

A Myersonian Framework for Optimal Liquidity Provision in Automated Market Makers

Jason Milionis

Department of Computer Science
Columbia University
jm@cs.columbia.edu

Ciamac C. Moallemi

Graduate School of Business
Columbia University
ciamac@gsb.columbia.edu

Tim Roughgarden

Department of Computer Science
Columbia University
a16z Crypto
tim.roughgarden@gmail.com

Initial version: October 19, 2022
Current version: February 28, 2023

Abstract

In decentralized finance (“DeFi”), automated market makers (AMMs) enable traders to programmatically exchange one asset for another. Such trades are enabled by the assets deposited by liquidity providers (LPs). The goal of this paper is to characterize and interpret the optimal (i.e., profit-maximizing) strategy of a monopolist liquidity provider, as a function of that LP’s beliefs about asset prices and trader behavior.

We introduce a general framework for reasoning about AMMs. In this model, the market maker (i.e., LP) chooses a demand curve that specifies the quantity of a risky asset (such as BTC or ETH) to be held at each dollar price. Traders arrive sequentially and submit a price bid that can be interpreted as their estimate of the risky asset price; the AMM responds to this submitted bid with an allocation of the risky asset to the trader, a payment that the trader must pay, and a revised internal estimate for the true asset price. We define an incentive-compatible (IC) AMM as one in which a trader’s optimal strategy is to submit its true estimate of the asset price, and characterize the IC AMMs as those with downward-sloping demand curves and payments defined by a formula familiar from Myerson’s optimal auction theory. We characterize the profit-maximizing IC AMM via a generalization of Myerson’s virtual values. The optimal demand curve generally has a jump that can be interpreted as a “bid-ask spread,” which we show is caused by a combination of adverse selection risk (dominant when the degree of information asymmetry is large) and monopoly pricing (dominant when asymmetry is small).

1 Introduction

1.1 Exchanges and Market Makers

The purpose of an exchange is to enable the trade of two or more assets. At a stock exchange, for example, shares of a stock might be exchanged for US dollars via a mechanism known as a “central limit order book” (CLOB), in which buyers and sellers post “limit orders” (offers to buy or sell a specified quantity at a specified price) that are then matched greedily.

Every trade has two sides: for one user to buy a risky asset at a given price, there must be a corresponding seller willing to sell at that same price. Such trades can occur organically due to a coincidence of wants from two traders, but in practice many trades are enabled by professional *market makers* who continually match buy and sell orders whenever the price is right. This paper is about the profit-maximization problem faced by such a market maker.

Digital assets secured by blockchains, such as cryptocurrencies, are also traded via exchanges. In centralized exchanges such as Coinbase or Binance, users hand control of their assets over to the exchange and must accept the credit risk of not getting them back (as in the FTX debacle). Such exchanges are typically based on the CLOB design.

Decentralized exchanges (DEXs), meanwhile, operate purely programmatically (i.e., “on-chain”), are typically non-custodial (meaning that traders at all times have direct control of their assets in the sense that assets are not entrusted to a third party), and often depart from the CLOB model. There are several reasons for this departure: on-chain computation and storage can be expensive, and CLOBs need a lot of both; the markets for many digital assets (especially long-tail ones) are thin and have no dedicated market makers, in which case CLOBs may fail to provide an acceptable level of liquidity; and a general openness in the blockchain world to experimental designs. The DEX design space is large, and many different designs have been deployed (on Ethereum and other general-purpose smart contract platforms) over the past five years.

For example, Uniswap (especially in its earlier iterations, v1 and v2) is a canonical example of an *automated market maker (AMM)*. In an AMM (unlike in a CLOB), the roles of traders and market makers are clearly separated. A market maker, which in this context is called a *liquidity provider (LP)*, deposits into the AMM some amounts of the assets being traded (and is, by default, passive thereafter). At all times, the AMM defines a spot price as a function of its current reserves of those assets. (In the case of Uniswap v1 and v2, the spot price of an AMM is simply the ratio of the quantities of its two assets.) The AMM is willing to take either side of a trade (buy or sell) at its current spot price, and in this sense is always providing liquidity to traders. The latest iteration of Uniswap (v3) blurs the line between AMMs and CLOBs by allowing a much wider range of LP demand curves than in v2 (to some extent approximating what can be expressed by limit orders in a CLOB), thereby encouraging an LP to more actively manage its demand curve as market conditions evolve.

The goal of this paper is to characterize the solution to the following basic question: What is the optimal liquidity demand curve for a profit-maximizing market maker? This question is relevant for a market maker in a CLOB, an LP of an expressive AMM such as Uniswap v3, or even the designer of a new AMM (in which case the market making liquidity demand curve is defined programmatically rather than submitted by a third-party LP). This paper offers a new model for reasoning about this question and a quite general answer, which resembles in many respects Myerson’s theory of optimal auctions (Myerson 1981). To get a feel for our model and results, we next discuss a highly restricted but nonetheless illuminating special case.

1.2 The Pure “Noise Trading” Case

The theory of market microstructure (see e.g., O’Hara (2011)) differentiates between “informed traders,” who have better information about the fundamental value of an asset than the market maker does, and “noise traders,” who trade for idiosyncratic reasons that do not reflect any private information about the fundamental value of the asset. Generally speaking, market makers lose money to informed traders due to adverse selection: by virtue of such a trader being willing to trade, it expects a profit, which, due to the zero-sum nature of common-value assets, comes at the expense of the market maker. Market makers can, however, generally profit from noise traders. This section considers, as an instructive warm-up, the special case in which *all* traders are noise traders (and hence adverse selection plays no role). Our general model in Section 3 accommodates arbitrary mixtures of informed and noise traders, and our general results must therefore reflect adverse selection effects in addition to the aspects identified in this section.

Fundamentally, a market maker must decide, as a function of the information available to it, how much of a risky asset they are willing to buy or sell at different prices. For this example, we will assume that the market maker knows the true value p_0 of the risky asset that it is trading.

An obvious strategy for such a market maker is to accept exactly the orders on which it makes money: all (and only) the buy orders with offered price more than p_0 , and all (and only) the sell orders with offered price less than p_0 . However, as we will see, this is not in general the profit-maximizing strategy for a market maker.

It will be convenient to encode a market maker’s strategy via the choice of *demand curve* g , which maps prices to quantities (i.e., an amount of risky asset that will be held by the market maker). The interpretation is that the market maker would be interested in selling $g(p_0) - g(p)$ units of the risky asset at the price p . (Note that this may be a negative number of units, in which case the market maker will be buying the risky asset from the trader.) A choice of p_0 and g induces the *allocation rule* $x(p) = g(p_0) - g(p)$, the amount of the risky asset allocated to a trader who reports a price of p .

