Automated Market Making with Continuity: Liquidity, Price Discovery, and Adverse Selection

Cesare Fracassi, Thomas J. George and Moazzam A. Khoja*

March, 2026

Abstract

Constant-Product Automated Market Makers (CPAMMs) are an established mechanism for trading crypto-assets and may soon support trading of tokenized equities. We develop a minimal-structure model in which traders' strategies and valuations are continuous. Despite its simplicity, the model delivers a rich set of sharp predictions. CPAMM markets provide incentives for perfectly competitive liquidity provision. Equilibrium prices are invariant to liquidity supply. Traders' valuations can be recovered linearly from reserves and transaction prices, allowing variation in traders' motivations to be decomposed into information-driven and liquidity-driven trading components. A simple equilibrium relation links liquidity provision, price spreads, and price impact; and the model provides a closed-form measure of expected adverse selection costs. We test these predictions using Uniswap V2 as the CPAMM market and Binance as the alternative market. Consistent with the theory, Uniswap reserves Granger-predict trading volume, estimated relationships among reserves, price spreads, and slippage have the predicted signs, and adverse selection costs are positively related to the variance of the permanent (information-driven) component of valuations.

*Fracassi is from the McCombs School of Business, University of Texas at Austin, Austin, TX 78705. George is from the C. T. Bauer College of Business, University of Houston, Houston, TX 77210. Khoja is from the University of Texas at San Antonio, San Antonio, TX 78249. Email: Cesare.Fracassi@mccombs.utexas.edu, tom-george@uh.edu, and moazzam.khoja@utsa.edu.

Automated Market Making with Continuity: Liquidity, Price Discovery, and Adverse Selection

Constant-Product Automated Market Makers (CPAMMs) are an established mechanism for trading crypto-assets and may soon support trading of tokenized equities. We develop a minimal-structure model in which traders' strategies and valuations are continuous. Despite its simplicity, the model delivers a rich set of sharp predictions. CPAMM markets provide incentives for perfectly competitive liquidity provision. Equilibrium prices are invariant to liquidity supply. Traders' valuations can be recovered linearly from reserves and transaction prices, allowing variation in traders' motivations to be decomposed into information-driven and liquidity-driven trading components. A simple equilibrium relation links liquidity provision, price spreads, and price impact; and the model provides a closed-form measure of expected adverse selection costs. We test these predictions using Uniswap V2 as the CPAMM market and Binance as the alternative market. Consistent with the theory, Uniswap reserves Granger-predict trading volume, estimated relationships among reserves, price spreads, and slippage have the predicted signs, and adverse selection costs are positively related to the variance of the permanent (information-driven) component of valuations.

Tokenized equities are digital representations of a company's ownership claims recorded on a blockchain, allowing issuance, transfer, and settlement to occur through smart contracts rather than centralized clearing systems. Markets for tokenized equities have moved from speculation to imminent reality, supported by recent regulatory developments.[1] Automated market makers (AMMs) are smart-contract protocols used today to trade cryptocurrency tokens in decentralized markets and, in the future, are expected to support trading of tokenized equities. AMMs pool liquidity and execute trades algorithmically rather than via a traditional order book. Uniswap, the leading AMM, has surpassed $3 trillion in cumulative trading volume.[2] As these developments blur the boundary between traditional limit-order book (LOB) markets and decentralized AMM venues, understanding the liquidity and price discovery properties of equity-type instruments across these environments becomes critical. This paper investigates, both theoretically and empirically, how efficiently AMM markets provide liquidity and price discovery in the presence of potentially informed traders and adverse selection, issues that become central as tokenized equities begin to trade alongside traditional markets.

In LOB markets, humans have historically made the decisions about supplying and demanding liquidity. Algorithmic and high-frequency trading strategies introduce automation into these markets by using complex statistical models to determine trading decisions. However, these models interact with the market in the same way as human traders, by submitting market and limit orders. Their strategies are not publicly known, as informed traders try to disguise their trade intentions, and dealers try to infer the informativeness of trades. Automated market makers (AMMs), in contrast, alter the fundamental structure of markets in three ways: First, the rules governing trade execution and the terms of trade are fixed a priori, encoded in open-source smart contracts such that the strategic interaction among market participants is transparent and can be analyzed without strong assumptions about hidden behavior. Second, anyone can become a liquidity provider, as participation does not require specialized market-making expertise. Third, liquidity is pooled, where individual liquidity providers hold a pro-rata claim on the pool, and their profits or losses arise through changes in the pool's value, rather than by competing to capture bid–ask spreads. These design features make AMMs a fundamentally different market environment from traditional LOBs, with distinct implications for price formation, liquidity provision, and strategic behavior.

In this paper, we model a constant-product automated market maker (CPAMM), meaning one in which quantities traded are constrained to maintain liquidity pool reserves at levels whose product is a constant (i.e., reserves a and b such that $a \times b = k$). This design or variations of it are among the most popular on-chain platforms for trading digital currencies. In such markets, a hyperbolic price function mediates the interaction between liquidity providers and demanders in a manner analogous to how a linear price function governs trading in the theoretical model of Kyle (1985). However, the price function in a CPAMM market is exogenously imposed a priori rather than emerging endogenously as a feature of the equilibrium.

Our model differs from prior theoretical models of AMMs by combining four key elements. First, while existing theoretical models typically adapt stylized assumptions from limit-order book analyses, such as restricting (some) agents to fixed trade sizes, we allow trade sizes to be continuous for all agents. Allowing

[1] See, for example, `https://www.sec.gov/files/rules/sro/nasdaq/2025/34-103989.pdf` and `https://www.reuters.com/business/coinbase-seeking-us-sec-approval-offer-blockchain-based-stocks-2025-06-17/`.

[2] `https://coincentral.com/uniswap-reaches-3-trillion-milestone-as-first-dex-to-hit-all-time-volume-mark/?utm_source=chatgpt.com`

all agents to optimize over a continuous range of quantities simplifies rather than complicates the analysis. The simplicity follows from the very special way in which the CPAMM price function provides incentives to liquidity demanders who have the freedom to choose from a continuum of quantities. Continuity also adds realism to the model because even traders who are motivated primarily by a liquidity need will adjust their quantities to the terms offered for their trade. Second, instead of modeling uncertainty about asset value or uninformed trading as draws from specific distributions (e.g., binomial), we make no parametric assumptions about the value process. Third, we do not require particular traders to be informed or uninformed. We model strategic liquidity providers as anticipating the actions of strategic liquidity demanders, either or both of whom may be informed or uninformed. Finally, since CPAMM markets exist alongside traditional LOB markets, we model the existence of an LOB market as an alternative to trading on a CPAMM.

Our model shows that CPAMMs exhibit properties that are desirable from a traditional market microstructure perspective, yet these properties arise without strong stylized assumptions. First, CPAMMs provide robust value discovery as a stand-alone platform. In equilibrium, transaction prices are invariant to the scale of liquidity provided, and post-transaction reserve prices closely mirror liquidity demanders' valuations. As a result, the unobservable sequence of traders' valuations can be inferred directly from transaction and reserve prices. Second, CPAMMs have the special property of promoting competitive liquidity provision. Although liquidity providers are finite and strategic, they are incentivized to act as if they are perfect competitors. Because transaction prices are invariant to liquidity providers' choices, liquidity providers' expected payoffs do not depend on how their choices might enlarge or shrink the pool of reserves. Therefore, liquidity providers do not benefit from holding back in providing liquidity to the pool in order to keep prices high, as Cournot-style competition would imply for markets without a constant-product constraint on quantities. In essence, the CPAMM synthesizes incentives for liquidity providers to behave as if they are perfect competitors.

These results follow from only the profit maximizing choices of liquidity demanders whose trades are executed by a CPAMM. They do not depend on expectations or distributional assumptions of any kind, and they hold regardless of the liquidity providers' choices. Consequently, they hold regardless of how liquidity providers perceive their options outside of the CPAMM, including the value of the token in LOB markets. In other words, a CPAMM is capable of discovering prices that closely reflect traders' valuations on its own, without further assumptions about price signals or arbitrage trading from other markets with more traditional designs.

The options outside the CPAMM are relevant to the quantity of liquidity supplied, however. Potential liquidity providers naturally compare participating in the CPAMM liquidity pool to their outside option (e.g., a LOB market for the same token pair). These outside-option perceptions are key inputs to the decisions of liquidity providers because the tokens they possess can be sold in either market, or additional tokens can be purchased in the outside-option market to increase the liquidity provider's participation in the CPAMM liquidity pool. Thus, prices and liquidity in the outside-option market are important to the scale of liquidity supplied in the CPAMM. In fact, the expected payoff functions of liquidity providers follow the simple form $p \cdot x - c(x)$ of price-taking firms with convex costs choosing quantities x under perfect competition, where the cost of trading in the outside-option market defines $c(x)$.

The simplicity of the liquidity provider's optimization problem leads to a first-order condition that

relates liquidity provided in a CPAMM to price spreads and the marginal cost of trading in the outside-option/LOB market. This equilibrium relation can be aggregated across liquidity providers and written in the form of a linear regression involving stationary variables that can be observed or estimated. A test of this relation using trading cost estimates from a LOB market is a joint test of the model's predictions and whether a particular LOB market is regarded as the relevant outside option by liquidity providers.

Using data from the Uniswap V2 ETH-USDT smart contract as the CPAMM and Binance as the LOB alternative spot market, we test and find strong empirical support for the model's implications regarding absorption capacity, the strategic equivalence of liquidity providers to competitive market makers, and price efficiency. Our empirical work focuses on three questions. The first is whether the scale of liquidity-demanding trades varies to absorb the scale of liquidity supplied. This feature of liquidity demand is what generates transaction prices that are invariant to liquidity supply in the model. To examine the plausibility of this absorption hypothesis, we estimate a vector autoregression involving trading volume and the scale of liquidity supply. Granger causality tests and the significance of individual VAR coefficients suggest that the data are consistent with the absorption hypothesis.

Second, we test whether the equilibrium relation that the model predicts for Uniswap liquidity supply is consistent with the data. This is the main prediction of the model, and follows from the incentives for perfectly competitive liquidity provision synthesized by the CPAMM. The regression tests are generally supportive, but with a caveat that Binance loses importance as the relevant outside option as the Uniswap market matures. Interestingly, the values of two key fundamental model parameters can be solved from the regression coefficients. The parameter identified with the number of potential liquidity providers is in line with its real-world value, but the parameter identified with the aggregate supply of the token is larger than the actual supply. The divergence of the latter value suggests that the overall level of liquidity provided on Uniswap was greater than the model would predict. Either the model is missing a benefit to providing liquidity on Uniswap, or market participants overestimate the profitability of doing so.

Finally, we investigate the relationship between the adverse selection faced by Uniswap dealers and the trading motivations of liquidity demanders. The valuations that motivate liquidity-demanding trades can be recovered from the sequence of equilibrium reserve and transaction prices. We identify a metric that is implied by the model for expected adverse selection costs borne by liquidity providers. We then estimate a regression of the adverse selection costs metric on the variances of the temporary and permanent components of changes in liquidity demanders' valuations. The results indicate that adverse selection costs are sensitive to both components and more strongly related to the permanent component. This finding suggests that liquidity providers suffer greater losses to adverse selection when motivations for trading are driven to a greater degree by information.

Our paper contributes to the AMM literature along three lines. It develops a model with continuous strategies and valuations, a finite number of strategic liquidity providers acting as Stackelberg leaders, and no parametric assumptions on the value process, thereby relaxing several common restrictions in existing frameworks. Within this setting, the paper shows that the constant-product bonding curve produces perfectly competitive dealer incentives, even among a finite number of strategic dealers. It also introduces an adverse selection measure that is intrinsic to the CPAMM mechanism and shows that traders' latent valuations can be recovered from observable reserve and transaction prices. Variation in valuations can then be decomposed

into permanent and temporary components. Finally, the paper tests the model's predictions using Uniswap V2 and Binance data, and the results are mostly supportive of the model's implications.

Our work is complementary to other models and empirical analyses of liquidity provision on AMMs. As AMM markets gained traction, Lehar and Parlour (2025) provided the first comprehensive review of the Uniswap V2 CPAMM, focusing on its design and comparing its performance to that of a LOB. They examine equilibrium liquidity supply under structural assumptions that enable a direct comparison between the performance of an AMM and a LOB. They show that the AMM can provide liquidity cheaper when asset volatility and adverse selection are low. Under similar assumptions, Capponi and Jia (2025) compare centralized (LOB) and decentralized (AMM) exchanges with a focus on the costs to liquidity providers of stale quotes. While their framework emphasizes stale quotes as a driver of adverse selection, our approach recovers adverse selection costs directly from the CPAMM's reserve and transaction prices, offering a complementary measurement strategy. Aoyagi and Ito (2023) develop a model to analyze how traders and liquidity providers navigate between limit order books and AMMs. Our model shares their insight that the LOB serves as an outside option shaping AMM participation, but focuses on the liquidity supply side of that relationship rather than on a full equilibrium between venues. Because our paper does not undertake a comparison between AMMs and LOBs, our model requires fewer structural assumptions and isolates some features of market quality of the AMM setting more directly, yielding results that complement those of these earlier studies.

To focus on the implications of the constant-product design feature of the smart contract, we assume that each transaction executes against reserves that are visible to market participants upon the order's submission. Our model does not address transaction-ordering strategies in block production. Transaction ordering frictions are analyzed by Park (2023) who shows that the block production protocol creates undesirable features of execution quality. In a similar vein, Capponi et al. (2024) show that the block-production auction mechanism used by some blockchains allows rents to be accrued by miners at the expense of liquidity providers. Both Park (2023) and Routledge et al. (2025) critique aspects of CPAMMs that are more fundamental, such as the shapes of the bonding curve and the requirement to post reserves in proportion to their values in the reserve pool. These critiques reflect a broader skepticism toward non-LOB mechanisms. Indeed, Glosten (1994) famously argues that the open limit order book is the inevitable endpoint of competition among trading venues. Our results on the positive market quality attributes of CPAMMs suggest that the constant-product design has virtues that complement, and in some respects offset, these concerns.

The role of fees in decentralized exchanges has also been an area of interest. Swap fees are transfers from liquidity demanders to liquidity providers. Hasbrouck et al. (2022) show that increases in such fees can lead to higher trading volumes by attracting more liquidity provision that reduces price impact. This feature is evident in our model also, as higher fees increase the return to liquidity provision. In our model, swap fees are added directly to the liquidity pool rather than distributed separately. Barbon and Ranaldo (2021) conduct an empirical analysis of market quality and price efficiency, finding that one of the major impediments to price efficiency is high gas fees. While gas fees are a distinct friction, swap fees create a wedge between reserve prices and traders' valuations in our model's equilibrium and are built into our empirical specification.

