

The Case of FBA as a DEX Processing Model

Tiantian Gong^{1*}, Zeyu Liu¹, and Aniket Kate^{2,3}

¹ Yale University, New Haven, USA,
{tiantian.gong, zeyu.liu}@yale.edu

² Purdue University, West Lafayette, USA,
aniket@purdue.edu

³ Supra Research

Abstract. We investigate the welfare loss of *continuous* and *discrete* order matching models in blockchain-based decentralized exchanges (DEX) that utilize order books to record outstanding orders. *Continuous processing* matches each incoming transaction against the current order book. The *discrete processing* model, i.e., *frequent batch auction* (FBA), executes transactions discretely in batches with a uniform price double auction: Orders are first matched according to *price*, then the exact transaction order if competing orders specify the same price.

We find that FBA imposes less **welfare loss** and provides better **liquidity** than continuous processing in typical scenarios, e.g., when few parties are *privately informed* about asset valuations. Even otherwise, it achieves better social welfare and liquidity provision in the following settings: when price takers and public information reflecting asset value changes arrive sufficiently frequently compared to private information, when the priority fees (for faster transaction inclusion into blockchains) are small, or when the market is more balanced on both buy and sell sides. Our empirical analysis on the BTC-USD and ETH-USD transactions on a DEX named dYdX indicates that FBA can reduce transaction costs by 21% – 37%.

1 Introduction

Blockchain-based decentralized exchanges (DEX) enable users to trade assets without relying on intermediary authorities through blockchain transactions. They are already valued at several billion USD and are expected to grow multi-fold in the coming years. DEXes typically take one of the two following forms. (1) In a limit order book-based design (e.g., Penumbra [23, 29], Uniswap limit orders [24], CoW Swap [12]), each transaction is a limit order (which specifies the direction⁴, price, and quantity of trade), and orders are matched individually. A set of outstanding limit orders is called a limit order book. (2) An automated market maker design (e.g., Uniswap [1]) features automatic algorithmic trading

* This work is mostly done while the author was a Ph.D. candidate at Purdue University.

⁴ Sell or buy

where the settling price ensures certain pre-defined function behaves in a pre-determined way.

In this paper, we analyze order book-based DEX designs with a focus on comparing two applicable *transaction processing models*, i.e., *continuous* processing (CLOB, continuous limit order book) where transactions are handled one by one in order or *discrete* frequent batch auction (FBA) where transactions are aggregated at a certain frequency and settled in a double auction (i.e., bids are from both sellers and buyers). Similar to the traditional centralized exchanges (CEXs), many order book-based DEX designs follow the CLOB model. Two main metrics to compare the performance of the two designs are (1) *welfare loss*, which measures the extra transaction costs paid by a common user and is determined by the adverse selection risk (explained below) and the adopted processing model (either CLOB or FBA); and (2) *liquidity provision*, which measures the extent to which an on-demand user’s transaction is settled (against the order book)⁵. We evaluate the suitability of FBA in the blockchain setting by uncovering the scenarios where FBA provides better welfare and liquidity.

Why consider FBA. While FBA is less considered, it can potentially fit the DEX setting better. In addition to the discussions in the context of centralized exchanges [8, 16], there are several other intuitive reasons. First, inherently, blockchains already treat time as discrete when ordering and assembling transactions into blocks, the same as in FBA. Second, when transactions are executed continuously (via CLOB), one can observe and act on the “future” if the pending transaction pool is not fully hidden, thus rendering one of CLOB’s main advantages moot. Similarly, latencies in block generation and message transmission along with priority fees allow even more space for latency arbitrage rents in DEXes, making CLOB even less favorable. Lastly, a common critique of FBA, the non-execution risks caused by nontransparent pre-trade order book, is alleviated in public blockchains where at least part of the transactions are overt.⁶

Analysis framework. In the DEX we consider three types of players: (1) *common investors* who submit inelastic trading orders, (2) *informed traders* who submit trades after observing publicly or privately available information about asset valuation jumps, and (3) (many) *arbitrageurs* who can apply market-making (i.e., placing orders to facilitate trades) and front-running strategies. Arbitrageurs are addressed as *liquidity providers* (LPs) when applying the market-making strategy, and as *front-runners* when front-running investors, traders, or other LPs. Our focus is then on how LPs set up their quotes, which decide the amount by which the asking price exceeds the bid price, i.e., the bid-ask spread.

Putting ourselves in LPs’ shoes, common investors are price takers and thus the *source of profits*. Informed traders and front-runners cause *losses* to LPs by reacting to information faster, which is also called the *adverse selection risk* [8, 17]. Hence, the arrival rates of common investors and the changes in public or

⁵ The bid-ask spread is usually viewed as a proxy for measuring liquidity.

⁶ In centralized exchanges employing FBA [11], quotes are hidden until being matched. For CEX, displaying orders directly could violate certain regulations [7], e.g., Rule 610 in Regulation National Market System.

private information about asset valuation affect how LPs determine the bid-ask spread. An LP does not set it to be too high, considering that other arbitrageurs can improve on it (i.e., setting a lower spread) to trade with common investors and earn a profit. The LP does not make it too small either, to absorb the potential loss from adverse selection. Here, the *markup* set aside to tolerate this adverse selection risk is the *welfare loss* of our interest. It is considered a “loss” of the welfare of common investors and traders because this portion of the spreads does not need to be paid when there is no adverse selection. Aside from the welfare loss, we are also interested in *liquidity provision*, which is measured with the bid-ask spread itself. Overall, the spread consists of the markup (i.e., the welfare loss) and the price impact of new incoming orders (because it may reflect public or private information about the underlying asset’s price changes).

1.1 Contributions

In summary, we make two contributions. First, we analyze the theoretical benefits of employing FBA (frequent batch auction) as the transaction processing model for DEX compared to CLOB (continuous limit order book). Specifically, we uncover typical market conditions where FBA provides better welfare and liquidity. Second, we conduct empirical analysis to straightforwardly compare the transaction costs under the two processing approaches.

Welfare and liquidity provision. For welfare loss, if there is *no* private information in DEXes, then FBA always imposes a *zero* welfare loss while CLOB has a positive welfare loss. If a specific DEX design allows privately informed parties, e.g., a board member or a developer capable of dictating protocol updates, adverse selection is present under FBA. In this case, FBA inflicts positive welfare loss. In this setting with private information, we find that first, FBA imposes less welfare loss than CLOB if common investors and public information concerning asset value changes arrive sufficiently often as compared to private information. This is because, under FBA, LPs have more time to respond to public information, and less private information reduces adverse selection. Second, the smaller the transaction priority fees [15] are, the higher the markups under CLOB are, as compared to under FBA.⁷ This is because the profits from front-running increase as costs decline, resulting in the LPs charging higher markups to counter adverse selection. Finally, a more balanced market decreases the welfare loss for FBA, making it typically less than the welfare loss under CLOB. In essence, when sufficiently many transactions are accumulated and matched among themselves, the welfare loss in FBA decreases.

In terms of liquidity provision, FBA has smaller bid-ask spreads (thus better liquidity provision) under similar conditions.

Empirical analysis. Overall, if the public blockchain that performs the DEX function has much less private information than public information, FBA provides lower welfare loss and better liquidity. More specifically, assuming that the

⁷ Average fee per transaction is \$1.18 for Ethereum between September 11-October 10 in 2024, during which the average daily price changes of ETH is \$49.37.

system has little private information, using real-world transactions acquired from the dYdX exchange [20] from January 2023 to October 2024, we find that empirically there is a 21%-37% increase in transaction costs when the transactions are executed with CLOB compared with using FBA.

2 Model and definitions

In this section, we specify the system and trading system models, and then introduce game theoretic concepts.

2.1 System model

A set of n processes called validators run a secure blockchain system that mitigates *transaction order manipulation attacks* with order-fair atomic broadcast (of-ABC) [10] (Definition 1). Up to f validators are Byzantine and behave arbitrarily. They communicate via reliable authenticated point-to-point channels. More specifically, they are connected with Byzantine fault-tolerant first-in-first-out (FIFO) *consistent broadcast links* [10] that securely deliver messages. The network is *partially synchronous* where the network imposes some known bounded message delay after a global stabilization time (GST).

Definition 1 (Differentially of-ABC [10]). *A secure κ -differentially of-ABC protocol satisfies:*

1. **Agreement:** *If a message m is delivered by some correct process, then m is eventually delivered by every correct process.*
2. **Integrity:** *No message is delivered more than once.*
3. **Weak validity:** *If all processes are correct and broadcast a finite number of messages, every correct process eventually delivers these broadcast messages.*
4. **Total order:** *Let m and m_0 be two messages, P_i and P_j be correct processes that deliver m, m_0 . If P_i delivers m before m_0 , P_j also delivers m before m_0 .*
5. **κ -differential order fairness:** *If $b(m; m_0) > b(m_0; m) + 2f + \kappa$ (where $b(x; y)$ counts the number of processes that consistent-broadcast x before y), then no correct process delivers m_0 before m .*

Note that we can adopt alternative order manipulation techniques. We only need to re-calibrate the front-running success probability.

