

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

October 23, 2025

Abstract:

Central limit order books, standard in traditional exchanges, are impractical on blockchains. Decentralized exchanges like Uniswap and Curve instead use Automated Market Makers (AMMs) that set prices via functions of token quantities. To study price discovery and market impact in this setting, we characterize the optimal liquidity provision of AMM participants. Theoretically and empirically, we show that price impact depends on both trade size and the dynamics of liquidity provision. Using data from 31 large Uniswap v2 pools, we find that liquidity providers adjust positions in ways that offset prior trades, especially when those trades are likely uninformed.

*We thank Bruno Biais and Cyril Monnet and participants at conferences and seminar audiences at the American Finance Association 2025 Meeting, CMU Secure Blockchain Summit, Ripple UBRI Partners Conference, HEC Paris Blockchain and Digital Assets Conference, the Richmond Fed, the Finance Theory Group Spring 2024 Meeting, the Federal Reserve Bank of Philadelphia, the Rice-Lemma Monetary Conference, the 2024 SFS Cavalcade, the 2024 Winter Meeting of the Society for Economic Dynamics, the 2025 Annual Meeting of the Society for the Advancement of Economic Theory, and the 2025 Wharton Conference on Liquidity and Financial Intermediation for useful comments and discussions. We are especially grateful to Nicolas Bonneton, Chuck Fang, Ji Hee Yoon, Julien Prat and Michael Sockin for insightful discussions. We are also grateful for the financial support of the Ripple Foundation. Contact authors at: azj@andrew.cmu.edu.

1 Introduction

Centralized exchanges for cryptocurrencies like Binance, Coinbase, FTX, and Kraken intermediate trade with a central limit order book. A central limit order book is constructed from participants' posts of quantity and price pairs they are willing to trade. This facilitates price discovery through matching buy and sell orders. Like similar mechanisms we see for trading equities (NASDAQ for example), settlement of cryptocurrency trades at limit order books happens later and not directly on the cryptocurrency's ledger. For example, Bitcoin trades at Coinbase are recorded only by Coinbase. Updates to the Bitcoin ledger happen only infrequently when traders deposit or withdraw from the exchange.

In contrast, decentralized finance (DeFi) exchanges facilitate trade directly using a blockchain. The computer codes that control the DeFi exchange, called smart contracts, and the messages traders use to execute transactions are recorded "on-chain" in the blockchain by the decentralized network of ledger validators ("miners"). This technology is currently incapable of replicating a central limit order book. The volume and speed of messages needed to implement a limit order book is not practical and is prohibitively expensive with most blockchains. As a consequence, decentralized exchanges, such as those developed by Uniswap or Curve, have created Automated Market Makers (AMMs) to intermediate trade. These alternatives to centralized markets now account for a large volume of cryptocurrency trade. From January 2024 to May 2025, decentralized exchanges processed an average of roughly \$210 billion of cryptocurrency spot transactions per month. As of May 2025, decentralized exchange spot transaction volume is approximately 20% of the cryptocurrency spot transaction volume on centralized exchanges.

To intermediate trade on blockchain with (far) fewer messages than a centralized exchange, AMMs limit traders to posting only quantities. The smart contract code defines trade in a liquidity pool with functions for adding (minting), removing (burning), or exchanging (swapping) the two coins (or tokens) that constitute the pool. Liquidity providers (LPs) supply a portfolio of two tokens to the AMM pool using the mint function. Liquidity takers (LTs) may swap one token for the other. The rate of this swap—effectively, the relative price of the two tokens—is a coded function of relative quantities of the tokens posted and swapped to date. For example, the constant-product market maker (CPMM) implies the marginal relative price of the two tokens is the ratio of the current balance of tokens.¹

¹The CPMM, defined below in equation 1, keeps the product of the two token quantities to be constant. The origin

To understand how DeFi markets facilitate price discovery and to measure price impact of trades, we study the dynamic provision of cryptocurrency quantities to a liquidity pool. Specifically, we focus on liquidity providers as they play an active and strategic role in price discovery by choosing the liquidity they supply in response both to information about coin values and to information about the distribution of future liquidity trades. Additionally, through our model we see that liquidity providers play an important role in understanding the price impact of AMM trades as their liquidity supply and thus prices respond directly to trades conducted by liquidity takers.

To motivate our study of the strategic behavior of liquidity providers, we use data from Uniswap v2 ("version 2") where the liquidity provision functions in the Uniswap smart contract are particularly stark and limited to adding ("mint") or removing ("burn") coins at the current ratio of coins in the pool. Here, liquidity mints or burns do not change the marginal price of the coins in the pool. We document that many liquidity providers are active in the price-setting process. Trades by LPs do constitute a small percent of the overall trade as most of the trade is swaps by LTs. However, the majority of the liquidity providers are active in the sense that a sizable proportion of their transactions involve the swap transaction (liquidity taking). When liquidity providers swap against their own pool, they directly impact prices faced by other liquidity takers at the AMM exchange.

Uniswap has augmented their pricing functions to offer liquidity providers additional control over their liquidity in Uniswap v3 and most recently v4. We view the data from v2 as particularly insightful since the limited choices liquidity providers allows us to measure their degree of activeness with their swap transaction behavior. The marginal tradeoffs we consider in our model are also relevant in the v3 and v4 settings.

To better understand how liquidity providers set prices and therefore explore the impact of trades on liquidity providers' price setting behavior, we build a dynamic model of AMM liquidity provision. The basic tradeoff in our model for the liquidity provider is familiar. We assume liquidity takers may be "informed" or "uninformed" giving rise to a classic form of adverse selection in asset markets (as in [Glosten and Milgrom \(1985\)](#) or [Kyle \(1985\)](#)). With AMMs, what creates this structure is timing. We posit that liquidity providers are "slow." They post their coins to an AMM and then a liquidity taker trades. Liquidity takers are "fast" and able to attain priority

of the CPMM is a blog post of Vitalik Buterin and Martin Koppelman in 2016. See https://www.reddit.com/r/ethereum/comments/55m04x/lets_run_onchain_decentralized_exchanges_the_way/.

for blockchain execution. The liquidity taker may be trading for reasons that are orthogonal to public information (akin to a private value). Alternatively, the trader may be trading having seen an update to public information about coin prices (a common value)—sometimes referred to as “impermanent loss” in AMM documentation. In either case, the liquidity provider cannot avoid the “fast” liquidity taker.

Introducing this conventional friction allows us to explore how adverse selection distorts the 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 AMM 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* directly to protect themselves from informed trading losses, on AMMs such distortions may only manifest in the quantities of deposits liquidity providers post.

The risk of trading against a better informed liquidity taker is an important consideration for the liquidity provider. However, our model does not assume perfect and continuous “arbitrageurs” whose AMM trades reflect a single agreed-upon external “market price.” In our setting, some traders may trade for reasons orthogonal to current market prices should they require liquidity—i.e., the very reason the market might exist. Our setting allows for traders’ beliefs to be updated by new information that, say, reflects information from a posted price on a centralized exchange. If the risk of that happening is sufficiently high, the liquidity provider will choose not to post liquidity. However, we do not model that event as an arbitrage. The costs to moving tokens between exchanges and decentralized exchanges (an on-chain transaction) is not trivial. Perhaps more importantly, a coin owned indirectly on an exchange is not a perfect substitute for owning the same coin directly on a blockchain. The bankruptcy of an exchange (FTX and others) or the loss of a private key are distinct risks. We are agnostic as to whether ownership through an exchange is better or worse than ownership on-chain. Security, convenience, and liquidity properties can all differ. Our point here is simply that they are not identical. Lastly, there is a large number of tokens, some of which have a sizable implied market capitalization, that are not listed on any centralized exchange. For these tokens, there is no applicable “market price” from a centralized exchange and our theory provides guidance on the price discovery process for such tokens.

Our model has interesting implications for the dynamics of liquidity provision. We show that

our model generates endogenous inactivity by liquidity providers. More precisely, even though the liquidity providers have the option to re-balance their liquidity on deposit after each trade by a liquidity taker, often they optimally do not do so. This endogenous inactivity arises because of our assumption that the LP is risk neutral and (despite the convexity arising from the exogenous pricing function) finds it optimal to supply her entire endowment of tokens for a range of relative prices—a maximal supply region. If an LT trade leaves the LP in this maximal supply region, then she will not re-balance her liquidity deposit. If an LT trade results in LP (ex post) balances outside of this region, then the LP will re-balance back to the boundary of the maximal supply region.

While this inaction region does not respond to uninformed trades at the AMM, it does shift when informed trades arrive. As a result, the extent of endogenous inaction depends on the relative proportion of informed versus uninformed trading. Indeed, these model dynamics give rise to predictions for the behavior of LPs. First, liquidity providers typically trade against liquidity takers; LP trades (when they happen) tend to move relative prices at the AMM in the opposite direction of trades completed by LTs. Second, liquidity providers are more likely to be active—are more likely to re-balance deposits—in markets with more uninformed trade.

The empirical dynamics of liquidity provider behavior in Uniswap v2 are consistent with these model predictions. Swaps—price setting behavior—completed by liquidity providers tend to impact relative prices in the opposite direction of (cumulative) trades completed by liquidity takers. Further, we adopt an empirical strategy used to identify informed trading in the Ethereum ecosystem as suggested by [Capponi, Jia and Yu \(2024\)](#) and aligned with ideas from the high-frequency trading literature.² The idea is that informed traders are more likely to be price sensitive. As a result, they should demand more blockchain priority and be more willing to pay (gas fees to) validators to prioritize their transactions. In AMM periods and markets where liquidity takers pay lower average gas fees—suggestive of less price sensitivity of traders and hence more uninformed trading—we find liquidity providers transact more frequently, consistent with our model.

We then use our model to explore how the shape of the pricing function impacts gains to trade and liquidity provider’s profits. Analogous to results in [Milionis, Moallemi and Roughgarden \(2023b\)](#), 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

²See [Aquilina, Budish and O’neill \(2022\)](#) and [Brugler and Hendershott \(2023\)](#) for recent papers that exploit the timing of trades or orders on centralized exchanges to identify high-frequency, informed trading.

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 the 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 constant-product market maker (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.

1.1 Related Literature

Much of the research on AMMs has focused on examining how AMMs perform alongside the presence of deep, liquid, centralized exchanges. One of the earliest examples is [Angeris and Chitra \(2020\)](#) who obtain conditions under which 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 after accounting for fees. [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 \(2023\)](#) demonstrate 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.

Similar to our model, [Aquilina et al. \(2024\)](#) considers heterogeneity among liquidity providers using size or external information to classify liquidity providers and study their empirical behavior on UniSwap V3 data. They classify liquidity providers with exceptionally large token positions or identified as VCs, asset managers, etc., as “sophisticated” and find that they provide majority of

the liquidity, actively manage their positions, and interact with multiple pools. In contrast, “unsophisticated” liquidity providers earn significantly smaller fees, and struggle to adapt their liquidity strategies during periods of high volatility. [Lehar, Parlour and Zoican \(2023\)](#), who also focus on UniSwap V3, find that larger liquidity providers dominate low-fee pools, while small liquidity providers dominate high-fee pools. As in our model, liquidity providers in their model adjust their pool positions after trades as well but only because the structure of contracts in Uniswap V3 prevents informed traders from fully arbitraging prices from centralized exchanges.

Directly supporting our assumption that trades can be categorized as informed and uninformed, [Capponi, Jia and Yu \(2024\)](#) provide empirical evidence showing that high-fee DEX trades contain more private information. Informed traders bid high fees both to mitigate execution risks from blockchain congestion and to secure execution priority. We build on these important papers by showing how liquidity providers directly impact decentralized exchange prices and then by building a dynamic, quantitative model where there is a role for liquidity providers to set prices. One of our contributions is to relax the assumption of perfect arbitrage with centralized exchanges and examine optimal liquidity provision when price discovery in some other market is not perfect.

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 \(2023a\)](#) extend the model to involve trading fees and provide results on the arbitrageur’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. [Cao et al. \(2023\)](#) develop a structural model where a platform sets the fee level to maximize liquidity in the pool with one token as a stablecoin. Like much of the earlier literature, they study this problem under an assumption that a true price is known and the fee

is designed to maximize rents from uninformed trades. They find that the optimal AMM fee structure dynamically adjusts to volatility, leading to better trade.

In terms of the design and efficiency of the price function, [Park \(2023\)](#) demonstrates that constant function market makers may cause economically meaningless and costly trading, such as front running. Front-running is a substantial concern that liquidity takers manage in practice by encoding a range of prices they are willing to trade at, known as “slippage”, but we abstract from front-running in our model as we focus on the interaction between liquidity providers and (an aggregate of) liquidity takers. [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 a mean-variance perspective. [Goyal et al. \(2023\)](#) focus on the design of convex pricing functions that maximize the fraction of trades with only uninformed trades. [Milionis, Moallemi and Roughgarden \(2023b\)](#) 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 \(2023b\)](#) but hold under a wider set of assumptions on traders’ beliefs about the token valuations.

The remainder of the paper is organized as follows. In [Section 2](#), we look at the empirical behavior of liquidity providers and document their active role in price setting on AMMs. We describe our model in [Section 3](#). In [Section 4](#), we focus on the one period model and analyze the resulting equilibrium and the welfare properties. [Section 5](#) looks at the dynamics of liquidity provision in our model and compares these dynamics to the empirical evidence. Finally, [Section 6](#) concludes.

2 Liquidity Provider Behavior


An AMM uses blockchain-based smart contracts so individuals can exchange cryptocurrencies (or tokens). Smart contracts are computer code stored on the blockchain. A feature of the Ethereum blockchain is that the functions in the code are transparent, verifiable, and immutable.³ Traders post transactions, calls to functions in the smart contracts, that are then executed by a decentralized network of validators (or “miners”). The typical AMM smart contract for a pool is specific to

³You can see the functions for a Uniswap contract at <https://etherscan.io/address/0x0d4a11d5EEaAC28EC3F61d100daF4d40471f1852#code>.

two coins. To characterize the empirical behavior of liquidity providers, we pull data for Uniswap v2 pools. Specifically, we look at v2 pools created prior to 2020-07-01 that have more than 100,000 transactions. The sample period is from 2020-07-01 to 2024-06-30 and contains 19.2 million transactions across 31 pools.⁴

For context, Figure 1 shows the evolution of trading volume for the Uniswap v2 and v3 contract pools (v1 and v4 both have negligible volumes). Note the volumes are denominated in Bit-

Figure 1: Uniswap v2 and v3 Trading Volume



Volume of trade on Uniswap across all pools measured in BTC. Uniswap v2 has mint and burn liquidity functions to add or remove coins from a pool. Uniswap v3 allows LPs to specify custom price ranges in which their liquidity is active. Volume for Uniswap v1 and v4 is negligible. Source: <https://dune.com/>. Monthly trading volume in dollars is converted to BTC using the average monthly USD/BTC exchange rate, computed from daily price data provided by the Federal Reserve Economic Data (FRED).

coin to help control for the large variation in the dollar-denominated value of cryptocurrencies over this period. The data from v2 are particularly relevant since these smart contracts limit the transaction space to a swap, a mint, or a burn. This limited smart contract functionality allows us to explore the behavior of liquidity providers in a straightforward and tractable fashion. Given the evidence we present below of active liquidity management, the smart contract modifications that follow in v3 (and now v4) that offer more active control for liquidity providers are understandable.