The constant-product design of AMMs has also been extended to allow for concentrated liquidity. Has-

brouck et al. (2025) provide a framework for analyzing liquidity provision in such markets, with applications to Uniswap versions 3 and 4. Their model highlights the trade-offs between liquidity concentration and price impact. Under a bonding curve with concentrated liquidity, liquidity provision improves within concentrated bands at the cost of reduced liquidity and thus greater price impact when prices are volatile and move outside the band. Also in the context of Uniswap V3, Lehar et al. (2023) show that fee differentials across pools with different tick sizes lead to efficient fragmentation of trading between fee pools, where large liquidity providers are dominant in pools with lower fees and smaller tick sizes that require aggressive rebalancing of reserves. While our model does not immediately extend to the setting of Uniswap V3, its tractability may be useful as a building block for extensions to that richer setting.

Empirical studies have also explored the evolutionary dynamics of decentralized exchanges. Han et al. (2021) investigate the role of trust in decentralized finance (DeFi) and find that user behavior on decentralized exchanges is influenced by perceptions of security and transparency. Their findings highlight the importance of trust in the adoption and growth of decentralized trading platforms. The emergence of Uniswap as a leading decentralized exchange has been documented by Lo and Medda (2021), who analyze the platform's growth and evolution. They highlight the role of AMMs in enabling decentralized trading and discuss the challenges associated with scaling and maintaining liquidity in a decentralized environment. Similarly, Aspris et al. (2021) describe decentralized exchanges as the "wild west" of cryptocurrency trading, emphasizing the need for regulatory oversight and improved market infrastructure. While our work does not venture to provide suggestions for regulation, the insights from our model could be useful in guiding the deployment of empirical microstructure tools to assess the quality of AMMs.

1 Model

We briefly describe the mechanics of a CPAMM as implemented in protocols such as Uniswap V2 to clarify how our model maps to an actual decentralized exchange.[3] A CPAMM is a smart contract that holds reserves of two tokens and exchanges them at prices determined by a deterministic rule requiring the product of reserves to remain constant before fees. Liquidity demanders specify only the input token and quantity; the contract mechanically computes the output, adds fees to reserves, and updates reserves, which are observable. There is no order book and no discretionary price setting: liquidity providers supply reserves in advance and own a pro rata share of the reserve pool. They affect the price schedule only through the quantity of liquidity they post. As a result, the CPAMM implements a transparent, continuously defined, and rigid price schedule.

This design contrasts sharply with both theoretical and institutional benchmarks. In canonical theoretical models such as Kyle (1985, 1989), smooth price schedules emerge endogenously from the optimizing behavior of liquidity providers with full pricing discretion. In limit order book markets, liquidity providers can post arbitrary quotes at each quantity, price schedules need not be continuous or conform to a particular functional form, and price schedules are often only partially observable. By contrast, AMM prices are fully

[3]There are AMMs that do not use the constant-product design. For example, Uniswap V3 uses constant-product within narrow price bands, but allows one-sided liquidity supply outside those bands. Curve mixes constant-product with constant-sum for trading stablecoin pairs. We focus on the constant-product design because it is a common building block for several smart contracts and because its economics are easy to analyze without imposing additional structure for the sake of tractability.

mechanical and do not allow liquidity providers to adjust price quotes in response to information, adverse selection, or competition. Nevertheless, we show that the constant-product AMM exhibits several desirable properties typically associated with market structures in which prices are flexible.

Several institutional features of contemporary AMM markets are not modeled here. We abstract from gas fees and transaction ordering strategies such as top-of-the-block arbitrage, front-running, and sandwiching by assuming immediate trade execution at the posted price schedule. We also abstract from just-in-time (JIT) liquidity provision and instead model reserve choices as being made prior to trader arrival. Later protocol innovations, such as concentrated liquidity in Uniswap V3 and customizable execution logic via V4 hooks, are likewise omitted. These features expand the design space of AMMs but do not alter the core mechanism we study: a deterministic, publicly known mapping from reserves to prices against which traders optimally choose quantities. Abstracting from these institutional details allows us to focus on the fundamental economics of liquidity provision and price formation in constant-product markets.

For concreteness, we label the token pair in the model as Ethereum (ETH) and Tether (USD). Trades that demand liquidity arrive one at a time and are indexed by t. Liquidity demanders are assumed only to be non-satiated in wealth. The symbol $v_t > 0$ denotes the per-unit value in USD that liquidity demander t attaches to ETH. By allowing continuous strategies, it is not necessary to specify v_t further. Nevertheless, we can imagine that v_t has a fundamental component that is common to all market participants and possibly a private component that is unique to the individual; i.e., valuations might be a blend of fundamental information that is common and persistent, and "noise" resulting from a liquidity shock that is unique to the individual and transitory. Trades are executed by the CPAMM sequentially, so we often refer to the transaction of liquidity demander t simply as the transaction at time t.

Liquidity demanders trade against a pool of reserves of ETH and USD. If a liquidity demander purchases ETH, the trade will draw down the reserves of ETH and add to the reserves of USD; and vice-versa if the liquidity demander sells ETH. CPAMM contracts charge liquidity demanders a transaction fee that is proportional to the units of the token added to the liquidity pool. The fee is deducted from the input token amount, and then the swap occurs according to the AMM's price function. For simplicity, we model the fee as being paid in USD (the numeraire) as happens when liquidity demanders buy ETH for USD. For brevity, we henceforth refer to liquidity demanders as traders and liquidity providers as dealers.

The reserves are posted by dealers. When tokens are posted as reserves, they are no longer accessible to the dealer, which eliminates the possibility that dealers default on a promise to deliver. Usually, AMM contracts do not charge fees for posting or withdrawing reserves, so we model dealers as choosing reserve levels afresh rather than modeling changes in reserves from the prior period. We assume there are J risk-neutral dealers and that dealer j has an endowment of $K_j > 0$ units of ETH. We assume that dealers have access to USD at a zero interest rate.[4]

After the transaction at time $t-1$ and prior to the arrival of trader t, each dealer j posts reserves of ETH and USD, which we denote by R_{otj} and R_{1tj}, respectively. We assume that dealers regard the current value of the remaining $K_j - R_{otj}$ units of ETH in their possession as its value if sold on an alternative platform such as a LOB market. Thus, the dealer's opportunity cost of posting ETH on reserve in the AMM

[4]A cost to raising cash could be added, which would affect liquidity provision in a manner similar to the friction associated with buying and selling ETH on a LOB market described below.

market is forgoing the proceeds available from selling the ETH on the LOB market. The outside option of selling ETH on the LOB market is what can prevent the dealer from committing the entire endowment to reserves on the CPAMM. The dealer can also buy ETH in the LOB market to post as reserves on the CPAMM.

In our model, agents are identified by their valuations. Because of the unique permissionless nature of blockchain-based AMMs, an individual agent may supply or demand liquidity (or both). Focusing on valuations frees us from specifying a stochastic process for trader arrivals, from labeling traders as informed or uninformed, and from restricting trade sizes. Instead, agents simply arrive with valuations. Whether traders adversely select dealers depends on the strength of the correlation between the traders' valuations and the future value of the token to dealers. We discuss such correlation in greater detail later in connection with estimating adverse selection in a CPAMM.

1.1 Reserve and Transaction Prices and Equilibrium Concept

Dealers face an important constraint when posting reserves on a CPAMM. Namely, dealers' individual postings and withdrawals of reserves are required to *preserve the ratio* of ETH and USD already in the pool. Thus, if q_{t-1} is the ratio of USD to ETH in the pool following transaction $t-1$, then all dealers must post reserves available for transaction t such that $R_{1tj}/R_{otj} = q_{t-1}$. This ratio is determined by the terms of the liquidity-demanding trade at time $t-1$. The ratio of USD to ETH in the pool is referred to as the *reserve price* because it is the dollar value per unit of ETH implied by the reserve levels of USD and ETH in the pool. This constraint allows dealers only to scale their participation or exposure to the pool, and not to change exposures to the tokens independently, which in turn would change the reserve price. In a CPAMM, reserve prices are changed by liquidity-demanding trades only.

The terms of each liquidity-demanding transaction in a CPAMM are governed by a time-invariant, deterministic formula that requires the transaction to *preserve the product* of the quantities of ETH and USD in the reserve pool. Whereas dealers are constrained to preserve the ratio of quantities when posting and withdrawing reserves, liquidity-demanding transactions are constrained to preserve the *product*. The relation that equates the product of pre-transaction reserves and post-transaction reserves is called the *bonding curve.* The particular constant-product shape of the bonding curve used by a CPAMM is crucial to the economically interesting results that follow. Indeed, a CPAMM's bonding curve design appears to be quite special in ways that support it having desirable attributes found in markets with traditional and more flexible structures.

Before transaction t, the reserve pool contains quantities $R_{ot} \equiv \sum_{j=1}^{J} R_{otj}$ of ETH and $R_{1t} \equiv \sum_{j=1}^{J} R_{1tj}$ of USD. The constant-product bonding curve requires that the quantity Q_{ot} of ETH that liquidity-demanding transaction t draws from the pool is offset by a deposit of Q_{1t} USD such that

$$R_{1t}R_{ot} = (R_{1t} + (1-f)Q_{1t})(R_{ot} - Q_{ot}), \tag{1}$$

where f is the proportional transaction fee collected by the CPAMM contract.[5] The bonding curve in

[5]The fee is equal to 0.30% (30bp) in the markets from which our data are drawn.

Equation (1) determines the transaction price. Solving for Q_{1t} yields

$$Q_{1t} = \left(\frac{1}{1-f}\right)\left(\frac{R_{1t}}{R_{ot} - Q_{ot}}\right) Q_{ot} \tag{2}$$

as the *total* USD the liquidity demander must add to reserves in exchange for receiving Q_{ot} units of ETH. Dividing by Q_{ot} yields the transaction price in USD *per unit* of ETH for a transaction to buy Q_{ot} total units of ETH:

$$P(Q_{ot}) = \left(\frac{1}{1-f}\right)\left(\frac{R_{1t}}{R_{ot} - Q_{ot}}\right). \tag{3}$$

The liquidity demander can be regarded as paying a uniform price for the Q_{ot} ETH units traded, where that price is a hyperbolically increasing function of the traded quantity. This deterministic price function, rather than the human or algorithmic placement of limit orders, mediates every interaction between dealers and traders in a CPAMM.

We assume that each liquidity-demanding transaction executes against the current state of reserves visible to market participants at the time the trade is chosen—i.e., liquidity demander t chooses quantities Q_{ot} and Q_{1t} knowing the current quantities of ETH and USD in the reserve pool. We solve first for the trader's optimal choices conditional on the reserve quantities. Then we solve for the dealers' choices of reserves that anticipate the trader's optimal choices; i.e., we model the dealers as Stackelberg leaders. An equilibrium consists of dealers' and traders' choices that are sequential in the manner just described and rational, meaning that the trader maximizes their payoff and dealers maximize their expected payoffs.

1.2 Trader's Optimization Problem and Post-Transaction Reserve Prices

Since the bonding curve is common knowledge and we assume the current state of reserves is observable, the price the trader must pay for each possible quantity is known, and the trader's payoff is a deterministic function of the quantity chosen.[6] The trader's optimization problem has the canonical form of a liquidity demander's problem in the microstructure literature, adapted here to the nonlinear price schedule of a CPAMM. The trader chooses to buy Q_{ot} ETH to maximize

$$\max_{Q_{ot}} \pi_t = (v_t - P(Q_{ot}))Q_{ot}, \tag{4}$$

where v_t is the value in USD that liquidity demander t places on each unit of ETH and $P(x)$ is the per-unit price function of ETH in USD implied by the bonding curve and observable to trader t. Substituting from Equation (3) and differentiating with respect to Q_{ot} yields a first-order condition for the trader's optimal quantity of ETH to buy:[7]

$$Q^*_{ot} = R_{ot} - \left(\frac{R_{ot}R_{1t}}{(1-f)v_t}\right)^{\frac{1}{2}}, \tag{5}$$

[6] This is an abstraction from the actual process of block building that happens between the submission of a trade and its posting on the blockchain. In reality, the trade will be blocked with other transactions from all the smart contracts that are active on the blockchain. It is possible that one of those other transactions is a scaling up or down of reserves by a dealer in the CPAMM that, if placed before the CPAMM trade, would change the reserves from what the trader observed as the state of reserves when submitting the trade.

[7] The negative square root is selected so that the trader's demand is increasing in v_t and the first-order condition identifies a maximum.

where R_{ot} and R_{1t} are the pre-transaction reserves of ETH and USD. The fee reduces demand the same as an equivalent reduction in the trader's valuation. Substituting from Equation (5) into Equation (2) yields the number Q^*_{1t} of USD the trader pays.[8]

The trader's quantity choice determines how the reserves of ETH and USD are changed by the transaction. The post-transaction ETH and USD reserves are given by

$$\hat{R}_{ot} = R_{ot} - Q^*_{ot} \qquad \text{and} \qquad \hat{R}_{1t} = R_{1t} + Q^*_{1t}. \tag{6}$$

An expression for the post-transaction reserves of ETH follows immediately from Equations (5) and (6)

$$\hat{R}_{ot} = \left(\frac{R_{ot}R_{1t}}{(1-f)v_t}\right)^{\frac{1}{2}} = \left(\frac{q_{t-1}}{(1-f)v_t}\right)^{\frac{1}{2}} R_{ot}, \tag{7}$$

where the second equality recognizes that because all dealers' reserve choices are constrained to satisfy $R_{1tj} = q_{t-1}R_{otj}$, the reserve pool as a whole likewise satisfies: $R_{1t} = q_{t-1}R_{ot}$. An expression for the post-transaction USD reserves is obtained by substituting Equation (5) into (2) to get Q^*_{1t} then adding R_{1t} as in Equation (6). The result is

$$\hat{R}_{1t} = \left\{\left(\frac{1}{1-f}\right)^{\frac{1}{2}} v_t^{\frac{1}{2}} - \left(\frac{f}{1-f}\right) q_{t-1}^{\frac{1}{2}}\right\} R_{ot}\ q_{t-1}^{\frac{1}{2}}. \tag{8}$$

The ratio of Equations (8) and (7) is the post-transaction reserve price, which after simplifying is

$$q_t = v_t - f \cdot \left(\frac{v_t\, q_{t-1}}{1-f}\right)^{\frac{1}{2}}. \tag{9}$$

This expression results only from the liquidity demander having optimized, and it is *independent* of the dealers' choices. Regardless of the dealers' choices, the optimal choice of the trader, *by itself,* drives the post-transaction reserve price to the liquidity demander's valuation, minus a small distortion attributable to the existence of the fee.

This result is a market efficiency feature of the CPAMM in the sense that post-transaction prices are informative about traders' valuations. This follows directly from the trader optimizing against the constant-product bonding curve as a monopolist. The first-order condition equates the trader's valuation v_t to the marginal cost of the trade, which is the post-transaction reserve price q_t plus the fee. As a result, the post-transaction reserve price reveals the liquidity demander's valuation up to a fee adjustment, regardless of the dealers' reserve choices and without any assumptions about information, distributions, or the behavior of other market participants. This would not be true in a competitive LOB environment such as Glosten (1994) because prices are set to protect liquidity providers from the possibility that a trade of any given size could have been part of a larger order. In that environment, the trader pays more than their valuation on the marginal unit traded, and the post-trade midquote is equal to the liquidity *provider's* expected value given the quantity traded.

