockchain–over the past five years.[7] Around this benchmark, we explore how changes in the variance of beliefs about ETH prices change the threshold probability for liquidity provision to be profitable, $\underline{\pi}$. We plot how this threshold varies with the variance of the LT's beliefs in Figure 4, which shows that increases in variance typically decrease this threshold. [8] In other words, liquidity provision becomes more profitable (LPs can tolerate more informed trading) as ETH price risk increases.

4 Efficiency Losses from Constant Product Market Making

In this section we examine how the shape of the AMM pricing function impacts gains to trade realized by liquidity providers. We focus on the (local) convexity of the CPMM price function and leave a full mechanism design perspective for future work (see Milionis, Moallemi and Roughgarden (2023*b*) for such an approach applied in an environment with only one risky token and limit pocket for the traders.) Specifically, we consider perturbing the CPMM price formula and study a class of pricing functions given by

$$(e_A + (1-\tau)q_A)(e_B - (1-\tau)q_B) = e_A e_B \tag{14}$$

where $\tau \in [0, 1)$. Notice that this class of price functions admits the CPMM function when $\tau = 0$. For values of q_i close to zero, an increase in τ reduces the convexity of the price

[6]In practice, the value of stablecoins do fluctuate at specific points in time, such as when USDC depegged for a short window in April 2023. For our example, we assume liquidity providers and takers believe the stablecoin peg will hold with certainty.

[7]Based on the Coinbase ETH index price obtained from fred.stlouis.org.

[8]We experimented with several other distributional assumptions for $\frac{\hat{\mu}_A}{\mu_A}$ and found similar results. Details are available upon request.

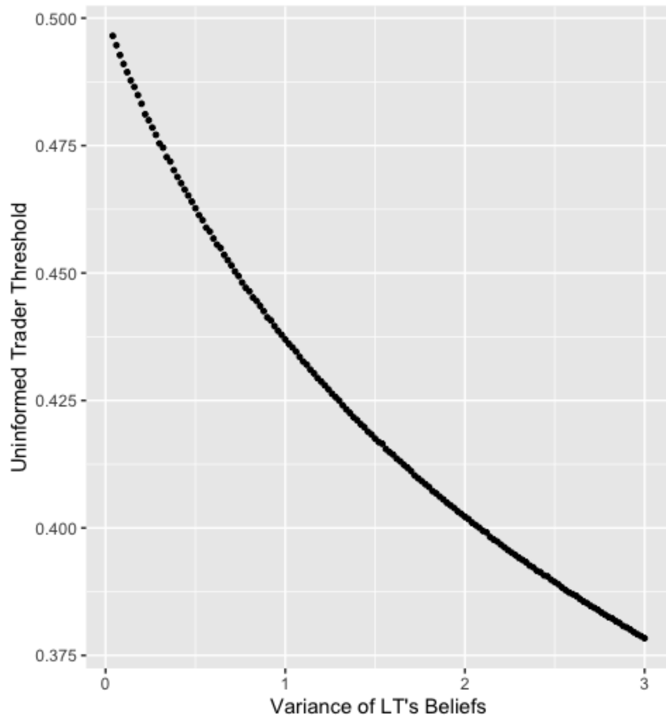

Figure 4: $\underline{\pi}$ against variance of $\hat{\mu}_A$

function. For larger values of q_i, it is possible the price function becomes more convex. Moreover, for any $\tau > 0$, there exist values of q_i such that the implied ex-post portfolio of the LP would have a negative amount of token A or B so we must impose the boundary conditions, $e_A \geqslant q_A$ and $e_B \geqslant q_B$. Such boundary conditions also tend to increase the global convexity of the price function.

We illustrate how an increase in τ impacts the price function locally in Figure 5 below. The solid curve represents the standard CPMM with $\tau = 0$. Around a given deposit point, (e_A, e_B), the dashed curve represents how the CPMM function changes when τ increases.

