

Just-In-Time Resale in an Ahead-Of-Time Auction for Faster Execution

Burak Öz ✉ 

Flashbots, Berlin, Germany

Akaki Mamageishvili ✉ 

Offchain Labs, Zurich, Switzerland

Christoph Schlegel ✉ 

Flashbots, Zurich, Switzerland

Ali Taslimi ✉

Entropy Advisors, London, UK

Abstract

We study Arbitrum’s Timeboost auction, an ahead-of-time mechanism that sells a 200 ms ordering advantage in an otherwise first-come, first-served transaction ordering policy. The market naturally divides into two phases: a *competition phase*, in which the dominant searchers compete directly in the primary auction, and a *coordination phase*, in which they source fast-lane access through Kairos, a Just-in-Time resale intermediary. We use auction bids, time-boosted transactions, and on-chain payment traces to study how well ahead-of-time bids predict realized CEX–DEX arbitrage profits and how the emergence of resale changes surplus allocation.

We find that ahead-of-time bids are noisy predictors of short-horizon arbitrage profits. During the competition phase, bid–profit correlations are statistically significant but economically modest at the round level, while correlations increase when profits are aggregated over longer horizons. This suggests that bidders can identify favorable market conditions but face substantial uncertainty about the realized value of any individual one-minute fast-lane interval. After the Kairos transition, competition in the primary auction weakens sharply: in our focused transition-analysis window, paid bids fall from 62.7% of the top bid before the transition to 14.8% thereafter, while total searcher profits remain broadly similar. The resulting dynamics are most consistent with coordination through a common intermediary rather than direct competition in the primary auction. More broadly, our findings suggest that ahead-of-time allocation mechanisms can be vulnerable to secondary-market intermediation when competition among dominant participants is weak.

2012 ACM Subject Classification Applied computing → Online auctions

Keywords and phrases ahead-of-time auctions, secondary markets, Timeboost, arbitrage

Digital Object Identifier 10.4230/LIPIcs.AFT.2026.20

Related Version *Full Version*: <https://arxiv.org/abs/2603.20175> [26]

1 Introduction

Many blockchains use auctions as part of their transaction ordering mechanism. In Ethereum Layer-1 (L1), *Proposer Builder Separation* (PBS) implements a builder auction in which specialized parties called *block builders* bid for the right to determine the content of the next proposed block. In most Layer-2 (L2) rollups, transactions are sorted in decreasing order of priority fees per gas. This mechanism is commonly called *Priority Ordering* or *Priority Gas Auction*, since it amounts to an auction for earlier execution. Despite significant differences – in Ethereum, builders bid for the right to order an entire block while in rollups the auction concerns relative ordering of individual transactions, and Ethereum has a public mempool while rollups do not – these formats share a common feature: they are *Just-in-Time* (JIT), meaning participants can submit bids until immediately before block content is finalized.



© Burak Öz, Akaki Mamageishvili, Christoph Schlegel, and Ali Taslimi;
licensed under Creative Commons License CC-BY 4.0


8th Conference on Advances in Financial Technologies (AFT 2026).

Editors: Aggelos Kiayias and Maria Kyropoulou; Article No. 20; pp. 20:1–20:23

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** Daily shares of Timeboost auction rounds won by each participant (April 2025–March 2026). The header strip delineates the *competition phase*, in which Wintermute and Selini win the majority of rounds, from the *coordination phase*, in which Kairos takes over as the dominant winner. Hatched regions indicate rounds with no active bidders.

An alternative design, *ahead-of-time* (AOT) auctions, introduces a substantial time gap between the auction and the determination of transaction order. In the Ethereum context, such mechanisms have been discussed as slot auctions: a bidder wins the right to propose the next block at auction time but commits to no specific block content until later (typically 12 seconds later) [14]. In the rollup ecosystem, Arbitrum’s Timeboost [2] is the first major deployed instance of a transaction ordering policy with an AOT component.