Next, we assume that a noise trader reports a price $\hat{p} \sim \mathcal{D}$ that is drawn according to a distribution \mathcal{D} known to the market maker. We assume that the trader is willing to buy or sell at most a bounded amount of the risky asset (normalized to be 1), and that the demand curve g and price p_0 accordingly satisfy $g(p_0) - g(\hat{p}) \in [-1, 1]$ for all \hat{p} . When the noise trader reports their price \hat{p} , they receive an allocation $x(\hat{p}) = g(p_0) - g(\hat{p})$ of the risky asset and must remit a payment $y(\hat{p})$ to the AMM for this provision. As we will see in Section 2, the natural notion of incentive-compatibility in this setting — that a utility-maximizing trader who values the risky asset at a price \hat{p} will in fact report \hat{p} to the AMM — uniquely pins down the payments (the $y(\hat{p})$ ’s) as a function of the allocation rule (the $x(\hat{p})$ ’s).¹

With informed traders, a rational market maker will update its estimate p_0 of the risky asset’s true price following a trade (generally increasing it after a trader buys the risky asset from the market maker and decreasing it after a sell). Our general model in Section 3 accommodates this through a generic update rule $\pi(p_0, \hat{p})$ that specifies the market maker’s new estimate as a function of its old estimate p_0 and the new information conveyed by a trader reporting the price \hat{p} . (For example, π might correspond to a Bayesian update with respect to some assumed information structure.) In

¹As the reader might guess, it is also essential that the allocation rule x is monotone in \hat{p} , or equivalently that the demand curve g is downward-sloping, with the market maker selling more and more of the risky asset as the going price gets higher and higher. See Section 2 for details.

the present example, because the behavior of noise traders is (by assumption) independent of the true value of the risky asset, the market maker has no cause for updating its internal price (i.e., $\pi(p_0, \hat{p}) = p_0$ for all p_0 and \hat{p}).

What is the optimal strategy (i.e., choice of demand curve g) for the market maker? That is, which choice of g maximizes

$$\mathbb{E}[\text{Revenue}] = \mathbb{E}[y(\hat{p}) - p_0 \cdot x(\hat{p})], \quad (1)$$

where the expectation is with respect to $\hat{p} \sim \mathcal{D}$, $x(\hat{p}) = g(p_0) - g(\hat{p})$ is the allocation rule, and y is the uniquely defined payment rule mentioned above? (In the general case, the “ p_0 ” term in (1) is replaced by $\pi(p_0, \hat{p})$, with the market maker evaluating its portfolio according to its new belief about the value of the risky asset, conditioned on observing the trader’s report of \hat{p} .)

The results in this paper imply that the answer to this question is always defined by two prices (depending on \mathcal{D}), $p_l \leq p_0$ and $p_h \geq p_0$. These two prices split the price range into three intervals, and the optimal strategy is to always buy the maximum amount in the lowest price interval, always sell in the highest price interval, and refuse to trade in the middle price interval.

We will be interested in the length of this middle price interval, which we call the *no-trade gap*. Bigger no-trade gaps are generally worse from a welfare perspective, with welfare-increasing trades sacrificed in the name of higher profit to the market maker. The no-trade gap is reminiscent of the “bid-ask spread” in CLOBs, which occupies a central position in the study of traditional market-making in CLOBs.² Traditionally, theoretical justifications for non-trivial bid-ask spreads in CLOBs have relied on frictions (like transaction fees or inventory costs) or adverse selection (with a market maker needing to exploit noise traders to cover the losses to informed traders). In the current example, with neither frictions nor adverse selection, the no-trade gap arises for a different reason, namely the monopolist position of the market-maker. In our general model, with the possibility of informed traders, monopoly pricing and adverse selection both contribute to the no-trade gap. For example, no matter what the distribution \mathcal{D} is, if traders have perfect information (i.e., extreme adverse selection), then the no-trade gap encompasses the full price range.

The prices p_l and p_h that split the price range into three intervals, and hence the no-trade gap, are distribution-dependent. For example, intuitively, one might expect a smaller gap for distributions \mathcal{D} that are tightly concentrated around p_0 . In general, our analysis shows that these two prices can be characterized as roots of two functions that resemble the virtual valuation functions used in optimal auction theory (Myerson 1981). For example, in the pure noise-trading case, if \mathcal{D} is the uniform distribution on $[p_{\min}, p_{\max}]$, then $p_l = (p_{\min} + p_0)/2$ and $p_h = (p_0 + p_{\max})/2$. If \mathcal{D} is the exponential distribution with parameter λ (on $[0, \infty)$), then $p_h = p_0 + \frac{1}{\lambda}$. The lower price p_l does not have a closed-form expression in this case, but can be easily approximated numerically (given p_0 and λ); the following table enumerates some examples.

²In a CLOB, the “bid” and “ask” refer to the highest price of an outstanding buy order and the lowest price of an outstanding sell order, respectively. Because limit orders in a CLOB are always matched greedily, the bid-ask spread is always non-negative. Small spreads are generally viewed as a good thing, indicating a “highly liquid” market (with lower trading costs, at least for small quantities).

$p_0 \setminus \lambda$	0.5	1	2
0.25	0.123	0.121	0.118
0.5	0.242	0.235	0.221
0.75	0.358	0.342	0.314
1	0.470	0.443	0.396
1.5	0.684	0.627	0.537
2	0.886	0.792	0.653

The rest of this paper develops a similarly precise understanding of optimal market-maker strategies in a more general setting.

1.3 Related Work

Automated market makers (AMMs) are objects familiar from prediction markets, which are used to perform belief aggregation around uncertain future events. An AMM is then used to automatically price — based on strategic agents’ behavior — a set of securities that represent the likelihoods of these particular events happening. For a survey of considerations around prediction markets, see D. N. Pennock and Sami (2007). More specifically, the framework that introduced AMMs where a function of the possible future states of the world remains constant is the utility-based market making framework of Chen and D. Pennock (2007). In that framework, utility indifference conditions define a bonding function for binary payoff Arrow-Debreu securities.

Recently, AMMs have found a new use in the environment of blockchains: to facilitate trading and pricing of arbitrary assets. Such use as a decentralized exchange mechanism was first suggested by Buterin (2016) and Lu and Köppelmann (2017). The latter authors suggested a constant product market maker, which was analyzed by Angeris, Kao, et al. (2019). Angeris and Chitra (2020) and Angeris, Evans, et al. (2021a) and Angeris, Evans, et al. (2021b) apply tools from convex analysis to study the more general case of constant function market makers.

In particular, constant function market makers (CFMMs) have arisen from the idea of holding some function constant across “trading states” (from states of the world, previously, on the utility-based framework of AMMs for prediction markets), and are characterized by such an invariant called a *bonding function*. This means that the AMM is willing to take the other side of any trade that corresponds to remaining on a constant level curve of the particular bonding function chosen.

This paper critically relates to revenue maximizing auction design, initiated under the pioneering work of Myerson (1981). Since then, various versions of virtual values have been used in a variety of settings among others for specifying optimal auctions and their approximations (e.g., Bulow and Roberts (1989), Chawla, Jason D Hartline, and Kleinberg (2007), Jason D Hartline (2012), Haghpanah and J. Hartline (2015), Chawla, Jason D. Hartline, et al. (2019), Chawla, Jason D Hartline, Malec, et al. (2010), Alaei et al. (2013), Roughgarden and Talgam-Cohen (2019), Roughgarden and Talgam-Cohen (2016), and Jason D Hartline and Roughgarden (2009)); in a similar spirit, we define and use functions that can be interpreted as virtual valuations. For a survey of results around revenue-optimal mechanisms, virtual values, optimal auctions, and their approximations, see the book by Jason D Hartline (2021).

The description of the framework that we use for AMMs — inspired by a reparameterization of a CFMM curve (established by Angeris and Chitra (2020)) in terms of portfolio holdings of the pool with respect to the price — resembles that in Milionis et al. (2022), who used a similar AMM framework for a different purpose, namely unifying constant function market makers (CFMMs)

and limit order books (LOBs) and complexity-approximation trade-offs between them. That paper considers neither incentive-compatibility constraints nor the problem of optimal liquidity provision.