Front-running success probability. Suppose after observing the transaction m_i of a user i , arbitrageur j decides to front-run m_i with transaction m_j . To compute j 's success probability in of-ABC, we assume the distribution of latencies on different communication links is known. The links connecting the validators have latency distribution F_v ; the links between common users (i.e., investors and traders) and validators have latency distribution F_{it} ; the links between arbitrageurs and validators have latency distribution F_a .

Let D be the random variable for the latency differences to independently deliver two messages on a communication link between validators. Let D follow distribution \bar{F}_{lt} . For example, if F_v is instantiated with normal distribution

$N(1, 1)$, then \tilde{F}_{lt} follows distribution $N(0, 2)$. Let \tilde{F}_{lt}^C be the CDF of \tilde{F}_{lt} . Assuming the worst case where j sees m_i immediately after it arrives at some validator, j then immediately sends m_j to a validator, and j always wins when neither transaction is broadcast first by more than $(2f + \kappa)$ validators. We denote j 's winning probability as p^* and compute it as follows:

$$p^* = 1 - \mathbb{P}[b(m_i; m_j) > b(m_j; m_i) + 2f + \kappa] = 1 - \sum_{s=2f+\kappa-1}^{n-1} \binom{n-1}{s} [\tilde{F}_{lt}^C(0)]^s [1 - \tilde{F}_{lt}^C(0)]^{n-1-s}$$

If Byzantine validators do not relay messages, we replace $(n-1)$ with $(n-f-1)$. In the above example where $F_v \sim N(1, 1)$ and $\tilde{F}_{lt} \sim N(0, 2)$, we have $p^* = 0.75$ for $n = 10, f = 3, \kappa = 1$ and $p^* = 0.95$ for $n = 31, f = 10, \kappa = 1$.

2.2 Trading system

We adapt the dynamic trading models in Eibelshäuser and Smetak [14] and Budish et al. [8] in traditional centralized exchanges to the blockchain-based DEX. The frequency of the batch auction is every one or multiple blocks. We denote the length of this period as I .

We consider an asset X with changing fundamental value V_t at time t ($t \in [0, T]$, for some $T > 0$). We assume an observable signal that equals the fundamental value of the asset and evolves according to a compound Poisson jump process, drawn from a symmetric distribution F_{jp} with arrival rate λ_{jp} , bounded support, and zero mean. The fundamental value jump can be observed as both *public* and *private* information. We capture the absolute value of the jump with the random variable J .

We model three types of risk-neutral⁸ trading players:

- **Investors** who arrive stochastically with probability λ_i and have inelastic demand, with buying and selling being equally likely.
- **Informed traders** who can observe both the public information arriving stochastically with probability λ_{pb} and the private information arriving stochastically with probability λ_{pr} . Both types of information affect V_t positively or negatively with equal probability.
- **r arbitrageurs** who can apply market-making or front-running strategy. They can also observe and respond to the public information mentioned above (arriving with probability λ_{pb}). At time t , they intend to sell at a price higher than V_t or buy at a price lower than V_t .

The arbitrageurs here can be trading firms and arbitrage bots, among others. As mentioned in §1, a front-running arbitrageur can front-run LPs, investors, and traders. Unlike in the modeling in centralized exchanges, the front-runner now can act after others have already submitted transactions.

Trading system states. We capture the history information at time t with \mathcal{H}_t , and it is observable by all. The state of the trading system at time t is a tuple $S_t = (P_t, J_t, (b_t, a_t), \vec{g})$, where $P_t = \mathbb{E}[V_t | \mathcal{H}_t]$ is the expected value of the fundamental value of the asset given the observed history, J_t is sampled from distribution J (i.e., $J_t \sim J$), (b_t, a_t) are quotes of the bid and ask prices, and

⁸ i.e., obtaining the same amount in expectation generates the same utility.

\vec{g} records the probabilities that arbitrageurs respond to investors' and traders' orders with front-running at each level of the order book.

Fees. The exchange collects fees from both sides of each settled trade. There are a base fee and a priority fee F which can promote the order of a transaction during tie-breaking. In the analysis, we disregard the base fee since it is the same for all transactions and is burnt in some systems [15].

2.3 Definitions - solution concepts

In the trading game, r arbitrageurs are the strategic players and aim to maximize their utilities by playing the best strategy. A *strategy* is a probabilistic distribution over possible actions, with all mass condensed at one action for a pure strategy. A *solution concept* then describes a profile or snapshot of all players' strategies with certain desired properties, e.g., Nash equilibrium discourages unilateral deviation. We naturally adopt stationary Markov Perfect Equilibrium (MPE) [26] for CLOB since the parameters of player and information arrival rates in the stochastic game are time-independent.

Definition 2 (Stationary MPE [26]). *A Nash equilibrium (NE) is a strategy profile \vec{s} where no player increases utility by unilaterally deviating from \vec{s} . A subgame perfect equilibrium (SPE) is a strategy profile \vec{s} that forms an NE for any subgame of the original game. An MPE is an SPE in which all players play Markov strategies, i.e., strategies that depend only on the current state of the game. A stationary MPE is an MPE where strategies are time-independent.*

We utilize a weaker notion, Order Book Equilibrium (OBE) for FBA since stationary MPE does not always exist for FBA [9]. Intuitively, this is because under FBA, when an arbitrageur provides liquidity at the bid-ask spread that equals costs from market-making, other arbitrageurs do not have the incentive to undercut (by improving the current quotes). Since others do not undercut, the arbitrageur has the incentive to widen the spread to increase profits, which is a unilateral action to increase the player's utility. But this would result in others undercutting the widened quotes. We first define OBE and discuss why adopting the weaker notion still gives a fair comparison.

Definition 3 (OBE [9]). *Given state S_t , an OBE at time t is a set of orders submitted by all arbitrageurs such that the following hold: There exist (1) no safe profitable price improvements and (2) no other safe profitable deviations.*

Here, a profitable price improvement is safe if it remains strictly profitable after other arbitrageurs take profitable responses, e.g., liquidity withdrawals, after the improvement. A profitable deviation is safe if it remains strictly profitable after other arbitrageurs react with safe profitable price improvements or liquidity withdrawals. Price improvements are liquidity provisions improving current quotes, and liquidity withdrawals are cancellations of limit orders. See [9] for formal justifications for OBE. OBE is weaker than MPE in the sense that it allows the existence of unilateral deviations that increase utility as long as they

can be made unprofitable by others' reactions. This is reasonable as FBA leaves time for others to respond to one's actions, and thus any reasonable player would not. Therefore, the players in CLOB and FBA would follow the corresponding equilibrium (MPE and OBE, the strongest equilibrium respectively), and the induced welfare losses are comparable.

3 DEX under CLOB and FBA

We now model DEXes under the CLOB and FBA. We then detail the events and provide the equilibrium analysis focusing on welfare loss and liquidity provision.

3.1 Order of events

We abstract the DEX as run by validators maintaining a secure blockchain. Arbitrageurs, investors, and traders submit transactions by sending messages to one or more validators, who then broadcast received messages to each other. In the following stochastic trading game \mathcal{G} , we capture the arrival of the three types of players and how the validators order and execute the transactions.

Trading game \mathcal{G} : Repeat the following for asset X with initial value V_0 .

- Ⓐ Arbitrageurs place orders on the exchange.
- Ⓑ One of the following events is then triggered:
 - (1) An investor arrives with probability λ_i and submits an order to the exchange. This may change order book quotes (\vec{b}_t, \vec{a}_t) at time t but not to the fundamental value.
 - (2) A private information event occurs with probability λ_{pr} , resulting in a jump in fundamental value, and an informed trader submits an order to the exchange. This may change the order book quotes.
 - (3) A public information event occurs with probability λ_{pb} , resulting in a jump in fundamental value, and arbitrageurs can submit withdrawals of existing orders, place new orders, or front-run stale quotes.
 - (4) Null event. The fundamental value and order books do not change.
- Ⓒ Each validator accumulates transactions and steps Ⓐ and Ⓑ are repeated until the validators can order transactions and output a block according to the underlying blockchain protocol.
- Ⓓ CLOB. After outputting a block, the validators running the exchange execute the transactions sequentially in the proposed order.
- FBA. Depending on the auction frequency, after outputting one or multiple blocks, the validators execute the transactions by clearing the market with a uniform price double auction: (1) the validators first gather the orders in the current batch and all previous outstanding orders; (2) they aggregate the bids and asks; (3) if there are intersections, the market clears where supply meets demand, at a uniform price. When there are conflicting ties, priority fees lift the execution order of transactions, and further ties are broken uniformly at random. Unmatched orders that are not canceled enter into the next trading game as outstanding orders.

3.2 MPE under CLOB

As described before, we aim to solve for the markup of LPs. It is contingent on the profits from satisfying investors' orders and the costs of defending against adverse selection. We denote the expected fundamental value of asset X conditioned on history and a new sell or buy order respectively as $\mathbb{E}[V_t|\mathcal{H}_t, \text{buy}]$ and $\mathbb{E}[V_t|\mathcal{H}_t, \text{sell}]$. For clarity, without loss of generality, we let each level of the book be one unit order. Let Q be the maximum number of units investors need and p_j be the probability of an investor transacting j units ($j = 1, \dots, Q$). We state the results for the quotes on the limit order book (LOB) under CLOB as follows.