Prior work on Ethereum slot auctions has argued [15, 22] that selling the right to propose blocks in advance may lead to secondary markets for proposal rights that benefit certain searchers. It has also been argued that predicting block value far in advance – as AOT bidders must – is inherently difficult and may lead to inefficient value capture, which can itself be the catalyst for such secondary markets to emerge. As the only major deployed AOT arbitrage auction for transaction ordering, Arbitrum’s Timeboost provides a natural empirical case study for evaluating these claims.

Since its launch in April 2025, the Timeboost auction has been dominated by two searchers, Wintermute and Selini Capital, who specialize in non-atomic arbitrage and have won the vast majority of auction rounds (see Figure 1). Kairos [10], a service that participates in the primary auction and resells fast-lane access at the transaction level JIT, attracted little traction for most of Timeboost’s existence. This changed sharply in mid-February 2026, when Wintermute and Selini began sourcing fast-lane access through Kairos’ resale market and reduced or ceased primary-auction bidding, after which Kairos began winning the majority of auction rounds.

More precisely, we observe the following sequence of events. On February 12, Wintermute and Selini stopped participating in the primary auction or significantly reduced their bids, leaving Kairos winning nearly all auction rounds at prices close to the reserve. In response, Arbitrum substantially increased the reserve price on February 18, after which no participant was willing to pay it for some period. Bidding later resumed in a small fraction of rounds before the market settled into a new stable regime when Arbitrum lowered the reserve price on February 25. In this regime, Kairos wins approximately 78% of auction rounds, with occasional but limited participation from Wintermute and Selini.

This naturally divides our study into two phases: a *competition phase* prior to February 12, which lets us assess AOT auction effectiveness in the absence of significant JIT resale, and a *coordination phase* thereafter, in which dominant searchers source fast-lane access through a

common JIT intermediary rather than competing directly in the primary auction. We use *coordination* in this descriptive sense and do not mean to imply that we observe an explicit agreement among the participants.

We find that, during the competition phase, bids are only weakly correlated with estimated markout arbitrage profit and loss (PnL), consistent with the difficulty of predicting short-horizon value in an AOT setting. We observe that the two dominant searchers route through Kairos’s JIT resale market, allowing them to obtain fast-lane access without directly committing to price-uncertain forward bids. Because the associated off-chain payments are unobserved, however, we cannot determine whether this shift increased their net payoffs. In our focused transition-analysis window, paid bids in the primary auction fall to 14.8% of the top bid after the transition, down from 62.7% in the February 1–12 pre-transition segment, while estimated arbitrage PnL remain broadly similar across regimes. This indicates that the primary auction captures a smaller share of the value reflected in submitted bids, although the final allocation of surplus between Kairos and the searchers remains unobserved. The observed dynamics are consistent with a shift from direct competition in the primary auction toward routing through a common intermediary. More broadly, our findings suggest that AOT allocation mechanisms may be particularly susceptible to secondary-market intermediation when competition among dominant participants is weak.

1.1 Related Work

In Ethereum, the next block proposer sorts transactions unilaterally. This allows the proposer to extract value from the power to order transactions, commonly referred to as Maximum Extractable Value (MEV) [5]. Transaction ordering through auctions is theoretically analyzed in [23]. An early version of Timeboost without an AOT auction was proposed in [16]. Spam levels before and after Timeboost are studied in [4], which finds that spam decreased after its introduction. The time advantage of centralized exchange (CEX) to decentralized exchange (DEX) arbitrage under a First-Come-First-Serve (FCFS) ordering policy is studied in [9], with comparisons to PGA auctions and pure FCFS. Latency advantage in a PBS setting are examined in [21, 24]. Possible downsides of AOT auctions are identified in [22] for settings where a secondary market allows JIT resale of AOT-acquired rights.

For methodology, we follow [25] in using markouts at CEX mid-prices to estimate CEX–DEX arbitrage profits. Theoretical predictions of such profits are developed in [20]. Our theoretical background for auctions with common-value component draws on the classical work of [19].