Goyal et al. (2022) also examine beliefs of LPs (specifically on CFMMs) around future asset prices; however, in their model, these beliefs are static (i.e., similar to a single prior that remains unchanged throughout the trades' occurrence) or modeled by a fixed price dynamic discounted to the present (i.e., still not dynamically revised in response to incoming trades). The focus of Goyal et al. (2022) is the design of CFMMs that maximize the fraction of trades that the CFMM can satisfy with small slippage under the stationary distribution of the Markov chain induced by their model. Our work focuses instead on LP revenue-maximization with informed traders (and consequent endogenous price updates).

There is also work specific to Uniswap v3 from the LP perspective around how beliefs about future prices should guide the choice of an LP's demand curve (Fan et al. 2022; Yin and Ren 2021; Neuder et al. 2021).

Finally, there is a large literature on the market microstructure of limit order books; see the textbook by O'Hara (2011) and references therein. Closest to the present work is the paper of Glosten (1989). Focusing on the role of asymmetric information, Glosten (1989) considers the problem of a monopolist optimizing their liquidity demand curve, but in the presence of risk averse traders (in contrast to our setting, where the traders can be viewed as risk neutral). The risk averse setting is more technically challenging, and hence Glosten (1989) is required to make very strong and highly specific technical assumptions (normally distributed values and observation noise) that are not necessary in our setting. Moreover, our risk neutrality assumption unlocks a fruitful connection between optimal market-making and the Myersonian theory of optimal auction design. In a slightly more general setting, Glosten (1994) considers the related problem of the liquidity demand curve in a competitive equilibrium.

2 Automated market makers

We begin by presenting a general framework of exchange for automated market makers (AMMs) that allows us to consider market making as conducting an auction, and is based on market participants (traders) trading with the AMM on the basis of future price beliefs.

More specifically, suppose there are two assets, a risky asset and a numéraire asset. An AMM according to this framework is defined by a **demand curve** $g : (0, \infty) \rightarrow \mathbb{R}^+$ such that the function $g(\cdot)$ represents the *current* demand of the market maker for the risky asset as a function of price, i.e., the amount of risky asset that the market maker wishes to hold at each possible price, along with a payment rule $y : (0, \infty) \rightarrow \mathbb{R}$ such that the function $y(\cdot)$ specifies the quantity of numéraire that a trader needs to pay depending on the *current* AMM price p_0 (which instantiates a specific state of this AMM's demand curve, i.e., a current demand of the AMM for $g(p_0)$ of risky asset) and what they declare to the AMM that their price estimate is (see the following paragraph on trading for the precise definition). The market maker is assumed to have no inventory costs or budget constraints.

The demand curve g along with an appropriately defined payment rule can arise through *bonding curves* of traditional constant function market makers (CFMMs, i.e., functions f such that the holdings of the joint pool (x, y) satisfy $f(x, y) = c$ for some c) but this is not necessary: the exchange mechanisms defined by this framework strictly generalize CFMMs (see also Milionis et al. (2022)). Additionally, a non-continuous demand curve g can also be used to express limit orders, with the discontinuities representing limit buy or limit sell orders at the corresponding price, of quantity of

risky asset equal to the jump at that price (Milionis et al. 2022).

Trading A trader who wants to trade with the AMM has some (hidden) belief about the asset price. The trader will come to the mechanism and specify a price bid \hat{p} . Based on their bid, the trader will get a quantity $g(p_0) - g(\hat{p})$ of risky asset. Finally, the AMM revises its current price (which was previously p_0) according to some update rule (we will see more about the update rule from Section 3 onwards; it will play no role in this section). We say that the trader was *allocated* a quantity $x(\hat{p}) \triangleq g(p_0) - g(\hat{p})$ of asset. This interpretation makes it natural to consider the transaction with the trader **as an auction** with one bidder; the trader has a price estimate (valuation) for the risky asset, will be allocated a portion of the risky asset (provided by the market maker) based on their submitted price bid, and will need to pay the market maker an amount $y(\hat{p})$ of numéraire for it. Notice that it may be the case that $x(\hat{p}) < 0$, which means that the trader will sell risky asset to the AMM, and the market maker will now need to come up with the corresponding payment to the trader for it.

Example: Uniswap v2 Let's consider the simple case of a constant product market maker (CPMM), such as Uniswap v2, to show how it emerges from the framework above. Consider the demand curve of the AMM to be of the form $g(p) = \frac{c}{\sqrt{p}}$, for some $c > 0$, and the payment rule to be of the form $y(\hat{p}) = c(\sqrt{\hat{p}} - \sqrt{p_0})$, for the same $c > 0$. A trader who will trade with this exchange at a current price p_0 with a price bid \hat{p} will obtain a quantity $x(\hat{p}) = g(p_0) - g(\hat{p}) = c\left(\frac{1}{\sqrt{p_0}} - \frac{1}{\sqrt{\hat{p}}}\right)$ of risky asset. Since the curve $g(p)$ above is just a reparameterization (in terms of prices) of the risky asset holdings of the CPMM curve³ $xy = k$ where $k = c^2$ (Angeris, Evans, et al. 2021b), by comparing the expression of the quantity of risky asset that the trader gets, we conclude that they get the same quantity as in the CPMM.

2.1 Incentive-compatible AMMs

We will now consider that traders interacting with AMMs have utility functions that emerge naturally from their *true* belief about the asset price versus what they bid to obtain a quantity of risky asset from the AMM. In particular, a trader who has a *true* belief of p for the asset's price, and submits a bid \hat{p} to the AMM seeks to optimize (maximize) their utility defined by

$$u(p, \hat{p}) \triangleq p \cdot x(\hat{p}) - y(\hat{p}), \tag{2}$$

because they consider that the asset's true value to them is p per unit, they obtain a quantity $x(\hat{p})$ of it, and pay $y(\hat{p})$ for this trade.

Definition 1. [Incentive-compatible AMM] An AMM defined according to a demand curve $g(p)$ and a payment rule $y(\hat{p})$ is called *incentive compatible* (IC), if any trader whose utility follows Eq. (2) has as an optimal strategy for interacting with the AMM to submit their *true* belief about the asset price as their bid to the AMM.

Is an AMM defined by an arbitrary demand curve $g(p)$ and an arbitrary payment rule $y(\hat{p})$ incentive compatible? The answer is no, and the characterization of IC AMMs follows standard results of dominant-strategy incentive-compatible (DSIC) single-parameter auctions.

³Note that here x and y — without the parentheses — are not the functions above, but rather the current total reserves of the AMM pool.

Proposition 2.1 (Characterization of IC AMMs). *An AMM defined according to a demand curve $g(p)$ can be paired with a payment rule $y(\hat{p})$ to obtain an IC AMM, if and only if $g(p)$ is a non-increasing function.*

Proof. First, we observe the analogy drawn to *allocations* of asset for trades: $x(\hat{p}) = g(p_0) - g(\hat{p})$. Since the utility function for the auction with a single bidder (the trader) is of the form of Eq. (2), the proposition follows from the characterization of dominant-strategy incentive-compatible (DSIC) single-parameter auctions' allocation rules $x(\hat{p})$ as only those functions which are monotone, and specifically non-decreasing functions, because $x(\hat{p})$ is non-decreasing precisely if and only if $g(\hat{p})$ is non-increasing.