If we impose the LP's ex-post token holdings ($x_A = e_A + q_A$ and $x_B = e_B - q_B$) then we may re-write (14) as

$$((1-\tau)x_A + \tau e_A)\,((1-\tau)x_B + \tau e_B) = e_A e_B. \tag{15}$$

The price function (14) is convex and smoothly decreasing when $x > 0$. The convexity of

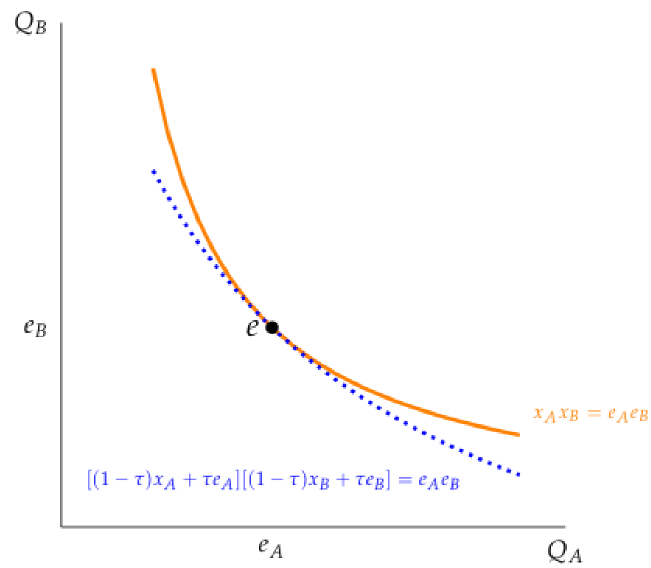

Figure 5: CPMM prices with $\tau = 0$ (the solid, orange curve) and with $\tau > 0$ (the dashed, blue curve).

the function is decreasing in τ. The boundary conditions on q_i simply imply $x_i \geqslant 0$.

For a given the realization of the LT's beliefs, $(\hat{\mu}_A, \hat{\mu}_B)$, the LP's net proceeds from trade satisfy

$$x - e_A = \frac{1}{1-\tau}\left[\sqrt{\frac{\hat{\mu}_B}{\hat{\mu}_A} e_A e_B} - e_A\right], \quad y - e_B = \frac{1}{1-\tau}\left[\sqrt{\frac{\hat{\mu}_A}{\hat{\mu}_B} e_A e_B} - e_B\right]. \tag{16}$$

Since net proceeds for both tokens scale by the same factor $1/(1-\tau)$, the LP's expected returns also scale by $\frac{1}{1-\tau}$. Moreover, gains from uninformed trading and losses from informed trading scale by the same ratio so that the break-even proportion $\underline{\pi}$ does not change with τ. As a result, increased (local) convexity of the CPMM hinders trading volume and reduces gains to trade for both the LP and the LT.

However, eliminating (global) convexity of the CPMM is not costless. When $\tau > 0$, equation 15 has finite positive intercepts: $(0, \frac{1+\tau}{\tau} e_B)$ and $(\frac{1+\tau}{\tau} e_A, 0)$. For such values of τ, trading volume cannot increase beyond the two intercepts, even for more extreme beliefs of the LT. Holding the LP's choice of liquidity fixed, we argue that relaxing the local convexity of the pricing function may be detrimental to the LP's ex-ante profits.

To illustrate this, it is simplest to consider a piece-wise linear approximation to the convex pricing function that runs through the LP's (fixed) choice of liquidity deposit. With piece-wise linear prices, liquidity takers either do not trade or trade up to one of the intercept points. For example, suppose p_h represents the (minus the) slope of the price function for values of x_A between 0 and e_A the amount of token A deposited by the LP. If the beliefs of the LT are more optimistic than p_h (so if $\hat{\mu}_A/\mu_A > p_h$), then the LT will trade up to the intercept where $x_A = 0$—the LT will buy all of token A in the pool at the prevailing price, p_h. Otherwise, for $p_h > \hat{\mu}_A/\mu_A > 1$, the LT will not trade.

Consider a marginal increase in p_h (in absolute value). Such a change increases the region of no trade by the LT and thus reduces trading volume on the extensive margin. Recall that the LP only loses expected value from informed trades (and earns exactly zero losses on the marginal informed LT who is just indifferent between trading at p_h and not trading). Therefore, decreasing the volume of trade reduces the LP's expected losses from informed trading. Among uninformed trades, reducing volume is costly on the extensive margin, but raising the intercept implies the LP realizes increased gains to trade for all beliefs where the LT continues to trade. An analogous argument occurs if beliefs of the LT are sufficiently low so that the LT trades to the point where $x_B = 0$. Consequently, it is possible that the gains from increasing the global convexity of a piece-wise linear price function outweigh the costs, implying some degree of convexity is desirable. We show this result both for piece-wise linear prices as well as for the continuously differentiable price function in (14) in Appendix B.