1.2 Background

Timeboost. Timeboost has been deployed on the Arbitrum rollup since April 17, 2025. Under this mechanism, the protocol sells fast-lane access in a permissionless second-price auction with a reserve price. The fast-lane integrates with Arbitrum’s FCFS transaction ordering policy and gives the auction winner a 200 ms advantage for one minute. The auction for the next minute begins at second 0 of each round and lasts until second 45¹. The remaining 15 seconds are used for on-chain auction settlement. The rule is enforced by the sequencer, which executes regular transactions with a 200 ms delay, while fast-lane transactions from the Timeboost winner are executed immediately upon arrival. In effect,

¹ Each round corresponds to a minute shifted by 9 seconds relative to the wall-clock minute.

the sequencer merges the two FCFS queues: one for time-boosted transactions and one for regular transactions. Empirically, most fast-lane activity and extracted value can be attributed to CEX–DEX arbitrage [18]. In particular, Wintermute exclusively conducts such arbitrage, while Selini occasionally uses fast-lane for other forms of value extraction methods.

Kairos. After winning a Timeboost round, Kairos organizes fast-lane “sub auctions” with bidding phases of approximately 100 ms. Transactions within each 100 ms batch are ordered by profitability, and the batch is submitted to the Arbitrum sequencer via the fast-lane, thereby giving all transactions in the batch a time advantage of at most 100 ms over regular transactions. The exact advantage depends on the latency between Kairos’s servers and the sequencer. As long as this latency remains below 100 ms, a competing transaction in the regular FCFS queue cannot be executed before the Kairos batch.

Payments to Kairos are typically made on-chain as explicit ETH transfers embedded in time-boosted transactions and sent to a designated contract address². An important exception is Kairos’s parallel lane [11], which is designed for CEX–DEX arbitrage searchers and operates via a subscription model with off-chain payments. This design bypasses the simulation step in which searchers’ transactions are continuously re-submitted under FCFS, thereby reducing latency and lowering the risk of losing priority to regular transactions. Since these transactions constitute a substantial share of all time-boosted transactions in our dataset, the absence of observable on-chain payments prevents us from tracking transfers to Kairos by the two main searchers, as we discuss further below.

2 Conceptual Framework and Empirical Predictions

We develop a conceptual framework to organize and guide our empirical analysis. We first use a standard interdependent-values auction model to discuss how common- and bidder-specific value components affect AOT bidding. We then relate expected arbitrage value to short-run price volatility. Finally, we identify several mechanisms through which a JIT resale market such as Kairos may emerge and derive their distinct empirical implications. The framework is intended to provide qualitative benchmarks rather than a complete structural model of the Timeboost market.

2.1 Common-Value Second-Price Auctions

The Timeboost auction has a clear common-value component: the market-wide arbitrage opportunities arising during a Timeboost interval are, in principle, available to all bidders. At the same time, valuations may contain bidder-specific components. Bidders may face private inventory constraints, hold accumulated positions that they want to unwind, or incur heterogeneous trading costs on external markets. For bidders that are also market makers on external exchanges, the trades and inventory accumulated through their market-making activity may further affect their value for winning the auction.

A standard framework for auctions with common- and private-value components is the interdependent-values model of [19]. Bidders $i = 1, \dots, n$ receive private signals X_i that are i.i.d., and the value of winning the auction for bidder i is a function of all signals $v_i(X_1, \dots, X_n)$, which is assumed to be symmetric in opponents’ signals, strictly increasing in

² Kairos payment address: 0x60E6a31591392f926e627ED871e670C3e81f1AB8

own signal, and nondecreasing in the others'. Let $Y_i := \max_{j \neq i} X_j$ denote the highest signal among bidder i 's competitors. Under risk neutrality, the symmetric equilibrium bidding strategy in a second-price auction is

$$b(x) = \mathbb{E}[v_i(X_i, X_{-i}) \mid X_i = x, Y_i = x],$$

the expected value of winning conditional on the pivotal event in which the bidder's signal is exactly tied with that of the strongest competitor.

In contrast to a pure private-value auction, bidders should not simply bid their expected value given their own signal, $\mathbb{E}[v_i(X_i, X_{-i}) \mid X_i = x]$: because the value of winning depends on rivals' signals, winning itself is informative – it reveals that all competing signals were lower – and a bidder who ignores this would systematically overpay. Equilibrium bidding corrects for this winner's curse by evaluating the expected value in the pivotal event where the bidder is exactly tied with the strongest rival.