[8] We need not assume that the trades at time t and $s \neq t$ are different individuals. However, a single trade receives better execution than multiple smaller trades against a constant-product bonding curve [see Angeris et al. (2021)].

1.3 Dealer's Optimization Problem and Transaction Prices

Liquidity is provided by the pool of reserves, and dealers' profit or loss from providing liquidity in the AMM arises from changes in the value of the pool. Each dealer's payoff from providing liquidity is the post-transaction USD value of that dealer's share of the reserve pool. Each dealer's ownership share of the reserve pool equals the pre-transaction contribution to total reserves, R_{otj}/R_{ot}, in ETH. This ratio will be the same as the ratio R_{1tj}/R_{1t} in USD because all dealers are required to post reserves of ETH and USD in the same proportion.

The objective of dealer j is to maximize the expected USD payoff by allocating the endowment K_j between liquidity provision in the AMM reserve pool against which transaction t will execute or to its sale on the alternative LOB market:

$$\tilde{\pi}_j = \left[\hat{R}_{1t} + \tilde{v}_j \hat{R}_{ot}\right] \frac{R_{otj}}{R_{ot}} - R_{1tj} + (K_j - R_{otj})(v_o - \lambda(K_j - R_{otj})). \tag{10}$$

The first term is the post-transaction USD value of dealer j's share of the reserve pool, where R_{otj} is the quantity of ETH dealer j commits to the pool, and $\tilde{v}_j$ is dealer j's post-transaction per-unit valuation of ETH in USD. We make no particular assumptions about the stochastic properties of $\tilde{v}_j$. The second term, $-R_{1tj}$ is the dollar cost of posting R_{1tj} USD to the reserve pool, which is required to balance the commitment of R_{ojt} to ETH reserves. The last term is the proceeds in dollars the dealer receives from selling on the LOB market the ETH endowment that is not posted as AMM reserves, $K_j - R_{otj}$ (or, if $K_j - R_{otj} < 0$, buying additional units on the LOB market to add to the AMM pool). Regarded more generally, the last term is the dealer's outside option that serves as an alternative to posting the entire endowment as reserves.

To model LOB market transaction costs in the simplest possible way, we assume that the LOB price per-unit of ETH traded is linear in order flow as in Kyle (1985), with an intercept of v_o, and a constant slope of $\lambda > 0$ that is the per-unit price impact of trades. Note that λ is in units of USD per ETH. We assume that each dealer trades independently as a liquidity demander on the LOB market and receives a price of $v_o - \lambda(K_j - R_{otj})$ per unit of ETH sold on the LOB market. A negative sign precedes λ in Equation (10) because the net sale of $K_j - R_{otj}$ units of ETH depresses the price that the AMM dealer receives. For convenience, we refer to v_o as the LOB midquote.

Substituting for $\hat{R}_{ot}$ and $\hat{R}_{1t}$ from Equations (7) and (8) yields

$$\tilde{\pi}_j = \left[\left\{\left(\frac{1}{1-f}\right)^{\frac{1}{2}} v_t^{\frac{1}{2}} - \left(\frac{f}{1-f}\right) q_{t-1}^{\frac{1}{2}}\right\} R_{ot}\ q_{t-1}^{\frac{1}{2}} + \tilde{v}_j \left(\frac{q_{t-1}}{(1-f)v_t}\right)^{\frac{1}{2}} R_{ot}\right] \frac{R_{otj}}{R_{ot}}$$
$$- R_{1tj} + (K_j - R_{otj})(v_o - \lambda(K_j - R_{otj})). \tag{11}$$

Inspection of Equation (11) reveals another special feature of the CPAMM design. The combination of the constant-product bonding curve and the pooling of reserves to provide liquidity implies that the R_{ot} terms in Equation (11) cancel. This cancelation leaves an expression that depends only on the reserves of dealer j who is making the choice. In arriving at Equation (11), we have substituted using Equations (7) and (8) that depend on the liquidity demander having optimized. Thus, when the liquidity demander is foreseen to optimize, each dealer's payoff depends only on that dealer's own reserve choice and not on the

reserve choices of other dealers or the size of the overall reserve pool. This is a remarkable, and possibly unique, property of the CPAMM architecture. [9]

The CPAMM price function generates a decision problem for the dealer that resembles that of perfect competition among firms; i.e., one in which the dealer optimizes a function of the simple form $p \cdot q_j - c_j(q_j)$. In other words, the pooling of reserves combined with the particular bonding curve specified by a CPAMM market generates incentives for individual dealers that resemble those of perfect competition, despite dealers being aware that they are not atomistic competitors. This special feature delivers more liquidity into the AMM market than would occur under Cournot style choices, regardless of the outside option. Nevertheless, the terms of the outside option are still important to the quantity of liquidity provided.

Using the dealer's constraint, $R_{1tj} = q_{t-1}R_{otj}$, and simplifying, Equation (11) becomes

$$\tilde{\pi}_j = \left[\left(\frac{q_{t-1}}{1-f}\right)^{\frac{1}{2}}\left(\frac{v_t + \tilde{v}_j}{v_t^{\frac{1}{2}}}\right) - \left(\frac{1}{1-f}\right)q_{t-1}\right]R_{otj} + (K_j - R_{otj})(v_o - \lambda(K_j - R_{otj})). \tag{12}$$

The dealer's payoff is uncertain because the next trader's valuation, v_t, is uncertain at the time reserves are chosen, and the dealer's future valuation v_j may also be regarded by the dealer as the realization of a random variable. To emphasize that these are sources of uncertainty, we place tildes on both variables when writing the dealer's expected payoff equation:

$$\mathrm{E}_j[\tilde{\pi}_j] = \left[\left(\frac{q_{t-1}}{1-f}\right)^{\frac{1}{2}}\mathrm{E}_j\left[\frac{\tilde{v}_t + \tilde{v}_j}{\tilde{v}_t^{\frac{1}{2}}}\right] - \left(\frac{1}{1-f}\right)q_{t-1}\right]R_{otj} + v_o(K_j - R_{otj}) - \lambda(K_j - R_{otj})^2, \tag{13}$$

The subscript j on the expectation operators emphasizes that dealers may have heterogeneous beliefs. Equation (13) is the dealer's objective function for choosing R_{otj}. The first-order condition is linear in R_{otj}. Solving the first-order condition for R_{otj} yields an expression for dealer j's optimal choice of ETH reserves to post. The choice of USD to post is constrained to be proportional to ETH, where the constant of proportionality is q_{t-1}. The resulting choices are:

$$R^*_{otj} = K_j - \frac{1}{2\lambda}\left[v_o - \left(\frac{q_{t-1}}{1-f}\right)^{\frac{1}{2}}\left\{\mathrm{E}_j\left[\frac{\tilde{v}_t + \tilde{v}_j}{\tilde{v}_t^{\frac{1}{2}}}\right] - \left(\frac{q_{t-1}}{1-f}\right)^{\frac{1}{2}}\right\}\right] \qquad \text{and} \qquad R^*_{1tj} = q_{t-1}R^*_{otj}. \tag{14}$$

The choices in Equation (14) anticipate how the liquidity demander's upcoming optimal trading decision will depend on the reserves—i.e., the dealer behaves as a Stackelberg leader. Summing across dealers yields the equilibrium quantities of ETH and USD in the reserve pool prior to the transaction at time t; i.e., $R^*_{ot} = \sum_{j=1}^{J} R^*_{otj}$ and $R^*_{1t} = q_{t-1}R^*_{ot}$, respectively. The AMM makes these choices visible to all market participants. When trader t arrives, the trade will execute at the terms defined by the bonding curve whose shape depends on the dealers' choices in Equation (14).

[9] This property is not evident in settings where liquidity is provided by agents strategically placing supply schedules in a limit order book [e.g., Glosten (1989), Kyle (1989), Back and Baruch (2013), Boulatov and George (2013)]. This feature is also not evident even when orders are pooled under a different bonding curve. Consider the alternative of a linear price function, $P(Q_o) = (\frac{1}{1-f})(R_{1t}/R_{ot} + bQ_{ot})$. This specification generates a profit function that depends on both the dealer's own choice, R_{otj}, and the dealer's share of total reserves R_{ojt}/R_{ot}. The dealer's optimization in this case involves choosing reserves while accounting for the impact on the total liquidity supplied.

The key endogenous variable in the interaction between the trader and the dealers is the price of ETH in the transaction at time t. The equilibrium transaction price is given by Equation (3) where the generic quantity and reserves variables are replaced by the trader's and dealers' optimal choices as given in Equations (5) and (14):

$$P^*(Q^*_{ot}) = \left(\frac{1}{1-f}\right)\left(\frac{R^*_{1t}}{R^*_{ot} - Q^*_{ot}}\right). \tag{15}$$

This equation illustrates how the dealers' choices affect the shape of the price function. But it turns out that the equilibrium realized transaction price is *invariant* to the dealers' choices (provided that reserves are non-zero). To see why, note that Equation (5) holds for any choices that the dealers make, and dealers are constrained to choices that satisfy $R_{1t} = q_{t-1}R_{ot}$. Substituting from (5) for the denominator in (15), and substituting from the constraint in (14) into the numerator, yields

$$P(Q^*_{ot}) = \left(\frac{1}{1-f}\right)\frac{q_{t-1}R_{ot}}{\left(\frac{q_{t-1}}{(1-f)v_t}\right)^{\frac{1}{2}}R_{ot}} = \left(\frac{v_t q_{t-1}}{1-f}\right)^{\frac{1}{2}}. \tag{16}$$

The R_{ot} terms cancel, implying that the transaction price does not depend on the dealers' choices. This result mirrors the invariance of the post-transaction reserve price to the dealers' choices in Equation (9). The invariance of the transaction price also is a special feature of the constant-product bonding curve.

In Appendix C, we show that in order for transaction prices to be invariant to liquidity providers' choices, a bonding curve has to generate a price function that satisfies a very specific partial differential equation. We then show that the price function associated with the constant-product bonding curve satisfies that pde *uniformly* in the domain of possible transaction quantities *and* reserve choices. The constant-product bonding curve is therefore very special in managing the expansion and contraction of the liquidity pool available to traders while isolating prices from dealers' choices that alter the size of the pool. Reserve prices and transaction prices depend only on traders' choices, and reserve prices reflect traders' valuations exactly, up to an adjustment for the fee.

In CPAMM markets, the particular shape of the bonding curve is such that the optimal trades of liquidity demanders absorb changes in liquidity supply in a manner that renders transaction prices invariant to the scale of reserves posted by liquidity providers. Any variation in that scale is (optimally) absorbed by liquidity demanders changing the scale of their orders. The pde in Appendix C is a precise statement of the condition that changes in liquidity supply are (optimally) absorbed by the liquidity demander. The observation that the CPAMM price function has this feature clarifies the intuition for why liquidity providers do not exhibit Cournot behavior despite being strategic. Since the *overall* supply of liquidity does not affect the price, the *contribution* of individual liquidity providers to the overall supply does not affect prices either. Consequently, individual dealer profit depends only on the dealer's exposure to the expected profit of the pool, so dealers optimize without regard to their impact on the size of the pool.

By Equation (16), when the liquidity demander optimizes, the transaction price is the geometric mean of the pre-transaction reserve price and the trader's valuation of ETH, grossed up by the fee. The reason the transaction price is geometrically only "half way" between the pre-trade reserve price and the trader's valuation is because each liquidity demander trades as a monopolist, resisting more aggressive trading that

would move the price too much against themselves. The transaction price is quite different from the value implied by the quantities that remain on reserve following the transaction, however. The post-transaction reserve price in Equation (9) updates the "rest of the way" to the trader's valuation (minus an adjustment attributable to the fee) as is evident from combining Equations (16) and (9):

$$q_t = v_t - f \cdot P(Q^*_{ot}). \tag{17}$$

This simple expression links equilibrium reserve and transaction prices to liquidity demander valuations, independent of the dealers' choices.

Equations (16) and (17) illustrate how *equilibrium* pricing in a CPAMM differs from that in a LOB. Glosten (1994) shows that in a LOB with competitive liquidity provision, the price quote at each quantity is the expected security value conditional on a transaction of at least that quantity. Liquidity providers choose bids and offers that reflect the inference they draw about the security's value at each potential quantity traded. The schedule of price-quantity pairs chosen in this way protects dealers from adverse selection. Such a schedule breaks even in expectation, and any schedule that undercuts this one earns expected losses for dealers.

Offering liquidity at prices that correspond to each potential quantity traded is not possible in a CPAMM because the bonding curve sets the schedule of quoted prices algorithmically according to a function that has only the two reserves as parameters. Ignoring the fee, if traders optimize, the transaction price per unit is geometrically half way between the prior reserve price and the trader's valuation, by Equation (16). Traders capture, as a surplus, the spread between their valuation and that half-way point, multiplied by the number of units they buy or sell. If the trader's valuation is the fundamental common value, the dealer loses the trader's surplus as the cost of adverse selection.

The mechanism by which dealers incur this loss is that the USD deposit and ETH withdrawal from reserves associated with the trade lead to a post-transaction reserve price that matches the trader's valuation; i.e., Equation (17). The CPAMM reserve price successfully "discovers" the trader's valuation, and liquidity providers can liquidate their share of the pool at the trader's valuation. However, the quantities of ETH and USD remaining in the pool are different after the transaction than before in a way that is adverse to the dealers if traders are informed. If the trader's valuation is the common fundamental value and the trader buys ETH, the number of units of ETH that remain in the reserve pool falls by enough that the surplus captured by the trader exactly equals the value lost by the dealers. The dealers sold ETH at a price below its fundamental value. Unlike prior analyses of adverse selection in AMMs, our model characterizes the dealer's loss precisely. In the common values case, the loss is the trader's surplus, and our model provides an exact expression for the post-transaction reserve price as a function of the trader's valuation. This exact characterization is what enables the empirical measurement of adverse selection costs developed in Subsection 1.5 under weaker assumptions than common values, requiring only that traders' and dealers' valuations satisfy the common-outlook assumption introduced above.

Dealers cannot use choices of price-quantity pairs to protect themselves from adverse selection in a CPAMM. They can only control their exposure to it by choosing how much to participate in the reserve pool. If adverse selection is severe enough, liquidity providers in the CPAMM can fail to break even (even

with the fee) and would therefore refuse to participate. In order for the CPAMM to be a viable trading platform, variation in traders' valuations, v_t, must be at least partly noise, meaning that v_t is less than perfectly correlated with the post-trade value of the token to dealers, v_j. This observation is an application of Glosten (1994)'s proposition that no market is more resilient to adverse selection than a competitive open LOB.

The relation in Equation (17) holds at each t, and therefore also describes the joint dynamics between the three variables. The unobservable valuation v_t is exactly identified by a sum of the two observable prices, with the transaction price weighted by the fee, i.e.,

$$v_t = q_t + fP_t. \tag{18}$$

Traders' unobserved valuations can be recovered from the observable reserve and transaction prices.