A transaction in a Uniswap v2 pool is a call to one of three functions defined by the pool contract. The functions are a swap, a mint, or a burn. The most commonly called AMM function is the swap transaction. Here, a trader deposits a quantity of one coin, say A, and withdraws a quantity of the other coin, say B. The rate (or price) of this exchange is calculated by the smart contract based on the reserve balance of coins currently on deposit at the pool. To calculate the rate Uniswap v2 uses the constant product market maker (CPMM). The rate of exchange is determined

⁴We use the Etherscan API Pro Services to collect the data.

so that the product of the quantity of coins A and B before and after the exchange is constant. In the absence of fees, this implies the marginal rate for the exchange depends only on the relative quantities of coin A and coin B in the pool. The swap is “taking” liquidity from the pool in the sense that the swap necessarily changes the relative quantities of coin A and coin B in the pool and thus distorts the marginal exchange rate faced by subsequent traders. We refer to a trader who only uses the swap function as a Liquidity Taker (denoted LT).

The liquidity provision portions of the smart contract happen through the mint and burn functions in the code. In a mint transaction, the trader deposits both coins A and B. Here, the mint refers to the creation of pool tokens that denominate the traders proportional claim to liquidity in the pool. Since this transaction provides more coins to the pool for use in swap transactions, we call this trader a Liquidity Provider (LP). In a burn transaction, the LP uses some or all of their pool tokens to withdraw some or all of their share of coins from the pool’s liquidity reserves. Both of these functions, by design, increase or decrease the size of the pool proportionally. This feature is hardcoded into the smart contract: mints and burns do not change the ratio of the quantity of coin A to B and so do not change the implied marginal price on offer at the pool.

Given the trading environment defined by these smart contract functions, how do liquidity providers behave? The general view is that LPs are passive in that they trade (post a transaction to the smart contract) infrequently and use only the liquidity provision functions mint and burn. They are akin to “buy-and-hold” investors. Uniswap in their documentation for v2, for example, highlights that the passive aspect is a feature that may increase participation of liquidity providers by removing the need for the sophisticated infrastructure and algorithms of a liquidity provider—a market maker—in a limit order book market.⁵

Table 1 characterizes traders’ behavior with transaction counts. (In Section 5.3 below we also explore the price impact of these trades once we have used our model to highlight relevant measures of price impact in the data.) To construct the table, we tag each transaction as coming from a liquidity provider (LP) or a liquidity taker (LT). We tag a transaction as coming from an LP if the trader, at the time of the transaction, owns a pool token. That is, the transaction is by someone who owns a proportionate claim to the pool.⁶ LTs own no pool tokens at the time of the trans-

⁵See <https://docs.uniswap.org/contracts/v2/concepts/core-concepts/pools>. See also the discussions in [Malinova and Park \(2024\)](#) and [Lehar and Parlour \(2023\)](#).

⁶More precisely, we classify a transaction as belonging to an LP if: (a) it is a mint or burn; (b) any of the addresses involved in the transaction have a positive balance of that pool’s tokens at the time of the transaction; or (c) the swap transaction is paired with a mint transaction. We describe the process we use to measure paired transactions below.

action. Since mints and burns both imply pool token ownership, LT transactions are exclusively swaps.

Table 1: Transaction Counts

	burn	mint	swap	Total	Percent
Full Sample: 2020-07-01 - 2024-06-30					
LP	149,911	196,128	108,802	454,841	2.4%
LT	0	0	18,742,272	18,742,272	97.6%
Total	149,911	196,128	18,851,074	19,197,113	100.0%
v2 Dominant Contract: 2020-07-01 - 2021-05-31					
LP	112,762	157,363	80,696	350,821	3.7%
LT	0	0	9,142,338	9,142,338	96.3%
Total	112,762	157,363	9,223,034	9,493,159	100.0%
v3 Dominant Contract: 2021-06-01 - 2024-06-30					
LP	37,149	38,765	28,106	104,020	1.1%
LT	0	0	9,599,934	9,599,934	98.9%
Total	37,149	38,765	9,628,040	9,703,954	100.0%

A transaction is classified as belonging to an LP if: (a) it is a mint or burn; (b) any of the addresses involved in the transaction have a positive balance of that pool’s tokens at the time of the transaction; or (c) the swap transaction is paired with a mint transaction. Each transaction can involve several addresses (both wallets and contracts). A transaction is paired if all the addresses on both transactions match and the transactions both occur within a three-minute interval. Data is pulled from all Uniswap v2 pools that were created prior to 2020-07-01 and have more than 100,000 transactions. There are 31 pools. The sample period is from 2020-07-01 to 2024-06-30. The total number of transactions is 19.2 million. A transaction is defined as a unique call to a Uniswap pool contract as a swap, mint, or burn, and involves multiple addresses (wallets and contracts) and token transfers.

Focusing on the full sample for a moment (in the top panel of Table 1), we can see that most of the transactions are from LTs. LPs are passive in that across all transactions, they transact infrequently. Trades by LPs are 2.4% of the total transactions. However, in contrast to the “buy-and-hold” passive characterization, liquidity providers are not completely passive. A significant proportion of liquidity provider transactions are swaps (108,802/454,841). Since every swap directly impacts the token exchange rate offered by the pool, when liquidity providers interact with their pools, 23.9% of the time they take actions that directly impact the pool’s exchange rate. We view these swaps by liquidity providers as evidence that at least some LPs play an active role in price setting and price discovery in AMM markets. Finally, the different panels in Table 1 show the transaction counts across subsample periods. The impact of the introduction of v3 contracts

on v2 volume and use is clear. However, the activeness of the LPs (the percentage of transactions as swaps) is similar across the subsamples.

Table 2: Uniswap LP Counts

	Unique ID Trader	Total Transactions	Liquidity Provisions	Liquidity Takings
Full Sample: 2020-07-01 - 2024-06-30				
LP Active	44,636	201,162	45.9%	54.1%
LP Passive	64,127	253,679	100.0%	0.0%
v2 Dominant Contract: 2020-07-01 - 2021-05-31				
LP Active	37,267	148,115	45.5%	54.5%
LP Passive	53,519	202,706	100.0%	0.0%
v3 Dominant Contract: 2021-06-01 - 2024-06-30				
LP Active	7,743	51,330	45.2%	54.8%
LP Passive	12,318	52,690	100.0%	0.0%

A transaction is classified as belonging to an LP if: (a) it is a mint or burn; (b) any of the addresses involved in the transaction have a positive balance of that pool's tokens at the time of the transaction; or (c) the swap transaction is paired with a mint transaction. Active liquidity providers are defined as having more than 1 percent of their trades as swaps. Data is pulled from all Uniswap v2 pools that were created prior to 2020-07-01 and have more than 100,000 transactions. There are 31 pools. The sample period is from 2020-07-01 to 2024-06-30. The total number of transactions is 19.2 million. A transaction is defined as a unique call to a Uniswap pool contract as a swap, mint, or burn, and involves multiple addresses (wallets and contracts) and token transfers.

Table 2 highlights that the percentage of LP transactions that are swaps differs across traders and shows that 41.0% of LPs in our sample have swap transactions while the remainder of traders are completely passive. The number of LPs does decline after the introduction of Uniswap v3, but the percentage of active LPs remains about the same at 38.6%. For the active LPs, about half of their trades are swaps. Figure 7(a) highlights that the swap percentage for LPs also differs across pools. (We will return to 7(b) below.)

While we describe behavior at the level of “traders” (LP or LT), we do not directly observe individuals. On the blockchain, activity is recorded at the level of addresses (public keys), and a given transaction typically involves multiple addresses. Addresses on Ethereum can be wallets or smart contracts. For example, an individual might trade by connecting her wallet to the Uniswap web-app. This creates a trade involving her wallet ID and a Uniswap router contract.

Larger traders often trade using their own smart contracts.⁷ We make the simplifying (and conservative) assumption of treating the set of addresses involved in a transaction on a specific pool as a single unit (effectively concatenating pool address and all addresses in the transaction to define a unique trader ID). This approach will under-count liquidity providers who also use swap transactions in cases where one trader (person) uses multiple wallets or modalities for different transaction types.

Since mint and burn transactions do not impact the token exchange rate by the pool, an LP who wishes to supply liquidity at a different price than the current price offered in the pool is likely to engage in a paired transaction that either swaps tokens (to change the price) and then mint (or vice-versa). Using our conservative definition of a unique “trader” as the concatenation of all addresses used in the trade, we count how often an LP actively changes marginal pool prices (with a swap) just prior to adding liquidity (with a mint). Specifically, we pair LP transactions when a swap and a mint (or burn) occur in a three minute window.⁸ Table 6 shows that almost all the paired transactions are swaps connected to mints. Focusing on the subsample where Uniswap v2 was the dominant contract, 10.7% of the mints were preceded by a swap transaction (by the same LP) and 21.0% of all swaps conducted by LPs were in support of a subsequent mint transaction (again, by the same LP).⁹ Figure 7(b) shows that the percentage of mints paired with a swap differs across pools similarly to the raw frequency of swap transactions by LP. Many of the summary statistics we have calculated here are similar across the subsample periods. Interestingly, this is not the case for paired transactions. Figure 8 shows the percentage of mints paired with a swap is declining over the Uniswap v2 era. After the introduction of v3 with more fine grained liquidity choices, the frequency of mints paired with a swap goes to near zero for most pools.¹⁰

In summary, we have presented evidence from Uniswap v2 that some liquidity providers play an active role in setting (marginal) relative prices by swapping against their own liquidity. Next, we build a dynamic model with frictions that provides liquidity providers incentives to actively set prices at an AMM. We use this model to develop insights into optimal active liquidity provision and to establish a lens to study the dynamics of the swap behavior of liquidity providers. We explore this dynamic behavior in Section 5.3 below.

⁷This is done for many reasons. For example, this can add a layer of security to the trading where first coins are transferred to a trader’s smart contract and then the smart contract is called in a second transaction.

⁸Defining pairs as happening in a three-minute interval is arbitrary. The specific numerical results change with different windows, say 0.5 to 5.0 minutes, but the general proportions are quantitatively similar.

⁹While not immediate from Table 6, almost all the paired transactions are where the swap precedes the mint.

¹⁰Oddly, in the brief window 2023/0/01 to 2022/06/30, in WETH-USDC and USDC-USDT the swap-then-mint pair accounts for about 75% of the mint transactions. As we saw, there are also fewer overall transactions in this period.

3 Model

We develop a model where traders’ relative valuation of coins consists of both a private and a common value component. The private value component motivates gains to trade. The common value component is public information that evolves over time. We model the arrival of trading opportunities as sequential and so some traders will be “informed” in that they have arrived at the same time as new information. The public, common value component creates the potential for an “adverse selection” cost. This cost is sometimes called “impermanent loss” in the AMM setting.

Our model is in discrete time, $t = 0, 1, 2, \dots$ and features two types of agents: liquidity takers (traders) and liquidity providers (or market makers). Liquidity takers are short-lived, have deep pockets, and care about net trading profits. Liquidity providers are long-lived, discount the future at rate $\delta \in (0, 1)$, and begin in period $t = 0$ with a fixed endowment of tokens or coin balances. We focus on a representative liquidity taker (in each period) and a representative liquidity provider.

Information. We study a model with two coins (or tokens), $i \in \{A, B\}$. Each coin i has a value at (the end of) date t given by $\exp(p_{i,t})$. We interpret the common value component, $\exp(p_{i,t})$, as either the “price” of token i at time t or possibly the service flow attainable by spending 1 unit of coin 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 ([Circle \(2023\)](#)). We assume the “price” or common payoff of coin i at time t evolves according to

$$p_{i,t} = \sum_{s=0}^t d_{i,s}$$

with the public information at each date $d_{i,s}$.

In particular, assume public information $\{d_{i,t}\}$ arrives independently across time and across tokens. For each token i , with probability $\hat{\pi}$, $d_{i,t} = 0$. With probability $1 - \hat{\pi}$, $d_{i,t} \in \{-\Delta_l, +\Delta_h\}$ 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 after the realization of public information is the same as it is before this information is realized. At the beginning of each period t , both LPs and LTs have beliefs about the common

value component of each token given by $\mu_{i,t-1}$ where

$$\mu_{i,t-1} = E[\exp(p_{i,t})|d_0, \dots, d_{t-1}] \equiv E_{t-1}[\exp(p_{i,t})].$$

Timing. At the beginning of each period, with probability $\beta \in (0, 1)$, the LP exogenously exits the game and realizes the current payoff of her endowment of tokens. If the LP does not exit, then, before the arrival of any public information, 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

$$v_{i,t} = E[\exp(p_{i,t})|d_0, \dots, d_t] \exp(\eta_{i,t}) \equiv E_t[\exp(p_{i,t})] \exp(\eta_{i,t})$$

where $\eta_{i,t}$ reflects a private value component of owning token i realized by the LT that trades in period t . The important timing assumption is that the LT trades before the LP can withdraw or adjust their deposits to the AMM. Once the LT trades, a new period begins and the LP may re-balance the liquidity supplied to the AMM.

We further specialize the information setting of our model to highlight the key forces at play. Note that in any period, before the arrival of public information, both the liquidity providers and liquidity takers have the same beliefs given by $\mu_{i,t-1}$. Once public information arrives, the LT who trades in period t has valuation $v_{i,t}$ distinct from $\mu_{i,t-1}$ because she has more public information and because of her private value shock. We impose a particular correlation between the public information and the LT's private value shocks. Since there are two coins, it is with probability $1 - \hat{\pi}^2$ that the LT has superior public information since $d_{i,t} \in \{-\Delta_l, \Delta_h\}$ for some token i . For such realizations, we impose $\eta_A = \eta_B = 0$.