This framework motivates an empirical comparison between bids and realized arbitrage value. In equilibrium, the price paid equals b evaluated at the second-highest signal, while the top bid reflects the highest signal. If the paid bid is more closely associated with our measure of realized common arbitrage value than the top bid, this is consistent with the common-value component being an important driver of bidder valuations, since the paid bid reflects the signal of the strongest losing bidder, which is informative about realized value only to the extent that values are interdependent. Conversely, if valuations are dominated by private idiosyncratic components, we would expect bids to be less closely associated with realized common arbitrage value. We emphasize that this comparison is suggestive rather than a sharp test: the theory does not imply in general that the second-highest signal must predict realized common value better than the highest, as the relative predictive power of the two order statistics depends on the joint distribution of signals and values.

2.2 Arbitrage Value and Price Volatility

The theory of arbitrage in automated market makers (AMM) [20] can inform how much arbitrage value searchers can expect to extract over a fast-lane interval. While the theory applies to the case where bidders do not face uncertainty about the amount of arbitrage their individual trades will extract³, quantitatively, we should expect similar results. The theory predicts that arbitrage value is proportional to price variance, i.e., expected arbitrage gains over a window of length T scale with price variance over the same window.

Cryptocurrency price volatility is typically non-stationary and is best modeled by a mean-reverting stochastic volatility process [13]. Therefore, this framework motivates the following empirical expectations. First, if price variance is linear in time, then arbitrage gains are also linear in the length of the Timeboost interval. Second, longer run averages of arbitrage gains should be easier to predict than realized gains in any short interval. This makes value prediction particularly difficult in an AOT auction, where bidders must price access before observing the exact opportunities that will arise.

³ Under the Timeboost policy, fast-lane transactions do not know the exact state of the AMMs, as unobserved slow-lane transaction might change it compared to the state that the arbitrageur last observed.

2.3 Resale Market

We identify three main mechanisms through which a JIT resale market such as Kairos could be successful.

1. **Unbundling:** The Timeboost auction wholesales a 200 ms time advantage for an entire minute. Reselling this advantage over shorter time intervals or on a per transaction basis could, in principle, enhance efficiency by allowing different searchers to obtain time-boosts within the same minute. Ideally, this would allocate the time advantage to the searcher that can extract the most value at a given moment, rather than to a single party that extracts the most value on average over the full minute. Similar dynamics arise in L1 block-building auctions [8], where builders resell blockspace to many different searchers. In this case, additional surplus from resale would come from improved market efficiency.
2. **Non-Competitive Behavior:** In the Timeboost setting, the presence of only two dominant searchers makes the market vulnerable to tacit coordination or collusive behavior. Kairos may effectively, though not necessarily deliberately, operate as a hub in a hub-and-spoke coordination structure: the two major searchers may abstain from bidding, or bid less aggressively than they otherwise would, allowing Kairos to acquire the fast-lane at a lower price than under competitive bidding. In this case, the combined surplus of the two searchers and the reseller could exceed that in the competitive benchmark, not because more value is created, but because less revenue is paid to the seller in the primary auction, namely Arbitrum.
3. **Information:** Bidding AOT for the right to use the fast-lane over an entire minute gives bidders less information about the value they will ultimately be able to extract than bidding JIT for an immediately submitted transaction. This matters when some opportunities are searcher-specific, e.g., because a searcher identifies or can exploit opportunities that others cannot. In such cases, a near-JIT auction can generate higher expected revenue⁴. A related but separate issue arises if the reseller is also a searcher themselves, for which there is no indication that Kairos qualifies⁵. In that case, the reseller could strategically choose to resell the time-boost only after receiving an unfavorable signal about extractable value, thereby capturing even more value from operating the resale market. In this setting, additional surplus from resale would again arise through improved allocation efficiency.