If the distribution of the LT's beliefs has bounded support, then the potential losses from reduced (global) convexity for extremal beliefs may be limited with an appropriate choice of τ. In other words, when the LT's beliefs have bounded support, then there exists $\tau > 0$ that increases the LP's expected returns. In fact, we generalize these results beyond the CPMM formula in the next Proposition (proved in Appendix A).

Proposition 4: Pareto Improvement. Consider a convex and smoothly decreasing price function $y = G(x)$. Assume the distributions of the LT's valuations of the tokens $(\hat{\mu}_A, \hat{\mu}_B)$ have bounded support such that a trade that exhausts one token never happens under the price function $G(x)$. Then there exists $\tau = \hat{\tau} \in (0, 1)$ such that the new price function $(1-\hat{\tau})\,y + \hat{\tau}e_B = G((1-\hat{\tau})\,x + \hat{\tau}e_A)$ is less convex at (e_A, e_B), the LP's optimal deposit is the same at $\tau = \hat{\tau}$ as at $\tau = 0$, and $\tau = \hat{\tau}$ increases both the LP's and the LT's expected returns proportionally by $\frac{\tau}{1-\tau}$.

In particular, if $G(x)$ is the CPMM function and if $\left[\underline{\mu_i}, \overline{\mu_i}\right]$ is the support of the distribution of $\hat{\mu}_i$, then the result of Proposition 4 hold for all $\tau \leqslant \bar{\tau} = \min\left\{\sqrt{\frac{\underline{\mu_B} e_B}{\overline{\mu_A} e_A}}, \sqrt{\frac{\underline{\mu_A} e_A}{\overline{\mu_B} e_B}}\right\}$ with $\bar{\tau} > 0$.

We see that with bounded beliefs, convexity hurts the LP's expected returns. In fact, with some additional conditions, the optimal price function for the LP is the linear price function:

$$\begin{cases} p_l x_A + x_B = p_l e_A + e_B, & x \geqslant e_A \\ p_h x_A + x_B = p_h e_A + e_B, & x < e_A \end{cases} \tag{17}$$

where again e_i are the LP's deposit and x_i are the tokens left in the pool after the LT's trading. Similar to the results in Milionis, Moallemi and Roughgarden (2023*b*), we have the following proposition (proved in Appendix C).

Proposition 5: LP's Optimal Pricing Function Assume the distributions of the values of the tokens have bounded support and the LT has a budget limit on at least one token, i.e. x or y can't go to infinite. Given the LP's deposit (e_A, e_B), the optimal pricing formula is the linear pricing formula is one of the following conditions is satisfied:

1. All trades are uninformed trading, i.e. $\pi = 1$;

2. The LT's value $(\hat{\mu}_A, \hat{\mu}_B)$ follows the same distribution for both informed and uninformed trading. And one of the two tokens is a stablecoin. In the case of token A is stable, it implies $\hat{\mu}_A = \mu_A$ for sure. Also, there exists some uninformed trading, i.e. $\pi \neq 0$.

5 Conclusion

Blockchain technology has spawned a very large variety of cryptocurrency tokens. Given the large disagreement about their speculative value and heterogeneity about any utility of the tokens, trading the tokens is important. Over the past decade, a large number of new centralized exchanges have been successful (and unsuccessful) at both generating large volumes and innovating. The perpetual futures contract is one example of innovation (Soska et al. (2021), Christin et al. (2023)). Similarly, Automated Market Makers (AMM) have innovated trade by designing smart contracts (automated code on the blockchain) to conduct trade directly on a blockchain.