Under this specification, our model features two types of information events as in [Glosten and Milgrom \(1985\)](#). The first type of information event—analogue to uninformed trading in [Glosten and Milgrom \(1985\)](#)—occurs when $d_{A,t} = d_{B,t} = 0$ and represents a case where the LT's new beliefs of the tokens' values, $v_{i,t}$ are uncorrelated with the LP's beliefs. That is, the LP believes the value of each token i will yield terminal value according to $E_t[\exp(p_{i,t})] = E_{t-1}[\exp(p_{i,t})]$ while the LT believes the value of each token i is distributed according to $v_{i,t} = E_t[\exp(p_{i,t})] \exp(\eta_{i,t})$. When $\eta_{i,t} \neq 0$, 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 denote $\pi = \hat{\pi}^2 \in [0, 1]$ as 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—analogue to informed trading in [Glosten and Milgrom \(1985\)](#)—occurs when $d_{i,t} \in \{-\Delta_l, \Delta_h\}$ (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 $v_{i,t} = E_t[\exp(p_{i,t})]$ and hence there are no gains to trade between the LT and the LP. For notational simplicity, we assume $v_{i,t}$ follow the same distributions under the two events.¹¹ Following the literature, we interpret such an event as pure information event that we describe as an *informed trade*. The correlation between the information arrival and private values of the LT that we impose allows us isolate the idea that liquidity takers may trade for “information” or may trade for “tastes.”

Denote $(E_{A,t-1}, E_{B,t-1})$ as the amount of tokens the LP owns at the end of each period $t - 1$. If the LP does not exit in period t , she chooses the amount of tokens $(e_{A,t}, e_{B,t})$ to deposit in the AMM. After her deposit, public information is realized and the LT trades in the AMM. Let $(x_{A,t}, x_{B,t})$ denote the amount of tokens remaining in the AMM after the LT's trade. Post-trade, the LP owns $(E_{A,t}, E_{B,t}) = (E_{A,t-1} - e_{A,t} + x_{A,t}, E_{B,t-1} - e_{B,t} + x_{B,t})$ tokens.

With probability π , the LT's trade is uninformed and the LP's valuation of each token remains unchanged. Alternatively, with probability $1 - \pi$, the trade is informed and the LP's valuation of each token updates to that of the LT.

Suppose that the LP has deposited a portfolio $(e_{A,t}, e_{B,t})$ with the smart contract of the AMM. We let $G(\cdot)$ be 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 $q_B = G(q_A | e_A, e_B)$. The most common implementation of automated markets imposes the constant product market maker (CPMM):

$$(e_A + q_A)(e_B - q_B) = e_A e_B \quad (1)$$

The simple function, although ad-hoc, has the attractive property that marginal prices are convex. The more you withdraw, the higher the marginal price. It also imposes that the contract pool will “run out” of either token, as marginal prices approach infinity as the withdrawal amount

¹¹Allowing $v_{i,t}$ to follow different distributions does not substantively change our theoretical results. We do allow for different distributions in our numerical results in Section 5.2 below.

approaches the token balance in the pool. In practice, the CPMM function includes a transaction fee paid to the pool. In our analysis, we abstract from transaction fees.

Next, we define the problem of the liquidity provider and liquidity taker working backwards from the LT's problem in each period. We maintain the Constant Product Market Making rule specified in Equation (1) through Section 3.1, 3.2, 4, and 5 below.

3.1 The Liquidity Taker's Problem

In each period, the LT—whether uninformed or informed—observes liquidity on deposit at the AMM as well as the realization of v_i . From their perspective, the LT perceives a favorable trading opportunity as prices in the AMM do not automatically adjust to their own valuation. Since the LT is short-lived, we omit time subscripts when describing the LT's behavior.

The LT maximizes the expected value of her tokens:

$$\begin{aligned} \max_{q_A, q_B} \quad & -v_A q_A + v_B q_B \\ \text{s.t.} \quad & (e_A + q_A)(e_B - q_B) = e_A e_B. \end{aligned} \tag{2}$$

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 that she faces in 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 $q_A \rightarrow -\infty$) and hence the implicit capacity constraints are slack under such a rule.


The solution to the LT's problem is straightforward, and, in terms of ex-post reserves remaining in the pool after the LT's trade implies

$$e_A + q_A = \sqrt{\frac{v_B}{v_A} e_A e_B}, \quad e_B - q_B = \sqrt{\frac{v_A}{v_B} e_A e_B}. \tag{3}$$

More succinctly, for any beliefs v_i , the LT will trade up until the relative price at the AMM equals her relative valuation of the tokens or

$$\frac{v_B}{v_A} = \frac{e_A + q_A}{e_B - q_B}. \tag{4}$$

Figure 2: Liquidity Taker's Optimal Trade



The liquidity taker's optimal trade is characterized by the tangency of their relative valuations (red line) and the constant product market making curve (orange line).

Notice that $x_A = e_A + q_A$ and $x_B = e_B - q_B$ are the post-trade AMM positions so that (1) and (4) imply that the post-trade positions satisfy

$$x_A x_B = e_A e_B \quad (5)$$

$$v_A x_A = v_B x_B. \quad (6)$$

The liquidity provider internalizes that for any realization of beliefs of the LT, v_i , her ex-post portfolio will satisfy (5)–(6). We may represent this behavior graphically as in Figure 2.

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 v_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 $-v_A/v_B$).

3.2 The Liquidity Provider's Problem

Given $\mu_{A,0}, E_{A,0}, \mu_{B,0}, E_{B,0}$, the problem of the LP at time 0 can be written as:

$$\max_{\{e_{A,t}, e_{B,t}\}_{t=1}^{\infty}} \sum_{t=0}^{\infty} \delta^t \beta (1 - \beta)^t \mathbb{E} [\mu_{A,t} E_{A,t} + \mu_{B,t} E_{B,t}] \quad (7)$$

where for $t = 1, 2, 3, \dots$

$$\begin{aligned}
x_{A,t} x_{B,t} &= e_{A,t} e_{B,t} \\
v_{A,t} x_{A,t} &= v_{B,t} x_{B,t} \\
E_{i,t} &= x_{i,t} + (E_{i,t-1} - e_{i,t}) \\
\mu_{i,t} &= \begin{cases} \mu_{i,t-1} & \text{with prob } \pi \\ v_{i,t} & \text{with prob } 1 - \pi \end{cases} \\
0 &\leq e_{i,t} \leq E_{i,t-1}
\end{aligned}$$

The first two constraints embed the LT's behavior in each period. The third and fourth constraints reflect the law of motion for the LP's endowments and her beliefs. The final set of constraints reflect feasibility constraints for the LP.

We now formulate the LP's problem as a stationary, dynamic program. Suppose the LP starts a given period with endowments (E_A, E_B) and beliefs about these token's values, (μ_A, μ_B) . Then the Bellman equation is given by

$$\begin{aligned}
V(E_A, E_B; \mu_A, \mu_B) &= \beta [\mu_A E_A + \mu_B E_B] \\
&+ (1 - \beta) \delta \max_{e_A, e_B} (\pi \mathbb{E}V(E'_A, E'_B; \mu_A, \mu_B) + (1 - \pi) \mathbb{E}V(E'_A, E'_B; v_A, v_B)) \quad (8)
\end{aligned}$$

subject to

$$0 \leq e_i \leq E_i \quad (9)$$

$$x_A x_B = e_A e_B \quad (10)$$

$$v_A x_A = v_B x_B \quad (11)$$

$$E'_i = x_i + (E_i - e_i). \quad (12)$$

With probability β , the LP exits and enjoys the expected utility of her endowment. Should the LP not exit, then she chooses the quantities of tokens to deposit on the exchange, e_i , subject to the feasibility constraints (9). The constraints (10) and (11) summarize the behavior of the liquidity taker for any realization of the public information or the private taste shock (v_A, v_B) which in turn dictate how the LP's endowment will evolve into the subsequent period as summarized in (12). Given her liquidity deposits (e_A, e_B) , with probability π there is no public information event so that the LP's beliefs remain constant at (μ_A, μ_B) . Alternatively, if there is a public information

event, which occurs with probability $(1 - \pi)$, then the LP's beliefs evolve and are consistent with those of the LT given by (v_A, v_B) .

4 AMM Liquidity: Insights from a One-Shot Model

In this section, we focus on a one-period model to highlight the key results that emerge from our dynamic model. Exactly as in the dynamic model, at the beginning of the period, the LP deposits a portfolio (e_A, e_B) with the AMM given a pricing function $G(\cdot)$ and 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 (v_A, v_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 (v_A, v_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. Here we set $\delta = 1$ (no discounting) and $\beta = 1$ (the LP exits and enjoys the terminal value of the tokens after one period.).

In the one-shot game, we solve for the LP's optimal liquidity supply and show how it depends on the LP's beliefs about the probability of informed versus uninformed trading. We use this simple model to 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 typically inconsistent with profit-maximizing behavior. It is optimal for the representative liquidity provider only when informed trading is so severe that the liquidity provider prefers to supply no liquidity. We demonstrate how adverse selection distorts the quantities of liquidity deposited by providers on automated exchanges. Finally, we examine how the shape of the AMM pricing function impacts gains to trade realized by liquidity providers.

4.1 The Liquidity Provider's Problem in the One-Shot Model

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

$$\begin{aligned}
& \max_{e_A, e_B} \pi(\mu_A \mathbb{E}[x_A - e_A] + \mu_B \mathbb{E}[x_B - e_B]) + \\
& \quad (1 - \pi)(\mathbb{E}v_A[x_A - e_A] + \mathbb{E}v_B[x_B - e_B]) \\
& \text{s.t. (5)-(6),} \\
& \quad 0 \leq e_i \leq E_i, \quad \forall i
\end{aligned} \tag{13}$$

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, v_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.

4.2 Liquidity Provision with Uninformed Trade Only

Suppose first that $\pi = 1$ so that there are only uninformed trades. The LP's problem (13) simplifies to

$$\begin{aligned}
& \max_{e_A, e_B} \mu_A \left(\mathbb{E} \sqrt{\frac{v_B}{v_A}} e_A e_B - e_A \right) + \mu_B \left(\mathbb{E} \sqrt{\frac{v_A}{v_B}} e_A e_B - e_B \right) \\
& \text{s.t. } 0 \leq e_i \leq E_i, \quad \forall i.
\end{aligned}$$

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


$$\left(\mathbb{E}\omega + \mathbb{E}\frac{1}{\omega} - 2 \right) \sqrt{\mu_A e_A} \sqrt{\mu_B e_B} - (\sqrt{\mu_A e_A} - \sqrt{\mu_B e_B})^2 \tag{14}$$

where $\omega = \sqrt{\frac{\nu_A/\mu_A}{\nu_B/\mu_B}}$. Equation (14) 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 (14), 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.

Panel (a) of Figure 3 illustrates this result. 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.

Figure 3: **Liquidity Provider's Deposit Choice**



Panel (a) illustrates a deposit choice of the LP that ensures zero losses if she only faces uninformed liquidity takers. The straight (blue) line shows the indifference curve an LP assuming her beliefs are given by (μ_A, μ_B) and do not change when trade takes place. The (orange) curve shows the set of possible ex-post allocations. Panel (b) shows an LP's optimal deposit is (typically) not a no loss deposit.

Should the LP provide liquidity different from an equal value deposit, then for small differences in beliefs from her own, the constant product 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 (14)—that depends on how her portfolio differs from an equal value ($\mu_A e_A = \mu_B e_B$) portfolio.

To the extent v_i differs from μ_i , there are gains to trade. The value of these gains depend on the term $\mathbb{E}\omega + \mathbb{E}\frac{1}{\omega} - 2 \geq 0$. (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 by supporting more trading with uninformed LTs at typically favorable pricing. 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.

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

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

Generically, 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 Panel (b) of Figure 3 where, according to Proposition 1 typically, we expect either $e_A = E_A$ or $e_B = E_B$.

4.3 Liquidity Provision with Informed Trade Only

Suppose next that $\pi = 0$ so that there are only informed trades. The LP's problem (13) simplifies to

$$\max_{e_A, e_B} \mathbb{E}v_A \left(\sqrt{\frac{v_B}{v_A}} e_A e_B - e_A \right) + \mathbb{E}v_B \left(\sqrt{\frac{v_A}{v_B}} e_A e_B - e_B \right) \quad (15)$$

$$\text{s.t. } 0 \leq e_i \leq E_i, \quad \forall i \quad (16)$$

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

$$(2\mathbb{E}\psi - 2) \sqrt{\mu_A e_A} \sqrt{\mu_B e_B} - (\sqrt{\mu_A e_A} - \sqrt{\mu_B e_B})^2 \quad (17)$$

where $\psi = \sqrt{\frac{v_A}{\mu_A} \frac{v_B}{\mu_B}}$. Equation (17) 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.¹²

The Cauchy-Schwarz inequality implies $\mathbb{E}\psi \leq \sqrt{\mathbb{E} \frac{v_A}{\mu_A} \mathbb{E} \frac{v_B}{\mu_B}}$ and holds with equality only when v_A and v_B are perfectly correlated. Since we impose $\mathbb{E} v_i / \mu_i = 1$, the above inequality implies $\mathbb{E}\psi \leq 1$. 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: $e_A^* = e_B^* = 0$.

Proposition 2 is consistent with the conventional view in the nascent literature on AMMs: if arbitrageurs have frictionless access to a centralized exchange where price discovery for the tokens takes place as well as the AMM, then in the absence of fees LPs cannot profit from supplying liquidity to the AMM. Fees must then be imposed by the AMM to make liquidity provision sustainable. In contrast, when arbitrageurs face enough frictions, which occurs in our model when π is not too small, then the market may be sustainable even in the absence of fees.

4.4 Liquidity Provision with Uninformed and Informed Trading

We now use these results to understand better the general problem (13) with arbitrary π . We once again simplify the LP's objective function as

$$\left[\pi \left(\mathbb{E}\omega + \mathbb{E} \frac{1}{\omega} \right) + (1 - \pi) 2\mathbb{E}\psi - 2 \right] \sqrt{\mu_A e_A} \sqrt{\mu_B e_B} - (\sqrt{\mu_A e_A} - \sqrt{\mu_B e_B})^2 \quad (18)$$

¹²Since 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. $\pi(\mathbb{E}\omega + \mathbb{E}\frac{1}{\omega}) + (1 - \pi)2\mathbb{E}\psi \geq 2$, 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

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

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

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


otherwise.