Equation (18) can be used to characterize traders' valuations empirically, such as decomposing the variance of the $q_t + fP_t$ series (as a perfect proxy for the v_t series) into temporary and permanent components, for example. Permanent changes in v_t would be regarded as information-based, while temporary changes would be attributable to other motivations, such as liquidating ETH to use in transactions for goods and services, portfolio rebalancing, or buying ETH as a digital store of value for excess cash [e.g., Cong et al. (2020)].

1.4 Optimal Liquidity Provision in CPAMM Markets

The results above show that the key endogenous variables of reserve price and transaction price are independent of the dealers' choices. As noted above, the dealers' choices determine the scale of reserves, and the liquidity demanders' trades adjust to absorb whatever scale is offered. Nevertheless, there is an equilibrium association between prices and liquidity provision because dealers' liquidity provision decisions depend on the prices that dealers *expect* in the CPAMM. The commitment of resources to liquidity provision is an investment, which depends on the value that the dealers expect the investment will return.

To make this connection easier to see, we make an assumption that we refer to as dealers and the trader as having a "common outlook" on the value of ETH, meaning that dealers' valuations are centered on the valuation of the liquidity demander, but are otherwise conditionally independently distributed. This assumption simplifies equations by eliminating terms that would relate to dealers and traders betting against biases they perceive in each other's valuations. Specifically, the common-outlook assumption is that $\tilde{v}_j = \tilde{v}_t \tilde{\epsilon}_{jt}$ where $\tilde{\epsilon}_{jt}$ and $\tilde{v}_t$ are mutually independent for all dealers j, and $\mathrm{E}_j[\tilde{\epsilon}_{jt}] = 1$ for all j. Under this assumption,

$$\mathrm{E}_j\left[\frac{\tilde{v}_t + \tilde{v}_j}{\tilde{v}_t^{\frac{1}{2}}}\right] = \mathrm{E}_j\left[\tilde{v}_t^{\frac{1}{2}}\right] + \mathrm{E}_j\left[\tilde{v}_t^{\frac{1}{2}}\tilde{\epsilon}_{jt}\right] = \mathrm{E}_j\left[\tilde{v}_t^{\frac{1}{2}}\right] + \mathrm{E}_j\left[\tilde{v}_t^{\frac{1}{2}}\right]\mathrm{E}_j\left[\tilde{\epsilon}_{jt}\right] = 2\mathrm{E}_j\left[\tilde{v}_t^{\frac{1}{2}}\right]. \tag{19}$$

Since Equation (16) is a deterministic relationship that holds for any choices that the dealers make, it

must hold at their optimal choices. Rearranging (16) with $v_t^{\frac{1}{2}}$ on the left we have

$$v_t^{\frac{1}{2}} = \left(\frac{1-f}{q_{t-1}}\right)^{\frac{1}{2}} P^*(Q_{ot}^*), \qquad \text{with probability 1.} \tag{20}$$

The deterministic nature of Equation (20) implies that, following the transaction at time $t-1$ and prior to the transaction at time t (i.e., while v_t and the transaction price are uncertain to dealers), the two sides of Equation (20) are equal in expectation:

$$\mathrm{E}_j\left[\tilde{v}_t^{\frac{1}{2}}\right] = \left(\frac{1-f}{q_{t-1}}\right)^{\frac{1}{2}} \mathrm{E}_j\left[\tilde{P}^*(Q_{ot}^*)\right]. \tag{21}$$

Substituting two times the right-hand side of Equation (21) for the expectation term in Equation (14) and simplifying yields the following equation for dealer j's optimal choice of ETH reserves:

$$\begin{aligned} R_{otj}^* &= K_j - \frac{1}{2\lambda}\left[v_o - \left\{2\mathrm{E}_j[\tilde{P}_t^*] - \frac{q_{t-1}}{1-f}\right\}\right] &(22)\\ &= K_j - \frac{1}{2\lambda}\left[\left(v_o - \mathrm{E}_j[\tilde{P}_t^*]\right) + \left(\frac{q_{t-1}}{1-f} - \mathrm{E}_j[\tilde{P}_t^*]\right)\right] &(23)\\ &= K_j - \frac{1}{2\lambda}\left[\mathrm{E}_j[\tilde{S}_t^*] - \left(\mathrm{E}_j[\tilde{P}_t^*] - v_o\right)\right] &(24) \end{aligned}$$

where $S_t^* \equiv \frac{q_{t-1}}{1-f} - P_t^*$ is a fee-adjusted signed "slippage" in value, from the pre-trade reserve price at which the dealer posted reserves to the transaction price the dealer receives. Positive slippage is costly to the dealer because the transaction price is lower than the pre-trade reserve price at which the dealer bought into the pool.

Equation (24) says that the dealer commits the full endowment, less an adjustment, to reserves in the AMM pool. The downward adjustment is larger if the dealer expects more slippage and smaller if the dealer expects the transaction price on the AMM to be higher than the LOB midquote. The sensitivity of the dealer's choice to these price spreads is greater when the LOB market is more liquid (i.e., $1/\lambda$ is larger). The lower is the price impact of trading on the LOB market, the more sensitive are dealers' reserve choices in AMM to expected price spreads.

It is interesting to note that reserves are not necessarily increasing in the fee rate because both adjusted slippage and the transaction price are higher when fees are higher. Since slippage and the transaction price affect the right-hand side of Equation (25) in opposite ways, the net effect of a change in the fee depends on whether the fee change elevates transaction prices more than it dampens demand [see Hasbrouck et al. (2022)].

Summing Equation (24) across dealers yields an expression for total ETH reserves:

$$R_{ot}^* = K + \frac{J}{2\lambda}\left[\left(\overline{\mathrm{E}}[\tilde{P}_t^*] - v_o\right) - \overline{\mathrm{E}}[\tilde{S}_t^*]\right] \tag{25}$$

where K is the aggregate endowment of ETH, and $\overline{\mathrm{E}}[\cdot]$ is the average of expectations across dealers. The first term in the square brackets is an expected AMM transaction price premium over the LOB market midquote,

and the minus slippage term is an AMM transaction price premium over the $t-1$ reserve price. Thus, the total reserves of ETH will be larger when dealers on average expect higher transaction prices on the AMM than on the LOB market and less slippage from the prior reserve price to the AMM transaction price. Note that R^*_{ot} remains positive (as it must) as long as the term in square brackets is not too negative. If the term in square brackets is positive, then dealers buy ETH on the LOB market to post as additional reserves on the AMM market to capture high expected AMM transaction prices.

1.5 Measuring Adverse Selection on CPAMM Markets

In this subsection, we examine how to measure adverse selection in a CPAMM market. We begin by identifying how adverse selection affects dealers' profit, which is given by Equation (11) and repeated here for convenience:

$$\tilde{\pi}_j = \left[\left\{ \left(\frac{1}{1-f} \right)^{\frac{1}{2}} v_t^{\frac{1}{2}} - \left(\frac{f}{1-f} \right) q_{t-1}^{\frac{1}{2}} \right\} R_{ot}\ q_{t-1}^{\frac{1}{2}} + \tilde{v}_j \left(\frac{q_{t-1}}{(1-f)v_t} \right)^{\frac{1}{2}} R_{ot} \right] \frac{R_{otj}}{R_{ot}} - R_{1tj} + (K_j - R_{otj})(v_o - \lambda(K_j - R_{otj})). \quad (26)$$

The term that begins with $\tilde{v}_j$ is the dollar value that dealer j attributes to the quantity of ETH that remains in the pool after the liquidity demander's trade. This term is the component of dealer profit that is affected by adverse selection.

The ratio q_{t-1}/v_t inside the parentheses following $\tilde{v}_j$ indicates that the pool is left with less ETH when the liquidity demander values ETH highly in relation to the reserve price, and vice-versa. Dealer j is adversely selected if liquidity demanders' trades systematically leave the pool with less (more) ETH when its value to dealer j is high (low). Thus, the degree of adverse selection depends on the relation between v_j and $(q_{t-1}/v_t)^{\frac{1}{2}}$. The stronger and more negative the relation, the more dealer j's profit is reduced by adverse selection. To construct a precise measure of expected losses due to adverse selection we compare the liquidity demander's actual trade with a hypothetical "informationless" trade; i.e., a trade based upon a valuation u_t that is independent of v_j and whose stochastic properties are otherwise identical to those of v_t. Holding all else equal, the realized loss perceived by dealer j in dollar profit (per unit of ETH reserves) associated with the hypothetical versus actual trade is proportional to

$$AS_t = \left[v_j \left(\frac{q_{t-1}}{u_t} \right)^{\frac{1}{2}} - v_j \left(\frac{q_{t-1}}{v_t} \right)^{\frac{1}{2}} \right]. \quad (27)$$

A stationary version of this measure can be obtained by normalizing by q_{t-1} to compute a *percentage* loss per unit of ETH reserves associated with adverse selection

$$AS^*_t = \left[\frac{v_j}{q_{t-1}} \left(\frac{q_{t-1}}{u_t} \right)^{\frac{1}{2}} - \frac{v_j}{q_{t-1}} \left(\frac{q_{t-1}}{v_t} \right)^{\frac{1}{2}} \right]. \quad (28)$$

Taking expectations at $t-1$ with respect to the objective distribution of the random variables (i.e., the

distribution that is relevant for empirical estimation) yields

$$\mathrm{E}\left[AS_t^*\right] = \mathrm{E}\left[\frac{v_j}{q_{t-1}}\right]\mathrm{E}\left[\left(\frac{q_{t-1}}{u_t}\right)^{\frac{1}{2}}\right] - \left(\mathrm{Cov}\left[\frac{v_j}{q_{t-1}}, \left(\frac{q_{t-1}}{v_t}\right)^{\frac{1}{2}}\right] + \mathrm{E}\left[\frac{v_j}{q_{t-1}}\right]\mathrm{E}\left[\left(\frac{q_{t-1}}{v_t}\right)^{\frac{1}{2}}\right]\right) \tag{29}$$

$$= -\mathrm{Cov}\left[\frac{v_j}{q_{t-1}}, \left(\frac{q_{t-1}}{v_t}\right)^{\frac{1}{2}}\right]. \tag{30}$$

The first equality follows from independence of $\tilde{v}_j$ and $\tilde{u}_t$. The non-covariance terms cancel because the univariate stochastic properties of $\tilde{v}_t$ and $\tilde{u}_t$ are identical. The expected loss to dealer j due to adverse selection is therefore proportional to the negative of the covariance between $\tilde{v}_j/q_{t-1}$ and $(q_{t-1}/\tilde{v}_t)^{\frac{1}{2}}$. The covariance will be larger and negative if the realizations of $\tilde{v}_t$ and $\tilde{v}_j$ are both high or both low in relation to the pre-trade reserve price. Thus, losses to adverse selection are greater if the liquidity demanders' valuations differ from the reserve price in the *same* way as the dealers' valuation differs from the reserve price.

This formula captures the idea of positive co-movement between $\tilde{v}_j/q_{t-1}$ and $\tilde{v}_t/q_{t-1}$. However, its specific functional form as implied by the design of the CPAMM protocol is not as simple as that. Instead, the correct formula (i.e., Equation (30)) is the negative of the covariance between the dealer's (relative) valuation and the square root of the inverse of the liquidity demander's corresponding valuation. We develop this formula further below.

1.6 Empirical Implications

1.6.1 Regression Tests

Equation (25) identifies the equilibrium supply of liquidity in the CPAMM. It connects liquidity provision on the CPAMM market (R_{ot}^*) to the liquidity environment on the LOB market (λ). The relation between those elements depends on all three prices—the LOB midquote and the expectation of the difference between the transaction and reserve prices on the AMM, and the relation is linear.

To specify a regression to test this association, rearrange Equation (25) with the variables in square brackets on the left-hand side

$$\overline{\mathrm{E}}[\tilde{S}_t^*] - \left(\overline{\mathrm{E}}[\tilde{P}_t^*] - v_o\right) = (2K/J)\ \lambda - (2/J)\ \lambda\ R_{ot}^*. \tag{31}$$

This equation is analogous to a production-based asset pricing model as in Cochrane (1991), where capital investment decisions are made based on agents' beliefs about expected return and the marginal cost of investment. Agents' first-order conditions can be rearranged with expected return on the left, and investment and marginal cost on the right. The same logic applies here: price spreads take the place of expected returns, reserve choices take the place of investment, and λ is the marginal cost of reserves.

Assuming that, on average, dealers have rational expectations, this equation can be regarded as making predictions about the coefficients of a regression involving variables that are observable to researchers:

$$S_t - (P_t - v_{ot}) = \beta_o + \beta_1\lambda_t + \beta_2 R_{ot} + \beta_3\lambda_t R_{ot} + \eta_t, \tag{32}$$

where Equation (31) predicts that $\beta_o = 0, \beta_1 > 0, \beta_2 = 0, \beta_3 < 0$. This regression is a joint test of the model and the hypothesis that the LOB market from which λ_t is estimated is the outside option that is relevant to dealers' decisions. If we estimate λ_t from Binance, and Binance is not the relevant outside option, then the price spreads will be unrelated to the λ_t estimates even if the model is generally correct in characterizing how dealers make decisions.

1.6.2 Estimating Adverse Selection Costs

Equation (30) provides a theoretical measure of adverse selection costs based on an unobservable v_j. This measure can be operationalized easily if the common-outlook assumption holds for the true data generating process—specifically, if $\mathrm{E}\left[\epsilon_{jt}\right] = 1$ for all j for the expectation operator of the empirical data generating process. This assumption serves a similar purpose as that of rational expectations of dealers for the regression specification above. Under this extension of the common-outlook assumption, Equation (30) can be written (see Appendix A) as

$$\mathrm{E}\left[AS_t^*\right] = -\mathrm{Cov}\left[\frac{v_t \epsilon_{jt}}{q_{t-1}}, \left(\frac{q_{t-1}}{v_t}\right)^{\frac{1}{2}}\right] \tag{33}$$

$$= \mathrm{Var}\left[\left(\frac{v_t}{q_{t-1}}\right)^{\frac{1}{2}}\right] \mathrm{E}\left[\left(\frac{q_{t-1}}{v_t}\right)^{\frac{1}{2}}\right] + \mathrm{E}\left[\left(\frac{v_t}{q_{t-1}}\right)^{\frac{1}{2}}\right] \left\{\mathrm{E}\left[\left(\frac{v_t}{q_{t-1}}\right)^{\frac{1}{2}}\right] \mathrm{E}\left[\left(\frac{q_{t-1}}{v_t}\right)^{\frac{1}{2}}\right] - 1\right\} \tag{34}$$

$$\approx \mathrm{Var}\left[\left(\frac{v_t}{q_{t-1}}\right)^{\frac{1}{2}}\right] \mathrm{E}\left[\left(\frac{q_{t-1}}{v_t}\right)^{\frac{1}{2}}\right]. \tag{35}$$

Equation (35) says that the cost of adverse selection can be approximated by the product of a variance of the square root of a return multiplied by the expectation of the square root of the inverse of that return. The justification for using Equation (35) to approximate the exact expression in Equation (34) is based on the observation that the trailing term in curly brackets is of the form $(\mathrm{E}[x]\mathrm{E}[1/x] - 1)$ where $x = v_t/q_{t-1}$. That term is non-zero (positive) only because of Jensen's inequality. The convexity in $1/x$ is stronger the closer is x to the origin, meaning when v_t/q_{t-1} is small or its inverse is large. Since the expectation of the inverse already multiplies the variance in the leading term, the leading term alone (i.e., Equation (35)) is likely to be a good simple empirical measure of adverse selection costs in CPAMM markets. Observations of v_t constructed from q_t and P_t via Equation (18) can be used to estimate the sample variance and average required in Equation (35); i.e.,

$$\widehat{\mathrm{E}\left[AS_t^*\right]} \approx \widehat{\mathrm{Var}}\left[\left(\frac{q_t + fP_t}{q_{t-1}}\right)^{\frac{1}{2}}\right] \widehat{\mathrm{E}}\left[\left(\frac{q_{t-1}}{q_t + fP_t}\right)^{\frac{1}{2}}\right]. \tag{36}$$

Nevertheless, in the tests that follow, we estimate the exact expression in Equation (34) using sample moments as just described for estimating Equation (36).