Theorem 1. *There exists a stationary MPE in trading game \mathcal{G} under CLOB processing. Bid and ask prices at the k -th level of the LOB in equilibrium satisfy*

$$b_t^k = \mathbb{E}[V_t|\mathcal{H}_t, \text{sell}] = P_t - \frac{s_k}{2}, a_t^k = \mathbb{E}[V_t|\mathcal{H}_t, \text{buy}] = P_t + \frac{s_k}{2}$$

where s_k satisfies

$$\lambda_i \sum_{j=k}^Q p_j \frac{s_k}{2} - \lambda_{pr} (1 + p^* \frac{\bar{g}_k}{\bar{g}_k(r-2)+1}) \bar{J}_k - \lambda_{pb} \bar{J}_k + (\lambda_i + \lambda_{pr}) \bar{g}_k F = 0$$

Here, $\bar{J}_k = \mathbb{P}[J > \frac{s_k}{2}] \mathbb{E}[J - \frac{s_k}{2} | J > \frac{s_k}{2}]$, and \bar{g}_k is the probability that an arbitrageur front-runs traders and investors at the k -th level. In equilibrium, \bar{g}_k is then updated to take the value that maximizes profits from front-running traders and investors, i.e., $\lambda_{pr} p^* \frac{\bar{g}_k}{\bar{g}_k(r-2)+1} \bar{J}_k - (\lambda_i + \lambda_{pr}) \bar{g}_k F$. The expected price impact of an incoming k -unit order is $\Delta_k = \tilde{J}_k \frac{\lambda_{pr}}{\lambda_{pr} + \lambda_i}$ with $\tilde{J}_k = \mathbb{P}[J > \frac{s_k}{2}] \mathbb{E}[J | J > \frac{s_k}{2}]$; the expected markup from liquidity providers is $(\lambda_{pr} + \lambda_i)(\frac{s_k}{2} - \Delta_k)$.

The key to the proof (which resides in Appendix A) is that in equilibrium, market-making and front-running strategies yield the same profits because arbitrageurs can choose strategies freely.

Interpret Theorem 1. The spread s at each level of the order book consists of (1) profits for market-making and (2) the part to absorb the price impact Δ from orders driven by private information. The price impact exists because arbitrageurs cannot distinguish common trades that do not carry information (which affects asset fundamental values) from privately informed trades. The expected welfare loss of investors and traders is contained in the profits, i.e., $(\lambda_{pr} + \lambda_i)(\frac{s}{2} - \Delta)$, which we address as markups. They are affected by the arrival rates of investors (λ_i), private information (λ_{pr}), and public information (λ_{pb}), the front-running probability (\bar{g}) and its success probability (p^*), priority fees F , jump size, the number of arbitrageurs (r), and transaction sizes. Below, we discuss how the markup within the spread is affected intuitively. In §4.3, we show how these formulas are reflected using real-life data, and in §4.4, we give an example demonstrating this theorem.

What affects the markups under CLOB. From the theorem, overall, the markup from LPs in equilibrium increases with the jump size, public information releases, the front-running probability, and its success probability. It decreases

with private information arrivals, fees, and the number of arbitrageurs. Fix a specific front-running probability \bar{g} , the markup first increases as the investor arrival rate rises and then decreases as it continues to grow. This is because λ_i raises market-making profits from common investors. In the meantime, it also lowers the profits from front-running traders due to increased costs (i.e., front-running arbitrageurs attack unprofitable common trades more often). Since arbitrageurs can choose freely between market-making and front-running, more arbitrageurs adopt the market-making strategy when market-making is more profitable. The competition reduces the markup until equilibrium is reached, where the profits from both strategies are equal.

3.3 OBE under FBA

Under FBA, we only need to process excessive demand or supply since only those transactions need to be fulfilled by LPs. We assume the withdrawals and updates of existing quotes can be completed in time during the batch auction interval. Let Z_I denote the excess demand (demand minus supply) during a batch interval I . $Z_I \leq Q + 1$ ($Q \geq 0$ is an integer) is bounded. We borrow \tilde{J}_k and \tilde{J}_k from Theorem 1 and formally state the result for FBA in Theorem 2.

Theorem 2. *There exists an OBE in pure strategies in trading game \mathcal{G} under FBA processing. If $|Z_I| \leq Q + 1$, bid and ask quotes in equilibrium satisfy*

$$\begin{cases} b_{(I)}^k = P_I - \Delta_k - M_k & Z_I < 0 \\ a_{(I)}^k = P_I + \Delta_k + M_k & Z_I > 0 \end{cases}, \quad k = 1, \dots, Q + 1$$

where $\Delta_k = \tilde{J}_k \frac{\lambda_{pr}}{\lambda_{pr} + \lambda_i}$ is the expected price impact from informed traders, $M_k = \sum_{u=k}^Q \Delta_u \prod_{v=k}^u \alpha_v$ for $k = 1, \dots, Q$, and $M_{Q+1} = 0$ are the markups from LPs, $\alpha_k = \frac{q_{k+1}}{q_k + q_{k+1}}$ and $q_k = \mathbb{P}[Z_I = k \mid |Z_I| \leq Q + 1]$. The expected markup per unit time is $\frac{2}{I} \sum_{k=1}^Q k q_k M_k$.

The proof utilizes induction on k . Intuitively, first, the ask quote for the $(Q+1)$ -th level LOB has 0 markup as otherwise, another arbitrageur can improve on it to settle the $(Q+1)$ -th transaction. This means that the last quote reflects only the price impact. In the induction process, given the $(k+1)$ -th quote ($k \leq Q$), we can determine the k -th markup by considering safe price improvements or other deviations in different cases of Z_I , which can take values from k (with probability q_k), $k+1$ (with probability q_{k+1}), to $Q+1$ (with probability q_{Q+1}), and different cases of an LP, who may already own a lower-level quote, an upper-level quote, or no existing quote. The full proof can be found in Appendix A.

Interpret Theorem 2. Same as before, the spread consists of the markup and the cushion for absorbing the price impact from adverse selection. The expected markup or welfare loss is influenced by the arrival rates of investors (λ_i), privately informed traders (λ_{pr}), the jump size, and the excessive demand. Again, we discuss these relationships below, show real-world data-based analysis in §4.3, and give an example in §4.4.

What affects the markups under FBA. Overall, the welfare loss in equilibrium increases with the jump size, private information releases, and excessive demand, and decreases with investor arrivals. First, when the jump size or private information arrival rates are larger or the investor arrival rates are smaller, the adverse selection is more severe, and the expected price impact rises. This in turn results in higher markups. Second, when there is more excessive demand (or supply), then the markups on the ask (or bid) quotes grow to incorporate more levels of price impacts on the LOB.

4 Welfare loss and liquidity provision

We compare the welfare loss and liquidity provision under CLOB and FBA in equilibrium in §4.1 and §4.2, give an empirical analysis with real world data in §4.3, and show an example with small parameters in §4.4.

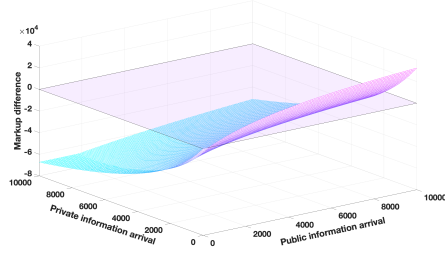
4.1 Welfare loss

This section compares FBA and CLOB under special cases and general conditions and ends with the main takeaways.

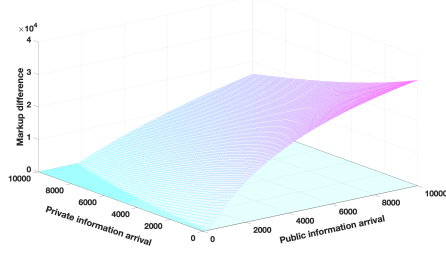
Special cases. When there are no privately informed traders ($\lambda_{pr} = 0$) but there exists public information ($\lambda_{pb} > 0$), FBA is strictly better than CLOB because the markup in FBA becomes zero while the markup in CLOB is still positive. As depicted in Fig. 1 and 3, FBA always suffers less welfare loss when $\lambda_{pr} = 0$. When there is no public information ($\lambda_{pb} = 0$) but there exists private information ($\lambda_{pr} > 0$), FBA has positive markups that increase with λ_{pr} . For CLOB, the markups tend to zero for small λ_{pr} but remain positive and also increase with λ_{pr} otherwise. This is also portrayed in Fig. 1 where the markups for both processing models approach zero.