In this paper, we have explored the key design characteristic of AMM technology, the pricing curve. Specifically, we look at two aspects related to the pricing curve, G. First, what is the optimal ratio for deposits? Contrary to conventional AMM wisdom, depositing tokens in equal value (measured through the lens of the liquidity provider) is not optimal. Second, we explore the convexity of G and its impact on the liquidity provider profits. The tradeoff is subtle since convexity impacts the profits from trading with both informed and uninformed liquidity takers.

There are, of course, several important areas we have left for future research. Our model treats the G function as given. This, along with the "deep pockets" assumption for the liquidity takers, means the liquidity provider's decision can be made in isolation (i.e., atomistic with respect to liquidity takers). In practice, there are multiple AMM exchanges. So, thinking about competition across the design of the G function is interesting. Second, our model takes a simplified view of the timing of transactions – first, the LP posts and then the LT trades. Again, in practice, the timing of transactions in a decentralized blockchain is complicated and potentially strategic.

References

Angeris, Guillermo, and Tarun Chitra. 2020. "Improved Price Oracles." ACM. 4

Angeris, Guillermo, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra. 2021. "An analysis of Uniswap markets." *arXiv 1911.03380*. 4

Aoyagi, Jun. 2022. "Liquidity Provision by Automated Market Makers." 5

Bergault, Philippe, Louis Bertucci, David Bouba, and Olivier Guéant. 2023. "Automated Market Makers: Mean-Variance Analysis of LPs Payoffs and Design of Pricing Functions." *arXiv 2212.00336*. 6, 10

Capponi, Agostino, and Ruizhe Jia. 2021. "The Adoption of Blockchain-based Decentralized Exchanges." *arXiv 2103.08842*. 5

Christin, Nicolas, Bryan Routledge, Kyle Soska, and Ariel Zetlin-Jones. 2023. "Crypto Carry Trade." Carnegie Mellon University, Working Paper. 25

Cicle. 2023. "Cicle Transparency & Stability." *url = https://www.circle.com/en/transparency*, Accessed on Sep 18th, 2023. 7

Fabi, Michele, and Julien Prat. 2023. "The Economics of Constant Function Market Makers." Blockchain at Polytechnique, Working Paper. 5

Glosten, Lawrence R., and Paul R. Milgrom. 1985. "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders." *Journal of Financial Economics*, 14(1): 71–100. 1, 3, 9

Goyal, Mohak, Geoffrey Ramseyer, Ashish Goel, and David MaziÃšres. 2023. "Finding the Right Curve: Optimal Design of Constant Function Market Makers." 6

Harrison, J. Michael, and David M. Kreps. 1978. "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations*." *The Quarterly Journal of Economics*, 92(2): 323–336. 2

Hasbrouck, Joel, Thomas Rivera, and Fahad Saleh. 2023. "The Need for Fees at a DEX: How Increases in Fees Can Increase DEX Trading Volume." *Available at SSRN: https://ssrn.com/abstract=4192925 or http://dx.doi.org/10.2139/ssrn.4192925*. 5

Lehar, Alfred, and Christine A. Parlour. 2023. "Decentralized Exchange: The Uniswap Automated Market Maker." Working Paper. 5

Milionis, Jason, Ciamac C. Moallemi, and Tim Roughgarden. 2023*a*. "Automated Market Making and Arbitrage Profits in the Presence of Fees." 5

Milionis, Jason, Ciamac C. Moallemi, and Tim Roughgarden. 2023*b*. "A Myersonian Framework for Optimal Liquidity Provision in Automated Market Makers." *arXiv 2303.00208*. 4, 6, 20, 24, 33

Milionis, Jason, Ciamac C. Moallemi, Tim Roughgarden, and Anthony Lee Zhang. 2022. "Quantifying Loss in Automated Market Makers." *DeFi'22*, 71–74. New York, NY, USA:Association for Computing Machinery. 2, 5, 17

Park, Andreas. 2023. "Conceptual Flaws of Decentralized Automated Market Making." *Available at SSRN: https://ssrn.com/abstract=3805750 or http://dx.doi.org/10.2139/ssrn.3805750*. 6

Soska, Kyle, Jin-Dong Dong, Alex Khodaverdian, Ariel Zetlin-Jones, Bryan Routledge, and Nicolas Christin. 2021. "Towards Understanding Cryptocurrency Derivatives: A Case Study of BitMEX." 25