1.6.3 Estimating the Components of Liquidity Demander Valuations

The key state variable in the model is the valuation of the liquidity demander, v_t. It determines the transaction price, the subsequent reserve price, and under the common-outlook assumption, it is also an

unbiased forecast of the dealers' long-run valuations, v_j. An interesting feature of the CPAMM is that when traders optimize, the liquidity demander valuations are revealed by observable prices. By Equation (18), $v_t = q_t + fP_t$, with probability one, which means that the sum of the reserve price and the fee rate times the transaction price is a perfect proxy for v_t.

We can use this connection to estimate interesting features of the stochastic properties of v_t. In particular, we can decompose variation in v_t into permanent and temporary components to assess the degree to which liquidity demanders' valuations are driven by information or non-informational motives for trading. The magnitudes of these components will capture variation through time in traders' motives, and we would expect that larger permanent components are associated with greater adverse selection costs borne by dealers. That is, dealers are more likely to be on the losing sides of trades when information is the motive for liquidity demanders' trading.

Our estimation is based on the Hasbrouck (1991) and Hasbrouck (1993) approach to estimating the permanent and temporary components of equity price changes. We begin by assuming that liquidity demander valuations are described as follows:

$$v_t = \theta_t + s_t \tag{37}$$

$$\theta_t = \theta_{t-1} + e_t \tag{38}$$

where the e_t are iid draws from a distribution with mean zero and finite variance, and s_t is a mean-zero covariance stationary stochastic process. Although e_t and s_t are unobservable, Hasbrouck (1993) shows that their variances are identified by the parameters of a vector autoregression (VAR) involving first differences in v_t and observable variables that capture non-price dynamics in the trading process that are related to e_t and s_t. Signed order flow is conventionally used as such a variable because larger buy (sell) trades are associated with important positive (negative) news and tend also to temporarily drive prices above (below) security fundamental values.

The VAR that we estimate uses the series $q_t + fP_t$ as a perfect proxy for v_t and the AMM signed order quantities Q_{ot} to capture non-price dynamics in the trading process. Estimates of $\text{Var}[e_t]$ measure the information content of liquidity demanders' valuations and therefore the degree to which information motivates their trading. Estimates of $\text{Var}[s_t]$ measure the extent to which liquidity demanders' valuations vary for reasons that are transitory. Appendix B describes the details of the calculations using the one-step closed-form procedure given in George and Hwang (2001) for estimating $\text{Var}[e_t]$ and $\text{Var}[s_t]$ from the VAR parameters.

2 Data

To test the empirical predictions of our model, we use the most liquid crypto trading pair (ETH-USDT) and the most liquid CPAMM (Uniswap V2) and LOB (Binance) markets. All data start on May 20, 2020 and end on April 30, 2024. We obtain Binance trade data for the ETH-USDT pair through the Binance API. The trade data includes a list of all transactions for the pair, accompanied by a millisecond-level time stamp, with transactions signed from the liquidity demander's (trader's) perspective: a positive quantity if

the trader buys ETH with USDT, and negative if it sells ETH. The Binance price impact λ is defined as the sum of the permanent and temporary price impacts estimated from transactions data as in George and Khoja (2023). The advantage of their approach over fitting a simple regression of returns on signed trades is that it estimates λ while controlling for serial and cross-autocorrelations between returns and trades. Uniswap data are obtained from Dune Analytics.[10] The data include all swap transactions and liquidity pool additions (mint) and withdrawals (burn) and the order in which they are added to the chain.[11]

First, we cumulate the reserve pool additions and withdrawals through time, to construct a transaction-level time series of reserve levels.[12] For each swap transaction, we compute the transaction price as the ratio of the change in the quantities of USDT and ETH in the pool, net of fees, as depicted on the left-hand side of Equation (3). Because we use the Uniswap V2 smart contract, fees are always set at 30bps of the input token. Fee-adjusted signed slippage is computed as the difference between the fee-adjusted reserve price before the transaction, and the transaction price, as in Equation (24). We use the price of ETH on Binance at the time of the block production as the outside-option LOB value v_{ot}.

We remove anomalous transactions where we observe both inflows and outflows of the same token. We also remove transactions where the quantity is less than the allowed minimum fraction of ETH (1 Gwei = 1e-18 ETH) and USD (1e-6 dollars). Finally, we remove anomalous transactions where ETH transaction prices are below \$100, as ETH prices during the sample period ranged between \$231 and \$4,856. These filters drop about 1.2% of the transaction-level observations.

The model predicts that $v_t = q_t + fP_t$. We estimate the variances of the permanent shocks (e) and the stationary (s) component of v_t using the observable variable $q_t + fP_t$ and the quantity of ETH traded Q_{ot} (see Appendix B for details). The Uniswap adverse selection cost measure is estimated by the expression in Equation (34) using sample moments as described before Equation (36).

Finally, we aggregate the transaction-level data to block (12s) or hourly levels by taking averages of the fee-adjusted slippage, Uniswap transaction price, and Binance price; and sums of trading volume and signed quantities. For beginning and end ETH reserves, we use the reserves at the start and the end of the block or hour. The hourly aggregation matches the hourly frequency at which we estimate λ from Binance and the adverse selection cost measure from Uniswap using transaction-level data.

3 Descriptive Statistics

Table 1 reports descriptive statistics of the variables used in the empirical analysis. Each hour, an average of 177 transactions occur on Uniswap V2, trading 902 ETH, worth approximately \$1.8M. On average, about 20 times as many tokens trade on Binance. Uniswap ETH trading volume is typically small as a fraction of the reserves available to trade. Average volume for a *full hour* is only 1.6% (=902/56,413) of average ETH

[10] The SQL query can be found in Appendix D.

[11] For the purpose of these tests, we follow the approach in the model and abstract away from the fact that the order of transactions in a block is determined by the amount of fees paid by the traders. Thus, we treat the transactions as chronologically ordered, where each trader/dealer has visibility of the state of the chain before making a trading decision. In reality, arbitrageurs and searchers often pay higher fees to ex-post move up in the ordering within a block to front-run and sandwich transactions, or use other Maximal Extractable Value (MEV) strategies like liquidation and top of the block arbitrage.

[12] We reconcile calculated reserves against the "reserve" query from Dune Analytics which reports reserves at the block-level for each contract, ensuring that no liquidity flows are missed in calculating reserves.

reserves. The fraction is even smaller at 0.77% based on medians. By these measures, typical individual liquidity-demanding transactions draw upon only a tiny fraction of available liquidity.

The Uniswap reserve and transaction price distributions are nearly identical. Furthermore, Binance prices are similar to Uniswap prices, albeit higher by $8 (or 40bps) on average, with a 95th percentile price difference of less than $9.

The average price impact is larger on Uniswap than on Binance. The average hourly estimate of Binance λ is $0.0057 per ETH. The price impact estimate for Uniswap is $0.0352 per ETH.[13] Since the average value of a single ETH is approximately $2,000, the price impacts in both markets are very small as a percentage of the dollar values of trades. Both are less than one-half basis point, much smaller than equity trading price impacts, which range from one to 15 bps (see George and Khoja (2023)).

The estimates of $\sqrt{Var(e)}/\overline{q}$ imply that within an average hour, the standard deviation of the permanent component of transaction-level changes in liquidity-demander valuations is only 4bp of the hour's average reserve price. The standard deviation of the transitory component is only 2bp of the hour's average reserve price. The distributions of these estimates are right-skewed. The *medians* are smaller at 3bp for the permanent component, and 1bp for the temporary component. These estimates are larger than the price impact, suggesting that liquidity demanders capture a surplus to trading. However, the small magnitudes of these numbers suggest that liquidity providers' exposure to adverse selection is small when trading ETH and USDT.

The interpretation of the raw Uniswap adverse selection metric in Equation (34) is the expected dollar cost, as a percentage of the reserve price, per unit of ETH exposure that the dealer has to the reserve pool. The raw estimates in Table 1 are computed from transactions data and reported for a hypothetical exposure to 10,000 ETH of the reserve pool. Note that the average total reserve pool has 56,413 ETH. Therefore, the total expected cost of adverse selection in USDT, as a percentage of the reserve price, would be the number in the table multiplied by 5.641, or $0.0027 \times 5.641 = 0.0152$ percent (or 1.52bp) of the reserve price. In comparison, adverse selection costs average 2.42bp and 10.2bp respectively for large- and small-cap quintiles for NYSE stocks (see Hendershott et al. (2011)). The interpretation of the Uniswap estimate is that dealers lose very little to adverse selection, even in the hour in which the losses are at the maximum reported in the table, $0.0264 \times 5.641 = 14.9$bp. This is consistent with the observation above that liquidity providers' exposure to adverse selection is small when trading ETH and USDT.

Figure 1 displays the distributions over time of hourly (net) signed quantities traded on Binance and Uniswap, and Uniswap reserves. All three variables are highly volatile at the beginning of the sample. Uniswap reserves are larger than hourly net quantities traded across most of the sample, and Binance net quantities traded are larger and more volatile than Uniswap quantities for most of the sample period. Because the model's predictions should hold regardless of such changes over the sample period, we examine the full sample and three subperiods in the regression tests below.

Table 2 reports correlations between some of the most important variables. Most of the correlations are

[13]The price impact on Uniswap is the first derivative of Equation (3), which is equal to $P(Q_{ot})/(R_{ot} - Q_{ot})$. The ratio of average hourly Uniswap volume to the number of transactions per hour is an estimate of the typical transaction size Q_{ot} of 902/177 = 5.10 ETH. The average Uniswap price and average reserves are $1,987 and 56,413 ETH, respectively. Using these inputs, the Uniswap price impact at these averages is 1,987/(56,413 - 5.10) = 0.0352 USD per ETH.

below 0.5. There are some notable exceptions. The positive correlation of 0.608 between the square roots of Var(e) and Var(s) indicates that shocks to the permanent and temporary components of traders' valuations vary together over time. The correlations between those variables and the adverse selection measure indicate that the adverse selection is more strongly associated with the permanent component of changes in traders' valuations (0.738) than the temporary component (0.608). We examine this relation in regressions below. Finally, the correlation of 0.77 between the price impact on Binance (i.e., Lambda) and Uniswap signed slippage suggests that transaction prices on Binance are more sensitive to order flow during hours when Uniswap transactions cause prices to fall farther from the prevailing reserve price. This correlation indicates that liquidity is low in both markets in hours with greater selling pressure, suggesting that liquidity providers back away when selling dominates buying by liquidity demanders.

4 Empirical Tests

4.1 Liquidity Provision and Demand

The model predicts that liquidity providers set the size of reserves as a function of anticipated transaction characteristics (slippage, transaction price, and outside option price), which elicits trading that absorbs the supply of liquidity and leads to a post-transaction reserve price that is near to the trader's valuation v_t regardless of the size of reserves. In other words, liquidity demand absorbs variation in liquidity supply. Thus, we would expect that variation in liquidity supply predicts future variation in quantities traded, and not the reverse. We conduct two analyses to test this prediction: a vector autoregression and Granger causality tests.

We estimate the vector auto-regression with three lags, where the variables are Uniswap hourly trading volume (TV) and beginning reserves (BR).

$$TV_t = c_1 + a_{11,1}TV_{t-1} + a_{12,1}BR_{t-1} + a_{11,2}TV_{t-2} + a_{12,2}BR_{t-2} + a_{11,3}TV_{t-3} + a_{12,3}BR_{t-3} + \varepsilon_{TV,t} \quad (39)$$

$$BR_t = c_2 + a_{21,1}TV_{t-1} + a_{22,1}BR_{t-1} + a_{21,2}TV_{t-2} + a_{22,2}BR_{t-2} + a_{21,3}TV_{t-3} + a_{22,3}BR_{t-3} + \varepsilon_{BR,t} \quad (40)$$

The theory leads us to expect that $a_{12,1}$, $a_{12,2}$, and $a_{12,3}$ are positive and significant, and $a_{21,1}$, $a_{21,2}$, and $a_{21,3}$ are not statistically significant.

Table 3 presents the VAR results, fitted with dummies for each week to capture low-frequency temporal fixed effects. The coefficient estimates on own lags indicate that both variables exhibit significant persistence. The cross coefficient estimates suggest that trading volume during an hour is positively related to beginning-of-hour reserves at lags 1 and 3, while reserves are negatively related to volume at lag 2, and unrelated to volume at lags 1 and 3. Thus, the pattern in the coefficients is more consistent with volume responding positively to changes in reserves than reserves responding positively to changes in trading volume, a pattern that the Granger causality tests below confirm more formally.

The Granger causality tests use three lags to test whether Uniswap beginning reserves predict trading volume, and vice-versa. First, we test whether the time series of beginning reserves and trading volume are non-stationary, because the Granger causality test requires stationary time series. We use the Augmented

Dickey-Fuller test:

$$\Delta y_t = \alpha + \beta t + \gamma y_{t-1} + \sum_{i=1}^{p} \delta_i \Delta y_{t-i} + \varepsilon_t \tag{41}$$

$$H_0 : \gamma = 0 \quad \text{(unit root)} \qquad H_1 : \gamma < 0 \quad \text{(stationary)}$$

Panel A of Table 4 reports that γ is negative and significant for both beginning reserves and trading volume, rejecting the null hypothesis that the time series are non-stationary. For the Granger causality tests, we use the following equations:

$$TV_t = \alpha_0 + \sum_{i=1}^{p} \alpha_i TV_{t-i} + \sum_{i=0}^{p} \gamma_{BR,i} BR_{t-i} + \varepsilon_t \qquad BR_t = \beta_0 + \sum_{i=1}^{p} \beta_i BR_{t-i} + \sum_{i=1}^{p} \gamma_{TV,i} TV_{t-i} + \varepsilon_t \tag{42}$$

The null hypothesis of no Granger causality from BR_t to TV_t is:

$$H_0 : \gamma_{BR,0} = \gamma_{BR,1} = \cdots = \gamma_{BR,p} = 0,$$

and the null hypothesis of no Granger causality from TV_t to BR_t is:

$$H_0 : \gamma_{TV,1} = \gamma_{TV,2} = \cdots = \gamma_{TV,p} = 0.$$

The contemporaneous BR is included in the equation for TV because BR precedes TV in time; BR_t is reserves at the start of the period over which TV_t is measured. Panel B of Table 4 reports an F-value for the test of the joint significance of the $\gamma_{BR,i}$ that is large (119.7) and highly significant (p-value < 0.001). Thus, beginning reserves Granger cause trading volume in Uniswap. On the other hand, the F-value for the joint significance of the $\gamma_{TV,i}$ is 0.067, with a p-value of 0.796, indicating that we cannot reject the null of no Granger causality from trading volume to beginning reserves. The tests in both Panel A and Panel B are conducted with $p = 3$ lags.