General comparison. The parameters affecting markups in both models are $\lambda_i, \lambda_{pr}, \lambda_{pb}$. We examine how they affect the markup differences in CLOB and FBA in four different settings (a)-(d). In setting (a), we pick the bad case for the remaining effective parameters for both CLOB and FBA. This means a high front-running success probability, small priority fees for CLOB, and high excessive demand for FBA. Parameters in settings (b)-(d) are set similarly to reflect the combinations of good- and bad-case scenarios for the two models. Additionally, in FBA, the excess demand Z_I follows a truncated Skellam distribution as the arrival rates of investors and traders follow the compound Poisson jump process. Therefore, when computing the markup differences, we consider the upper bound for the markup in FBA. This means that FBA can perform *strictly better* than the predictions.

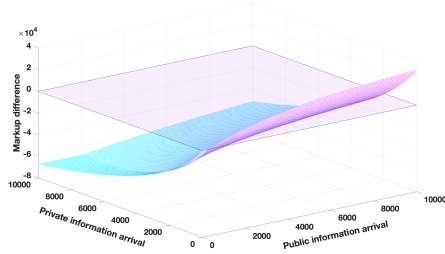
Public versus private information. Overall, as shown in Fig. 1, more public information renders FBA appealing since markups in CLOB increase due to front-running. As a result, FBA performs better when λ_{pb} is high compared to λ_{pr} . More private information gives CLOB an advantage by increasing the spread in CLOB. As for the state of the market, if it is one-sided and there is a large



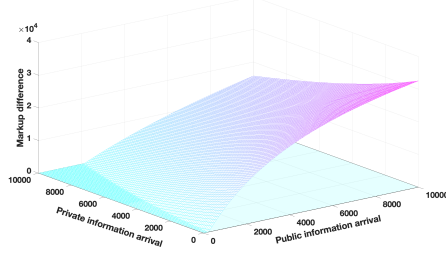
(a) Bad case for both CLOB and FBA.



(b) Bad case for CLOB and good case for FBA.

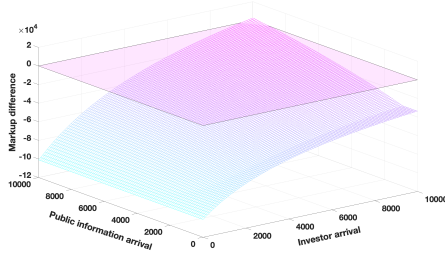


(c) Good case for CLOB and bad case for FBA.

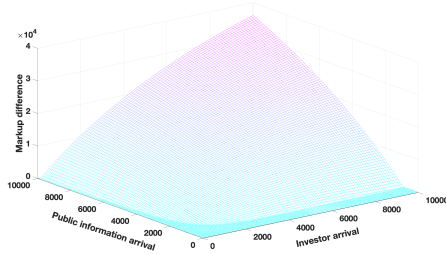


(d) Good case for both CLOB and FBA.

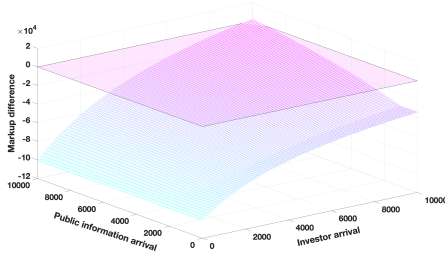
Fig. 1: The markup difference between CLOB and FBA with respect to public and private information arrivals. Positive regions are where FBA has fewer markups, i.e., less welfare loss. The investor arrival rate λ_i is set to 5000. Increasing (decreasing) λ_i pushes the surface up (down).



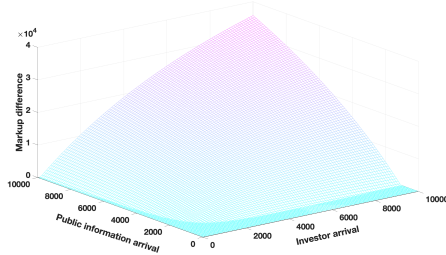
(a) Bad case for both CLOB and FBA.



(b) Bad case for CLOB and good case for FBA.



(c) Good case for CLOB and bad case for FBA.



(d) Good case for both CLOB and FBA.

Fig. 2: The markup difference between CLOB and FBA with respect to public information and investor arrivals. The private information arrival rate λ_{pr} is set to be 5000. Increasing (decreasing) λ_{pb} pushes the surface down (up).

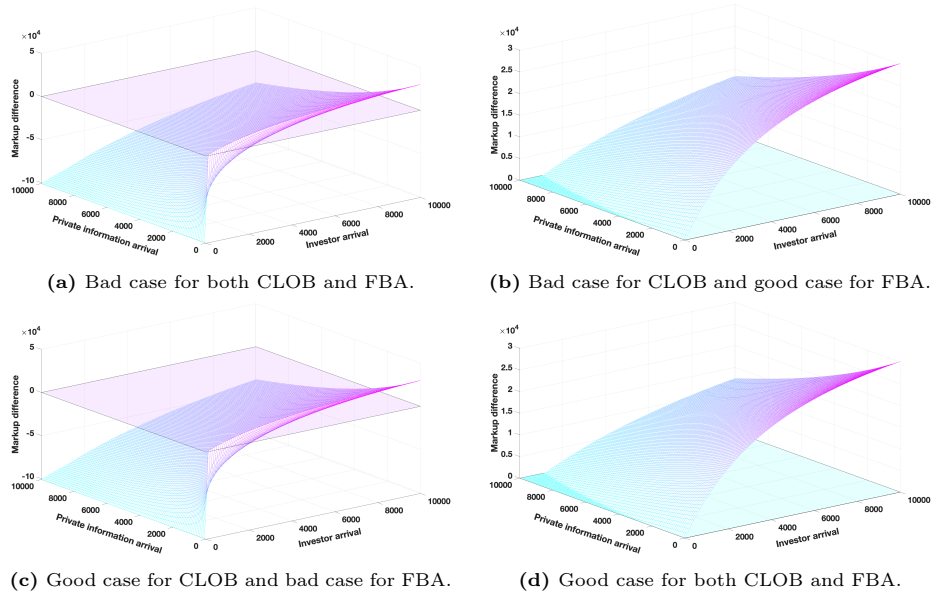


Fig. 3: The markup difference between CLOB and FBA with respect to private information and investor arrivals. The public information arrival rate λ_{pb} is set to be 5000. Increasing (decreasing) λ_{pb} pushes the surface up (down).

excessive demand, CLOB outperforms FBA and FBA realizes lower welfare loss only when λ_{pr} is small. When the market is balanced and submitted orders are mostly settled among themselves, FBA has lower welfare loss when λ_{pb} surpasses λ_{pr} by a smaller amount. This can also be observed in Fig. 2 and 3: in the good cases for FBA where the market is more balanced, FBA outperforms CLOB in the majority of the parameter regions.

Investor arrivals versus public information. We know from Theorem 2 that the markups for FBA do not depend on public information while markups under CLOB increase with λ_{pb} . As a result, FBA performs better in terms of welfare as λ_{pb} increases, which is also evident in Fig. 2. Investor arrivals have a mixed effect. FBA experiences a smaller price impact as λ_i grows larger, diminishing its markup. CLOB also faces a smaller price impact. And larger λ_i decreases the profits from front-running and more arbitrageurs compete as LPs, reducing the markups. Overall, FBA realizes better social welfare if sufficiently many market participants are price takers or are publicly informed.

Investor arrivals versus private information. As described in §3.3, FBA's welfare loss decreases with the investor arrival rate and increases with private information. CLOB's induced welfare loss decreases with λ_{pr} and also eventually decreases with λ_i . As shown in Figure 3, in the bad case for FBA, its markups are amplified by the excessive demand and only larger λ_i absorbs this effect. In the good case for FBA, its markups are already small even when λ_i is minute.

Main takeaways. In a nutshell, we believe FBA is more suitable for the setting of blockchain systems, since we believe blockchain markets tend to have (a lot) more public information than private information. For example, for Bitcoin and Ethereum, all the core codes and projects are open-sourced, and the transactions are overt on-chain. Furthermore, the main goal of blockchain systems is to be decentralized and to have less hidden information. Under this belief, FBA tend to have less welfare loss, as also shown in §4.3.

4.2 Liquidity comparisons from bid-ask spread

We now turn to compare the two execution models from a liquidity provision perspective, with the bid-ask spread being the proxy quantity that we examine. The spread consists of the price impact from orders driven by information about assets’ valuation changes, and markups from LPs.

We include the figures depicting how $\lambda_i, \lambda_{pb}, \lambda_{pr}$ affect the spreads in FBA and CLOB in Figures 5 to 7 (in Appendix B). Similar to the analysis with welfare loss, λ_{pr} approaching 0 always benefits liquidity provision in FBA. For a larger λ_{pr} , FBA has better liquidity provision only when λ_i, λ_{pb} are sufficiently high with respect to λ_{pr} . When the market is thin or one-sided and more orders are fulfilled with LPs’ orders, this requirement on λ_i and λ_{pb} is even more demanding. When the DEX has a more balanced market and a larger proportion of the orders can be filled among themselves, the requirement is looser.

Therefore, similar to the previous section, we conclude that FBA provides better liquidity in the blockchain settings, under the assumption that the blockchain market has much more public information than private information.