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

$$\frac{\mu_A e_A^*}{\mu_B e_B^*} = \begin{cases} \frac{1}{\Pi^2} & \text{if } E_A \leq \frac{1}{\Pi^2} \frac{\mu_B}{\mu_A} E_B \\ \frac{\mu_A E_A}{\mu_B E_B} & \text{if } \frac{1}{\Pi^2} \frac{\mu_B}{\mu_A} E_B < E_A < \Pi^2 \frac{\mu_B}{\mu_A} E_B \\ \Pi^2 & \text{if } \Pi^2 \frac{\mu_B}{\mu_A} E_B \leq E_A \end{cases} . \quad (19)$$

We illustrate Proposition 3 in Figure 4 for cases where the optimal deposit is strictly positive (so $\Pi > 1$). We slowly vary the LP's endowments of tokens (E_A, E_B) making token A relatively more plentiful as we move from Panel A to Panel C. Initially, when the LP's endowment of token A is relatively scarce (shown in Panel A), she deposits all of token A and an interior amount of token B. When the LP's endowments are tokens are relatively balanced (near $\mu_A E_A = \mu_B E_B$ shown

Figure 4: Liquidity Provider's Optimal Deposit Choice



The figures illustrate different cases of the optimal liquidity deposit (e_A^*, e_B^*) described in Proposition 3 for various possible initial endowment points (E_A, E_B) . Moving from Panel A to Panel C, we slowly increase the LP's initial endowment of token A making it relatively more plentiful and highlighting how according to Proposition 3, differences in her relative endowments impacts which token she supplies completely.

in Panel B), she deposits all of both tokens. When the LP's endowment of token A is relatively abundant (shown in Panel C), she deposits all of token B and an interior amount of token A.

Next, we use Proposition 3 to explore the optimality of the conventional wisdom that liquidity providers should deposit portfolios with equal values. Notice that when $\Pi > 1$, the optimal deposit ratio, $\mu_A E_A / \mu_B E_B$ is only 1 if the LP's endowments are relatively balanced (as in Panel B of Figure 4) and her endowment satisfies $\mu_A E_A = \mu_B E_B$. This suggests that the conventional wisdom the liquidity providers should deposit portfolios with equal values is typically not profit maximizing for liquidity providers. Furthermore, Proposition 3 reveals that as $\Pi \rightarrow 1$ then $\mu_A e_A^* \rightarrow \mu_B e_B^*$ 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.

Proposition 3 also suggests an important feature of optimal liquidity provision in our model. The liquidity provider optimally supplies her entire endowment to the AMM when the ratio of her endowments (in token quantities) E_A/E_B lies in a region that depends on her relative value of tokens μ_A/μ_B as well as the profit margin Π . If we interpret the LP's endowment in the static model as having arisen from past trades by liquidity takers, then should her endowment lie in this region, she will not "re-balance" her deposits—she will simply leave her entire endowment on the AMM. We return to this point in Section 5 when we study dynamics of liquidity provision below.

Note also 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 $\mu_A E_A \neq \mu_B E_B$.¹³ The equal value deposit $\mu_A e_A = \mu_B e_B$ is optimal only when $\pi = \underline{\pi}$.

4.5 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 v_i , a mean preserving spread of the LT's beliefs v_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.

To simplify the numerical analysis, consider a special case where one token is a stablecoin whose value does not fluctuate over time such as USDC or Tether.¹⁴ We let token B represent the stable coin and set $v_B = \mu_B = 1$ and $h_B(v_B) = 1$ if $v_B = 1$. Then we have $\omega = \psi = \sqrt{\frac{v_A}{\mu_A}}$. We assume $\frac{v_A}{\mu_A}$ is a log-normally distributed random variable with $E[v_A/\mu_A] = 1$ and $\text{Var}[v_A/\mu_A] = \sigma_A^2$. 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—from 2018 October to 2023 October.¹⁵ 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 9, which shows that increases in variance typically decrease this threshold.¹⁶ In other words, liquidity provision becomes more profitable (LPs can tolerate more informed trading) as ETH price risk increases.

¹³If 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.

¹⁴In practice, the value of stablecoins do fluctuate at specific points in time, such as when USDC de-pegged for a short window in April 2023. For our example, we assume liquidity providers and takers believe the stablecoin peg will hold with certainty.

¹⁵Based on the Coinbase ETH index price obtained from fred.stlouisfed.org.

¹⁶We experimented with several other distributional assumptions for $\frac{v_A}{\mu_A}$ and found similar results. Details are available upon request.

4.6 Efficiency Losses from Constant Product Market Making

Finally 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 \(2023b\)](#) for such an approach applied in an environment with only one risky token and endowment limits for the LTs.) Specifically, we consider perturbing the CPMM price formula and study a class of pricing functions given by

$$(e_A + (1 - \tau)q_A)(e_B - (1 - \tau)q_B) = e_A e_B \quad (20)$$

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 function. For larger values of q_i , it is possible that 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, $e_A \geq q_A$ and $e_B \geq q_B$. 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 (20) as

$$((1 - \tau)x_A + \tau e_A)((1 - \tau)x_B + \tau e_B) = e_A e_B. \quad (21)$$


The price function (20) is convex and smoothly decreasing when $x > 0$. The convexity of the function is decreasing in τ . The boundary conditions on q_i simply imply $x_i \geq 0$.

For a given the realization of the LT's beliefs, (v_A, v_B) , (21) implies that the LP's net proceeds from trade satisfy

$$x - e_A = \frac{1}{1 - \tau} \left[\sqrt{\frac{v_B}{v_A}} e_A e_B - e_A \right], \quad y - e_B = \frac{1}{1 - \tau} \left[\sqrt{\frac{v_A}{v_B}} e_A e_B - e_B \right]. \quad (22)$$

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 π does not change with τ . As a result,

Figure 5: CPMM with change in local convexity



The figure illustrates a perturbation of the CPMM curve that decreases the (local) convexity at the LP's deposit point. The CPMM curve with $\tau = 0$ is displayed as the solid, orange curve while the CPMM curve with $\tau > 0$ is displayed as the dashed, blue curve.

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 21 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 see this, 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 $v_A/\mu_A > p_h$), 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 $p_h > v_A/\mu_A > 1$, 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 (20) 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 (v_A, v_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 $[\underline{\mu}_i, \overline{\mu}_i]$ is the support of the distribution of v_i , then the result of Proposition 4 holds for all $\tau \leq \bar{\tau} = \min \left\{ \sqrt{\frac{\mu_B e_B}{\mu_A e_A}}, \sqrt{\frac{\mu_A e_A}{\mu_B e_B}} \right\}$ 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:

$$\begin{cases} p_l x_A + x_B = p_l e_A + e_B, & x \geq e_A \\ p_h x_A + x_B = p_h e_A + e_B, & x < e_A \end{cases} \quad (23)$$

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 \(2023b\)](#), 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. her endowment

of either token A or B is finite. Given the LP's deposit (e_A, e_B) , the optimal pricing formula is the linear pricing formula if one of the following conditions is satisfied:

1. All trades are uninformed trading, i.e. $\pi = 1$;
2. The LT's value (v_A, v_B) follows the same distribution for both informed and uninformed trading, one of the two tokens is a stablecoin (i.e. if token A is stable then $v_A = \mu_A$ for sure) and $\pi \neq 0$ (i.e. there exists some uninformed trading).

5 AMM Liquidity: Insights from the Dynamic Model

In this section, we solve for and simulate the optimal supply of liquidity in the dynamic model. We show how the LP's dynamic problem, which is a function of four state variables (her endowment of each token and her beliefs about the value of each token), may be simplified using an auxiliary dynamic problem with a single endogenous state variable and a single exogenous state variable. We use this approach to solve and simulate the dynamic supply of liquidity.

We use our simulations to study the dynamics of optimal liquidity—that is, we study how optimal liquidity responds to trading by liquidity takers. First, we show that the LP's responses typically feature action and inaction regions. For trades (informed or uninformed) that have little price impact, LPs typically do not re-balance their liquidity. For large trades, however, LPs typically re-balance. Moreover, LPs are more likely to re-balance when trading is uninformed. Finally, we explore our empirical evidence on AMM trades and liquidity provider behavior and demonstrate similar patterns exist in the data.

5.1 The Dynamic Model Solution

Although the liquidity provider has linear (risk-neutral) preferences, the evolution of her endowments is not immediately linear given the convex pricing curve G and the behavior of the liquidity takers in each period. Nonetheless, we are able to simplify our dynamic model which naturally has four state variables.

Towards this end, observe that $E_{A,t}, E_{B,t}, \mu_{A,t}, \mu_{B,t}$ only show up in the objective function (7) as products $\mu_{A,t} E_{A,t}$ and $\mu_{B,t} E_{B,t}$. We now show that the same holds in the LP's constraints in

each period. Define $(r_{A,t}, r_{B,t}) = \left(\frac{v_{A,t}}{\mu_{A,t-1}}, \frac{v_{B,t}}{\mu_{B,t-1}} \right)$ as the rates of change in the LT's valuations relative to the LP's. Let $X_{i,t} = \mu_{i,t} E_{i,t}$ represent the expected value of each token the LP has at the end of each period t and $Y_{i,t} = \mu_{i,t-1} e_{i,t}$ represent the expected value of each token the LP deposits into the pool at the beginning of period t .

Now, suppose that $(r_{A,t}, r_{B,t})$ follows the distribution H_t , which is independent of the current beliefs $(\mu_{A,t-1}, \mu_{B,t-1})$. Then, there is a one-to-one mapping from the initial values $\{E_{A,0}, E_{B,0}, \mu_{A,0}, \mu_{B,0}\}$ and the sequence of $\{E_{A,t}, E_{B,t}, \mu_{A,t}, \mu_{B,t}, e_{A,t}, e_{B,t}\}_{t=1}^{\infty}$, to the initial values $\{X_{A,0}, X_{B,0}, \mu_{A,0}, \mu_{B,0}\}$ and a sequence of $\{X_{A,t}, X_{B,t}, r_{A,t}, r_{B,t}, Y_{A,t}, Y_{B,t}\}_{t=1}^{\infty}$ where

$$E_{i,t} = \frac{X_{i,t}}{\mu_{i,t}}, \quad \mu_{i,t} = \mu_{i,0} \prod_{s=1}^t r_{i,s}, \quad e_{i,t} = \frac{Y_{i,t}}{\mu_{i,t-1}}.$$

In other words, assuming the change in beliefs is independent of the level of beliefs renders the LP's optimal liquidity supply in each period independent of the level of beliefs. Her payoffs, of course, depend on these levels so we must track their values, but they do not influence the LP's optimal supply.

Using notation similar to that from the one-shot model, $\omega_t = \sqrt{\frac{r_{A,t}}{r_{B,t}}}$ and $\psi_t = \sqrt{r_{A,t} r_{B,t}}$, we may re-write the LP's dynamic problem as

$$\max_{\{Y_{A,t}\}_{t=1}^{\infty}, \{Y_{B,t}\}_{t=1}^{\infty}} \sum_{t=0}^{\infty} \delta^t \beta (1 - \beta)^t \mathbb{E} [X_{A,t} + X_{B,t}]$$

where for $t = 1, 2, 3, \dots$

$$\begin{pmatrix} X_{A,t} \\ X_{B,t} \end{pmatrix} = \begin{cases} \begin{pmatrix} \frac{1}{\omega_t} \sqrt{Y_{A,t} Y_{B,t}} + (X_{A,t-1} - Y_{A,t}) \\ \omega_t \sqrt{Y_{A,t} Y_{B,t}} + (X_{B,t-1} - Y_{B,t}) \end{pmatrix} & \text{with prob } \pi \\ \begin{pmatrix} \psi_t \sqrt{Y_{A,t} Y_{B,t}} + r_{A,t} (X_{A,t-1} - Y_{A,t}) \\ \psi_t \sqrt{Y_{A,t} Y_{B,t}} + r_{B,t} (X_{B,t-1} - Y_{B,t}) \end{pmatrix} & \text{with prob } 1 - \pi \end{cases}$$

$$0 \leq Y_{i,t} \leq X_{i,t-1}$$

given $\mu_{A,0}, X_{A,0}, \mu_{B,0}, X_{B,0}$. We leave a detailed proof in Appendix D. This result implies that we may re-write the sequential problem as a dynamic program with value function $V(X_{A,t}, X_{B,t})$ (rather than as $V(E_{A,t}, E_{B,t} | \mu_{A,t}, \mu_{B,t})$).

Next, we argue that this value function is homogeneous with degree one (constant return to scale).

Proposition 7: Constant Return to Scale. For any $X_{A,0}^k = kX_{A,0}$ and $X_{B,0}^k = kX_{B,0}$, $V(X_{A,0}^k, X_{B,0}^k) = V(kX_{A,0}, kX_{B,0}) = kV(X_{A,0}, X_{B,0})$ for any $k > 0$.

Proposition 7 implies that we may use two one-dimensional functions $V^i : (0, 1] \rightarrow \mathbb{R}_+$ to represent the value function instead of one two-dimensional function $V : \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$.

$$V(X_{A,0}, X_{B,0}) = \begin{cases} X_{A,0} V\left(1, \frac{X_{B,0}}{X_{A,0}}\right) & X_{A,0} \geq X_{B,0} \\ X_{B,0} V\left(\frac{X_{A,0}}{X_{B,0}}, 1\right) & X_{A,0} < X_{B,0} \end{cases} \equiv \begin{cases} X_{A,0} V^B\left(\frac{X_{B,0}}{X_{A,0}}\right) & X_{A,0} \geq X_{B,0} \\ X_{B,0} V^A\left(\frac{X_{A,0}}{X_{B,0}}\right) & X_{A,0} < X_{B,0} \end{cases}. \quad (24)$$

Being able to reduce the value function into the form of a two-dimensional function with domain $(0, 1]$ allows us to numerically solve the value functions and policy functions through policy function iteration.

5.2 Simulated Dynamics of Liquidity Provider Behavior

We now explore the features of the dynamics of optimal liquidity using simulated data from our model. Specifically, we consider the case of one risky coin and one stable coin. We assume the rate of change of beliefs of the risky coin follows a truncated normal distribution. We first examine what fraction of each token the LP deposits as a function of the ratio of the expected value of each token the LP has at the beginning of each period. Numerical experiments across a wide range of parameters suggest that at least one corner constraint always binds as in our static model.¹⁷

Figure 10 displays an illustration of the LP's optimal deposit strategy for a typical numerical example (we provide the numerical details in the figure's description). In the example, token A is a risky coin and token B is a stable coin whose value never fluctuates. On the left-hand side of the panel, the LP owns risky coins with lower expected value relative to the stable coins she owns. In this case, she deposits all of her risky coins with the AMM and retains a portion of her stable coins. The fraction of stable coins she deposits on the AMM grows as the ratio of the value of her endowment of token tends towards one. We see the inverse behavior when instead the LP owns risky coins with greater expected value relative to the stable coins she owns. In this case, she deposits all of her stable coins and only a fraction of her risky coins. In this example, the full

¹⁷ Our simulation results below (Figures 6, 10, 11) assume $\pi = 0.8$, $r_{B,t} = 1$ (token B is a stable coin), $r_{A,t}$ follows a normal distribution where the distribution for uninformed trades has mean 1 and standard deviation 0.5 while for informed trades has mean 1 and standard deviation 0.25 (both distributions are truncated to a range of 0.5 to 1.5), $\delta = 0.99$, $\beta = 0.01$. We simulate conduct 20,000 simulations of activity at the AMM in our model with each simulation lasting for 500 periods.

deposit region—when the LP deposits all of her endowment of both tokens—is when the expected value of her risky coin endowment is slightly bigger than that of her stable coin endowment (when X_B/X_A is close to 1).