The VAR and the Granger causality test support the prediction that trading volume responds to absorb changes in reserves (and not the reverse). This absorption effect is the mechanism by which prices can be invariant to liquidity supply in a CPAMM.

4.2 Equilibrium Reserves

In the model, equilibrium reserves are set by dealers optimizing their use of Uniswap and an alternative (LOB) market to offer their endowment for sale (and possibly to buy more in the alternative market to offer for sale in Uniswap). Equation (32) is the equilibrium implication of the dealers' optimization for market-level variables. We estimate this equation using OLS, where the dependent variable is the difference between the Uniswap fee-adjusted slippage, and the price difference between the Uniswap transaction price and the price of ETH on Binance. The independent variables are Binance price impact λ, Uniswap beginning reserves, and the interaction between the two. The model predicts a positive coefficient for λ, and a negative coefficient for the interaction.

Table 5 reports the coefficients of the OLS regression estimated at the hourly level for the full sample (column 1), and for three subsamples identified with early (column 2), middle (column 3), and late years (column 4) in our data. Our full sample covers May 20, 2020 to April 30, 2024. The middle years are identified with the "crypto winter" between November 20, 2021 and December 21, 2022 when crypto prices declined significantly.

For the full sample, the coefficients are consistent with the model's predictions: the coefficient on the Binance λ is positive and the coefficient on the interaction variable is negative. These results are generally robust in the sub-periods before, during, and after the 2022 crypto winter. Interestingly, the interaction coefficient is not significant in the latest sample period, possibly because the prominence of Binance as a reference outside option for Uniswap liquidity providers diminished. Binance's global market share of the spot market declined significantly from a peak of 60.3% at the end of the crypto winter to 38.3% at the end of the sample period.[14] An intercept and beginning reserves as a raw variable are included to be sure the interaction is properly captured by its coefficient. The significance of the intercept and beginning reserves in the regression are not predicted by the model, and their signs change across the subperiods.

By Equations (31) and (32), the regression coefficients on λ and the interaction variable have literal interpretations from the model as $\beta_1 = 2K/J$ and $\beta_3 = -2/J$, where J is the number of agents who potentially act as dealers and K is the aggregate endowment of ETH available for sale. We can solve for the K and J implied by the coefficient estimates as an additional check on the model's realism. We focus on the full-sample regression in Table 5. Because the variables in the regression are normalized by their standard deviations (and reserves are also divided by 10,000), we have to re-scale the estimated coefficients to recover β_1 and β_3 in Equation (32). In particular, $\beta_1 = (\sigma_y/\sigma_\lambda)\hat{\beta}_1$, where σ_y is the standard deviation of the dependent variable involving price spreads, σ_λ is the standard deviation of the λ regressor, and $\hat{\beta}_1$ is the estimated coefficient in Table 5. Likewise, $\beta_3 = (\sigma_y/(10,000\,\sigma_\lambda\sigma_R))\hat{\beta}_3$. Using the estimates of $\hat{\beta}_1$ and $\hat{\beta}_3$ from column 1 of Table 5, and standard deviations from Table 1, we have $\beta_1 = 502.43$ and β_3 = -1.96E-07. Solving for $J = -2/\beta_3 = 10.2$ million potential dealers, and $K = J\beta_1/2 = 2.56$ billion ETH available for sale.

The implied value of J is empirically plausible. In early 2024, the number of addresses with an ETH balance reached 100 million, with estimates that these represent about 50 million distinct individual holders. The implied value of K is too high, however. There were only 120 million ETH in existence in early 2024. This discrepancy implies that the overall level of liquidity provided during the sample period is much greater than our model would have predicted. This suggests that the model is missing a benefit that market participants perceive to providing liquidity or market participants overestimate the profitability of providing liquidity on Uniswap.[15]

The regressions in Table 5 are based on hourly-level data for all variables. We repeat these regressions in Table 6 using data at the individual block level. The λ variables are held constant across blocks within each hour. The results are very similar to the hourly estimates, and the inferences and qualitative conclusions are the same as well. The difference is that the interaction is statistically significant in all three subperiods.

[14] Source: The Block `https://www.theblock.co/data/crypto-markets/spot/the-block-legitimate-index-market-share`

[15] This finding is consistent with that of Cartea et al. (2024) who compare the performance of an optimized liquidity-provision strategy they derive to actual liquidity providers in ETH/USDC Uniswap V3 during the period January 1 - August 18 of 2022. The authors' strategy is profitable, but actual liquidity providers earn losses.

4.3 Adverse Selection

One of the main contributions of the model is to provide an empirical measure of expected adverse selection costs for liquidity providers that is specific to CPAMM markets, as found in Equation (34). To test whether this measure is related to the permanent (information-based) and temporary (liquidity-driven) components of changes in liquidity-demander valuations, we estimate a time-series OLS regression, where the dependent variable is the expected adverse selection cost, and the independent variables are the square roots of $Var(e)$ and $Var(s)$ divided by the average reserve price within the measurement hour of the variances. Taking square roots and dividing by the average reserve price puts these variables into units of USDT as a fraction of the reserve price to match the units of the dependent variable.

Table 7 reports the coefficients of OLS regressions for the full sample and subsamples as in the prior tables. The dependent variable is the raw adverse selection measure (USDT expected cost per unit of ETH exposure as a fraction of the reserve price) scaled by the sample average of beginning ETH reserves (56,413 for the full sample). The regressions are estimated with and without lags of the adverse selection measure.

In regressions without lags and across all sample periods, the coefficients of the $Var(e)$ and $Var(s)$ variables are positive and statistically significant, and the magnitude of the coefficient of $Var(e)$ is three times or more that of the $Var(s)$ variable. The interpretation is that adverse selection is greater at times when the permanent (information) component of changes in traders' valuations is greater. The significance of the $Var(s)$ variable is not surprising given the high correlation between the $Var(e)$ and $Var(s)$ variables reported in Table 2. The lesser importance of the $Var(s)$ variable is consistent with the adverse selection measure being more sensitive to permanent changes in traders' valuations. The adjusted R-square for the full sample is 0.57, implying that a majority of the variation in the adverse selection measure is explained by variation in the components of traders' valuations.

In the regressions with lags, adverse selection costs are persistent, with a half life of about 4 hours in the full sample. The persistence is strongest in the early part of the sample, and weakens in the latter part of the sample. More importantly, adverse selection costs are positively and statistically significantly related to the permanent component of traders' valuations in the full sample, as well as across all sub-sample periods. The transitory component is also significant, but its coefficient is half or less in magnitude. Viewing the lags as capturing an autoregressive forecast of adverse selection, these regressions indicate that unexpected variation in adverse selection from one hour to the next is related to the permanent component of changes in traders' valuations more strongly than the temporary component.

Taken together, the regression results are strongly consistent with the notion that our new measure of adverse selection in AMM markets reflects dealers' costs to providing liquidity to informed traders.

5 Conclusion

We model the strategic interaction between agents who supply and demand liquidity via a constant-product automated market maker. Our model is free of parametric assumptions about the distributions of traders' valuations and also free of structural assumptions about the quantities they can trade. Our approach

leads to some new insights about market quality that are surprising given the simplicity of the CPAMM protocol. In particular, regardless of how many dealers participate, dealers act as if they are in a perfectly competitive market. We also show that the CPAMM generates equilibrium prices that are invariant to the quantity of liquidity supplied. A simple relation describes liquidity supply in terms of observable prices and the price impact in the outside-option market. Furthermore, the CPAMM discovers prices well in the sense that token values implied by reserves in the pool track closely traders' valuations. The sequence of traders' valuations can be recovered from the time series of reserve and transaction prices. This finding enables us to decompose the variance of valuations into permanent (information-related) and temporary (liquidity-related) components. Finally, the model also offers a simple measure of the adverse selection costs that dealers expect from traders. We then test the predictions. Consistent with the model, we find that beginning ETH reserves Granger predict trading volume on Uniswap, but not vice-versa. Furthermore, the levels of reserves on Uniswap are consistent with the modeled relationship between slippage in Uniswap, Binance price impact, and the price difference between the two markets. Finally, the novel measure of adverse selection is more strongly related to the permanent component of variation in traders' valuations than the temporary component, suggesting that information-driven trading reduces dealers' profits. As tokenized equities begin to trade on CPAMM platforms, the tools developed here may prove useful for assessing market quality in these emerging markets.

Table 1: Summary Statistics

Summary statistics computed across hours are reported for key variables aggregated within the hour. Trade volume represents the total unsigned trade quantities (buys and sells), whereas trade quantity is the sum of signed quantities, within the hour. Beginning ETH reserves are the starting reserves of each hour. Quantity variables are in units of ETH. The reserve price is calculated as the ratio of USDT reserves to ETH reserves at the beginning of each hour. Binance and Uniswap transaction prices are averaged within the hour. Price variables are in units of USDT. The reserve price return is the percentage change in hourly beginning reserve prices. Lambda is estimated as described in George and Khoja (2023) each hour for Binance using millisecond-level transactions data. Lambda is in units of USDT per ETH. Fee-adjusted slippage is the signed difference between the grossed-up pre-transaction reserve price and the transaction price for each Uniswap trade, averaged within the hour. Price difference is the spread between the Uniswap transaction price and the average of the Binance transaction prices with the same second-level time stamp, averaged within the hour. $Var(e)$ and $Var(s)$ are estimated each hour using Uniswap transactions data as described in Appendix B. The distributions reported below are of the square-root of the hourly estimate divided by the average of the pre-transaction reserve prices within the hour. The Uniswap adverse selection measure is estimated from Uniswap transactions data as in Equation (34) using sample moments as described above Equation (36). Variables labeled with (Win) are winsorized at the 5th and 95th percentiles. The sample period is May 20, 2020 to April 30, 2024.

Statistic	N	Mean	Median	SD	P5	P95	Min	Max
Number of Uniswap Transactions	34,271	177	122	140	42	461	1	1,600
Uniswap Trade Volume	34,271	902	175	2,435	19	3,899	0	154,442
Binance Trade Volume	34,607	19,980	15,072	17,129	4,542	51,776	117	214,298
Uniswap Trade Quantity	34,271	3	1	483	-323	342	-13,487	14,359
Binance Trade Quantity	34,607	-107	-55	3,407	-4,714	4,373	-54,758	38,873
Beginning ETH Reserves (Win)	34,271	56,413	22,592	120,703	11,018	224,565	11,018	631,516
Reserve Price	34,271	1,987	1,837	1,046	350	3,941	196	4,856
Binance Price	34,607	1,969	1,832	1,055	323	3,935	195	4,845
Uniswap Transaction Price	34,271	1,987	1,837	1,045	350	3,940	198	4,855
Reserve Price Return	34,270	0.0001	0.0002	0.0089	-0.0132	0.0132	-0.1778	0.0945
Lambda (Win)	33,902	0.0057	0.0035	0.0055	0.0006	0.0197	0.0006	0.0205
Fee Adjusted Slippage (Win)	34,271	5.6125	4.9518	3.2374	1.1627	11.6564	1.0741	11.8624
Price Difference (Win)	34,110	0.1719	0.1251	2.3436	-3.7405	4.1284	-8.3180	8.7676
Fee Adjusted slippage (Win) - Price Diff (Win)	34,110	5.4573	4.6529	3.9182	0.6016	13.1315	-5.8965	19.9531
$\sqrt{Var(e)}/\bar{q}$ (Win)	27,520	0.0004	0.0003	0.0004	0.0001	0.0011	0.0000	0.0106
$\sqrt{Var(s)}/\bar{q}$ (Win)	27,520	0.0002	0.0001	0.0004	0.0000	0.0009	0.0000	0.0084
Uniswap Adverse Selection (Win) x 10,000	34,045	0.0027	0.0005	0.0062	0.0000	0.0264	0.0000	0.0264
ETH Reserves (Win) x Lambda (Win)	33,902	199.2748	102.8546	213.6370	9.0723	605.6565	6.3852	1464.0210

Table 2: Correlations

Correlations are computed from hourly values of key variables. Beginning ETH reserves are the starting reserves of each hour. Lambda is estimated as described in George and Khoja (2023) each hour for Binance using millisecond-level transactions data. Fee-adjusted slippage is the signed difference between the grossed-up pre-transaction reserve price and the transaction price for each Uniswap trade, averaged within the hour. Price difference is the spread between the Uniswap transaction price and the average of the Binance transaction prices with the same second-level time stamp, averaged within the hour. $Var(e)$ and $Var(s)$ are estimated each hour using Uniswap transactions data as described in Appendix B then divided by the average of the pre-transaction reserve prices within the hour. The Uniswap adverse selection measure is estimated from Uniswap transactions data as in Equation (34) using sample moments as described above Equation (36). Variables labeled (Win) are winsorized at the 5th and 95th percentiles.The sample period is May 20, 2020 to April 30, 2024.

	ETH Beg Reserves	Lambda (Win)	Fee-Adj Signed Slippage (Win)	Price Difference (Win)	$\sqrt{Var(e)}/\bar{q}$ (Win)	$\sqrt{Var(s)}/\bar{q}$ (Win)	Adverse Selection x 10k (Win)
ETH Beg Reserves	1.000	-0.2179	-0.37	-0.0349	-0.106	-0.058	-0.090
Lambda (Win)	-0.218	1.0000	0.77	-0.0038	0.117	0.012	0.073
Fee-Adj Signed Slippage (Win)	-0.370	0.7711	1.00	0.0402	-0.115	-0.115	-0.109
Price Difference (Win)	-0.035	-0.0038	0.04	1.0000	-0.035	-0.024	-0.026
$\sqrt{Var(e)}/\bar{q}$ (Win)	-0.106	0.1165	-0.12	-0.0352	1.000	0.608	0.738
$\sqrt{Var(s)}/\bar{q}$ (Win)	-0.058	0.0117	-0.12	-0.0238	0.608	1.000	0.567
Adverse Selection x 10,000 (Win)	-0.090	0.0734	-0.11	-0.0264	0.738	0.567	1.000

Table 3: VAR Results - Reserves and Trade Volume

Coefficient estimates of a vector autoregression involving Uniswap trading volume (TV) and the beginning ETH reserves (BR) of each hour. Weekly dummies are included (but not reported) to account for temporal fixed effects. Uniswap volume is the sum of unsigned quantities traded (buys and sells) within the hour. The sample period is May 20, 2020 to April 30, 2024.