These mechanisms imply different empirical patterns. The unbundling mechanism predicts a broader and more differentiated set of active searchers under Kairos, together with higher total value extraction through time-boosted transactions. The weakened-competition mechanism predicts that the same dominant searchers remain active and total value extraction stays broadly similar, while revenue in the primary Timeboost auction declines. It may also imply that the dominant searchers stop bidding competitively against each other in the primary auction. The information mechanism predicts higher value extraction from time-boosted transactions, together with a significant spread between revenue captured in the primary auction and revenue generated through Kairos.

⁴ The mathematical intuition is simple: in the extreme case of a pure private-value auction, the mean of i.i.d. random variables is smaller than their expected maximum, so allocating closer to the point of realization can increase expected value extraction.

⁵ [22] focus particularly on this case and argue that it can be a driver of market centralization in the searcher/builder market.

As we show below, the strongest empirical support is for the weakened-competition mechanism. After the transition to Kairos, the same searchers remain central, total value extraction changes only modestly, and primary auction revenue declines. By contrast, we find more limited evidence for unbundling or information-based efficiency gains.

3 Methodology

We now describe the datasets used in our analysis and the methodology for estimating arbitrage profits. Our primary data source is Dune Analytics [7].⁶

3.1 Datasets

Our datasets cover two complementary windows. The main auction and time-boosted transaction datasets span August 2025 through March 12, 2026. For the analysis of the Kairos transition, regular non-time-boosted CEX–DEX transactions and Kairos payments are collected over a focused February 1–March 12, 2026 window.

Timeboost auctions. We collect Timeboost auction data from August 2025 until March 12, 2026. The dataset contains 298,530 auction rounds with at least one bid. We obtain detailed bidding data from public endpoints provided by Arbitrum [1]. For each round, we observe the start and end time of the one-minute interval during which the auction winner holds fast-lane access. We use this interval to assign time-boosted transactions to auction rounds based on block time.

Transactions. In total, we identify 30,310,245 time-boosted transactions executed during our study period. Of these, 27,399,487 (90.4%) target contracts belonging to Wintermute⁷ and Selini⁸. Wintermute accounts for 22,383,187 such transactions, while Selini accounts for 5,016,300. Since these contracts are associated with non-atomic CEX–DEX arbitrage, we apply a heuristic to identify transactions likely executing such strategies. We classify a transaction as CEX–DEX arbitrage if it targets these contracts, emits only one swap-related event, and trades a known liquid pair involving WETH, WBTC, USDC, or USDT. Using this heuristic, we label 11,320,741 Wintermute transactions and 3,399,555 Selini transactions as CEX–DEX arbitrage. These correspond to 50.6% and 67.8% of their respective time-boosted contract interactions. Applying the same heuristic to regular, non-time-boosted transactions targeting these contracts, we identify an additional 1,290,594 Wintermute and 316,927 Selini CEX–DEX arbitrage transactions during the February 1–March 12 analysis window.

Kairos Payments. Using Arbitrum transaction traces, we identify all transactions during the February 1–March 12 analysis window in which Kairos’s payment contract receives an explicit ETH transfer with positive value. In total, we identify 39,521 such payments.

⁶ The queries used in the analysis are publicly accessible: Timeboost auctions, <https://dune.com/queries/6219197>; CEX–DEX transactions, <https://dune.com/queries/6219209>; and Kairos payments, <https://dune.com/queries/6813491>.

⁷ Wintermute CEX–DEX contracts: `0xcb43d843f6cadf4f4844f3f57032468aadd9b95c` and `0x27920e8039d2b6e93e36f5d5f53b998e2e631a70`.

⁸ Selini CEX–DEX contract: `0xee2e7bbb67676292af2e31dffd1fea2276d6c7ba`.

3.2 Markouts and Profits

To estimate the PnL of Wintermute and Selini, we measure the realized value of their CEX–DEX arbitrage trades using markouts. The searchers’ exact off-chain trading strategies are unobserved, including when and how positions accumulated on DEXes are offset on CEXes. Consider an arbitrage trade i executed on-chain at time t that purchases x units of token A for y units of token B on a DEX. Assuming the position is offloaded on a CEX after a markout horizon m at prices $P_A(t + m)$ and $P_B(t + m)$, the profit is defined as

$$\Pi(i)_{t,m} := x P_A(t + m) - y P_B(t + m) - \text{fees}.$$

Unless otherwise stated, our empirical analyses involving markouts restrict attention to trades with positive estimated PnL, thereby focusing on realized profitable arbitrage activity. For the markout-horizon sensitivity exercise, however, we report net PnL over a fixed sample of all identified trades, since applying the positive-PnL filter separately at each horizon would change the composition of the sample.