4.3 Empirical analysis

We sample 707,267 BTC-USD transactions and 786,727 ETH-USD transactions on dYdX [20] from January 2023 to October 2024 via Tardis [32]. We then simulate CLOB and FBA processing on the order book at the receiving time (measured in nanoseconds) of each transaction. The simulation code and result summary are available here [3]. We compute the realized spread [19] as the welfare loss. We trim outliers in the computed realized spread value with three median absolute deviations for a more robust comparison.

For BTC-USD, the realized spread of CLOB has a mean of 0.0864 and a median of 0.0643; the realized spread of FBA has mean values (0.0688, 0.0652, 0.0632) and median values (0.0520, 0.0490, 0.0476) for auction frequency of 5, 10, and 15 seconds. CLOB inflicts 24%-37% more transaction costs. As shown in Figure 4a, the realized spread for FBA, especially with 5-second auction frequency is more concentrated at the smaller end. The actual clearing price is on average 1.06, 1.74, and 2.20 better than the posted prices of settled trades under FBA with auction frequency of 5, 10, and 15 seconds. A longer auction period allows for the arrival of more transactions, bringing about more and at least not worse trading opportunities.

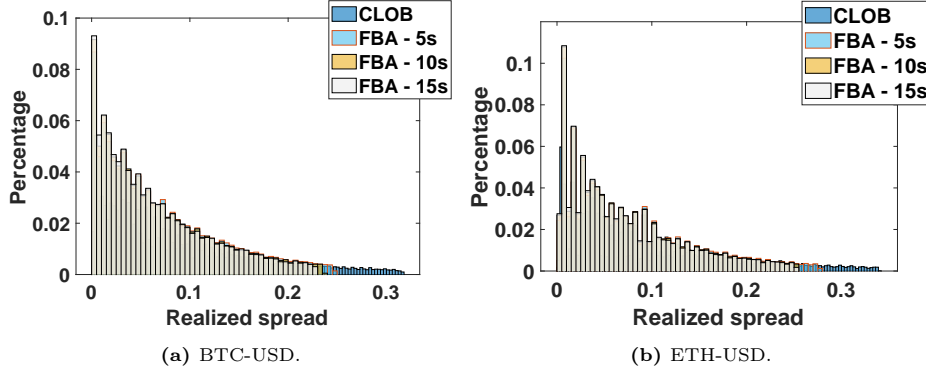


Fig. 4: Distribution of realized spread (with outliers trimmed) from simulations on sampled transactions. The bars for depicting the four execution methods are placed on top of each other. The bars for FBA with 15 second auction frequency concentrate on the lower end, while the rest stretch to higher realized spread, with CLOB having the rightmost bars.

For ETH-USD, the realized spread of CLOB has a mean value 0.0950 and a median value 0.0710. FBA has mean realized spreads of (0.0770, 0.0720, 0.0709) and median realized spreads of (0.0586, 0.0550, 0.0544) for auction frequency of 5, 10, 15 seconds. CLOB has 21%-34% more transaction costs. The clearing price is on average 0.05, 0.11, and 0.12 better than the posted prices of the settled trades under FBA with an auction frequency of 5, 10, and 15 seconds.

In summary, for the two trading targets during the sampled time window, FBA reduces welfare loss and saves transaction costs for users. Note that for the lack of more suitable datasets, we acquired the tick-level order book snapshots and transactions from an exchange that utilizes the CLOB processing model. If FBA were adopted, more transactions could have been settled since placed orders can be matched against each other, while the order books could have had different spreads.

4.4 An example with small parameters

While previous comparisons and empirical analysis are relatively thorough, one thing they lack is reflecting how all different parameters in Theorems 1 and 2 interact. To this end, we provide a small-parameter example as a demonstration.

Welfare loss. Without loss of generality, we consider a unit order for CLOB and a specific jump size J . Under CLOB, let $a = 1 + p^* \frac{g_1}{g_1(r-2)+1}$ and the spread in Equation (1) satisfies $\frac{s}{2} = \frac{(\lambda_{pb} + a\lambda_{pr})J - (\lambda_i + \lambda_{pr})g_1F}{\lambda_{pb} + a\lambda_{pr} + \lambda_i}$. And the price impact from informed trader is $\Delta = J \frac{\lambda_{pr}}{\lambda_{pr} + \lambda_i}$. Then the markup can be computed as $M_C = J \frac{\lambda_i \lambda_{pb} + (a-1)\lambda_i \lambda_{pr}}{(\lambda_i + a\lambda_{pr} + \lambda_{pb})(\lambda_{pr} + \lambda_i)} - \frac{(\lambda_i + \lambda_{pr})g_1F}{\lambda_{pb} + a\lambda_{pr} + \lambda_i}$ and the expected markup is $J \frac{\lambda_i \lambda_{pb} + (a-1)\lambda_i \lambda_{pr}}{\lambda_i + a\lambda_{pr} + \lambda_{pb}} - \frac{(\lambda_i + \lambda_{pr})^2 g_1F}{\lambda_{pb} + a\lambda_{pr} + \lambda_i}$. This increases with the jump size J and de-

creases with priority fee F . For simplicity, we can measure F relative to the jump, i.e., let $F = xJ$ for some $x < 1$.

Under FBA, we have that $\Delta_k = J \frac{\lambda_{pr}}{\lambda_{pr} + \lambda_i}$. Consider $Q = 2$. The expected markups become $\frac{2}{I} J \frac{\lambda_{pr}}{\lambda_{pr} + \lambda_i} (q_1(\alpha_1 + \alpha_1\alpha_2) + 2q_2\alpha_2)$. For simplicity and without loss of generality, we let $q_3 = q_2 = 1/16, q_1 = 1/8$, FBA is better if $\frac{\lambda_i \lambda_{pb} + (a-1)\lambda_i \lambda_{pr} - (\lambda_i + \lambda_{pr})^2 \mathbf{g}_1 x}{\lambda_i + a\lambda_{pr} + \lambda_{pb}} > \frac{\lambda_{pr}}{I(\lambda_{pr} + \lambda_i)}$, i.e.,

$$\lambda_{pb} > \frac{\lambda_{pr}(\lambda_i + a\lambda_{pr})}{I\lambda_i(\lambda_{pr} + \lambda_i) - \lambda_{pr}} - \frac{I(\lambda_i + \lambda_{pr})[(a-1)\lambda_i \lambda_{pr} - (\lambda_i + \lambda_{pr})^2 \mathbf{g}_1 x]}{I\lambda_i(\lambda_{pr} + \lambda_i) - \lambda_{pr}}$$

Note that $(q_1 + q_2 + q_3)$ is smaller than 1 because the excess demand follows the Skellam distribution (since investor and trader arrivals are Poisson processes, and they buy or sell equally likely). Intuitively, users' transactions can be executed against each other and the LPs only need to satisfy excessive demands. We depict the regions where FBA induces less welfare loss for common investors and traders in Figure 8a in Appendix C. We intentionally let x or F be not too small, i.e., $x = 0.15$, in the figure to make the regions where CLOB imposes fewer markups more visible. This means that if the jump $J = \$1$, the priority fee is $\$0.15$. When the priority fee F decreases, the region where FBA has fewer markups expands.

Spread. In the same setting with constant jump size, under CLOB, the bid-ask spread is $s = \frac{(\lambda_{pb} + a\lambda_{pr})J - (\lambda_i + \lambda_{pr})\mathbf{g}_1 F}{\lambda_{pb} + a\lambda_{pr} + \lambda_i}$. Under FBA, we denote the bid-ask spread as s_F and $s_F = 2(\Delta + M_1)$. Consider $Q \geq 1$. Same as before, although Z_I follows a Skellam distribution, for simplicity and without loss of generality, we consider $q_1 = \frac{1}{8}$ and $q_2 = \dots = q_{Q+1} = \frac{1}{8Q}$. Then

$$\begin{aligned} M_1 &= \sum_{u=1}^Q \Delta_u \prod_{v=1}^u \alpha_v = \Delta(\alpha_1 + \alpha_1\alpha_2 + \dots + \alpha_1 \dots \alpha_Q) \\ &= \Delta \frac{\frac{1}{8Q}}{\frac{1}{8} + \frac{1}{8Q}} \left(1 + \frac{1}{2} + \dots + \frac{1}{2^{Q-1}}\right) = \Delta \frac{2}{Q+1} \left(1 - \frac{1}{2^Q}\right) \end{aligned}$$

Let a transaction be fulfilled with transactions submitted by other common investors or traders with probability $q_0 = 1 - 2 \sum_{j=1}^{Q+1} q_j$. FBA provides a smaller bid-ask spread in expectation when $(1 - q_0)s > s_F$. This gives

$$2(1 - q_0) \frac{\lambda_{pr}}{\lambda_{pr} + \lambda_i} \left(1 + \frac{2}{Q+1} \left(1 - \frac{1}{2^Q}\right)\right) < \frac{\lambda_{pb} + a\lambda_{pr} - (\lambda_i + \lambda_{pr})\mathbf{g}_1 x}{\lambda_{pb} + a\lambda_{pr} + \lambda_i}$$

We depict the regions where FBA has a smaller expected bid-ask spread in Figure 8b. When there is no private information, the spread under FBA is constant while the spread under CLOB increases with λ_{pb} . We can also consider general parameterization for Z_I 's distribution. When $\sum_{u=1}^Q \prod_{v=1}^u \alpha_v$ decreases, the region where FBA has smaller spreads expands. In summary, when public information is released more often or common investors arrive more often compared with privately informed traders, FBA yields less welfare loss than CLOB.