Uniswap-V2. 2023. "Uniswap V2 Protocol." *url = https://docs.uniswap.org/contracts/v2/concepts/core-concepts/pools*, Accessed on Sep 18th, 2023. 3, 10

A Proof of Optimal Liquidity Provision

LP's optimal deposit problem is

$$\begin{aligned}\max_{e_A, e_B} \; & \pi(\mu_A \mathbb{E}[x_A - e_A] + \mu_B \mathbb{E}[x_B - e_B]) + \\ & \quad (1-\pi)(\mathbb{E}\hat{\mu}_A[x_A - e_A] + \mu_B \mathbb{E}\hat{\mu}_B[x_B - e_B]) \\ \text{s.t. } & (5)\text{–}(6), \\ & \quad 0 \leqslant e_i \leqslant E_i, \quad \forall i\end{aligned}$$

Based on equation (5)–(6), we can write down the post-trade portfolio of the LP as

$$x_A = \sqrt{\frac{\hat{\mu}_B}{\hat{\mu}_A} e_A e_B},\ x_B = \sqrt{\frac{\hat{\mu}_A}{\hat{\mu}_B} e_A e_B}$$

Then we can write the post-trade net value gains from each token in the hand of the LP by depositing as

$$\begin{aligned}
\mu_A(x_A - e_A) &= \sqrt{\frac{\hat{\mu}_B/\mu_B}{\hat{\mu}_A/\mu_A} \mu_A \mu_B e_A e_B} - \mu_A e_A \\
\mu_B(x_B - e_B) &= \sqrt{\frac{\hat{\mu}_A/\mu_A}{\hat{\mu}_B/\mu_B} \mu_A \mu_B e_A e_B} - \mu_B e_B
\end{aligned}$$

for uninformed trades and

$$\begin{aligned}
\hat{\mu}_A(x_A - e_A) &= \sqrt{\frac{\hat{\mu}_A \hat{\mu}_B}{\mu_A \mu_B} \mu_A \mu_B e_A e_B} - \hat{\mu}_A e_A \\
\hat{\mu}_B(x_B - e_B) &= \sqrt{\frac{\hat{\mu}_A \hat{\mu}_B}{\mu_A \mu_B} \mu_A \mu_B e_A e_B} - \hat{\mu}_B e_B
\end{aligned}$$

for informed trades.

Denote $\omega = \sqrt{\frac{\hat{\mu}_A/\mu_A}{\hat{\mu}_B/\mu_B}}$ and $\psi = \sqrt{\frac{\hat{\mu}_A}{\mu_A}\frac{\hat{\mu}_B}{\mu_B}}$. With assumptions that $\mathbb{E}\hat{\mu}_i = \mu_i$, LP's optimal deposit problem becomes

$$\begin{aligned}
&\max_{e_A, e_B} \left[\pi\left(\mathbb{E}\omega + \mathbb{E}\frac{1}{\omega}\right) + (1-\pi)2\mathbb{E}\psi\right] \sqrt{\mu_A e_A}\sqrt{\mu_B e_B} - \mu_A e_A - \mu_B e_B \\
&\text{s.t. } 0 \leqslant e_i \leqslant E_i, \quad \forall i
\end{aligned}$$

Further denote $\Pi = \pi\frac{\mathbb{E}\omega + \mathbb{E}\frac{1}{\omega}}{2} + (1-\pi)\mathbb{E}\psi$. We can use the standard Lagrangian method to solve the above constraint optimization problem. The FOCs are

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial e_A}: &\quad \mu_A\left(\Pi\sqrt{\frac{\mu_B e_B}{\mu_A e_A}} - 1\right) + \eta_A - \lambda_A = 0 \\
\frac{\partial \mathcal{L}}{\partial e_B}: &\quad \mu_B\left(\Pi\sqrt{\frac{\mu_A e_A}{\mu_B e_B}} - 1\right) + \eta_B - \lambda_B = 0
\end{aligned}$$

where η_i is the Lagrangian multiplier for $0 \leqslant e_i$ and λ_i is the Lagrangian multiplier for

$e_i \leqslant E_i$.

If $\Pi < 1$, the above FOCs only hold when $e_A = e_B = 0$. In this case $\eta_i > 0$ and $\lambda_i = 0$.