Variable	Estimate	StdError	t-stat	p Value
TV : TV Lag 1	0.2672***	0.0066	40.2915	0.0000
TV : BR Lag 1	0.0061**	0.0028	2.1663	0.0303
TV : TV Lag 2	0.0597***	0.0069	8.7138	0.0000
TV : BR Lag 2	-0.0053	0.0036	-1.4573	0.1450
TV : TV Lag3	0.0574***	0.0066	8.6537	0.0000
TV : BR Lag3	0.0062**	0.0028	2.2103	0.0271
TV : Constant	0.0013	12.8058	0.0001	0.9999
BR : TV Lag 1	0.0100	0.0158	0.6373	0.5239
BR : BR Lag 1	0.8451***	0.0066	127.2343	0.0000
BR : TV Lag 2	-0.0550***	0.0163	-3.3763	0.0007
BR : BR Lag 2	0.1401***	0.0086	16.2108	0.0000
BR : TV Lag3	0.0087	0.0158	0.5499	0.5824
BR : BR Lag3	-0.0121*	0.0066	-1.8239	0.0682
BR : Constant	0.0387	30.4538	0.0013	0.9990

Table 4: Augmented Dickey Fuller and Granger Causality Tests - Liquidity Demand and Supply

Uniswap trading volume (TV) is the sum of unsigned quantities traded (buys and sells) within the hour. ETH beginning reserves (BR) are the starting reserves for the hour. The Augmented Dickey-Fuller and Granger Causality tests are conducted using three lags of each variable. The $BR \rightarrow TV$ test includes contemporaneous BR as a predictor of hourly TV because BR is determined at the start of the hour.The sample period is May 20, 2020 to April 30, 2024.

Test	Test Statistic	P-Value	Comment
Panel A: Stationarity Tests			
BR ADF Test	-19.916	0.010	Non-Stationarity Rejected
TV ADF Test	-26.698	0.010	Non-Stationarity Rejected
Panel B: Granger Causality Tests			
$BR \rightarrow TV$ (Contemporaneous)	114.384	<0.001	Beg Reserves Granger Causes Uniswap Volume
$TV \rightarrow BR$ (Standard Lag)	0.067	0.796	Uniswap Volume Does not Granger Cause Reserves

Table 5: Equilibrium Reserves (Hourly Level)

The dependent variable is the difference between fee-adjusted slippage and the price difference between Binance and Uniswap, as depicted in Equation (32). Fee-adjusted slippage is the signed difference between the grossed-up beginning reserve price and the transaction price for each Uniswap trade, averaged within the hour. Price difference is the spread between the Uniswap transaction price and the Binance traded price at the time of the Uniswap transaction, averaged within the hour. ETH reserves are the starting reserves of each hour. Reserves are divided by 10,000 to represent a unit of 10,000 ETHs. Lambda is estimated as described in George and Khoja (2023) each hour for Binance using millisecond-level transactions data. All variables are winsorized at the 5th and 95th percentiles then divided by their respective standard deviations. The Pre-Winter period spans from the inception of ETH-USDT trading until November 10, 2021, when Bitcoin prices reached their peak. The Winter period begins at this peak and extends to Bitcoin's lowest point on December 21, 2022. This period includes major market events such as the Terra Luna crisis in May 2022 and the collapse of FTX in November 2022. The Post-Winter period starts from Bitcoin's low on December 21, 2022, and continues until the end of the sample period on April 30, 2024. Standard errors are robust to heteroskedasticity.

	(1)	(2)	(3)	(4)
	Full Sample	Pre Winter	Winter	Post Winter
	5/20/2020-4/30/2024	5/20/2020-11/10/2021	11/21/2021-12/21/2022	12/22/2022-4/30/2024
Lambda (Win)	0.7075***	0.7277***	0.5176***	0.7777***
	(0.0088)	(0.0119)	(0.0451)	(0.1013)
Beg Res (Win)	-0.0756***	-0.0421***	11.7549***	4.8874***
	(0.0037)	(0.0041)	(0.2585)	(0.3573)
Lambda (Win) × Beg Res (Win)	-0.3331***	-0.1998***	-1.2624***	-0.1455
	(0.0246)	(0.0323)	(0.2228)	(0.6208)
Constant	0.8005***	0.5499***	-0.6014***	0.0925*
	(0.0061)	(0.0099)	(0.0344)	(0.0545)
Observations	33,768	12,303	9,663	11,802
R^2	0.4324	0.6176	0.5875	0.2916

Standard Errors in parentheses are robust to heteroskedasticity.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Equilibrium Reserves (Block Level)

The dependent variable is the difference between fee-adjusted slippage and the price difference between Binance and Uniswap, as depicted in Equation (32). Fee-adjusted slippage is the signed difference between the grossed-up beginning reserve price and the transaction price for each Uniswap trade, averaged within the block. Price difference is the spread between the average Uniswap transaction prices within the block and the average Binance transaction prices that occur within the time period covered by transactions within the block. ETH reserves are the starting reserves of each block. Reserves are divided by 10,000 to represent a unit of 10,000 ETHs. Lambda is estimated as described in George and Khoja (2023) each hour for Binance using millisecond-level transactions data. The Lambda variable is constant across blocks within each hour. All variables are winsorized at the 5th and 95th percentiles then divided by their respective standard deviations. The Pre-Winter period spans from the inception of ETH-USDT trading until November 10, 2021, when Bitcoin prices reached their peak. The Winter period begins at this peak and extends to Bitcoin's lowest point on December 21, 2022. This period includes major market events such as the Terra Luna crisis in May 2022 and the collapse of FTX in November 2022. The Post-Winter period starts from Bitcoin's low on December 21, 2022, and continues until the end of the sample period on April 30, 2024. Standard errors are clustered by hour, allowing disturbances within each hour during which Lambda is constant to be correlated, but uncorrelated with disturbances in other hours.

	(1)	(2)	(3)	(4)
	Full Sample	Pre Winter	Winter	Post Winter
	5/20/2020-4/30/2024	5/20/2020-11/10/2021	11/21/2021-12/21/2022	12/22/2022-4/30/2024
Lambda (Win)	0.5167***	0.5065***	0.3773***	0.8233***
	(0.0066)	(0.0066)	(0.0066)	(0.0066)
Beg Res (Win)	-0.0835***	-0.0711***	11.5418***	5.6709***
	(0.0035)	(0.0035)	(0.0035)	(0.0035)
Lambda (Win) × Beg Res (Win)	-0.4603***	-0.3183***	-1.2657***	-1.9981***
	(0.0234)	(0.0234)	(0.0234)	(0.0234)
Constant	0.6718***	0.5616***	-0.4678***	-0.0394***
	(0.0053)	(0.0053)	(0.0053)	(0.0053)
Observations	3,203,039	1,654,231	586,948	961,860
R^2	0.2291	0.2934	0.3306	0.1509

Standard errors in parentheses are clustered by hour.

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 7: Adverse Selection Regression with Newey-West Standard Errors

The dependent variable is the Uniswap adverse selection measure, estimated within each hour from Uniswap transactions data as in Equation (34) using sample moments as described above Equation (36). $Var(e)$ and $Var(s)$ are estimated each hour using Uniswap transactions data as described in Appendix B. The square-root of the hourly estimate is divided by the average of the pre-transaction reserve prices within the hour to align its units with the units of the adverse selection measure. All variables are winsorized at the 5th and 95th percentiles. Pre-Winter period spans from the inception of ETH-USDT trading until November 10, 2021, when Bitcoin prices reached their peak. The Winter period begins at this peak and extends to Bitcoin's lowest point on December 21, 2022. This period includes major market events such as the Terra Luna crisis in May 2022 and the collapse of FTX in November 2022. The Post-Winter period starts from Bitcoin's low on December 21, 2022, and continues until the end of the sample period on April 30, 2024.

	Full Sample		Pre-Winter		Winter		Post-Winter	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\sqrt{Var(e)}/\overline{q}$ (Win)	6.908***	4.721***	7.630***	4.905***	6.590***	5.392***	4.416***	4.070***
	(0.209)	(0.155)	(0.399)	(0.328)	(0.242)	(0.243)	(0.327)	(0.297)
$\sqrt{Var(s)}/\overline{q}$ (Win)	2.163***	2.194***	2.995***	2.680***	1.367***	1.634***	1.191***	1.237***
	(0.165)	(0.137)	(0.264)	(0.238)	(0.262)	(0.251)	(0.269)	(0.263)
L.Adv Sel		0.1736***		0.1845***		0.1237***		0.1335***
		(0.0114)		(0.0171)		(0.0191)		(0.0376)
L2.Adv Sel		0.1108***		0.1164***		0.0826***		0.0257
		(0.0101)		(0.0154)		(0.0170)		(0.0229)
L3.Adv Sel		0.0843***		0.0937***		0.0459***		0.0226
		(0.0090)		(0.0144)		(0.0145)		(0.0203)
Constant	−0.001494***	−0.001281***	−0.001849***	−0.001482***	−0.001504***	−0.001553***	−0.000674***	−0.000703***
	(0.0000735)	(0.0000458)	(0.0001880)	(0.0001220)	(0.0000973)	(0.0000853)	(0.0000646)	(0.0000654)
R-squared	0.5670	0.6300	0.6660	0.7347	0.4229	0.4530	0.3461	0.3668
N	27,520	27,517	10,363	10,360	7,325	7,325	9,832	9,832

Newey-West standard errors (lag = 3) in parentheses. Odd columns (1,3,5,7): without lags. Even columns (2,4,6,8): with lags. * p<0.10, ** p<0.05, *** p<0.01

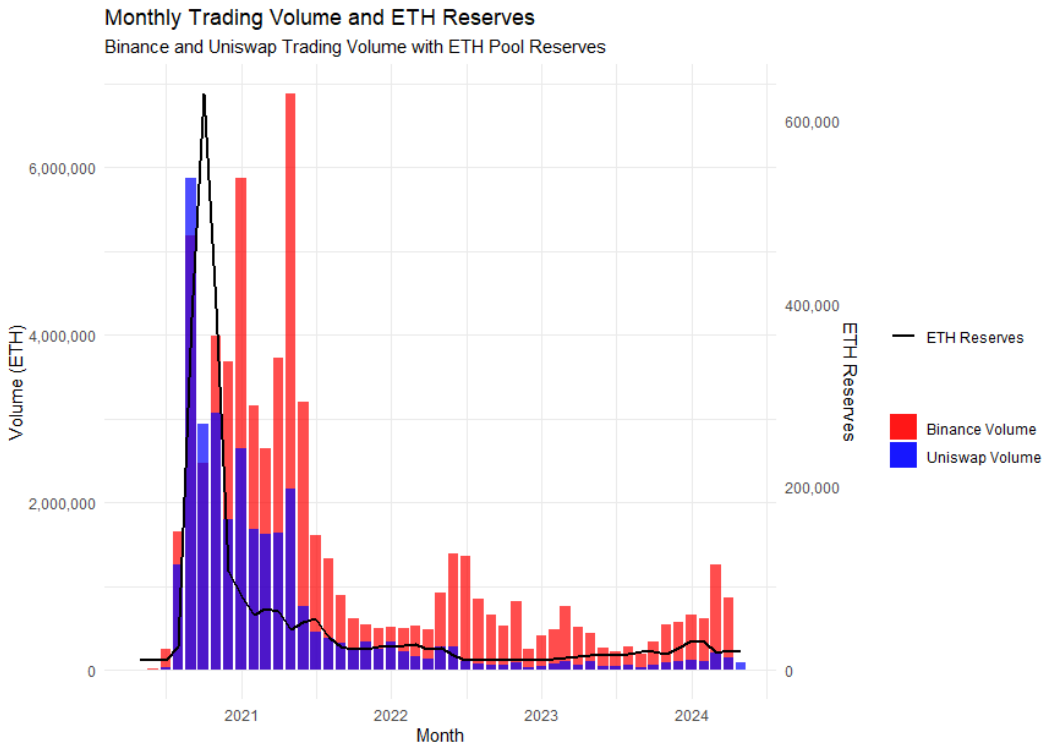

Figure 1: Hourly trading metrics showing Binance Volume, Uniswap Volume, and Uniswap beginning reserves over time.

6 Appendix A

Derivation of Equation (34): Substituting for $v_t\epsilon_{jt}$ for v_j in Equation (30) in the text, and recalling that $\mathrm{E}[\epsilon_{jt}] = 1$,

$$\mathrm{E}\left[AS_t^*\right] = -\mathrm{Cov}\left[\frac{v_t\epsilon_{jt}}{q_{t-1}}, \left(\frac{q_{t-1}}{v_t}\right)^{\frac{1}{2}}\right] \tag{43}$$

$$= -\mathrm{E}\left[\frac{v_t\epsilon_{jt}}{q_{t-1}^{1/2}v_t^{1/2}}\right] + \mathrm{E}\left[\frac{v_t\epsilon_{jt}}{q_{t-1}}\right]\mathrm{E}\left[\left(\frac{q_{t-1}}{v_t}\right)^{\frac{1}{2}}\right] \tag{44}$$

$$= -\mathrm{E}\left[\left(\frac{v_t}{q_{t-1}}\right)^{\frac{1}{2}}\right]\mathrm{E}\left[\epsilon_{jt}\right] + \mathrm{E}\left[\frac{v_t}{q_{t-1}}\right]\mathrm{E}\left[\epsilon_{jt}\right]\mathrm{E}\left[\left(\frac{q_{t-1}}{v_t}\right)^{\frac{1}{2}}\right] \tag{45}$$

$$= \mathrm{E}\left[\frac{v_t}{q_{t-1}}\right]\mathrm{E}\left[\left(\frac{q_{t-1}}{v_t}\right)^{\frac{1}{2}}\right] - \mathrm{E}\left[\left(\frac{v_t}{q_{t-1}}\right)^{\frac{1}{2}}\right] \tag{46}$$

$$= \left(\mathrm{Var}\left[\left(\frac{v_t}{q_{t-1}}\right)^{\frac{1}{2}}\right] + \mathrm{E}\left[\left(\frac{v_t}{q_{t-1}}\right)^{\frac{1}{2}}\right]^2\right)\mathrm{E}\left[\left(\frac{q_{t-1}}{v_t}\right)^{\frac{1}{2}}\right] - \mathrm{E}\left[\left(\frac{v_t}{q_{t-1}}\right)^{\frac{1}{2}}\right] \tag{47}$$

$$= \mathrm{Var}\left[\left(\frac{v_t}{q_{t-1}}\right)^{\frac{1}{2}}\right]\mathrm{E}\left[\left(\frac{q_{t-1}}{v_t}\right)^{\frac{1}{2}}\right] + \mathrm{E}\left[\left(\frac{v_t}{q_{t-1}}\right)^{\frac{1}{2}}\right]\left\{\mathrm{E}\left[\left(\frac{v_t}{q_{t-1}}\right)^{\frac{1}{2}}\right]\mathrm{E}\left[\left(\frac{q_{t-1}}{v_t}\right)^{\frac{1}{2}}\right] - 1\right\} \tag{48}$$

7 Appendix B

The model predicts that $v_t = q_t + fP_t$, so we estimate the components of v_t using the observable variable $q_t + fP_t$ and the quantity of ETH traded Q_{ot}. We define the first differences $r_t \equiv (q_t + fP_t) - (q_{t-1} + fP_{t-1})$, and we assume the vector $\mathbf{x}_t \equiv (r_t, Q_{ot})'$, has a finite (say, pth-order) auto-regressive representation:

$$A\mathbf{x}_t = B_1\mathbf{x}_{t-1} + \cdots + B_p\mathbf{x}_{t-p} + \mathbf{u}_t \qquad \mathbf{u}_t \sim iid(0\,,\,\Omega). \tag{49}$$

We stack the lags of $\mathbf{x}_t$ and redefine the parameter matrices in the standard way [see Appendix B in George and Hwang (2001)] so that the VAR(p) can be written as a VAR(1):

$$A_*\mathbf{x}_{*t} = B_*\mathbf{x}_{*t-1} + \mathbf{u}_{*t}, \qquad \text{where } \mathbf{u}_{*t} \sim iid(\mathbf{0}\,,\,\Omega_*) \tag{50}$$

A_*, B_* and Ω_* are $(2p \times 2p)$ matrices, and $\mathbf{x}_{*t}$ and $\mathbf{u}_{*t}$ are $(2p \times 1)$. Proposition 1 below adapts Theorem 1 in George and Hwang (2001) to our simpler setting. The proposition provides closed-form expressions for the variances of ϵ_t and s_t in Equations (37) and (38) in the text in terms of the auto-regressive parameters contained in the matrices A_*, B_* and Ω_* of equation (50).