Using numerically solved dynamic optimal policy functions, we conduct Monte Carlo experiments to simulate the dynamic behavior of the LP under different parameterizations. As shown in Section 2, liquidity providers transact at the AMM rarely compared to liquidity takers. Our model generates this inactivity endogenously since an LP who deposits their whole endowment may still do so even after a trade by the liquidity taker.

To understand this inaction, recall from our discussion following Proposition 3 that there is a range of relative token endowments (given the LP’s beliefs) where the LP finds it optimal to deposit her entire endowment. Suppose in some period t the LP deposits her entire endowment and there is subsequently an uninformed trade. If this trade is small—it shifts the LP’s ex-post endowment very little—then the LP is likely to remain in the region where she finds it optimal to supply her entire endowment. (With uninformed trades, the LP’s relative valuations $\mu_{A,t}$ and $\mu_{B,t}$ do not adjust.) When a large uninformed trade takes place, the LP is likely to find herself outside of the maximal supply region and will adjust her balance. However, when a large *informed* trade takes place, the LP’s maximal supply region also shifts (simulations suggest this shift is typically in the same direction as the LT’s trade) and so sometimes the LP will remain in the (new) maximal supply region. This result suggests that LP’s will trade more frequently when there is more *uninformed* trade.

Figure 11 explores this inactivity numerically by examining how the extent of the LP’s inaction depends on the severity of informed trading. This figure shows that an increase in π —the degree of uninformed trade—leads to a higher probability that the LP will re-balance her deposits in any given period.¹⁸ In other words, we should expect more LP transactions in pools that feature less informed trading (or less adverse selection).

Regardless of whether the LP faces informed or uninformed trades, when trades are sufficiently large we expect the liquidity provider to re-balance her deposits in such a way that adjusts relative prices in the opposite direction of the liquidity taker’s realized trades. Moreover, it is more likely for the LP to re-balance her deposits following an LT swap after an uninformed trade relative to an informed trade.

¹⁸While we display the percentage of periods the LP trades, the level of this value is not determined in our model as we may assume an arbitrary amount of trades by LTs in each period before the LP has the opportunity to re-balance.

To study how LPs adjust prices empirically, we construct measures of price impact by LTs and LPs at AMMs. Recall, that the ratio of the quantity of coin A to coin B defines the marginal price in the liquidity pool. Motivated by this feature of the AMM, we define the price impact of an LT trade at time t as

$$\lambda_t^{LT} = \log(x_{A,t}/x_{B,t}) - \log(e_{A,t}/e_{B,t}) \quad (25)$$

where $e_{i,t}$ are the quantities prior to the trade and $x_{i,t}$ are quantities immediately after the trade, as in 7. Similarly, we define the price impact of the LP swap after an LT trade at time t as

$$\lambda_t^{LP} = \log(e_{A,t+1}/e_{B,t+1}) - \log(x_{A,t}/x_{B,t}). \quad (26)$$


We expect these measures of price impact to be negatively correlated and that the correlation is more negative for uninformed trades than informed trades. Figure 6 illustrates this result in our numerical simulations. Each panel displays the scatter plot of λ_t^{LP} against λ_t^{LT} (panel (a) shows this relationship following informed trades and panel (b) shows this relationship following uninformed trades). In a given period, the LT experiences a shock to her valuation of token A relative to the LP's beliefs at the start of the period $r_{A,t}$ (recall token B in this simulation is a stable coin). When $r_{A,t}$ is larger than 1, the LT becomes relatively optimistic about token A and so withdraws token A from the pool and deposits token B. From (25), this swap induces a negative price impact ($\lambda_t^{LT} < 0$). Conversely, when $r_{A,t}$ is smaller than 1, the LT is relatively pessimistic about token A and will induce a positive price impact ($\lambda_t^{LT} > 0$).

In the figure, darker colors with larger positive LT price impact are associated with lower values of $r_{A,t}$ (more pessimistic views of the LT) while lighter colors with larger negative LT price impact are associated with higher values of $r_{A,t}$ (more optimistic views of the LT). There is variation in the implied price impact for a given level of $r_{A,t}$ because the state of the AMM at the time of the shock varies (and because the colors represent bins of the belief update distribution).

For a given level of $r_{A,t}$, whether trades are informed or uninformed, we observe a negative correlation between the LT's price impact and the LP's price impact caused through deposit re-balancing. (In both plots for a given color, we see a strong negative relationship.) However, for any level of $r_{A,t}$, when trades are informed, the LP is less likely to re-balance their deposits, and, as a result, the distribution of LP price impacts is (roughly) centered at no price impact. Instead, LPs respond more aggressively to large trades by uninformed LTs. When examining the overall correlation of the price impact of swaps by LTs and LPs, we expect a much weaker correlation for

informed trading than for uninformed trading as shown in Figure 6.

Figure 6: LP Swap Slope against LT Trade Slope



Scatterplot of price impact of LT swaps (λ^{LT}) with price impact of subsequent LP swaps (λ^{LP}) from numerical simulation where the value of token A follows a truncated normal distribution and the value of token B is fixed or stable. Different colors represent different rates of change in the LT's valuation of token A relative to the LP's at the beginning of each period ($r_{A,t}$). Panel (a) shows the cases of informed trades. Panel (b) shows the cases of uninformed trades. For detailed choices of parameters, see 17.

5.3 Empirical Dynamics of Liquidity Provider Behavior

In Section 2, we used data from Uniswap v2 contract to show that even with a limited set of actions, many liquidity providers are actively involved in price setting. Specifically, we observed that many liquidity providers execute swap transactions in pools where they had also provided liquidity. We also found that many mint transactions, which add liquidity to a pool, were preceded by swaps that adjust the marginal price in the pool.

In our dynamic model, building on the static framework shown in Figure 4, we characterized how liquidity providers “set prices” (by posting quantities) based on their beliefs. As liquidity takers post swap transactions, liquidity providers may want to adjust their positions. Whether they do so depends on the nature of the LT trade. Any LT swap, by definition, changes the relative quantities of pool tokens and, given the CPMM function, the relative marginal token price. If the swap came from a LT who is uninformed, the beliefs of the LP are unchanged. So, to move the pool back towards the LP optimal position, they would need to swap in the opposite direction. In

contrast, if the LT trade came from an informed trade, then the beliefs of the LP have updated. In this case, the LP may wish to swap to attain the optimal position. However, given the new beliefs centered at those motivating the last LT trade, there is no bias for the LP to swap in one direction or the other.

We can use the transaction data for the Uniswap v2 data to examine this behavior empirically. We begin by defining an empirical counterpart to our measures of price impact as

$$\lambda_t = \log(q_{A,t_1}/q_{B,t_1}) - \log(q_{A,t_0}/q_{B,t_0})$$

where t_0 are the quantities prior to a swap and t_1 are quantities immediately after the swap. We are interested in the price impact of the swap of an LP, λ_t^{LP} , and how that compares to LT swaps prior to t .

In Table 3 we present summary statistics on the typical size of LP and LT price impacts in our Uniswap data. Note in this table we focus on individual rather than cumulative trades.

Table 3: Price Impact

	Mean	Standard Deviation	Min	Max
LP	0.002958	0.010522	-0.526232	0.326240
LT	0.001911	0.014523	-7.496103	6.977194

The table presents the average and standard deviation of the absolute value of price impact by LPs and LTs. The table also includes the smallest and largest price impacts by LPs and LTs. Data is pulled from all Uniswap v2 pools that were created prior to 2020-07-01 and have more than 100,000 transactions. There are 31 pools. The sample period is from 2020-07-01 to 2024-06-30.

Given our focus on transaction counts in Section 2 above, it is interesting to note that a typical swap by an LP is of a comparable size (by price impact) to a typical swap by an LT. We view this as supporting evidence that LP swap behavior is meaningful for overall price discovery at AMMs.

Additionally, our model suggests that the response of the LP to LT trade as measured by the correlation between λ_t^{LP} and λ_t^{LT} should vary with the extent of informed trading in the market. We pursue an approach to identify the extent of informed trading behavior in AMMs motivated by ideas in [Capponi, Jia and Yu \(2024\)](#) who use blockchain gas fees (transaction costs) paid on swaps to categorize informed and uninformed trades. The logic is that a LT with information that

is short-lived will pay a higher mining (gas) fee to increase the priority of her transaction and ensure it is added to the blockchain quickly.¹⁹ Building on this idea, we categorize a transaction by an LT as an informed trade if the gas fee is high. We define “high” as a gas fee paid in a specific pool that is in the top 25% decile over the prior (rolling) seven day window in that pool.

Using this categorization, we find that the average (absolute) price impact of an informed trade is roughly 0.0029 and of the same order of magnitude as an LP swap as reported in Table 3 while the average (absolute) price impact of an uninformed trade is slightly smaller and roughly 0.0016. Focusing on one specific pool, Table 4 shows the distribution of the size of trades in the largest pool in our sample, the WETH–USDT pool where traders may swap WETH for the Tether stable coin USDT.²⁰

Table 4: Price Impact in the WETH-USDT Pool.

Trader Type	5%	50%	95%	Min	Max
LP	-0.000361	0.000000	0.000535	-0.278843	0.273400
LT	-0.000330	-0.000001	0.000327	-0.232827	0.322526
Informed	-0.000737	-0.000002	0.000700	-0.232827	0.322526
Noise	-0.000241	-0.000001	0.000246	-0.230243	0.234302

The table presents distributional statistics of price impact by LPs and LTs. The table also shows distributional statistics of price impact by LT trades classified as Informed trades compared to those classified as uninformed trades. Trades are classified as informed if the gas fee associated with the swap transaction is in the top 25% quartile of gas fees paid for swaps over the prior (rolling) seven day window in that pool. Data is pulled from all Uniswap v2 pools that were created prior to 2020-07-01 and have more than 100,000 transactions. There are 31 pools. The sample period is from 2020-07-01 to 2024-06-30.

Table 7 shows similar results for all pools in our sample. Consistently, we observe the price impact of a typical swap by an LP is comparable to the price impact of a typical swap by an LT. While this result holds across most pools for a wide range of the distribution, from the 5th to the 95th percentile, for some pools we do observe that the LT swaps with the largest price impact in absolute value tend to be larger than those for LP swaps. We view this evidence as consistent with our assumption that a typical liquidity provider is “slow”—they wish to trade less often and would obtain less transactional priority should they trade—compared to the typical trader. Table 7 shows similar results for all pools in our sample.

¹⁹To expedite the processing of an Ethereum transaction, an LT can increase the gas fee, effectively offering a higher amount of ETH as an incentive for miners to include it in the next block. This prioritization mechanism ensures that, during periods of network congestion, transactions with higher fees are processed ahead of those with lower fees.

²⁰WETH represents “wrapped” ETH which is a smart-contract based representation of the native cryptocurrency of the Ethereum blockchain, ETH.

Consistently, we observe the price impact of a typical swap by an LP is comparable to the price impact of a typical swap by an LT. While this result holds across most pools for a wide range of the distribution, from the 5th to the 95th percentile, for some pools we do observe that the LT swaps with the largest price impact in absolute value tend to be larger than those for LP swaps. We view this evidence as consistent with our assumption that a typical liquidity provider is “slow”—they wish to trade less often and would obtain less transactional priority should they trade—compared to the typical trader.

We now use this evidence to understand if the dynamics of LP behavior in the data are consistent with our theory. In our data, we typically have multiple LT swaps in a row (recall from Table 1 that LT transactions are, by far, the most common). Instead of focusing on the price impact of a single trade by a liquidity taker, then, we define λ_t^{LT} as the cumulative price impact of LT trades after the previous LP swap and prior to t . In other words, the impact of all the LT swaps between two LP swaps. More precisely, recall that we have labeled all swap transactions as coming from a LP or a LT. Define $\mathcal{T}_{LT} = \{\tau \mid \text{the trade at time } \tau \text{ is by an LT}\}$ and $\mathcal{T}_{LP} = \{t \mid \text{the trade at time } t \text{ is by an LP}\}$. So, the date of the last swap by an LP prior to t is $\rho(t) = \max\{\tau \in \mathcal{T}_{LP} \mid \tau < t\}$. Hence, the cumulative price impact of all intervening LT trades prior to date t is

$$\lambda_t^{LT} = \sum_{\tau \in \mathcal{T}_{LT} : \rho(t) < \tau < t} \lambda_\tau.$$

Then the price impact λ_t^{LT} is informed if any of the intervening LT swaps were informed. This is captured in the indicator function I_t as:

$$I_t = \begin{cases} 1 & \text{if } \exists \tau \in \mathcal{T}_{LT} : \rho(t) < \tau < t \text{ and } \text{gas}_\tau \text{ is high} \\ 0 & \text{otherwise} \end{cases}$$

Table 5 reports the results of the following regression

$$\lambda_t^{LP} = b_0 + b_1 I_t + b_2 \lambda_t^{LT} + b_3 \lambda_t^{LT} I_t + \epsilon_t.$$

The focus of our analysis is the reaction of the LP to prior trades of the LT. For LT trades classified as noise trades, this reaction is b_2 and for LT trades classified as informed, the reaction is $b_2 + b_3$. The finding that b_2 is negative implies the LT actively adjusts marginal prices in a direction opposite that of the noise trader LT. The finding that $b_2 + b_3$ is approximately zero implies that when the intervening swaps are informed, the LP swap is, on average, directionless. (See table 5

Table 5: LP swap reaction to LT trade

Sample Period	Full	v2 era	v3 era	Full	v2 era	v3 era	Full	Full
Intercept b_0	0.00010 (0.00008)	0.00009 (0.00008)	0.00030 (0.00030)	–	–	–	–	–
Intercept Informed b_1	0.00010 (0.00008)	0.00010 (0.00009)	–0.00008 (0.0003)	0.00010 (0.00008)	0.00020 (0.00009)	–0.00030 (0.0003)	0.00010 (0.00009)	0.00006 (0.00008)
LT Price Impact b_2	–0.2284 (0.0170)	–0.2302 (0.017)	–0.1725 (0.096)	–0.2204 (0.017)	–0.2227 (0.017)	–0.1422 (0.080)	–0.2200 (0.017)	–0.2063 (0.018)
LT Price Impact $\times$ has_informed b_3	0.2264 (0.017)	0.2053 (0.019)	0.1721 (0.096)	0.2187 (0.010)	0.1985 (0.019)	0.1422 (0.080)	0.2184 (0.017)	0.2039 (0.018)
Pool fixed effect	No	No	No	Yes	Yes	Yes	Yes	No
Month fixed effect	No	No	No	No	No	No	Yes	No
Pool–Month fixed effect	No	No	No	No	No	No	No	Yes
# of observations	75,047	64,312	10,735	75,047	64,312	10,735	75,047	75,047
R ²	0.025	0.036	0.007	0.045	0.054	0.081	0.046	0.156

λ_t^{LP} is the price impact of LP swap at t . λ_t^{LT} is the price impact of all LT swaps at between the last LP swap and prior to t . I_t is an indicator that is one if the LT swap is identified to be informed. LT swaps are tagged as informed if their gas fee (transaction cost) is high (top 25% decile of past seven days, by pool). Standard Errors (in parentheses are heteroscedasticity-robust (HC2). A transaction is classified as belonging to an LP if: (a) it is a mint or burn; (b) any of the addresses involved in the transaction have a positive balance of that pool’s tokens at the time of the transaction; or c the swap transaction is paired with a mint transaction. Each transaction can involve several addresses (both wallets and contracts). A transaction is paired if all the addresses on both transactions match and the transactions both occur within a three-minute interval. Data is pulled from all Uniswap v2 pools that were created prior to 2020-07-01 and have more than 100,000 transactions. There are 31 pools. The sample period is from 2020-07-01 to 2024-06-30. The total number of transactions is 19.2 million. A transaction is defined as a unique call to a Uniswap pool contract as a swap, mint, or burn, and involves multiple addresses (wallets and contracts) and token transfers.

for results using the full sample as well as on partial samples that include the periods before and after the introduction of Uniswap v3.) Additionally, we run the regression with fixed effects for pools, months, and the pool-month combination. The main result that $b_2 < 0$ and $b_2 + b_3 \approx 0$ is consistent across all these specifications.