If $\Pi > 1$, the solution is always at the corner, i.e. at least one of the $\lambda_i > 0$. To see this, consider the interior cases where $\eta_i = 0$ and $\lambda_i = 0$. For the FOCs to hold, we need $\Pi\sqrt{\frac{\mu_B e_B}{\mu_A e_A}} = \Pi\sqrt{\frac{\mu_A e_A}{\mu_B e_B}} = 1$, which is impossible. Since $\Pi > 1$, if one of $\Pi\sqrt{\frac{\mu_B e_B}{\mu_A e_A}}$ and $\Pi\sqrt{\frac{\mu_A e_A}{\mu_B e_B}}$ equals to 1, then the other one must be bigger than 1. And it needs the corresponding λ_i to be positive for the FOCs to hold.

Therefore, we have the following optimal deposit of the LP

$$\begin{cases} e_A^* = E_A,\ e_B^* = \min\left\{\Pi^2 \frac{\mu_A}{\mu_B} E_A, E_B\right\}, & \text{if } \mu_A E_A \leqslant \mu_B E_B \\ e_A^* = \min\left\{\Pi^2 \frac{\mu_B}{\mu_A} E_B, E_A\right\},\ e_B^* = E_B, & \text{if } \mu_A E_A > \mu_B E_B \end{cases}$$

if $\Pi > 1$ and

$$e_A^* = e_B^* = 0$$

if $\Pi < 1$.

B Proof of Pareto Improvement

Let $y = G(x)$ be a convex and smoothly decreasing price function where $e_B = G(e_A)$. Consider a uniform stretch of the function around the initial deposit point (e_A, e_B): $(1-\tau)y + \tau e_B = G((1-\tau)x + \tau e_A)$ where $\tau \in (0,1)$. Then the second order derivatives is $\frac{d^2y}{dx^2} = (1-\tau)^2 G''((1-\tau)x + \tau e_A)$. Therefore, the transformation is less convex around the initial deposit point (e_A, e_B) as τ increases.

Now we can write the LT's problem as:

$$\begin{aligned} \max_{e_A, e_B}\ & \hat{\mu}_A(e_A - x) + \hat{\mu}_B(e_B - y) \\ \text{s.t. } & (1-\tau)y + \tau e_B = G((1-\tau)x + \tau e_A) \end{aligned}$$

Assume the distributions of the LT's values of the tokens ($\hat{\mu}_A$, $\hat{\mu}_B$) have bounded support such that a trade that exhausts one token never happens. Then the first order condition becomes $G'((1-\tau)x+\tau e_A) = -\frac{\hat{\mu}_A}{\hat{\mu}_B}$. Similar to the CPMM case, the LP's post-trade portfolio satisfies

$$
\begin{aligned}
(1-\tau)y + \tau e_B &= G((1-\tau)x+\tau e_A) \\
G'((1-\tau)x+\tau e_A) &= -\frac{\hat{\mu}_A}{\hat{\mu}_B}
\end{aligned}
$$

Let (x_0, y_0) be the post-trade portfolio for the original function, i.e., when $\tau = 0$. Let (x_τ, y_τ) be the portfolio for some $\tau \in (0, 1)$. Then given $\frac{\hat{\mu}_A}{\hat{\mu}_B}$, the ex post portfolios satisfies

$$
\begin{aligned}
(1-\tau)x_\tau + \tau e_A &= x_0 \\
(1-\tau)y_\tau + \tau e_A &= y_0
\end{aligned}
$$

which can be written as

$$
\begin{aligned}
x_\tau - e_A &= \frac{1}{1-\tau}(x_0 - e_A) \\
y_\tau - e_B &= \frac{1}{1-\tau}(y_0 - e_B)
\end{aligned}
$$

Therefore, the trading volume is proportionally increased by $1 - \frac{1}{1-\tau} = \frac{\tau}{1-\tau}$ for every ex post scenario.