More specifically, we follow the standard analysis to showcase the exact correspondence of the settings but do not show the entire analysis, as this follows Myerson's. In order for the optimal strategy of any trader with true belief p to be to submit p to the AMM, we need to have that for any p, \hat{p} :

$$u(p, p) \geq u(p, \hat{p}) \Leftrightarrow p \cdot x(p) - y(p) \geq p \cdot x(\hat{p}) - y(\hat{p}) \Leftrightarrow p \cdot [x(p) - x(\hat{p})] \geq y(p) - y(\hat{p}).$$

Additionally, if the converse holds, and the true belief of the trader was \hat{p} while it is considering submitting p to the AMM, then we also need to have that for any p, \hat{p} :

$$u(\hat{p}, \hat{p}) \geq u(\hat{p}, p) \Leftrightarrow \hat{p} \cdot x(\hat{p}) - y(\hat{p}) \geq \hat{p} \cdot x(p) - y(p) \Leftrightarrow \hat{p} \cdot [x(p) - x(\hat{p})] \leq y(p) - y(\hat{p}).$$

Combining these two inequalities, we get that for any p, \hat{p} :

$$\hat{p} \cdot [x(p) - x(\hat{p})] \leq y(p) - y(\hat{p}) \leq p \cdot [x(p) - x(\hat{p})],$$

hence it is immediate (from the derived inequality $(p - \hat{p}) \cdot [x(p) - x(\hat{p})] \geq 0$) that $x(\hat{p})$ needs to be non-decreasing. \square

The next question to be answered is whether arbitrary payment rules $y(\hat{p})$ are allowed to be requested by an IC AMM. Not surprisingly, as is the case for payments defined by an allocation rule of a DSIC mechanism, the payments need to be precisely defined by the demand curve for an IC AMM. The below formula is familiar from Myerson's single-parameter DSIC auction design.

Corollary 2.2 (Payments of traders on IC AMMs). *Suppose that any trader who is allocated a quantity 0 of risky asset pays 0 amount of numéraire to an IC AMM according to Definition 1 with a demand curve $g(p)$. The IC AMM's payment rule $y(\hat{p})$ for a trader submitting a price bid \hat{p} needs to be exactly:*

$$y(\hat{p}) = \int_{p_0}^{\hat{p}} s \, dx(s) = - \int_{p_0}^{\hat{p}} p \, dg(p). \quad (3)$$

The above integrals are Riemann–Stieltjes integrals. In cases where $g(p)$ is differentiable, the differential takes the form $dg(p) = g'(p) \, dp$.

Proof. From Myerson's single-parameter auction design, we know that there is a unique payment rule (up to an additive constant) for DSIC single-parameter auctions; that is, $y(\hat{p}) = \int_{p_0}^{\hat{p}} s \, dx(s)$, where the arbitrary additive constant is no longer present, as we have integrated the required normalization that for an allocation of $x(p_0) = g(p_0) - g(p_0) = 0$ asset, the payment needs to be $y(p_0) = 0$. This directly translates to the formula given in the corollary's statement in terms of the demand curve $g(p)$ of the AMM. \square

Note that the integral that defines the payments is non-negative if $\hat{p} \geq p_0$, and non-positive if $\hat{p} \leq p_0$. This is compatible with the direction of the trade, as defined by the allocation rule and [Proposition 2.1](#), which is that $x(\hat{p}) \geq 0$ for $\hat{p} \geq p_0$, and $x(\hat{p}) \leq 0$ for $\hat{p} \leq p_0$.

3 Optimal liquidity provision via revenue maximization

3.1 Model for optimal allocation rule

First, we describe our model for how the market maker should optimize their allocation rule for the asset. From now on, the market maker will be assumed to be risk-neutral, i.e., be indifferent to “risk,” and only care to be profit-maximizing.

The Bayesian belief framework operates as follows: the market maker has a prior belief distribution on the asset price, and the (single) price it decides to use is p_0 ; for example, and most commonly, p_0 will be the mean of this prior distribution. A trader then comes to the AMM with an estimate \hat{p} (that is the true belief of the trader, if the AMM is IC according to [Definition 1](#)). Traders are on the market to buy or sell up to 1 unit of the asset;⁴ that is, the allocation rule should satisfy $-1 \leq x(\hat{p}) \leq 1$ for all \hat{p} . The trader’s estimate \hat{p} is perceived by the market maker as one observation of an asset price estimate coming from a distribution \mathcal{D} that is included in the market maker’s model. For this distribution \mathcal{D} , we assume that it has a continuous, positive density function $f(\hat{p})$ on a compact support $[p_{\min}, p_{\max}]$ and a cumulative density function (CDF) $F(\hat{p}) = \int_{p_{\min}}^{\hat{p}} f(s) ds$.⁵

Applying Bayes’ theorem allows the market maker to specify the posterior distribution of the asset price, conditional on the observation from the trader. The market maker then sets this posterior distribution as their new prior for future observations (trades). We will assume that the way this posterior is used by the AMM is through valuing the quantity of asset given or obtained at a price of $\pi(p_0, \hat{p})$ which can be, for instance, the mean of the posterior distribution according to the Bayesian framework.

In general, we will call $\pi(p_0, \hat{p})$ the “update rule” and — for generality — we will not impose further constraints arising from the Bayesian framework on what it may be other than that:

Assumption 1. *The following conditions have to hold for the update rule:*

- $\pi(p_0, \hat{p})$ lies in the interval between p_0 and \hat{p} .
- $\pi(p_0, \hat{p})$ is non-decreasing in \hat{p} .
- $\pi(p_0, p_0) = p_0$, i.e., if the trader’s estimate is the same as the market maker’s prior, then the updated valuation of the market maker for the asset should be the same as their prior.
- $\mathbb{E}_{\hat{p} \sim \mathcal{D}}[\pi(p_0, \hat{p})] = p_0$, i.e., the prior is consistent with the update rule.

Therefore, according to the above, the market maker’s next best estimate for the asset’s price is $\pi(p_0, \hat{p})$ conditioned on receiving the trader’s estimate \hat{p} . It is then natural to consider that, when selling (respectively, buying) some quantity $x(\hat{p})$ of the asset from the trader, the AMM will not

⁴This is just for notational convenience. In fact, it generalizes easily for demand up to any arbitrary finite bound, and a subsequent normalization enables using exactly our results here with unit demand.

⁵Our results can be extended to non-continuous density functions supported at possibly non-compact intervals, e.g., $(0, \infty)$, but for simplicity of exposition, we do not focus on handling these edge cases here, since they do not add to the general discussion points of the paper.

only consider the amount of money that they receive (respectively, give) but also their perceived best estimate of the value of the asset that they lost (respectively, gained). This, then, arises as an interesting contrast to traditional auction theory, where the revenue of the auction is only the payment that is made. Here, according to the above reasoning, it makes sense for the AMM to value their revenue according to

$$\text{Revenue} = y(\hat{p}) - \pi(p_0, \hat{p}) \cdot x(\hat{p}). \quad (4)$$

From the perspective of determining an optimal allocation rule, maximizing this revenue expression is equivalent to maximizing profit. Therefore, an AMM choosing an optimal strategy to provide liquidity will consider the strategy that maximizes their revenue, among all feasible, incentive-compatible (for the traders) strategies.