These findings on the dynamic behavior of liquidity providers in Uniswap v2 are consistent with our theoretical model’s predictions for this same behavior. The fact that there is systematic variation in how liquidity providers actively set prices on AMMs suggests the importance of better understanding the role liquidity providers in AMMs play in aiding price discovery for cryptocurrencies.

6 Conclusion

In this paper, we have explored how the innovative market structure featured in most decentralized exchanges transmits new information into prices and how prices respond when traders

demand liquidity. Our empirical approach is facilitated by the stark, limited message space for Uniswap v2. We can identify (define) LPs directly by their mint and burn transactions that supply or remove tokens to the pool. These transactions are constrained by the code to not change marginal token prices. We can also track LPs ownership tokens in the pool (the result of a mint). We exploit this feature of decentralized exchanges to document LPs' active price setting behavior.

Our model also guides our empirical measure of price impact. The dynamic strategy of an LP in our model features a cone of inaction and so LPs do not necessarily re-balance their position after each LT trade. When they do find it optimal to re-balance it is typically in the opposite direction as the LT trade. The direction and size of these LT trades by the LP is central to evaluating price impact in these markets. Importantly, the LP is more likely to re-balance her position when the prior trades were from uninformed LTs. We are able to confirm this in the data when we proxy for the informedness of LTs by their willingness to incur a higher transaction (gas) fee (a measure suggested [Capponi, Jia and Yu \(2024\)](#)).

DeFi markets and AMMs continue to evolve. Uniswap v3, released in May 2021, gives LPs the control of setting a range for allowable transactions. In v2, once liquidity is posted, all future trades follow constant-product curve across the entire domain. In v3, LPs instead choose bounds, so that the constant-product rule applies only inside that interval. The same tradeoffs we explore in our model still apply in this new setting since the marginal tradeoffs within the active range follow the same mechanics. What is new and interesting is how the ability to choose those bounds introduces an extensive-margin decision.

Our model also abstracts from execution risk such as sandwich attacks that impact LT profits and, presumably, behavior. If LTs anticipate higher execution costs, their demand to trade may fall or become more selective. That in turn alters the flow of fees to LPs and could change the incentives to post liquidity. This channel raises a further question about protocol design. In March 2023 Uniswap changed its default slippage tolerance from a static value to a dynamic setting. The goal of this policy was to balance the risk of failed trades against vulnerability to manipulation. Exploring how such design choices affect both sides of the market, the willingness of LTs to trade and the willingness of LPs to provide liquidity, would be a natural extension of our framework.

References

- Angeris, Guillermo, and Tarun Chitra. 2020. "Improved Price Oracles." ACM. 5
- Angeris, Guillermo, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra. 2021. "An analysis of Uniswap markets." *arXiv* 1911.03380. 5
- Aoyagi, Jun. 2022. "Liquidity Provision by Automated Market Makers." 5
- Aquilina, Matteo, Eric Budish, and Peter O’neill. 2022. "Quantifying the high-frequency trading "arms race"." *The Quarterly Journal of Economics*, 137(1): 493–564. 4
- Aquilina, Matteo, Sean Foley, Leonardo Gambacorta, and William Krekel. 2024. "Decentralised dealers? Examining liquidity provision in decentralised exchanges." Bank for International Settlements BIS Working Papers 1227. 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. 7
- Brugler, James, and Terrence Hendershott. 2023. "Measuring public and private information using quote and trade races." *Available at SSRN* 4674014. 4
- Cao, David, Leonid Kogan, Gerry Tsoukalas, and Brett Hemenway Falk. 2023. "A Structural Model of Automated Market Making." <https://ssrn.com/abstract=4591447>, The Wharton School Research Paper, SSRN Working Paper No. 4591447. 6
- Capponi, Agostino, and Ruizhe Jia. 2021. "Liquidity Provision on Blockchain-based Decentralized Exchanges." *Review of Financial Studies*. Forthcoming. Available at <https://ssrn.com/abstract=3805095> or <http://dx.doi.org/10.2139/ssrn.3805095>. 6
- Capponi, Agostino, Ruizhe Jia, and Shihao Yu. 2024. "Price Discovery on Decentralized Exchanges." <https://ssrn.com/abstract=4236993>, Available at SSRN: <http://dx.doi.org/10.2139/ssrn.4236993>. 4, 6, 36, 40
- Circle. 2023. "Circle Transparency & Stability." *url = https://www.circle.com/en/transparency*, Accessed on Sep 18th, 2023. 13
- 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. 2, 3, 14, 15
- Goyal, Mohak, Geoffrey Ramseyer, Ashish Goel, and David MaziÅšres.** 2023. "Finding the Right Curve: Optimal Design of Constant Function Market Makers." 7
- 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>. 6
- Kyle, Albert S.** 1985. "Continuous Auctions and Insider Trading." *Econometrica*, 53(6): 1315–1335. 2
- Lehar, Alfred, and Christine A. Parlour.** 2023. "Decentralized Exchange: The Uniswap Automated Market Maker." Working Paper. 5, 9
- Lehar, Alfred, Christine Parlour, and Marius Zoican.** 2023. "Liquidity fragmentation on decentralized exchanges." *arXiv preprint arXiv:2307.13772*. 6
- Malinova, Katya, and Andreas Park.** 2024. "Learning from DeFi: Would Automated Market Makers Improve Equity Trading?" Available at SSRN 4531670. 9
- Milionis, Jason, Ciamac C. Moallemi, and Tim Roughgarden.** 2023a. "Automated Market Making and Arbitrage Profits in the Presence of Fees." 6
- Milionis, Jason, Ciamac C. Moallemi, and Tim Roughgarden.** 2023b. "A Myersonian Framework for Optimal Liquidity Provision in Automated Market Makers." *arXiv 2303.00208*. 4, 7, 27, 29, 48
- 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. 6, 23
- 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>. 7

A Proof of Optimal Liquidity Provision

The LP's optimal deposit problem is

$$\begin{aligned} \max_{e_A, e_B} & \pi(\mu_A \mathbb{E}[x_A - e_A] + \mu_B \mathbb{E}[x_B - e_B]) + \\ & (1 - \pi)(\mathbb{E}v_A[x_A - e_A] + \mu_B \mathbb{E}v_B[x_B - e_B]) \\ \text{s.t. } & \text{(5)-(6),} \\ & 0 \leq e_i \leq E_i, \quad \forall i \end{aligned}$$

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

$$x_A = \sqrt{\frac{v_B}{v_A}} e_A e_B, \quad x_B = \sqrt{\frac{v_A}{v_B}} e_A e_B$$

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

$$\begin{aligned} \mu_A(x_A - e_A) &= \sqrt{\frac{v_B/\mu_B}{v_A/\mu_A}} \mu_A \mu_B e_A e_B - \mu_A e_A \\ \mu_B(x_B - e_B) &= \sqrt{\frac{v_A/\mu_A}{v_B/\mu_B}} \mu_A \mu_B e_A e_B - \mu_B e_B \end{aligned}$$

for uninformed trades and

$$\begin{aligned} v_A(x_A - e_A) &= \sqrt{\frac{v_A v_B}{\mu_A \mu_B}} \mu_A \mu_B e_A e_B - v_A e_A \\ v_B(x_B - e_B) &= \sqrt{\frac{v_A v_B}{\mu_A \mu_B}} \mu_A \mu_B e_A e_B - v_B e_B \end{aligned}$$

for informed trades.

Denote $\omega = \sqrt{\frac{v_A/\mu_A}{v_B/\mu_B}}$ and $\psi = \sqrt{\frac{v_A v_B}{\mu_A \mu_B}}$. With assumptions that $\mathbb{E}v_i = \mu_i$, LP's optimal deposit problem becomes

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

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

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

where η_i is the Lagrangian multiplier for $0 \leq e_i$ and ξ_i is the Lagrangian multiplier for $e_i \leq E_i$.

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

If $\Pi > 1$, the solution is always at the corner, i.e. at least one of the $\xi_i > 0$. To see this, consider the interior cases where $\eta_i = 0$ and $\xi_i = 0$. 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 ξ_i to be positive for the FOCs to hold.

Therefore, we have the following optimal deposit of the LP

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

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 $\frac{d^2 y}{dx^2} = (1 - \tau)^2 G''((1 - \tau)x + \tau e_A)$. 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:

$$\begin{aligned} \max_{e_A, e_B} \quad & v_A(e_A - x) + v_B(e_B - y) \\ \text{s.t.} \quad & (1 - \tau)y + \tau e_B = G((1 - \tau)x + \tau e_A) \end{aligned}$$

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

$$(1 - \tau)y + \tau e_B = G((1 - \tau)x + \tau e_A)$$

$$G'((1 - \tau)x + \tau e_A) = -\frac{v_A}{v_B}$$

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 $\frac{v_A}{v_B}$, the ex post portfolios satisfy

$$(1 - \tau)x_\tau + \tau e_A = x_0$$

$$(1 - \tau)y_\tau + \tau e_B = y_0$$

which can be written as

$$x_\tau - e_A = \frac{1}{1 - \tau}(x_0 - e_A)$$

$$y_\tau - e_B = \frac{1}{1 - \tau}(y_0 - e_B)$$

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

$$R_\tau = \mathbb{E}[(\pi\mu_A + (1 - \pi)v_A)(x_\tau - e_A) + (\pi\mu_B + (1 - \pi)v_B)(y_\tau - e_B)]$$

$$= \frac{1}{1 - \tau} \mathbb{E}[(\pi\mu_A + (1 - \pi)v_A)(x_0 - e_A) + (\pi\mu_B + (1 - \pi)v_B)(y_0 - e_B)]$$

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 $v_B = \mu_B = 1$ and $h_B(v_B) = 1$ if $v_B = 1$. Denote $r_A = v_A/\mu_A$. 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 23. The region of belief where a trade happens with price p_h is when $r_A \geq p_h$. 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 21. Similarly, the region of belief where a trade happens with price p_h is when $r_A \geq \frac{1}{c^2}$. From the LP's perspective, the trading volume in this region is $-e_A$ for token A and $\frac{1}{c} e_B$ 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 $c \geq \frac{\pi}{1-\pi}$. 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 (v_A, v_B) such that the LT is willing to participate (Individual Rational) and truthfully report the values (Incentive Compatible).

Assume the LT's value (v_A, v_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:

$$\begin{aligned} \max_{x,y} \mathbb{E}_{\{v_A, v_B\}} [& -(\pi\mu_A + (1 - \pi)v_A)t_A(v_A, v_B) - (\pi\mu_B + (1 - \pi)v_B)t_B(v_A, v_B)] \\ \text{s.t. } & v_A t_A(v_A, v_B) + v_B t_B(v_A, v_B) \geq v'_A t_A(v'_A, v'_B) + v'_B t_B(v'_A, v'_B) \\ & v_A t_A(v_A, v_B) + v_B t_B(v_A, v_B) \geq 0 \\ & t_A(v_A, v_B) \leq e_A, -l_B \leq t_B(v_A, v_B) \leq e_B \end{aligned}$$

Since only $p = \frac{v_B}{v_A} \frac{e_B}{e_A}$ matters in the constraints, the problem can be written as

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

Under one of the two conditions, i.e. $\pi = 0$ or $v_A = \mu_A$ for sure, we know $\pi\mu_A + (1-\pi)v_A$ 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 $\frac{(\pi\mu_B + (1-\pi)v_B)}{(\pi\mu_A + (1-\pi)v_A)} \frac{e_B}{e_A} = \pi(p_0, p)$. The problem then has the same expression as [Milionis, Moallemi and Roughgarden \(2023b\)](#).

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

E Dimension Reduction of the Dynamic Model

E.1 Proof of Total Value Equivalence

We see that $E_{A,t}, E_{B,t}, \mu_{A,t}, \mu_{B,t}$ only show up in the objective function (7) as products $\mu_{A,t}E_{A,t}$ and $\mu_{B,t}E_{B,t}$. We need to show that this is also the case for the constraints.