Given the probability of uninformed trading π, the LP's expected return with the transformed price function is

$$
\begin{aligned}
R_\tau =& \mathbb{E}[(\pi\mu_A + (1-\pi)\hat{\mu}_A)(x_\tau - e_A) + (\pi\mu_B + (1-\pi)\hat{\mu}_B)(y_\tau - e_B)] \\
=& \frac{1}{1-\tau}\mathbb{E}[(\pi\mu_A + (1-\pi)\hat{\mu}_A)(x_0 - e_A) + (\pi\mu_B + (1-\pi)\hat{\mu}_B)(y_0 - e_B)]
\end{aligned}
$$

Since the objective is just scaled up by a constant, the optimal deposit decision (e_A^*, e_B^*) shouldn't change as well.

C Cost of Convexity

Again let token B represent a stable coin and set $\hat{\mu}_B = \mu_B = 1$ and $h_B(\hat{\mu}_B) = 1$ if $\hat{\mu}_B = 1$. Denote $r_A = \hat{\mu}_A/\mu_A$. Assume r_A follows a distribution with CDF $F(r_A)$. For simplicity, assume $\frac{\mu_A e_A}{\mu_B e_B} = 1$. The results still go through when $\frac{\mu_A e_A}{\mu_B e_B}$ equals to some constant other than 1.

C.1 Piece-wise Linear

Consider the piece-wise linear prices 17. The region of belief where a trade happens with price p_h is when $r_A \geqslant p_h$. From the LP's perspective, the trading volume in this region is $-e_A$ for token A and $p_h e_A$ for token B. The expected return of the LP from uninformed trading is

$$\int_{p_h}^{\infty} (p_h - 1) dF(r_A) \mu_A e_A$$

with derivative as $[1 - F(p_h) - (p_h - 1) f(p_h)] \mu_A e_A$. The first term represents the increased gains to trade for all beliefs where the LT continues to trade. The second term represents the reduced trading volume on the margin.

On the other hand, the expected return (negative) of the LP from informed trading is

$$\int_{p_h}^{\infty} (p_h - r_A) dF(r_A) \mu_A e_A$$

with derivative as $(1 - F(p_h)) \mu_A e_A$. Since on the marginal informed LT is just indifferent between trading and not, the second term in the case of uninformed trades is not here.

Given the proportion of uninformed trades π, the marginal benefits of increasing p_h (increasing convexity) is

$$[1 - F(p_h) - \pi(p_h - 1) f(p_h)] \mu_A e_A$$

which has finite number of roots. It implies that some degree of convexity is desirable.

C.2 Continuously Differentiable Price

Now consider the continuously differentiable price function in 15. Similarly, the region of belief where a trade happens with price p_h is when $r_A \geqslant \frac{1}{\tau^2}$. From the LP's perspective, the trading volume in this region is $-e_A$ for token A and $\frac{1}{\tau}e_B$ for token B. Denote $c = \frac{1}{\tau} \in (1, \infty)$. So, increasing c increases the local convexity. The expected return of the LP from uninformed trading is

$$\int_{c^2}^{\infty}(c-1)dF(r_A)\mu_A e_A$$

with derivative as $[1 - F(c^2) - (c-1)f(c^2)]\mu_A e_A$. Again the first term represents the increased gains to trade for all beliefs where the LT continues to trade. The second term represents the reduced trading volume on the margin.

On the other hand, the expected return (negative) of the LP from informed trading is

$$\int_{c^2}^{\infty}(c-r_A)dF(r_A)\mu_A e_A$$

with derivative as $[1 - F(c^2) + c(c-1)f(c^2)]\mu_A e_A$. Since $c > 1$ there is an additional gain for the LP from reducing the trading volume further.

Given the proportion of uninformed trades π, the marginal benefits of increasing p_h (increasing convexity) is

$$[1 - F(c^2) + (c-1)((1-\pi)c - \pi)f(c^2)]\mu_A e_A$$

which is always positive for $c \geqslant \frac{\pi}{1-\pi}$. In these cases, increasing (local) convexity is always beneficial for trades induced by extremal beliefs. However, it reduces the trading volume and the returns from mild beliefs.

D Proof of Optimal Pricing Function

We can consider the optimal design problem as the LP post the ending position of the pool given the new valuation of the LT $(\hat{\mu}_A, \hat{\mu}_B)$ such that the LT is willing to participate

(Individual Rational) and truthfully report the values (Incentive Compatible).