3.2 Equivalence of expected revenue and virtual welfare, and virtual value functions

By re-working Myerson's optimal auction theory under the new objective that we want to optimize and the constraints that come with the definition of the allocation rule from an AMM, we get that:

Theorem 3.1. *[Expected Revenue = Expected Virtual Welfare] The expected revenue according to Eq. (4) of a market maker under the assumptions of Section 3.1 is equal to the expected virtual welfare, i.e.,*

$$\mathbb{E}_{\hat{p} \sim \mathcal{D}}[\text{Revenue}] = \mathbb{E}_{\hat{p} \sim \mathcal{D}}[|x(\hat{p})| \cdot (\phi_u(\hat{p}) \cdot \mathbf{1}_{\hat{p} \geq p_0} + \phi_l(\hat{p}) \cdot \mathbf{1}_{\hat{p} \leq p_0})]$$

where $\mathbf{1}_A$ is the indicator function that is 1 when A holds and 0 otherwise, and $\phi_u(s), \phi_l(s)$ are the virtual value functions defined by

$$\phi_u(s) \triangleq s - \frac{1 - F(s)}{f(s)} - \pi(p_0, s) \quad (5)$$

and

$$\phi_l(s) \triangleq \pi(p_0, s) - s - \frac{F(s)}{f(s)}. \quad (6)$$

Proof. By Eq. (3) and re-writing $x(\hat{p})$ using the normalization condition, we have that⁶

$$y(\hat{p}) = \int_{p_0}^{\hat{p}} sx'(s) ds \text{ and } x(\hat{p}) = \int_{p_0}^{\hat{p}} x'(s) ds,$$

⁶This proof assumes differentiability of $x(s)$ for simplicity of the technical exposition but carries through verbatim in terms of Riemann–Stieltjes integrals where $x'(s) ds = dx(s)$, which always exist because $x(s)$ is guaranteed to be monotonic by Proposition 2.1.

so that the expected revenue $\mathbb{E}_{\hat{p} \sim \mathcal{D}}[\text{Revenue}]$ is

$$\begin{aligned}
& \int_{p_{\min}}^{p_{\max}} \int_{p_0}^{\hat{p}} (s - \pi(p_0, \hat{p})) x'(s) f(\hat{p}) ds d\hat{p} \\
&= \int_{p_0}^{p_{\max}} \int_{p_0}^{\hat{p}} (s - \pi(p_0, \hat{p})) x'(s) f(\hat{p}) ds d\hat{p} - \int_{p_{\min}}^{p_0} \int_{\hat{p}}^{p_0} (s - \pi(p_0, \hat{p})) x'(s) f(\hat{p}) ds d\hat{p} \\
&= \int_{p_0}^{p_{\max}} \int_s^{p_{\max}} (s - \pi(p_0, \hat{p})) x'(s) f(\hat{p}) d\hat{p} ds - \int_{p_{\min}}^{p_0} \int_{p_{\min}}^s (s - \pi(p_0, \hat{p})) x'(s) f(\hat{p}) d\hat{p} ds \\
&= \int_{p_0}^{p_{\max}} x'(s) \left(s(1 - F(s)) - \int_s^{p_{\max}} \pi(p_0, \hat{p}) f(\hat{p}) d\hat{p} \right) ds \\
&\quad - \int_{p_{\min}}^{p_0} x'(s) \left(sF(s) - \int_{p_{\min}}^s \pi(p_0, \hat{p}) f(\hat{p}) d\hat{p} \right) ds \\
&= - \int_{p_0}^{p_{\max}} x(s) (1 - F(s) + (\pi(p_0, s) - s) f(s)) ds + \int_{p_{\min}}^{p_0} x(s) (F(s) + (s - \pi(p_0, s)) f(s)) ds \\
&= \int_{p_0}^{p_{\max}} |x(s)| \left(s - \pi(p_0, s) - \frac{1 - F(s)}{f(s)} \right) f(s) ds + \int_{p_{\min}}^{p_0} |x(s)| \left(\pi(p_0, s) - s - \frac{F(s)}{f(s)} \right) f(s) ds,
\end{aligned}$$

where the third line arises from switching the order of the integrals (and the correct switching depends on the relative order of p_0 and \hat{p} which is why the second line is needed; the first integral when $\hat{p} \geq p_0$ and the second vice-versa), the fifth line from integration by parts, and the sixth line because $x(s) \leq 0$ for $s \leq p_0$ and $x(s) \geq 0$ for $s \geq p_0$.

Therefore, the two virtual value functions are as in the theorem description. \square

3.3 Derivation of the optimal allocation rule and the no-trade gap

We now turn our attention to optimizing the expected revenue, i.e., finding the allocation rule (equivalently, the demand curve up to an arbitrary positive additive constant) which maximizes the expected revenue.

Theorem 3.2. *The optimal allocation rule $x^*(\hat{p})$ that maximizes the expected revenue according to Eq. (4) under the assumptions of Section 3.1 has the form*

$$x^*(s) = \begin{cases} 1, & \text{for } p_1 \leq \hat{p} \leq p_{\max} \\ 0, & \text{for } p_2 < \hat{p} < p_1 \\ -1, & \text{for } p_{\min} \leq \hat{p} \leq p_2 \end{cases},$$

where $p_1 \geq p_0$ and $p_2 \leq p_0$ are some roots of the upper and lower virtual value functions $\phi_u(s), \phi_l(s)$ of Eqs. (5) and (6) respectively.

Proof sketch. We will use the equivalence of expected revenue and expected virtual welfare established by Theorem 3.1. Because we operate under the constraints that $x(p_0) = 0$, $x(\hat{p}) \geq 0$ for $\hat{p} \geq p_0$, and $x(\hat{p}) \leq 0$ for $\hat{p} \leq p_0$, we can just equivalently optimize the two integrals arising from the right hand side of Theorem 3.1 separately to obtain the optimal allocation rule. Additionally, since $\phi_u(p_0) = -\frac{1-F(p_0)}{f(p_0)} \leq 0$, $\phi_u(p_{\max}) = p_{\max} - \pi(p_0, p_{\max}) \geq 0$, $\phi_l(p_{\min}) = \pi(p_0, p_{\min}) - p_{\min} \geq 0$, and $\phi_l(p_0) = -\frac{F(p_0)}{f(p_0)} \leq 0$, there always exists a non-trivial optimal solution to the optimization problem.