Remember that $(r_{A,t}, r_{B,t}) = \left(\frac{v_{A,t}}{\mu_{A,t-1}}, \frac{v_{B,t}}{\mu_{A,t-1}} \right)$ are the rate of belief change, $X_{i,t} = \mu_{i,t}E_{i,t}$ are the total value of each token LP has at the end of time t and $Y_{i,t} = \mu_{i,t-1}e_{i,t}$ are the total value of each token LP deposits into the pool at the beginning of time t . We also assume that $(r_{A,t}, r_{B,t})$ follows the distribution G_t , which is independent of the current belief of valuations $(\mu_{A,t-1}, \mu_{B,t-1})$. Then

we have the following mappings:

$$\begin{aligned} E_{i,t} &= \frac{X_{i,t}}{\mu_{i,t}} \\ \mu_{i,t} &= \mu_{i,0} \prod_{s=1}^t r_{i,s} \\ e_{i,t} &= \frac{Y_{i,t}}{\mu_{i,t-1}} \end{aligned}$$

First we can pin down $(x_{A,t}, x_{B,t})$ —the pool position after trade at time t —using Constant Product and LT's optimality:

$$\begin{pmatrix} x_{A,t} \\ x_{B,t} \end{pmatrix} = \begin{pmatrix} \sqrt{\frac{v_{B,t}}{v_{A,t}}} e_{A,t} e_{B,t} \\ \sqrt{\frac{v_{A,t}}{v_{B,t}}} e_{A,t} e_{B,t} \end{pmatrix} = \begin{pmatrix} \frac{1}{\mu_{A,t-1}} \sqrt{\frac{v_{B,t}/\mu_{B,t-1}}{v_{A,t}/\mu_{A,t-1}}} \sqrt{\mu_{A,t-1} e_{A,t} \mu_{B,t-1} e_{B,t}} \\ \frac{1}{\mu_{B,t-1}} \sqrt{\frac{v_{A,t}/\mu_{A,t-1}}{v_{B,t}/\mu_{B,t-1}}} \sqrt{\mu_{A,t-1} e_{A,t} \mu_{B,t-1} e_{B,t}} \end{pmatrix}$$

Combining with Accounting, we have the post-trade token in the hand of LP as

$$\begin{pmatrix} E_{A,t} \\ E_{B,t} \end{pmatrix} = \begin{pmatrix} \sqrt{\frac{v_{B,t}}{v_{A,t}}} e_{A,t} e_{B,t} + (E_{A,t-1} - e_{A,t}) \\ \sqrt{\frac{v_{A,t}}{v_{B,t}}} e_{A,t} e_{B,t} + (E_{B,t-1} - e_{B,t}) \end{pmatrix}$$

Similar to the static problem, we can write the post-trade values of each token in the hand of LP as

$$\begin{pmatrix} \mu_{A,t} E_{A,t} \\ \mu_{B,t} E_{B,t} \end{pmatrix} = \begin{pmatrix} \sqrt{\frac{v_{B,t}/\mu_{B,t-1}}{v_{A,t}/\mu_{A,t-1}}} \sqrt{\mu_{A,t-1} e_{A,t} \mu_{B,t-1} e_{B,t}} + \mu_{A,t-1} (E_{A,t-1} - e_{A,t}) \\ \sqrt{\frac{v_{A,t}/\mu_{A,t-1}}{v_{B,t}/\mu_{B,t-1}}} \sqrt{\mu_{A,t-1} e_{A,t} \mu_{B,t-1} e_{B,t}} + \mu_{B,t-1} (E_{B,t-1} - e_{B,t}) \end{pmatrix}$$

for uninformed trades, and

$$\begin{pmatrix} \mu_{A,t} E_{A,t} \\ \mu_{B,t} E_{B,t} \end{pmatrix} = \begin{pmatrix} \sqrt{\frac{v_{A,t}}{\mu_{A,t-1}} \frac{v_{B,t}}{\mu_{B,t-1}}} \sqrt{\mu_{A,t-1} e_{A,t} \mu_{B,t-1} e_{B,t}} + \frac{v_{A,t}}{\mu_{A,t-1}} \mu_{A,t-1} (E_{A,t-1} - e_{A,t}) \\ \sqrt{\frac{v_{A,t}}{\mu_{A,t-1}} \frac{v_{B,t}}{\mu_{B,t-1}}} \sqrt{\mu_{A,t-1} e_{A,t} \mu_{B,t-1} e_{B,t}} + \frac{v_{B,t}}{\mu_{B,t-1}} \mu_{B,t-1} (E_{B,t-1} - e_{B,t}) \end{pmatrix}$$

for informed trades.

Therefore, with the assumption that $(r_{A,t}, r_{B,t}) = \left(\frac{v_{A,t}}{\mu_{A,t-1}}, \frac{v_{B,t}}{\mu_{B,t-1}} \right)$ is independent of the state variables $E_{A,t-1}, E_{B,t-1}, \mu_{A,t-1}, \mu_{B,t-1}$, we can use a notation similar to that of the static model, $\omega_t = \sqrt{\frac{r_{A,t}}{r_{B,t}}}$ and $\psi_t = \sqrt{r_{A,t} r_{B,t}}$. And the constraints become:

$$\begin{pmatrix} \mu_{A,t} E_{A,t} \\ \mu_{B,t} E_{B,t} \end{pmatrix} = \begin{pmatrix} \frac{1}{\omega_t} \sqrt{\mu_{A,t-1} e_{A,t} \mu_{B,t-1} e_{B,t}} + (\mu_{A,t-1} E_{A,t-1} - \mu_{A,t-1} e_{A,t}) \\ \omega_t \sqrt{\mu_{A,t-1} e_{A,t} \mu_{B,t-1} e_{B,t}} + (\mu_{B,t-1} E_{B,t-1} - \mu_{B,t-1} e_{B,t}) \end{pmatrix}$$

for uninformed trades, and

$$\begin{pmatrix} \mu_{A,t} E_{A,t} \\ \mu_{B,t} E_{B,t} \end{pmatrix} = \begin{pmatrix} \psi_t \sqrt{\mu_{A,t-1} e_{A,t} \mu_{B,t-1} e_{B,t}} + r_{A,t} (\mu_{A,t-1} E_{A,t-1} - \mu_{A,t-1} e_{A,t}) \\ \psi_t \sqrt{\mu_{A,t-1} e_{A,t} \mu_{B,t-1} e_{B,t}} + r_{B,t} (\mu_{B,t-1} E_{B,t-1} - \mu_{B,t-1} e_{B,t}) \end{pmatrix}$$

for informed trades.

Now we have that $E_{A,t}, E_{B,t}, \mu_{A,t}, \mu_{B,t}$ only show up in both the objective function and the constraints as products $\mu_{A,t} E_{A,t}$ and $\mu_{B,t} E_{B,t}$. By changing the sequence of variables $\{E_{A,t}, E_{B,t}, \mu_{A,t}, \mu_{B,t}, e_{A,t}, e_{B,t}\}_{t=1}^{\infty}$ to $\{X_{A,t}, X_{B,t}, r_{A,t}, r_{B,t}, Y_{A,t}, Y_{B,t}\}_{t=1}^{\infty}$ as we define above, the problem can be written as:

$$\max_{\{Y_{A,t}\}_{t=1}^{\infty}, \{Y_{B,t}\}_{t=1}^{\infty}} \sum_{t=0}^{\infty} \delta^t \beta (1 - \beta)^t \mathbb{E} [X_{A,t} + X_{B,t}]$$

where for $t = 1, 2, 3, \dots$

$$\begin{pmatrix} X_{A,t} \\ X_{B,t} \end{pmatrix} = \begin{cases} \begin{pmatrix} \frac{1}{\omega_t} \sqrt{Y_{A,t} Y_{B,t}} + (X_{A,t-1} - Y_{A,t}) \\ \omega_t \sqrt{Y_{A,t} Y_{B,t}} + (X_{B,t-1} - Y_{B,t}) \end{pmatrix} & \text{with prob } \pi \\ \begin{pmatrix} \psi_t \sqrt{Y_{A,t} Y_{B,t}} + r_{A,t} (X_{A,t-1} - Y_{A,t}) \\ \psi_t \sqrt{Y_{A,t} Y_{B,t}} + r_{B,t} (X_{B,t-1} - Y_{B,t}) \end{pmatrix} & \text{with prob } 1 - \pi \end{cases}$$

$$0 \leq Y_{i,t} \leq X_{i,t-1}$$

given $\mu_{A,0}, X_{A,0}, \mu_{B,0}, X_{B,0}$.

E.2 Proof of Constant Returns to Scale

We need to show that for any $X_{A,0}^k = kX_{A,0}$ and $X_{B,0}^k = kX_{B,0}$, it must be $V(X_{A,0}^k, X_{B,0}^k) = V(kX_{A,0}, kX_{B,0}) = kV(X_{A,0}, X_{B,0})$ for any $k > 0$.

To do that, we can show that for any realization of $\{r_{A,t}, r_{B,t}\}_{t=1}^{\infty}$, we have $\{X_{A,t}^k, X_{B,t}^k\}_{t=1}^{\infty} = \{kX_{A,t}, kX_{B,t}\}_{t=1}^{\infty}$, then the objective function implies the above statement directly.

Notice that the assumption of $\{r_{A,t}, r_{B,t}\}_{t=1}^{\infty}$ is an independent process of the pool position and LP's move is needed here.

Let $\{Y_{A,t}^*, Y_{B,t}^*\}_{t=1}^{\infty}$ be the optimal deposits for $X_{A,0}, X_{B,0}, \{r_{A,t}, r_{B,t}\}_{t=1}^{\infty}$ and together they induce $\{X_{A,t}, X_{B,t}\}_{t=1}^{\infty}$.

We can do this by two steps.

1. The first step is to show that $\{kY_{A,t}^*, kY_{B,t}^*\}_{t=1}^{\infty}$ is a feasible sequence of deposit for $X_{A,0}^k, X_{B,0}^k$, and they induce $\{kX_{A,t}, kX_{B,t}\}_{t=1}^{\infty}$, which implies that $V(X_{A,0}^k, X_{B,0}^k) \geq kV(X_{A,0}, X_{B,0})$;
2. The second step is to show that there's no other deposit for $X_{A,0}^k, X_{B,0}^k$ that can achieve higher value than $kY_{A,t}^*, kY_{B,t}^*$ induce, i.e. $V(X_{A,0}^k, X_{B,0}^k) \leq kV(X_{A,0}, X_{B,0})$.

Let us first show that $V(X_{A,0}^k, X_{B,0}^k) \geq kV(X_{A,0}, X_{B,0})$.

Given $\{r_{A,t}, r_{B,t}\}_{t=1}^{\infty}$, it is easy to see that $kY_{A,t}^*, kY_{B,t}^*$ is feasible given $(X_{A,t-1}^k, X_{B,t-1}^k) = (kX_{A,t-1}, kX_{B,t-1})$.

$$0 \leq Y_{i,t}^* \leq X_{i,t-1} \Rightarrow 0 \leq kY_{i,t}^* \leq kX_{i,t-1}$$

then it implies for next period

$$\begin{aligned} \begin{pmatrix} X_{A,t}^k \\ X_{B,t}^k \end{pmatrix} &= \begin{cases} \begin{pmatrix} \sqrt{\frac{r_{B,t}}{r_{A,t}}} kY_{A,t}^* kY_{B,t}^* + (kX_{A,t-1} - kY_{A,t}^*) \\ \sqrt{\frac{r_{A,t}}{r_{B,t}}} kY_{A,t}^* kY_{B,t}^* + (kX_{B,t-1} - kY_{B,t}^*) \end{pmatrix} & \text{with prob } \pi \\ \begin{pmatrix} \sqrt{r_{A,t} r_{B,t}} kY_{A,t}^* kY_{B,t}^* + r_{A,t} (kX_{A,t-1} - kY_{A,t}^*) \\ \sqrt{r_{A,t} r_{B,t}} kY_{A,t}^* kY_{B,t}^* + r_{B,t} (kX_{B,t-1} - kY_{B,t}^*) \end{pmatrix} & \text{with prob } 1 - \pi \end{cases} \\ &= \begin{pmatrix} kX_{A,t} \\ kX_{B,t} \end{pmatrix} \end{aligned}$$

Therefore, we construct a feasible sequence $\{X_{A,t}^k, X_{B,t}^k\}_{t=1}^{\infty} = \{kX_{A,t}, kX_{B,t}\}_{t=1}^{\infty}$ by depositing $\{kY_{A,t}^*, kY_{B,t}^*\}_{t=1}^{\infty}$ for any realization $\{r_{A,t}, r_{B,t}\}_{t=1}^{\infty}$. We don't know if this is optimal for $X_{A,0}^k, X_{B,0}^k$. But at least it implies that $V(X_{A,0}^k, X_{B,0}^k) \geq kV(X_{A,0}, X_{B,0})$.

Next we show that $V(X_{A,0}^k, X_{B,0}^k) \leq kV(X_{A,0}, X_{B,0})$. In other words, no deposit can achieve higher value.

Suppose for some realization $\{r_{A,t}, r_{B,t}\}_{t=1}^{\infty}$, there exist $\{Y_{A,t}^{*k}, Y_{B,t}^{*k}\}_{t=1}^{\infty}$, such that the corresponding $X_{i,t}^{*k}$ yields

$$\sum_{t=0}^{\infty} \delta^t \beta (1 - \beta)^t (X_{A,t}^{*k} + X_{B,t}^{*k}) > \sum_{t=0}^{\infty} \delta^t \beta (1 - \beta)^t (kX_{A,t} + kX_{B,t})$$

Then for $(X_{A,0}, X_{B,0})$, we can use deposit $\frac{1}{k}Y_{i,t}^{*k}$, which by the same logic as in step one, is feasible and yields $\left(\frac{1}{k}X_{A,t}^{*k}, \frac{1}{k}X_{B,t}^{*k}\right)$. And it gives

$$\sum_{t=0}^{\infty} \delta^t \beta (1 - \beta)^t \left(\frac{1}{k}X_{A,t}^{*k} + \frac{1}{k}X_{B,t}^{*k} \right) > \sum_{t=0}^{\infty} \delta^t \beta (1 - \beta)^t (X_{A,t} + X_{B,t})$$

which contradicts the definition that $\left\{Y_{A,t}^*, Y_{B,t}^*\right\}_{t=1}^{\infty}$ is optimal.

Therefore, together we have $V(kX_{A,0}, kX_{B,0}) = kV(X_{A,0}, X_{B,0})$.

F Additional Tables and Figures


Table 6: LP Transactions - Paired

	burn	mint	swap
Full Sample: 2020-07-01 - 2024-06-30			
Naked	149,855	169,015	81,686
Paired	56	27,113	27,116
Total	149,911	196,128	108,802
v2 Dominant Contract: 2020-07-01 - 2021-05-31			
Naked	112,719	140,452	63,789
Paired	43	16,911	16,907
Total	112,762	157,363	80,696
v3 Dominant Contract: 2021-06-01 - 2024-06-30			
Naked	37,136	28,563	17,897
Paired	13	10,202	10,209
Total	37,149	38,765	28,106


A transaction is classified as belonging to an LP if: (a) it is a mint or burn; (b) any of the addresses involved in the transaction have a positive balance of that pool's tokens at the time of the transaction; or (c) the swap transaction is paired with a mint transaction. Each transaction can involve several addresses (both wallets and contracts). A transaction is paired if all the addresses on both transactions match and the transactions both occur within a three-minute interval. Data is pulled from all Uniswap v2 pools that were created prior to 2020-07-01 and have more than 100,000 transactions. There are 31 pools. The sample period is from 2020-07-01 to 2024-06-30. The total number of transactions is 19.2 million. A transaction is defined as a unique call to a Uniswap pool contract as a swap, mint, or burn, and involves multiple addresses (wallets and contracts) and token transfers.