Assume the LT's value $(\hat{\mu}_A, \hat{\mu}_B)$ follows the same distribution for both informed and uninformed trading. Also, assume the LT has at most l_B token B to trade in.

Let $t_A = e_A - x$ and $t_B = e_B - y$ be the net amount of token the LP loses by trading. With the percentage of uninformed trading π, the problem can be written as:

$$
\begin{aligned}
\max_{x,y} \ & \mathbb{E}_{\{\hat{\mu}_A,\hat{\mu}_B\}} \left[-\left(\pi\mu_A + (1-\pi)\,\hat{\mu}_A\right) t_A\left(\hat{\mu}_A, \hat{\mu}_B\right) - \left(\pi\mu_B + (1-\pi)\,\hat{\mu}_B\right) t_B\left(\hat{\mu}_A, \hat{\mu}_B\right) \right] \\
\text{s.t. } & \hat{\mu}_A t_A\left(\hat{\mu}_A, \hat{\mu}_B\right) + \hat{\mu}_B t_B\left(\hat{\mu}_A, \hat{\mu}_B\right) \geqslant \hat{\mu}_A t_A\left(\hat{\mu}'_A, \hat{\mu}'_B\right) + \hat{\mu}_B t_B\left(\hat{\mu}'_A, \hat{\mu}'_B\right) \\
& \hat{\mu}_A t_A\left(\hat{\mu}_A, \hat{\mu}_B\right) + \hat{\mu}_B t_B\left(\hat{\mu}_A, \hat{\mu}_B\right) \geqslant 0 \\
& t_A\left(\hat{\mu}_A, \hat{\mu}_B\right) \leqslant e_A, \ -l_B \leqslant t_B\left(\hat{\mu}_A, \hat{\mu}_B\right) \leqslant e_B
\end{aligned}
$$

Since only $p = \frac{\hat{\mu}_B}{\hat{\mu}_A}\frac{e_B}{e_A}$ matters in the constraints, the problem can be written as

$$
\begin{aligned}
\max_{t_A,t_B} \ & \mathbb{E}_p \left[\left(-\frac{t_A(p)}{e_A} - \frac{\left(\pi\mu_B + (1-\pi)\,\hat{\mu}_B\right)}{\left(\pi\mu_A + (1-\pi)\,\hat{\mu}_A\right)} \frac{e_B}{e_A} \frac{t_B(p)}{e_B} \right) \left(\pi\mu_A + (1-\pi)\,\hat{\mu}_A\right) \right] e_A \\
\text{s.t. } & \frac{t_A(p)}{e_A} + p\frac{t_B(p)}{e_B} \geqslant \frac{t_B(\hat{p})}{e_B} + p\frac{t_B(\hat{p})}{e_B} \\
& \frac{t_A(p)}{e_A} + p\frac{t_B(p)}{e_B} \geqslant 0 \\
& \frac{t_A(p)}{e_A} \leqslant 1, \ -\frac{l_B}{e_B} \leqslant \frac{t_B(p)}{e_B} \leqslant 1
\end{aligned}
$$

Under one of the two conditions, i.e. $\pi = 0$ or $\hat{\mu}_A = \mu_A$ for sure, we know $\pi\mu_A + (1-\pi)\,\hat{\mu}_A$ is a constant. So the objective can be simplified. Let $-\frac{t_A(p)}{e_A} + 1 = y(p)$, $\frac{t_B(p)}{e_B} = x(p)$ and $\frac{(\pi\mu_B + (1-\pi)\hat{\mu}_B)}{(\pi\mu_A + (1-\pi)\hat{\mu}_A)}\frac{e_B}{e_A} = \pi(p_0, p)$. The problem then has the same expression as Milionis, Moallemi and Roughgarden (2023*b*).

$$
\begin{aligned}
\max_{x,y} \ & \mathbb{E}_p \left[y(p) - \pi(p_0, p)\, x(p) \right] \\
\text{s.t. } & p x(p) - y(p) \geqslant p x(\hat{p}) - y(\hat{p}) \\
& p x(p) - y(p) \geqslant 0 \\
& y(p) \geqslant 0, \ -c \leqslant x(p) \leqslant 1
\end{aligned}
$$