If $\phi_u(s)$ was monotone (increasing), then by the above sign changes, there would exist a (unique) root p_1 such that $\phi_u(p_1) = 0$. A simple argument by contradiction then shows that the structure of

the optimal allocation rule is to allocate the maximum allowable quantity of asset (by the constraint of unit demand, $x(s) \leq 1$) to be sold from the price p_1 onwards; that is,

$$x^*(s) = \begin{cases} 1, & \text{for } p_1 \leq \hat{p} \leq p_{\max} \\ 0, & \text{for } p_0 \leq \hat{p} < p_1 \end{cases}. \quad (7)$$

For the lower virtual value function, a similar argument shows that, if $\phi_l(s)$ was monotone (decreasing), and p_2 is its unique root (which is guaranteed to exist by the aforementioned sign changes), then the optimal allocation rule obeying the constraint of unit demand $x(s) \geq -1$ would be

$$x^*(s) = \begin{cases} 0, & \text{for } p_2 < \hat{p} \leq p_0 \\ -1, & \text{for } p_{\min} \leq \hat{p} \leq p_2 \end{cases}. \quad (8)$$

In full generality, if the virtual value functions are not monotone, then there may now be multiple roots of those functions. In this case, and with an argument by contradiction that if the allocation rule were to switch in between roots then we could obtain better or equal expected revenue by modifying the allocation rule to remain constant in between roots, along with the argument that if the allocation rule were not to allocate the entire available demand (1, -1 , respectively) then we could obtain better or equal expected revenue by increasing (or decreasing, respectively) the allocated quantity of the asset until the limit (of 1 or -1 , respectively), we obtain that the optimal allocation rule is exactly characterized by Eqs. (7) and (8), where this time the choice of the roots as p_1 (and p_2 , respectively) needs to be the root of the upper (lower, respectively) virtual value function that gives the maximum expected revenue under the allocation rule that has the form of Eq. (7) (Eq. (8), respectively). Therefore, the optimal allocation rule is as per the theorem statement. \square

Therefore, in any case, we observe that the optimal solution of Eqs. (7) and (8) has a well-specified **no-trade gap** around p_0 where the market maker is not willing to buy or sell any quantity of the asset to the trader. More specifically, the length of this no-trade gap is precisely $p_1 - p_2$, where $p_1 \geq p_0$, $p_2 \leq p_0$ are as per the above analysis roots of the upper and lower virtual value functions respectively.

We can now specialize in the two interesting extreme cases that were examined in Section 1.2 to obtain the respective observations there: first, consider the pure noise trading case, where the traders are assumed to offer no information on the true asset price, and hence the market maker's prior is static, i.e., in the case that the update rule is $\pi(p_0, \hat{p}) = p_0$, $\forall \hat{p} \in [p_{\min}, p_{\max}]$. In this case, we see that there is a *non-trivial* no-trade gap, since we have for the extremal values that $\phi_u(p_0) < 0$, $\phi_u(p_{\max}) > 0$, $\phi_l(p_{\min}) > 0$, and $\phi_l(p_0) < 0$ (see the beginning of Section 3.3 for the expressions), hence the respective roots are contained strictly in between their corresponding intervals (by continuity of the virtual value functions). This no-trade gap is akin to a bid-ask spread observed in traditional CLOB-based markets, and is attributable to the monopoly pricing power of the market maker. In addition, it is *distribution-dependent* and its length can be intuitively understood as being smaller under distributions \mathcal{D} , according to which the traders' estimates \hat{p} are distributed, that are more "concentrated" around p_0 ⁷. To mention an example to illustrate

⁷With this statement, we intend to only give intuition around the length of the interval, because there can be pathological cases of pairs of distributions that do not follow this generic intuition, and hence it need not be true in full generality.

this intuition, the no-trade gap under pure noise trading and a uniform distribution supported on $[p_{\min}^{(1)}, p_{\max}^{(1)}]$ will be smaller than the one under another uniform distribution supported on $[p_{\min}^{(2)}, p_{\max}^{(2)}]$, where $p_{\min}^{(1)} > p_{\min}^{(2)}$ and $p_{\max}^{(1)} < p_{\max}^{(2)}$.

At the other extreme, consider the case where the traders always have perfect information about the true asset price, i.e., the market maker's update rule is $\pi(p_0, \hat{p}) = \hat{p}$, $\forall \hat{p} \in [p_{\min}, p_{\max}]$. In this case, we observe that the two virtual value functions become

$$\phi_u(s) = -\frac{1 - F(s)}{f(s)} \leq 0 \text{ and } \phi_l(s) = -\frac{F(s)}{f(s)} \leq 0,$$

thus the optimal allocation rule is everywhere zero: $x^*(\hat{p}) = 0$, $\forall \hat{p} \in [p_{\min}, p_{\max}]$. It is therefore the case that this corollary which may be phrased as a version of a *no-trade theorem* in the presence of perfect adverse selection is *distribution-independent*: that is, regardless of the distribution \mathcal{D} of the traders' asset price estimates $\hat{p} \sim \mathcal{D}$, if the market maker believes that the traders have perfect information, then their optimal allocation rule is to perform no-trade at all. More specifically:

Corollary 3.3 (No-trade theorem when traders have perfect information). *If the update rule $\pi(p_0, \hat{p}) = \hat{p}$, $\forall \hat{p} \in [p_{\min}, p_{\max}]$, then for every distribution \mathcal{D} , the optimal allocation rule is $x^*(\hat{p}) = 0$, $\forall \hat{p} \in [p_{\min}, p_{\max}]$.*

4 An important special case: linear instantiation of the update rule

We move on to consider an interesting special case of the update rule of [Section 3.1](#): the simplest specialization that allows us to capture an arbitrary mixture of both noise trading and adverse selection, and enables us to show exactly how they intermingle with one another. More specifically, we consider the case when the update rule is a convex combination of the two above extreme cases, controlled by a parameter $\lambda \in [0, 1]$:

$$\pi(p_0, \hat{p}) = \lambda p_0 + (1 - \lambda)\hat{p} \tag{9}$$

The extreme cases arise as follows: for $\lambda = 1$, we obtain that the update remains the same as the market maker's prior (p_0), i.e., the pure noise trading model, where the trader is assumed to provide no information on the asset price. For $\lambda = 0$, we obtain that the update always follows exactly the trader's estimate (\hat{p}), thus the trader is considered as having perfect information around the asset price, i.e., pure adverse selection is attained.

There are standard normality assumptions under which this linear update rule naturally emerges as the update rule for Bayesian updating. For example, if we consider that the asset value $V \sim \mathcal{N}(p_0, \sigma_0^2)$ where \mathcal{N} is the normal distribution, and the trader's estimate is $\hat{p} = V + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$, i.e., $\hat{p} \sim \mathcal{N}(p_0, \sigma_0^2 + \sigma_\epsilon^2)$, then the posterior according to a Bayesian update step would also be a normal distribution with mean $\pi(p_0, \hat{p}) = \lambda p_0 + (1 - \lambda)\hat{p}$, where $\lambda = \frac{\sigma_\epsilon^2}{\sigma_0^2 + \sigma_\epsilon^2}$ (Gelman et al. [2014](#)). Of course, this is just an illustration of the motivation behind the linear update rule; it ignores the fact that asset prices cannot obtain negative values, but a similar calculation can be performed for log-normal distributions, for instance, and also holds approximately for truncated Gaussian distributions in the case of the standard deviation being small compared to the mean value. However, in this section, we will not be necessarily making any such Gaussian assumptions;

we will only be assuming that the conducted update corresponds to a linear interpolation (in fact, convex combination) of the two cases of p_0 and \hat{p} that we have seen in [Section 1.2](#) as indicative of noise trading and adverse selection.

In order to proceed with a more precise characterization of the interval of no-trade around p_0 , we will make a similar assumption to the regularity of the distribution \mathcal{D} made by Myerson ([1981](#)):

Assumption 2. *The upper and lower virtual value functions, obtained by plugging in [Eq. \(9\)](#) to [Eqs. \(5\)](#) and [\(6\)](#),*

$$\phi_u(s; \lambda) = \lambda(s - p_0) - \frac{1 - F(s)}{f(s)} \text{ and } \phi_l(s; \lambda) = \lambda(p_0 - s) - \frac{F(s)}{f(s)}$$

are non-decreasing and non-increasing (respectively) with respect to s for all $\lambda \in [0, 1]$.