5 Related work

Batch auction-based DEX. Penumbra [29] adopts FBA as the processing model. CoW Swap [12] combines batch auction with trade intents where users specify the assets and the amounts to trade. Users’ intents are then settled via direct matching among intents for the same tokens and opposite sides.

Manipulation-aware DEX design. We mention mechanisms that update transaction *execution* rules to mitigate order manipulation attacks in DEX because they are closely related to this work. For general order manipulation mitigation techniques that update *ordering* algorithms, we refer the readers to comprehensive surveys like [18]. Fair-TraDEX [27] combines FBA with commit-then-reveal style transaction masking where users commit to transactions and later reveal the contents upon finalization. Each user needs to make deposits in an escrow service so that they are incentivized to post correct commitments. Similarly, Injective [22] and Penumbra [29] utilize FBA for executing transactions on proof-of-stake blockchains, with Penumbra being a private chain. It additionally hides transactions’ trading amounts. Masking transaction limits may not eliminate all manipulation opportunities. In P2DEX [6], servers run secure multi-party computation (MPC) to match orders. This introduces computation overhead and latencies and adds constraints inherent to the adopted MPC. SPEDEX [30] approximates the clearing price in a block in general equilibrium where demand meets supply given a set of static orders. The equilibrium does not capture participants’ strategies, acquired information, or sequential moves. Xavier et al. [33] alternatively consider a greedy sequencing rule in two-token constant product AMM. Since this sequencing rule ensures good properties only inside a block, sequencers can push submitted transactions to future blocks.

Comparison of CLOB and FBA in centralized exchanges. It has been shown that compared with CLOB, FBA leads to lower transaction costs [2], decreases adverse selection and spreads [28, 31], achieves an optimal trade-off between liquidity and price discovery [5], and increases market quality [13, 21, 25]. However, the severity of the inefficiency of liquidity provision under FBA can exceed the inefficiency from latency arbitrage under CLOB [14]. This liquidity provision inefficiency of FBA originates from bid shading⁹ in UPDA: every equilibrium in multi-unit uniform price auctions is inefficient due to bid shading [4].

References

1. Adams, H., Zinsmeister, N., Salem, M., Keefer, R., Robinson, D.: Uniswap v3 core. Tech. rep., Uniswap, Tech. Rep. (2021)
2. Aldrich, E.M., López Vargas, K.: Experiments in high-frequency trading: comparing two market institutions. *Experimental Economics* **23**(2), 322–352 (2020)
3. Anonymized: Simulation source code, https://drive.google.com/drive/folders/1XtnstLf5o0Y8Ktp_YM0wpNbHgs01JpGC?usp=share_link
4. Ausubel, L.M., Cramton, P.: Demand reduction and inefficiency in multi-unit auctions (2002)

⁹ (Non-unit-demand) bidders’ tendency to bid less than true valuations for later units.

5. Baldauf, M., Mollner, J.: High-frequency trading and market performance. *The Journal of Finance* **75**(3), 1495–1526 (2020)
6. Baum, C., David, B., Frederiksen, T.K.: P2dex: privacy-preserving decentralized cryptocurrency exchange. In: ACNS. pp. 163–194 (2021)
7. Budish, E., Cramton, P., Shim, J.: Implementation details for frequent batch auctions: Slowing down markets to the blink of an eye. *American Economic Review* **104**(5), 418–424 (2014)
8. Budish, E., Cramton, P., Shim, J.: The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics* **130**(4), 1547–1621 (2015)
9. Budish, E., Lee, R.S., Shim, J.J.: A theory of stock exchange competition and innovation: Will the market fix the market? Tech. rep., National Bureau of Economic Research (2019)
10. Cachin, C., Mičić, J., Steinhauer, N., Zanolini, L.: Quick order fairness. In: FC. pp. 316–333 (2022)
11. Corporation, T.S.E.: Taiwan stock exchange, <https://www.twse.com.tw/en/>
12. CoW: Cow protocol, <https://docs.cow.fi/cow-protocol>
13. Economides, N., Schwartz, R.A.: Electronic call market trading. In: *The Electronic Call Auction: Market Mechanism and Trading*, pp. 87–99. Springer (2001)
14. Eibelshäuser, S., Smetak, F.: Frequent batch auctions and informed trading (2022)
15. ethereum.org: Priority fees in ethereum, <https://ethereum.org/en/developers/docs/gas/#priority-fee>
16. Farmer, D., Skouras, S.: Review of the benefits of a continuous market vs. randomised stop auctions and of alternative priority rules (policy options 7 and 12). Manuscript, Foresight. Government Office for Science (2012)
17. Glosten, L.R., Milgrom, P.R.: Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of financial economics* **14**(1), 71–100 (1985)
18. Heimbach, L., Wattenhofer, R.: Sok: Preventing transaction reordering manipulations in decentralized finance. In: *Proceedings of the 4th ACM Conference on Advances in Financial Technologies*. pp. 47–60 (2022)
19. Indriawan, I., Pascual, R., Shkilko, A.: On the effects of continuous trading. Available at SSRN 3707154 (2020)
20. Juliano, A.: dydx. <https://github.com/dydxprotocol> (2023)
21. Kandel, E., Rindi, B., Bosetti, L.: The effect of a closing call auction on market quality and trading strategies. *Journal of Financial Intermediation* **21**(1), 23–49 (2012)
22. Labs, I.: Injective, <https://injective.com/about>
23. Labs, M.: sui. <https://github.com/MystenLabs/sui> (2023)
24. Labs, U.: Uniswap limit order book. <https://blog.uniswap.org/limit-orders> (2024)
25. Madhavan, A.: Trading mechanisms in securities markets. *the Journal of Finance* **47**(2), 607–641 (1992)
26. Maskin, E., Tirole, J.: Markov perfect equilibrium: I. observable actions. *Journal of Economic Theory* **100**(2), 191–219 (2001)
27. McMenamin, C., Daza, V., Fitzi, M., O’Donoghue, P.: Fairtradex: A decentralised exchange preventing value extraction. In: *ACM CCS DeFi Workshop’22*. pp. 39–46 (2022)
28. Menkveld, A.J., Zoican, M.A.: Need for speed? exchange latency and liquidity. *The Review of Financial Studies* **30**(4), 1188–1228 (2017)

29. Penumbra: Penumbra zswap, <https://protocol.penumbra.zone>
30. Ramseyer, G., Goel, A., Mazières, D.: Speedex: A scalable, parallelizable, and economically efficient digital exchange. arXiv preprint arXiv:2111.02719 (2021)
31. Riccò, R., Wang, K.: Frequent batch auctions vs. continuous trading: Evidence from taiwan. Continuous Trading: Evidence from Taiwan (November 19, 2020) (2020)
32. Tardis.dev: Tardis.dev node. <https://github.com/tardis-dev/tardis-node> (2023)
33. Xavier Ferreira, M.V., Parkes, D.C.: Credible decentralized exchange design via verifiable sequencing rules. In: Proceedings of the 55th Annual ACM Symposium on Theory of Computing. pp. 723–736 (2023)

A Omitted proofs

Theorem 3. *There exists a stationary MPE in trading game \mathcal{G} under CLOB processing. Bid and ask prices at the k -th level of the LOB in equilibrium satisfy*

$$b_t^k = \mathbb{E}[V_t | \mathcal{H}_t, \text{sell}] = P_t - \frac{s_k}{2}, a_t^k = \mathbb{E}[V_t | \mathcal{H}_t, \text{buy}] = P_t + \frac{s_k}{2}$$

where s_k satisfies

$$\lambda_i \sum_{j=k}^Q p_j \frac{s_k}{2} - \lambda_{pr} (1 + p^* \frac{\vec{g}_k}{\vec{g}_k(r-2)+1}) \bar{J}_k - \lambda_{pb} \bar{J}_k + (\lambda_i + \lambda_{pr}) \vec{g}_k F = 0$$

Here, $\bar{J}_k = \mathbb{P}[J > \frac{s_k}{2}] \mathbb{E}[J - \frac{s_k}{2} | J > \frac{s_k}{2}]$, and \vec{g}_k is the probability that an arbitrageur front-runs traders and investors at the k -th level. In equilibrium, \vec{g}_k is then updated to take the value that maximizes profits from front-running traders and investors, i.e., $\lambda_{pr} p^* \frac{\vec{g}_k}{\vec{g}_k(r-2)+1} \bar{J}_k - (\lambda_i + \lambda_{pr}) \vec{g}_k F$. The expected price impact of an incoming k -unit order is $\Delta_k = \tilde{J}_k \frac{\lambda_{pr}}{\lambda_{pr} + \lambda_i}$ with $\tilde{J}_k = \mathbb{P}[J > \frac{s_k}{2}] \mathbb{E}[J | J > \frac{s_k}{2}]$; the expected markup from liquidity providers is $(\lambda_{pr} + \lambda_i)(\frac{s_k}{2} - \Delta_k)$.