Proposition 1. *With reference to equation (50), suppose $\mathbf{x}_{*t}$ and $\mathbf{u}_{*t}$ are $(2p \times 1)$. Let e_i be the $(2p \times 1)$ vector with unity in the ith position and zeros elsewhere, and define the $(2p \times 2)$ matrix $H \equiv (e_1 \quad e_2)$. If A_*^{-1} exists and the eigenvalues of $A_*^{-1}B_*$ lie inside the unit circle, then the variance of the random-walk innovation (denoted by e_t in Equation (38) and ϵ_t here) can be written as*

$$Var[\epsilon_t] = e_1'\Lambda_{\{1;1\}}e_1 + 2e_1'\Lambda_{\{2;1\}}e_1 + e_1'\Lambda_{\{2;2\}}e_1, \tag{51}$$

and the variance of the stationary component can be written as

$$Var[s_t] = e_1'\left(\mathcal{S} - \Lambda_{\{12;12\}}\right)e_1, \tag{52}$$

where

$$Vec[\mathcal{S}] = \left[I - \Theta \otimes \Theta\right]^{-1} Vec[\Lambda_{\{12;12\}}], \qquad \Theta = A_*^{-1}B_*,$$

$$\Lambda_{\{\alpha_1;\alpha_2\}} = \left\{A_* - B_*\right\}^{-1} H \left(I_{\{\alpha_1\}}\,\Omega\, I_{\{\alpha_2\}}\right) H' \left\{(A_* - B_*)^{-1}\right\}',$$

$$\Omega = H'\Omega_* H, \qquad \Omega_* = E[\mathbf{u}_{*t}\mathbf{u}_{*t}'],$$

and $I_{\{\alpha_j\}}$ is the 2×2 matrix with unity in the combination of diagonal positions specified in the vector α_j and zeros elsewhere (e.g., $I_{\{2\}}$ has unity in the second diagonal, whereas $I_{\{12\}}$ is the 2×2 identity matrix).

8 Appendix C

This appendix contains two results. First, we derive a general condition that the price function associated with a bonding curve must satisfy in order for the equilbrium transaction price to be invariant to liquidity supply. This result follows from an application of the implicit function theorem to the liquidity demander's first-order condition. This condition is quite specific and is not likely to be satisfied by a large class of functions. Second, we show that the price function associated with the constant product bonding curve satisfies this condition uniformly in the range of possible quantities of liquidity supplied and demanded.

Let $P(Q_{ot}, R_{ot})$ denote the transaction price function associated with a bonding curve. It suffices to write this function in terms of these arguments only because this price is Q_{1t} divided by Q_{ot}, and R_{1t} is determined by R_{ot} by the constraint $R_{1t} = q_{t-1}R_{ot}$.

Liquidity demander t chooses Q_{ot} to solve

$$\max_{Q_{ot}} (v_t - P(Q_{ot}, R_{ot}))Q_{ot}. \tag{53}$$

The first-order condition is

$$v_t - P(Q^*_{ot}, R_{ot}) - \frac{\partial P(Q^*_{ot}, R_{ot})}{\partial Q_{ot}} \; Q^*_{ot} = 0, \tag{54}$$

where Q^*_{ot} denotes the liquidity demander's optimal choice. The second-order condition is

$$-2\frac{\partial P}{\partial Q_{ot}} - \frac{\partial^2 P}{\partial Q^2_{ot}} \; Q^*_{ot} < 0, \tag{55}$$

which says that the AMM price function cannot be too concave.

Proposition 2. *The equilibrium transaction price is invariant to liquidity supplied if and only if*

$$\frac{\frac{\partial}{\partial R_{ot}} \; P(Q_{ot}, R_{ot})}{\frac{\partial}{\partial R_{ot}} \left[\frac{\partial P(Q_{ot}, R_{ot})}{\partial Q_{ot}} \; Q_{o,t} \right]} = \frac{\frac{\partial}{\partial Q_{ot}} \; P(Q_{ot}, R_{ot})}{\frac{\partial}{\partial Q_{ot}} \left[\frac{\partial P(Q_{ot}, R_{ot})}{\partial Q_{ot}} \; Q_{o,t} \right]}, \tag{56}$$

where all functions are evaluated at the liquidity demander's optimal choice Q^*_{ot}.

Proof. The trader's optimal choice Q^*_{ot} depends on R_{ot} through the price function. We denote this dependence by writing $Q^*_{ot}(R_{ot})$. The equilibrium price is invariant to R_{ot} if and only if the total derivative of $P(Q^*_{ot}(R_{ot}), R_{ot})$ is equal to zero:

$$\frac{d}{dR_{ot}} P(\cdot) = \frac{\partial P}{\partial Q_{ot}} \frac{\partial Q^*_{ot}}{\partial R_{ot}} + \frac{\partial P}{\partial R_{ot}} = 0 \tag{57}$$

when evaluated at Q^*_{ot}. Or, equivalently,

$$\frac{\partial P(Q^*_{ot}, R_{ot})}{\partial R_{ot}} = -\frac{\partial P(Q^*_{ot}, R_{ot})}{\partial Q_{ot}} \; \frac{\partial Q^*_{ot}(R_{ot})}{\partial R_{ot}}. \tag{58}$$

The first-order condition, Equation (54), is an identity in R_{ot}. Totally differentiating it with respect to R_{ot}

and rearranging yields

$$\frac{\partial Q_{ot}^*}{\partial R_{ot}} \left\{ 2\frac{\partial P}{\partial Q_{ot}} + \frac{\partial^2 P}{\partial Q_{ot}^2}\ Q_{ot}^* \right\} + \frac{\partial P}{\partial R_{ot}} + \frac{\partial^2 P}{\partial R_{ot} \partial Q_{ot}}\ Q_{ot}^* = 0, \tag{59}$$

where the functions are evaluated at Q_{ot}^*. The quantity in curly brackets is the second-order condition. Substituting from Equation (59) into Equation (58) and rearranging yields

$$\frac{\frac{\partial P}{\partial R_{ot}}}{\frac{\partial P}{\partial Q_{ot}}} = \frac{\frac{\partial P}{\partial R_{ot}} + \frac{\partial^2 P}{\partial R_{ot} \partial Q_{ot}} Q_{ot}^*}{2\frac{\partial P}{\partial Q_{ot}} + \frac{\partial^2}{\partial Q_{ot}^2} Q_{ot}^*}. \tag{60}$$

To simplify further, let $K \equiv \frac{\frac{\partial P}{\partial R_{ot}}}{\frac{\partial P}{\partial Q_{ot}}}$. Then the preceding equation can be written as

$$\left\{ \frac{\partial P}{\partial Q_{ot}} + \frac{\partial P}{\partial Q_{ot}} + \frac{\partial^2 P}{\partial Q_{ot}^2} Q_{ot}^* \right\} K = \frac{\partial P}{\partial Q_{ot}}\ K + \frac{\partial^2 P}{\partial R_{ot} \partial Q_{ot}}\ Q_{ot}^* \tag{61}$$

or

$$K \equiv \frac{\frac{\partial P}{\partial R_{ot}}}{\frac{\partial P}{\partial Q_{ot}}} = \frac{\frac{\partial^2 P}{\partial R_{ot} \partial Q_{ot}}\ Q_{ot}^*}{\frac{\partial P}{\partial Q_{ot}} + \frac{\partial^2 P}{\partial Q_{ot}^2}\ Q_{ot}^*} = \frac{\frac{\partial}{\partial R_{ot}} \left[\frac{\partial P}{\partial Q_{ot}} Q_{ot} \right] \Big|_{Q_{ot}^*}}{\frac{\partial}{\partial Q_{ot}} \left[\frac{\partial P}{\partial Q_{ot}} Q_{ot} \right] \Big|_{Q_{ot}^*}}, \tag{62}$$

which is equivalent to the expression in the proposition. □

Proposition 3. *The price function associated with the constant product bonding curve satisfies the expression in Proposition 2 uniformly in Q_{ot} and R_{ot}.*

Proof. Straightforward to verify by direct calculation. □

9 Appendix D

9.1 Mint Transactions

```
SELECT *
FROM uniswap_v2_ethereum.Pair_evt_Mint
WHERE contract_address = 0x0d4a11d5eeaac28ec3f61d100daf4d40471f1852
AND evt_block_time > TIMESTAMP '2020-05-19 11:59'
ORDER BY evt_block_time ASC
;
```

9.2 Burn (Withdrawals) Transactions

```
SELECT *
FROM uniswap_v2_ethereum.Pair_evt_Burn
WHERE contract_address = 0x0d4a11d5eeaac28ec3f61d100daf4d40471f1852
AND evt_block_time > TIMESTAMP '2020-05-19 11:59'
ORDER BY evt_block_time ASC
;
```

9.3 Swap Transactions

```
SELECT *
FROM uniswap_v2_ethereum.Pair_evt_Swap
WHERE contract_address = 0x0d4a11d5eeaac28ec3f61d100daf4d40471f1852
AND evt_block_time > TIMESTAMP '2020-05-19 11:59'
ORDER BY evt_block_time ASC
;
```

9.4 Reserves (To Reconcile Transactions)

```
SELECT *
FROM uniswap_v2_ethereum.Pair_call_getReserves
WHERE contract_address = 0x0d4a11d5eeaac28ec3f61d100daf4d40471f1852
AND evt_block_time > TIMESTAMP '2020-05-19 11:59'
ORDER BY call_block_time ASC
;
```

References

Angeris, G., Kao, H.-T., Chiang, R., Noyes, C., and Chitra, T. (2021). An analysis of uniswap markets.

Aoyagi, M. and Ito, Y. (2023). Coexisting exchange platforms: Limit order books and automated market makers. *Journal of Political Economy Microeconomics*, 1(1):1–35.

Aspris, A., Foley, S., Svec, J., and Wang, L. (2021). Decentralized exchanges: The "wild west" of cryptocurrency trading. *International Review of Financial Analysis*, 77:101845.

Back, K. and Baruch, S. (2013). Strategic liquidity provision in limit order markets. *Econometrica*, 81(1):363–392.

Barbon, A. and Ranaldo, A. (2021). On the quality of cryptocurrency markets: Centralized versus decentralized exchanges. *arXiv preprint arXiv:2112.07386*.

Boulatov, A. and George, T. J. (2013). Hidden and displayed liquidity in securities markets with informed liquidity providers. *The Review of Financial Studies*, 26(8):2096–2137.

Capponi, A. and Jia, R. (2025). Liquidity provision on blockchain-based decentralized exchanges. *Forthcoming at the Review of Financial Studies*.

Capponi, A., Jia, R., and Yu, S. (2024). Price discovery on decentralized exchanges. *Available at SSRN 4236993*.

Cartea, Á., Drissi, F., and Monga, M. (2024). Decentralized finance and automated market making: Predictable loss and optimal liquidity provision. *SIAM Journal on Financial Mathematics*, 15(3):931–959.

Cochrane, J. H. (1991). Production-based asset pricing and the link between stock returns and economic fluctuations. *The journal of finance*, 46(1):209–237.

Cong, L. W., Li, Y., and Wang, N. (2020). Tokenomics: Dynamic adoption and valuation. *The Review of Financial Studies*, 34(3):1105–1155.

George, T. J. and Hwang, C.-Y. (2001). Information flow and pricing errors: A unified approach to estimation and testing. *The Review of Financial Studies*, 14(4):979–1020.

George, T. J. and Khoja, M. (2023). Estimating price impact and its components: Evidence from hft around earnings announcements. *Available at SSRN 4627622*.

Glosten, L. R. (1989). Insider trading, liquidity, and the role of the monopolist specialist. *Journal of business*, pages 211–235.

Glosten, L. R. (1994). Is the electronic open limit order book inevitable? *The Journal of Finance*, 49(4):1127–1161.

Han, J., Huang, S., and Zhong, Z. (2021). Trust in defi: an empirical study of the decentralized exchange. *Available at SSRN*, 3896461.

Hasbrouck, J. (1991). The summary informativeness of stock trades: An econometric analysis. *The Review of Financial Studies*, 4(3):571–595.

Hasbrouck, J. (1993). Assessing the quality of a security market: A new approach to transaction-cost measurement. *The Review of Financial Studies*, 6(1):191–212.

Hasbrouck, J., Rivera, T. J., and Saleh, F. (2022). The need for fees at a dex: How increases in fees can increase dex trading volume. *Available at SSRN*, 4192925.

Hasbrouck, J., Rivera, T. J., and Saleh, F. (2025). An economic model of a decentralized exchange with concentrated liquidity. *Management Science*.

Hendershott, T., Jones, C. M., and Menkveld, A. J. (2011). Does algorithmic trading improve liquidity? *The Journal of finance*, 66(1):1–33.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, pages 1315–1335.

Kyle, A. S. (1989). Informed speculation with imperfect competition. *The Review of Economic Studies*, 56(3):317–355.

Lehar, A. and Parlour, C. (2025). Decentralized exchange: The uniswap automated market maker. *The Journal of Finance*, 80(1):321–374.

Lehar, A., Parlour, C., and Zoican, M. (2023). Fragmentation and optimal liquidity supply on decentralized exchanges. *arXiv preprint arXiv:2307.13772*.

Lo, Y. C. and Medda, F. (2021). Uniswap and the emergence of the decentralized exchange. *Journal of financial market infrastructures*, 10(2):1–25.

Park, A. (2023). The conceptual flaws of decentralized automated market making. *Management Science*, 69(11):6731–6751.

Routledge, B. R., Shen, Y., and Zetlin-Jones, A. (2025). Automated exchange economies.