Discussion of Assumption 2 This assumption can be intuitively understood as stating that the distribution of \hat{p} (i.e., \mathcal{D}) does not have heavy tails away from p_0 on either direction (i.e., left or right tails). We now explain why this is the case. First, observe that the original regular distributional assumption of Myerson ([1981](#)) is equivalent to the assumption that $\phi_u(s; \lambda)$ is strictly increasing for $\lambda = 1$. Additionally, notice that an equivalent re-statement of [Assumption 2](#) is that the respective monotonicities hold for $\lambda = 0$, since this implies the monotonicities for all $\lambda \in [0, 1]$. In particular, for the monotonicity of the upper virtual value function, we observe that it exactly corresponds to having a monotone hazard rate, where the hazard rate of a distribution is defined as

$$h(s) \triangleq \frac{f(s)}{1 - F(s)}.$$

Distributions that have a monotone hazard rate (MHR) have been widely used before as common distributional assumptions in revenue maximizing settings ([Jason D Hartline 2021](#); [Cole and Rao 2017](#); [Cole and Roughgarden 2014](#); [Ewerhart 2013](#)). More generally, $\phi_u(s; \lambda)$ being non-decreasing for a specific λ corresponds to the notion of an α -strongly regular distribution ([Definition 1.1](#)) of [Cole and Roughgarden \(2014\)](#) where $\alpha = 1 - \lambda$. As for the monotonicity of the lower virtual value function assumption, we observe that if we make the change of variables $p' = p_0 - \hat{p}$ (where $\hat{p} \sim \mathcal{D}$ is our original random variable, and p' is the new one), then $\phi_l(s; 0)$ being non-increasing would be equivalent to the distribution of p' having a monotone hazard rate (MHR), because if we denote by $G(s)$ and $g(s)$ the CDF and the PDF of p' respectively, then from standard probabilities, we get that $G(s) = 1 - F(p_0 - s)$, and $g(s) = f(p_0 - s)$, hence

$$\frac{g(s)}{1 - G(s)} = \frac{f(p_0 - s)}{F(p_0 - s)},$$

and $\phi_l(s; 0) = -\frac{F(s)}{f(s)}$ being non-increasing is equivalent to $\frac{f(p_0 - s)}{F(p_0 - s)}$ being non-decreasing.

Intuitively, distributions satisfying [Assumption 2](#) are well-behaved, in the sense that they correspond to distributions with MHR on both of their tails, i.e., distributions that have both left and right tails that are no heavier than an exponential.

Many interesting distributions satisfy the “non-heavy-tailed” property that [Assumption 2](#) refers to. For example, all uniform and exponential distributions satisfy the assumption. More generally, all distributions with log-concave densities have a monotone hazard rate ([Ewerhart 2013](#)), and if centered around p_0 (for instance symmetrically on both sides), both tail ends would not be heavy

and they would satisfy the assumption. As with classical non-MHR distributions, [Assumption 2](#) is not satisfied by most heavy-tailed distributions; the difference is that this assumption includes both left and right tails of the distribution. A good example of a double-sided (left and right) heavy-tailed distribution that would fail to satisfy [Assumption 2](#) would be the Cauchy distribution with a location parameter of p_0 , for which the moment generating function does not exist anywhere: in this case, both of its left and right tails around p_0 are heavy ([Mantel 1976](#)).

With [Assumption 2](#) in mind, we move on to characterize the length of the no-trade interval present in the optimal solution of [Eqs. \(7\) and \(8\)](#) as λ varies.

Proposition 4.1. *The length of the no-trade gap in the optimal solution that obtains the maximum expected revenue under [Eq. \(9\)](#) is a decreasing function of λ , when the distribution \mathcal{D} satisfies [Assumption 2](#).*

Proof. First, observe that under [Assumption 2](#), the optimal allocation rule is exactly characterized by [Eqs. \(7\) and \(8\)](#), whereby there exists exactly one root for each of the corresponding upper and lower virtual value functions. We now need to show how the roots move, while λ varies. In particular, we will prove that the root of the upper virtual value function increases as λ decreases; the proof that the root of the lower virtual value function decreases is similar. These two combined facts mean that, as λ decreases, the length of the interval of no-trade increases (because each of the roots moves away from p_0), thus proving the desired result.

We finish the proof as follows: consider $\lambda_1 \leq \lambda_2$. The roots p_1, p_2 of the upper virtual value functions $\phi_u(s; \lambda_1), \phi_u(s; \lambda_2)$ for λ_1, λ_2 respectively satisfy

$$\phi_u(p_1; \lambda_1) = 0 \text{ and } \phi_u(p_2; \lambda_2) = (\lambda_2 - \lambda_1)(p_2 - p_0) + \phi_u(p_2; \lambda_1) = 0.$$

Additionally, $(\lambda_2 - \lambda_1)(p_2 - p_0) \geq 0$ since $\lambda_1 \leq \lambda_2$ and $p_2 \geq p_0$ (by definition of the upper virtual value function), so the latter equality means that $\phi_u(p_2; \lambda_1) \leq 0 = \phi_u(p_1; \lambda_1)$. Therefore, by the monotonicity of $\phi_u(s; \lambda_1)$, the latter is equivalent to the ordering $p_2 \leq p_1$ of the respective roots, as desired. \square

[Proposition 4.1](#) allows us to make precise the important interpretation of how the effects of monopoly pricing and adverse selection interplay in a model of trading that incorporates both noise trading and better-informed traders. More specifically, in the optimal solution to the revenue maximization problem of the market maker, there is a minimal interval of no-trade (corresponding to the pure noise trading case, $\lambda = 1$) that exhibits the monopoly pricing power being the sole driving force of the shape of the optimal mechanism. As λ gets smaller, and the market maker's model for the trader pre-supposes more and more traders with perfect information, according to [Proposition 4.1](#), the interval around p_0 on which the AMM is not willing to trade grows larger and larger, and this change is attributed due to adverse selection being added in the mixture because of the traders that are better informed than the market maker. At the other extreme, where pure adverse selection is observed ($\lambda = 0$), the optimal solution is to provide no liquidity, or equivalently, perform no trade. This is a distribution-independent result that is due to the corresponding virtual value functions (both upper and lower) being purely non-positive in that case (c.f., the expressions of [Assumption 2](#) for $\phi_u(s; 0)$ and $\phi_l(s; 0)$ where we substitute $\lambda = 0$), making the optimal solution that obtains the maximum revenue be exactly the no-trade-anywhere solution $x(\hat{p}) = 0, \forall p_{\min} \leq \hat{p} \leq p_{\max}$.

Acknowledgments

The first author is supported in part by NSF awards CNS-2212745, CCF-2212233, DMS-2134059, and CCF-1763970. The second author is supported by the Briger Family Digital Finance Lab at Columbia Business School. The third author’s research at Columbia University is supported in part by NSF awards CNS-2212745, and CCF-2006737.

Disclosures

The first author is a Research Fellow with automated market making protocols, including ones mentioned in this work. The second author is an advisor to fintech companies. The third author is Head of Research at a16z Crypto, a venture capital firm with investments in automated market making protocols.