Proof. Let s_k denote the spread at the k -th level. In a DEX, unlike in centralized exchanges, an arbitrageur can also front-run privately informed traders. Since a transaction may be submitted by a common investor or privately informed trader, the front-runner can expect premiums when privately informed traders arrive. This happens with probability $\frac{\lambda_{pr}}{\lambda_{pr} + \lambda_i}$ if the arbitrageur always front-runs ($\vec{g} = \vec{1}$). More meaningfully, let the arbitrageur front-run at the k -th level with probability $\vec{g}_k \in [0, 1]$. As described, let the front-running transaction beat a user's transaction with probability p^* .

Suppose in an incoming order, there are $q \leq Q$ ($Q \in N$) quantities needed. We need the k -th spread s_k ($k = 1, \dots, Q$) to satisfy that the market-making

strategy yields the same returns as the front-running strategy in equilibrium:

$$\begin{aligned}
& \underbrace{\lambda_i \sum_{j=k}^Q p_j \frac{s_k}{2}}_{\text{Profits from investor}} - \underbrace{\lambda_{pr} \bar{J}_k}_{\text{Private info-induced loss}} - \underbrace{(\lambda_{pb} \frac{r-1}{r} \bar{J}_k + \lambda_{pb} F)}_{\text{Public info-induced loss}} = \\
& \underbrace{\lambda_{pb} \frac{1}{r} \bar{J}_k - \lambda_{pb} F}_{\text{Profits from LPs}} + \underbrace{\lambda_{pr} p^* \frac{\bar{g}_k}{\bar{g}_k(r-2)+1} \bar{J}_k - (\lambda_i + \lambda_{pr}) \bar{g}_k F}_{\text{Profits from front-running investors and traders}} \quad (1)
\end{aligned}$$

where $\bar{J}_k = \mathbb{P}[J > \frac{s_k}{2}] \mathbb{E}[J - \frac{s_k}{2} | J > \frac{s_k}{2}]$ is the expected extra value change uncovered by the spread (induced by the asset fundamental value jump), and p_j is the probability of an investor transacting j units.

In Equation (1), the left-hand side is the profits from the market-making strategy and the right-hand side is the returns from front-running LPs and private traders. To update \bar{g}_k stored in the current state, we lay out the right-hand side as a function of \bar{g}_k and solve for \bar{g}_k that maximizes the function value in range $[0, 1]$ for a front-running arbitrageur, by taking the first and second derivatives of the function. One can solve it analytically or numerically if closed-form formulas do not exist for specific distributions of the jump J .

We now compute the additional transaction costs for common investors. When an arbitrageur front-runs common investors, the fundamental value does not change. This front-runner only earns premiums when private traders arrive, which happens with probability $\frac{\lambda_{pr}}{\lambda_{pr} + \lambda_i}$. To compute the welfare loss for investors, we can then treat front-running orders as part of private traders' orders. Depending on the size of the jump, they can place multi-unit orders profitably, and we first consider a single unit. Let Δ_1 be the expected price impact of an incoming unit order at time t . Since one cannot distinguish between orders from informed traders from investors, we have

$$\Delta_1 = \pm \mathbb{E}[V_t | \mathcal{H}_t, \lambda_i \vee \lambda_{pr}] \mp P_t = \pm \mathbb{E}[V_t | \mathcal{H}_t, \lambda_i \vee \lambda_{pr}] \mp \mathbb{E}[V_t | \mathcal{H}_t] = \tilde{J}_1 \frac{\lambda_{pr}}{\lambda_{pr} + \lambda_i}$$

where $+$, $-$ indicate buy, sell order and $\tilde{J}_k = \mathbb{P}[J > \frac{s_k}{2}] \mathbb{E}[J | J > \frac{s_k}{2}]$ (for $1 \leq k \leq Q$). After absorbing the price impact, the additional markup from the liquidity provider on the first level of the order book is then $M_C = \frac{s_1}{2} - \Delta_1$. The expected markup is then $(\lambda_{pr} + \lambda_i) M_C$. We can calculate multi-unit orders in the same way. Let Δ_k be the expected price impact of the k -unit order. We can compute $\Delta_k = \tilde{J}_k \frac{\lambda_{pr}}{\lambda_{pr} + \lambda_i}$. Then the markup on the k -th level of the order book is then $M_C^{(k)} = \frac{s_k}{2} - \Delta_k$, and its expected value is $(\lambda_{pr} + \lambda_i) M_C^{(k)}$.

Theorem 4. *There exists an OBE in pure strategies in trading game \mathcal{G} under FBA processing. If $|Z_I| \leq Q + 1$, bid and ask quotes in equilibrium satisfy*

$$\begin{cases} b_{(I)}^k = P_I - \Delta_k - M_k & Z_I < 0 \\ a_{(I)}^k = P_I + \Delta_k + M_k & Z_I > 0 \end{cases}, \quad k = 1, \dots, Q + 1$$

where $\Delta_k = \tilde{J}_k \frac{\lambda_{pr}}{\lambda_{pr} + \lambda_i}$ is the expected price impact from informed traders, $M_k = \sum_{u=k}^Q \Delta_u \prod_{v=k}^u \alpha_v$ for $k = 1, \dots, Q$, and $M_{Q+1} = 0$ are the markups from LPs,

$\alpha_k = \frac{q_{k+1}}{q_k + q_{k+1}}$ and $q_k = \mathbb{P}[Z_I = k \mid |Z_I| \leq Q + 1]$. The expected markup per unit time is $\frac{2}{I} \sum_{k=1}^Q k q_k M_k$.

Proof. We show by induction that the markups $\{M_k\}_{k=[Q+1]}$ are the largest such that there exist (i) neither strictly profitable safe price improvements (ii) nor strictly profitable robust deviations. More specifically, (i) means that for a targeted quote at the k -th level, an arbitrageur cannot safely insert a new profitable quote at this level, i.e., pushing the original k -th quote to be the new $(k + 1)$ -th quote and so forth, and the current k -th quote issuer cannot improve this quote (after being pushed) to safely increase profits. (ii) means that the k -th quote issuer cannot robustly increase profits by issuing or updating quotes deviating from the stated markups, considering that others can react with quote withdrawals and updates. We will focus on the ask quotes, but we only need to adapt signs for analyzing bid quotes.

$(Q + 1)$ -th level. We first show that $M_{Q+1} = 0$ is the OBE, i.e., $a_{(I)}^{Q+1} = P_I + \Delta_{Q+1}$. We write $a_{(I)}^{Q+1}$ as a^{Q+1} in the following analysis for visual simplicity. (i) We show that safe profitable price improvements on the last quote with zero markups do not exist. First, an arbitrageur inserting a new $(Q + 1)$ -th quote with a price lower than a^{Q+1} experiences a loss in expectation. Because the spread does not cover the potential jump in asset value. Similarly, the current $(Q + 1)$ -th quote issuer does not have the incentive to lower the quote. Second, the issuer is not incentivized to increase the quote either, i.e., to $\bar{a}^{Q+1} = P_I + \Delta_{Q+1} + \epsilon$ for even a small $\epsilon > 0$. Because this is not a safe price improvement: another arbitrageur can safely insert a quote at price a^{Q+1} and push this issuer's new quote to be the $(Q + 2)$ -th quote. Since $Z_I \leq Q + 1$, the expected profit after others' reaction is zero, which is less than the original profit. (ii) With similar reasoning, we can show that robust profitable deviation does not exist. This is because the owner of the $(Q + 1)$ -th quote is not incentivized to insert a new quote or update another quote to be the $(Q + 1)$ -th that deviates from zero markups.

$(k + 1)$ -th level. Suppose there do not exist safe profitable price improvements or robust profitable deviations on the $(k + 1)$ -th quote. The markups on the $(k + 1)$ -th level and onwards to the $(Q + 1)$ -th level are as stated in Theorem 2.

k -th level. We show that the stated M_k is the OBE given the quotes at the $(k + 1)$ -th level. (i) Suppose there exists a safe profitable price improvement on the k -th quote by ϵ , i.e., $\bar{a}^k = a^k - \epsilon$. First, consider an arbitrageur who does not own quotes and inserts a new k -th quote. If $Z_I = k$, then the arbitrageur expects positive profits, and if $Z_I > k$, then a loss is expected because the settlement price would be below \bar{a}^k . Let $\Delta_k = \tilde{J}_k \frac{\lambda_{pr}}{\lambda_{pr} + \lambda_i}$ and $q_k = \mathbb{P}[Z_I = k \mid |Z_I| \leq Q + 1]$. We can express the profits as:

$$A_k = q_k(M_k - \epsilon) + \sum_{i=k+1}^{Q+1} q_i(M_{i-1} - \Delta_{i-1})$$

For this improvement to be safe, the largest markup needs to be achieved, i.e., $\epsilon \rightarrow 0$, so that no other arbitrageurs have the incentive to undercut again. We

next demonstrate the conditions such that $A_k < 0$. We know that the analog equation for the $(k+1)$ -th level ask quote is

$$A_{k+1} = q_{k+1}(M_{k+1} - \epsilon) + \sum_{i=k+2}^{Q+1} q_i(M_{i-1} - \Delta_{i-1}) < 0$$

Observe that instead of directly requiring $A_k < 0$, we can locate conditions that let $A_k = A_{k+1}$ as $\epsilon \rightarrow 0$. This gives us $M_k = \alpha_k \Delta_k + \alpha_k M_{k+1}$ where $\alpha_k = \frac{q_{k+1}}{q_k + q_{k+1}}$. Applying this recursively, we obtain $M_k = \sum_{u=k}^Q \Delta_u \prod_{v=k}^u \alpha_v$. In other words, when the original markup of the k -th level quote satisfies this condition, $A_k = A_{k+1} < 0$.