Figure 7: Swap Trades by LPs

(a) Swaps




(b) Swaps Paired with Mints




Panel (a) shows the proportion of LP transactions that are swaps. Panel (b) shows the proportion of mint transactions that are paired with a swap transaction. A transaction is classified as belonging to an LP if: (a) it is a mint or burn; (b) any of the addresses involved in the transaction have a positive balance of that pool's tokens at the time of the transaction; or c the swap transaction is paired with a mint transaction. Each transaction can involve several addresses (both wallets and contracts). A transaction is paired if all the addresses on both transactions match and the transactions both occur within a three-minute interval. Data is pulled from all Uniswap v2 pools that were created prior to 2020-07-01 and have more than 100,000 transactions. This subsample is the period where v2 was the dominant contract. There are 31 pools. The sample period is from 2020-07-01 to 2021-05-31. The total number of transactions is 9.5 million. A transaction is defined as a unique call to a Uniswap pool contract as a swap, mint, or burn, and involves multiple addresses (wallets and contracts) and token transfers.

Figure 8: Share of Paired Transactions




The proportion of mint transactions that are paired with a swap transaction by month. A transaction is classified as belonging to an LP if: (a) it is a mint or burn; (b) any of the addresses involved in the transaction have a positive balance of that pool's tokens at the time of the transaction; or c the swap transaction is paired with a mint transaction. Each transaction can involve several addresses (both wallets and contracts). A transaction is paired if all the addresses on both transactions match and the transactions both occur within a three-minute interval. Data is pulled from all Uniswap v2 pools that were created prior to 2020-07-01 and have more than 100,000 transactions. There are 31 pools. The sample period is from 2020-07-01 to 2024-06-30. The total number of transactions is 19.2 million. A transaction is defined as a unique call to a Uniswap pool contract as a swap, mint, or burn, and involves multiple addresses (wallets and contracts) and token transfers.

Figure 9: π against variance of v_A




The break even proportion of uninformed trade for liquidity provision against the variance of the value change of token A, given token B is a stable coin. It shows that the break even level goes down with the variance, which suggests liquidity provision becomes more profitable as liquidity traders' beliefs become more disperse.

Figure 10: Optimal Policy Functions



The figure displays the LP's optimal deposit policy functions for various values of her state, represented as the ratio of the value of her endowment of token A to the value of her endowment of token B at the start of the period (or its inverse). Each panel displays the fraction of her endowment of tokens A and B that she deposits at the AMM as a function of this state variable. These policy functions were obtained for a numerical simulation. For detailed choices of parameters, see Footnote 17.

Figure 11: Inactivity of Liquidity Providers



The figure displays the percentage of all swaps conducted by LPs in our numerical simulation as a function of the extent of uninformed trading, π . When this percentage is smaller than 0.5, it reflects periods in the simulation where the LP does not adjust her portfolio following previous trade by the LT. In this numerical simulation, token B is set as a stable coin and the values of token A follow a truncated normal distribution. For detailed choices of parameters, see 17.

Table 7: Price Impact of LPs and LTs by Pool.

Pool	Type	5%	50%	95%	Min	Max
CEL-WETH	LP	-0.019008	-0.000088	0.019823	-0.124251	0.101654
	LT	-0.025152	-0.000095	0.025556	-1.222129	2.320158
	Informed	-0.034065	-0.000158	0.033186	-0.895435	2.320158
	Noise	-0.021682	-0.000083	0.022608	-1.222129	0.663037
COMP-WETH	LP	-0.005833	0.000174	0.022375	-0.076514	0.115606
	LT	-0.013155	0.000013	0.013184	-7.496103	2.783089
	Informed	-0.018981	-0.000000	0.018419	-4.216026	1.102276
	Noise	-0.010440	0.000018	0.010974	-7.496103	2.783089
DAI-MKR	LP	0.022853	0.024274	0.029212	0.022695	0.029761
	LT	-0.002773	-0.000147	0.002229	-0.684456	4.471649
	Informed	-0.003569	-0.000517	0.002788	-0.684456	0.499704
	Noise	-0.002496	-0.000069	0.002021	-0.277409	4.471649
DAI-USDC	LP	-0.002194	0.000003	0.006046	-0.007676	0.011746
	LT	-0.002277	0.000001	0.002241	-0.178995	0.221806
	Informed	-0.003420	0.000001	0.003305	-0.176501	0.221806
	Noise	-0.001842	0.000001	0.001835	-0.178995	0.135151
DAI-USDT	LP	-0.004179	-0.000004	0.003937	-0.011912	0.007690
	LT	-0.003262	0.000000	0.003272	-5.489722	1.176629
	Informed	-0.004558	-0.000000	0.004693	-1.560066	1.176629
	Noise	-0.002826	0.000000	0.002788	-5.489722	0.512787
DAI-WETH	LP	-0.001208	-0.000000	0.001294	-0.141629	0.144807
	LT	-0.001548	-0.000000	0.001594	-1.054344	1.545626
	Informed	-0.002662	-0.000000	0.002718	-0.682491	0.419199
	Noise	-0.001064	-0.000000	0.001103	-1.054344	1.545626
HEX-USDC	LP	-0.006601	0.001632	0.024926	-0.060493	0.140660
	LT	-0.016694	-0.000261	0.018381	-6.502482	6.977194
	Informed	-0.020806	-0.000296	0.027118	-6.498733	6.977194
	Noise	-0.015019	-0.000250	0.015279	-6.502482	2.701488
HEX-WETH	LP	-0.003425	-0.000030	0.005687	-0.055129	0.060230
	LT	-0.009227	-0.000206	0.010673	-3.839029	4.952318
	Informed	-0.013632	-0.000252	0.019413	-2.526518	4.952318
	Noise	-0.007484	-0.000195	0.008219	-3.839029	2.466702
LINK-WETH	LP	-0.001386	0.000022	0.001994	-0.051710	0.058549

Table 7 – continued from previous page

Pool	Type	5%	50%	95%	Min	Max
	LT	-0.004160	0.000002	0.004120	-2.645252	1.573756
	Informed	-0.005813	-0.000001	0.005716	-0.756233	0.891011
	Noise	-0.003555	0.000002	0.003497	-2.645252	1.573756
LRC-WETH	LP	-0.003021	0.000096	0.017321	-0.057972	0.057807
	LT	-0.023721	-0.000006	0.023445	-2.417452	1.441687
	Informed	-0.029764	0.000258	0.032368	-2.417452	1.441687
	Noise	-0.021543	-0.000015	0.019676	-2.345306	1.228012
MANA-WETH	LP	-0.018598	0.003434	0.032088	-0.053487	0.080270
	LT	-0.020260	-0.000000	0.020070	-2.960467	1.719663
	Informed	-0.028621	-0.000014	0.030003	-2.960467	1.719663
	Noise	-0.016974	-0.000000	0.016343	-2.687419	1.711054
MATIC-WETH	LP	-0.017385	-0.000073	0.012818	-0.099014	0.084713
	LT	-0.009910	-0.000025	0.010137	-3.689463	3.431068
	Informed	-0.015221	-0.000027	0.016651	-2.018141	3.431068
	Noise	-0.007794	-0.000025	0.007464	-3.689463	1.711428
MKR-WETH	LP	-0.006637	-0.000598	0.049773	-0.023171	0.080592
	LT	-0.007169	0.000000	0.006948	-0.431381	0.454065
	Informed	-0.009965	-0.000000	0.009685	-0.302546	0.454065
	Noise	-0.006113	0.000000	0.005965	-0.431381	0.290323
PAXG-WETH	LP	-0.022055	-0.000284	0.021186	-0.059943	0.059336
	LT	-0.009833	-0.000009	0.010149	-0.372182	0.386372
	Informed	-0.013534	-0.000016	0.015665	-0.312674	0.320812
	Noise	-0.008736	-0.000008	0.008392	-0.372182	0.386372
QNT-WETH	LP	-0.019505	0.001533	0.028131	-0.122191	0.218000
	LT	-0.028062	-0.000339	0.029270	-3.155449	4.539785
	Informed	-0.036108	-0.000500	0.039161	-2.856490	4.539785
	Noise	-0.024025	-0.000308	0.025175	-3.155449	1.209371
REN-WETH	LP	-0.011550	-0.000028	0.009614	-0.043950	0.026732
	LT	-0.015423	-0.000038	0.015180	-0.791739	0.834232
	Informed	-0.020305	-0.000120	0.020391	-0.532673	0.834232
	Noise	-0.013353	-0.000031	0.013068	-0.791739	0.765245
RNDR-WETH	LP	-0.045083	0.036478	0.157010	-0.193828	0.299502
	LT	-0.056394	-0.000467	0.057643	-4.619541	2.397623
	Informed	-0.072629	-0.000059	0.079192	-3.097615	1.758754

Table 7 – continued from previous page

Pool	Type	5%	50%	95%	Min	Max
	Noise	-0.049457	-0.000547	0.048709	-4.619541	2.397623
RSR-WETH	LP	-0.005970	-0.000041	0.008135	-0.024516	0.096186
	LT	-0.013662	-0.000151	0.014059	-0.505260	0.533010
	Informed	-0.019359	-0.000202	0.017951	-0.505260	0.533010
	Noise	-0.011294	-0.000139	0.012153	-0.399528	0.416920
SNX-WETH	LP	-0.009840	-0.000063	0.006957	-0.135156	0.141418
	LT	-0.011296	-0.000029	0.011529	-1.777572	2.181157
	Informed	-0.014557	-0.000031	0.015854	-0.739633	0.635834
	Noise	-0.009860	-0.000029	0.009867	-1.777572	2.181157
STAKE-WETH	LP	-0.006568	-0.000390	0.006901	-0.060975	0.102807
	LT	-0.015172	-0.000143	0.015630	-0.661681	1.390066
	Informed	-0.019871	-0.000228	0.020297	-0.661681	1.390066
	Noise	-0.012978	-0.000127	0.013479	-0.407096	0.374979
UBT-WETH	LP	-0.013436	-0.000046	0.017058	-0.056210	0.041197
	LT	-0.023418	-0.000272	0.024239	-2.173492	0.684236
	Informed	-0.028200	-0.000318	0.030909	-2.173492	0.530942
	Noise	-0.021631	-0.000260	0.021349	-0.611111	0.684236
USDC-USDT	LP	-0.003526	-0.000000	0.002146	-0.016443	0.010921
	LT	-0.001345	-0.000003	0.001394	-0.415557	0.396513
	Informed	-0.002029	-0.000007	0.002096	-0.415557	0.396513
	Noise	-0.001137	-0.000002	0.001157	-0.154831	0.145328
USDC-WETH	LP	-0.000452	-0.000001	0.000285	-0.013262	0.044963
	LT	-0.000362	-0.000000	0.000368	-0.445412	0.422483
	Informed	-0.000828	-0.000000	0.000936	-0.445412	0.422483
	Noise	-0.000253	-0.000000	0.000242	-0.351411	0.349653
WBTC-WETH	LP	-0.001434	-0.000002	0.001369	-0.038737	0.039957
	LT	-0.001605	-0.000002	0.001647	-2.622859	1.285458
	Informed	-0.002635	-0.000003	0.002659	-2.622859	1.285458
	Noise	-0.001204	-0.000002	0.001203	-1.056087	0.848669
WETH-DMG	LP	-0.015085	0.000169	0.010431	-0.050471	0.065274
	LT	-0.015825	0.000177	0.014510	-4.025062	0.618852
	Informed	-0.018714	0.000306	0.018980	-0.769612	0.618852
	Noise	-0.014540	0.000145	0.012645	-4.025062	0.564875

Table 7 – continued from previous page

Pool	Type	5%	50%	95%	Min	Max
WETH-ENJ	LP	-0.011665	-0.000202	0.010551	-0.045953	0.034753
	LT	-0.020387	0.000021	0.020311	-4.485019	2.492653
	Informed	-0.028719	0.000014	0.027724	-4.485019	1.763173
	Noise	-0.017228	0.000023	0.017054	-2.771008	2.492653
WETH-FOX	LP	-0.005538	-0.000014	0.000939	-0.192240	0.239064
	LT	-0.006630	-0.000054	0.006748	-0.426760	0.417105
	Informed	-0.008960	-0.000053	0.008903	-0.406457	0.357383
	Noise	-0.005380	-0.000054	0.005436	-0.426760	0.417105
WETH-HEX2T	LP	-0.048888	0.000292	0.051842	-0.526232	0.326240
	LT	-0.032260	0.000276	0.030335	-4.693004	4.824305
	Informed	-0.050251	0.000744	0.049012	-4.693004	4.824305
	Noise	-0.025834	0.000207	0.022002	-2.669469	2.209537
WETH-USDT	LP	-0.000361	0.000000	0.000535	-0.278843	0.273400
	LT	-0.000330	-0.000001	0.000327	-0.232827	0.322526
	Informed	-0.000737	-0.000002	0.000700	-0.232827	0.322526
	Noise	-0.000241	-0.000001	0.000246	-0.230243	0.234302
XOR-WETH	LP	-0.011513	-0.000186	0.009555	-0.103902	0.096408
	LT	-0.070346	-0.000247	0.083216	-6.086125	5.036696
	Informed	-0.055550	0.000024	0.119308	-6.086125	5.036696
	Noise	-0.075187	-0.000353	0.067829	-3.353420	1.164287
XRT-WETH	LP	-0.043845	0.001514	0.048237	-0.203818	0.209805
	LT	-0.025771	-0.000424	0.026596	-0.651707	0.575380
	Informed	-0.033454	-0.000695	0.033175	-0.651707	0.575380
	Noise	-0.021270	-0.000382	0.022934	-0.406637	0.455064

The table presents distributional statistics of price impact by LPs and LTs and by Informed and Uninformed LT swaps for each AMM pool in our sample. Trades are classified as informed if the gas fee associated with the swap transaction is in the top 25% quartile of gas fees paid for swaps over the prior (rolling) seven day window in that pool. Data is pulled from all Uniswap v2 pools that were created prior to 2020-07-01 and have more than 100,000 transactions. The sample period is from 2020-07-01 to 2024-06-30.