References

- Alaei, Saeed, Hu Fu, Nima Haghpanah, and Jason Hartline (2013). “The simple economics of approximately optimal auctions”. In: *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. IEEE, pp. 628–637.
- Angeris, Guillermo and Tarun Chitra (2020). “Improved price oracles: Constant function market makers”. In: *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pp. 80–91.
- Angeris, Guillermo, Alex Evans, and Tarun Chitra (2021a). “Replicating market makers”. In: *arXiv preprint arXiv:2103.14769*.
- (2021b). “Replicating monotonic payoffs without oracles”. In: *arXiv preprint arXiv:2111.13740*.
- Angeris, Guillermo, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra (2019). “An analysis of Uniswap markets”. In: *arXiv preprint arXiv:1911.03380*.
- Bulow, Jeremy and John Roberts (1989). “The Simple Economics of Optimal Auctions”. In: *Journal of Political Economy* 97.5, pp. 1060–1090. DOI: [10.1086/261643](https://doi.org/10.1086/261643). eprint: <https://doi.org/10.1086/261643>. URL: <https://doi.org/10.1086/261643>.
- Buterin, Vitalik (Oct. 2016). *Let’s run on-chain decentralized exchanges the way we run prediction markets*. Reddit Post. URL: www.reddit.com/r/ethereum/comments/55m04x/lets_run_onchain_decentralized_exchanges_the_way/ (visited on 07/31/2022).
- Chawla, Shuchi, Jason D Hartline, and Robert Kleinberg (2007). “Algorithmic pricing via virtual valuations”. In: *Proceedings of the 8th ACM Conference on Electronic Commerce*, pp. 243–251.
- Chawla, Shuchi, Jason D Hartline, David L Malec, and Balasubramanian Sivan (2010). “Multi-parameter mechanism design and sequential posted pricing”. In: *Proceedings of the forty-second ACM symposium on Theory of computing*, pp. 311–320.
- Chawla, Shuchi, Jason D. Hartline, and Balasubramanian Sivan (2019). “Optimal crowdsourcing contests”. In: *Games and Economic Behavior* 113, pp. 80–96. ISSN: 0899-8256. DOI: <https://doi.org/10.1016/j.geb.2015.09.001>. URL: <https://www.sciencedirect.com/science/article/pii/S089982561500127X>.
- Chen, Y. and D.M Pennock (2007). “A utility framework for bounded-loss market makers”. In: *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence (UAI 2007)*. Vancouver, BC, Canada, pp. 49–56.

- Cole, Richard and Shravas Rao (Dec. 2017). “Applications of α -Strongly Regular Distributions to Bayesian Auctions”. In: *ACM Trans. Econ. Comput.* 5.4. ISSN: 2167-8375. DOI: [10.1145/3157083](https://doi.org/10.1145/3157083). URL: <https://doi.org/10.1145/3157083>.
- Cole, Richard and Tim Roughgarden (2014). “The sample complexity of revenue maximization”. In: *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pp. 243–252.
- Ewerhart, Christian (2013). “Regular type distributions in mechanism design and ρ -concavity”. In: *Economic Theory* 53.3. The original publication is available at www.springerlink.com, pp. 591–603. ISSN: 0938-2259. DOI: [10.1007/s00199-012-0705-3](https://doi.org/10.1007/s00199-012-0705-3). URL: <https://doi.org/10.5167/uzh-63129>.
- Fan, Zhou, Francisco J. Marmolejo-Cossío, Ben Altschuler, He Sun, Xintong Wang, and David Parkes (2022). “Differential Liquidity Provision in Uniswap v3 and Implications for Contract Design”. In: *Proceedings of the Third ACM International Conference on AI in Finance. ICAIF ’22*. New York, NY, USA: Association for Computing Machinery, pp. 9–17. ISBN: 9781450393768. DOI: [10.1145/3533271.3561775](https://doi.org/10.1145/3533271.3561775). URL: <https://doi.org/10.1145/3533271.3561775>.
- Gelman, Andrew, John B. Carlin, Hal Steven Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin (2014). *Bayesian data analysis*. eng. Third edition. OCLC: 909477393. Boca Raton: CRC Press. ISBN: 9781439898208.
- Glosten, Lawrence R (1989). “Insider trading, liquidity, and the role of the monopolist specialist”. In: *Journal of business*, pp. 211–235.
- (1994). “Is the electronic open limit order book inevitable?” In: *The Journal of Finance* 49.4, pp. 1127–1161.
- Goyal, Mohak, Geoffrey Ramseyer, Ashish Goel, and David Mazières (2022). *Finding the Right Curve: Optimal Design of Constant Function Market Makers*. DOI: [10.48550/ARXIV.2212.03340](https://arxiv.org/abs/2212.03340). URL: <https://arxiv.org/abs/2212.03340>.
- Haghpanah, Nima and Jason Hartline (2015). “Reverse mechanism design”. In: *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pp. 757–758.
- Hartline, Jason D (2012). “Approximation in mechanism design”. In: *American Economic Review* 102.3, pp. 330–336.
- (2021). “Mechanism design and approximation”. In: *Book draft* 122.1.
- Hartline, Jason D and Tim Roughgarden (2009). “Simple versus optimal mechanisms”. In: *Proceedings of the 10th ACM conference on Electronic commerce*, pp. 225–234.
- Lu, Alan and Martin Köppelmann (Mar. 2017). *Building a Decentralized Exchange in Ethereum*. en. URL: <https://blog.gnosis.pm/building-a-decentralized-exchange-in-ethereum-eea4e7452d6e> (visited on 07/31/2022).
- Mantel, Nathan (1976). “Tails of Distributions”. In: *The American Statistician* 30.1, pp. 14–17. ISSN: 00031305. URL: <http://www.jstor.org/stable/2682880> (visited on 01/24/2023).
- Milionis, Jason, Ciamac C. Moallemi, and Tim Roughgarden (2022). *Complexity – Approximation Trade-offs in Exchange Mechanisms: AMMs vs. LOBs*. To Appear in FC 2023.
- Myerson, Roger B. (1981). “Optimal Auction Design”. In: *Mathematics of Operations Research* 6.1, pp. 58–73. ISSN: 0364765X, 15265471. URL: <http://www.jstor.org/stable/3689266> (visited on 01/22/2023).
- Neuder, Michael, Rithvik Rao, Daniel J Moroz, and David C Parkes (2021). “Strategic liquidity provision in uniswap v3”. In: *arXiv preprint arXiv:2106.12033*.
- O’Hara, Maureen (2011). *Market microstructure theory*. eng. repr. Malden, Mass.: Blackwell. ISBN: 9780631207610.

- Pennock, D. N. and R. Sami (2007). “Computational Aspects of Prediction Markets”. In: *Algorithmic Game Theory*. Ed. by N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani. Cambridge University Press. Chap. 26, pp. 651–676.
- Roughgarden, Tim and Inbal Talgam-Cohen (2016). “Optimal and robust mechanism design with interdependent values”. In: *ACM Transactions on Economics and Computation (TEAC)* 4.3, pp. 1–34.
- (2019). “Approximately optimal mechanism design”. In: *Annual Review of Economics* 11, pp. 355–381.
- Yin, Jimmy and Mac Ren (2021). “On Liquidity Mining for Uniswap v3”. In: *arXiv preprint arXiv:2108.05800*.