Now consider an arbitrageur who owns some existing quotes. To implement price improvements, she can insert a new k -th quote or adjust existing quotes. As discussed above, insertion results in a loss. If she moves up a lower quote at the k^* -th ($k^* < k$) level to the k -th, the expected profit is lower because some profits are lost for $k^* \leq Z_I < k$. If she moves down one of the quotes to be the new k -th quote, then the extra profit generated is bounded by A_k . Because A_k is the extra profits of an arbitrageur who originally expects 0 returns.

(ii) Suppose there exists a robust profitable deviation for the issuer of the k -th quote. To allow the issuer to deviate other than improving quotes (which we just discussed), we let the issuer own some quotes under the k -th level. The issuer can update the k -th quote to be the k' -th quote ($k' \geq k$) with a markup higher than the original $M_{k'}$ but still lower than $\Delta_{k'+1} + M_{k'+1}$. This is profitable but not robust: another arbitrageur can re-issue the original k -th level quote and push the updated k' -th quote to be the $(k'+1)$ -th. This arbitrageur's price improvement is safe as it restores the original k -th quote.

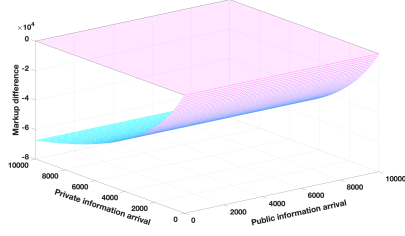
By induction, the markups from LPs at each level satisfy the established conditions. The expected markup of the liquidity providers per unit of time can then be calculated as $\frac{2}{I} \sum_{k=1}^Q k q_k M_k$.

B Bid-ask spread comparison figures

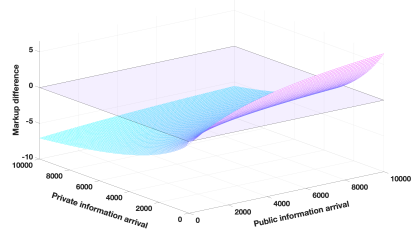
We include the figures (Figures 5 to 7) depicting how $\lambda_i, \lambda_{pb}, \lambda_{pr}$ affect the spreads in FBA and CLOB, as mentioned in §4.2.

C Figures for the small example

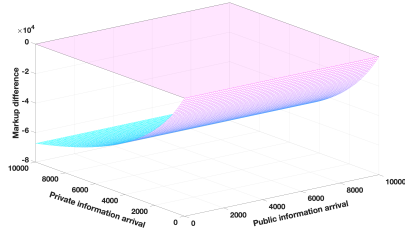
We provide the figures for the small example in §4.4 in Figure 8b.



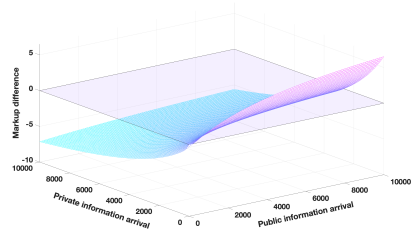
(a) Bad case for both CLOB and FBA.



(b) Bad case for CLOB and good case for FBA.

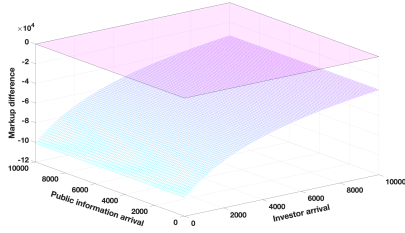


(c) Good case for CLOB and bad case for FBA.

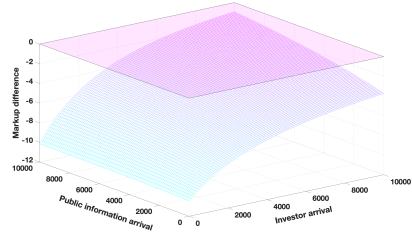


(d) Good case for both CLOB and FBA.

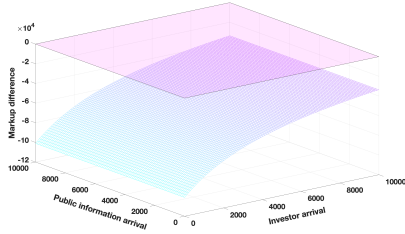
Fig. 5: The spread difference between CLOB and FBA with respect to public and private information arrivals. The investor arrival rate λ_i is set to be 5000. Increasing (decreasing) λ_i pushes the surface up (down).



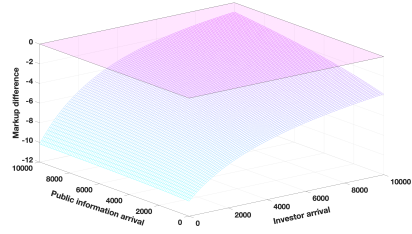
(a) Bad case for both CLOB and FBA.



(b) Bad case for CLOB and good case for FBA.

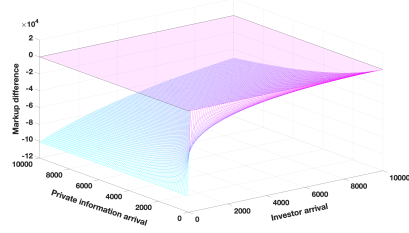


(c) Good case for CLOB and bad case for FBA.

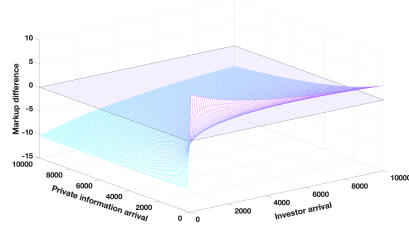


(d) Good case for both CLOB and FBA.

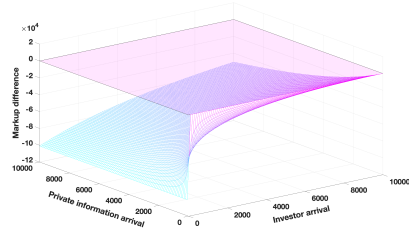
Fig. 6: The spread difference between CLOB and FBA with respect to public and private information arrivals. The private information arrival rate λ_{pr} is set to be 5000. Increasing (decreasing) λ_{pb} pushes the surface down (up).



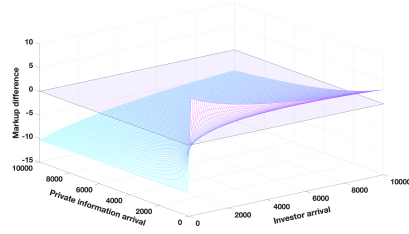
(a) Bad case for both CLOB and FBA.



(b) Bad case for CLOB and good case for FBA.

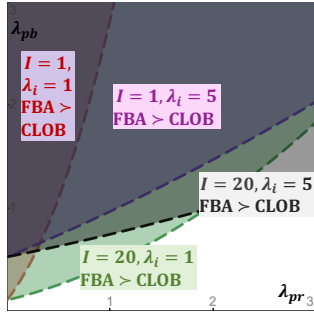


(c) Good case for CLOB and bad case for FBA.

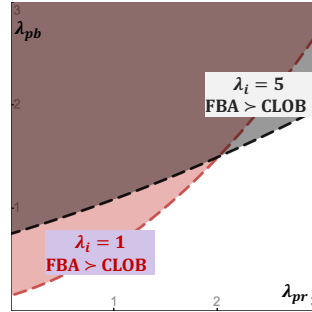


(d) Good case for both CLOB and FBA.

Fig. 7: The spread difference between CLOB and FBA with respect to public and private information arrivals. The public information arrival rate λ_{pb} is set to be 5000. Increasing (decreasing) λ_{pb} pushes the surface up (down).



(a) Example regions where FBA has less welfare loss when $F = 0.15J$, $p^* = 0.8$, $r = 35$, truncated at $\lambda_{pr} = 3$, $\lambda_{pb} = 3$.



(b) Example regions where FBA has smaller bid-ask spreads under the same parameters as Figure 8a and additionally, $q_0 = 0.5$, $Q = 100$.

Fig. 8: Examples with small parameters. The dashed lines represent that given certain I , λ_i , the region above the line means FBA performs better than CLOB, and vice versa.