For our baseline specification, we use a markout horizon of $m = 5$ seconds. This choice reflects a trade-off: very short horizons may value the DEX leg before prices have sufficiently adjusted across on-chain and off-chain markets, whereas longer horizons are increasingly likely to incorporate price movements unrelated to the original arbitrage opportunity.

To assess the sensitivity of the markout valuation to the chosen horizon, we recompute net markout PnL over a fixed sample of all identified time-boosted trades in August 2025 using horizons of $m \in \{1, 2, 5, 10, 30\}$ seconds. As shown in the appendix of the full version [26], estimated PnL increases between one and five seconds but is comparatively stable thereafter. Relative to the five-second baseline, total estimated PnL is only 1.2% higher at ten seconds and 5.0% higher at thirty seconds. At ten seconds, the corresponding differences are 1.8% for Wintermute and 0.7% for Selini. We therefore select five seconds as the shortest horizon in this relatively stable range, allowing time for cross-market price adjustment while limiting contamination from subsequent market movements.

We obtain price data from the Binance historical data repository [3], using USDT trading pairs sampled at one-second resolution granularity. Following [25], we compute markouts using CEX mid-prices.

3.3 Limitations

Our methodology has several limitations.

1. The heuristic used to detect CEX–DEX arbitrage captures only simple swaps involving a small set of liquid assets known to trade on CEXes. In practice, non-atomic arbitrage searchers may execute more complex strategies involving a long tail of assets [25], as well as over-the-counter markets or venues on other chains [27].
2. Profit estimation assumes that the entire inventory is unwound in a single trade at the chosen markout horizon. In reality, traders may unwind positions gradually to manage price risk, or may choose to retain inventory exposure.
3. Trade coverage is limited by the DEX protocol data available on Dune Analytics.
4. The observability of payments to Kairos is limited to on-chain transfers and does not include payments made through its CEX–DEX subscription model.
5. We assume $1\text{USDT} = 1\text{USDC} = 1\text{USD}$, which may not always hold in practice.

However, because the methodology is applied consistently across regimes, any resulting over- or underestimation of profits should affect all periods similarly and still allow meaningful comparisons across regimes.

■ **Table 1** Pearson and Spearman correlations between per-round Timeboost auction bids and total positive-PnL time-boosted searcher profits during the competition phase. Auction bids are denominated in ETH, while searcher profits are estimated using a 5 s markout. n denotes the number of auction rounds.

Winner	Pearson		Spearman		n
	Top Bid	Paid Bid	Top Bid	Paid Bid	
Wintermute	0.321	0.348	0.449	0.458	108,489
Selini	0.115	0.109	0.337	0.305	93,246

4 Empirical Results

We now present the empirical results. We first evaluate auction effectiveness during the pre-transition competition phase, before turning to the emergence of Kairos-mediated JIT resale and its effects on bidding behavior, auction revenue, and surplus allocation.

4.1 Auction Effectiveness in the Competition Phase

We first assess the effectiveness of the AOT auction during the competition phase, before the Kairos transition. To do so, we measure the relationship between auction bids and estimated searcher PnL at the round level. Table 1 shows that bids are noisy predictors of arbitrage profits. During the competition phase, the Pearson correlations between per-round profits and both the top bid and the paid bid are statistically significant but economically modest. Spearman correlations are higher for both searchers, suggesting that bidders could directionally identify higher-value intervals, but had limited ability to predict the magnitude of realized value.


One way to interpret this pattern is through the relationship between arbitrage profits and return variance. Return variance is a moderately informative predictor of estimated arbitrage profits for the relevant trading pairs, but the fit is substantially stronger over longer horizons than over the one-minute intervals used in the Timeboost auction. Consistent with this, both the top bid and the paid bid are noisy predictors of short-run markout profits, but become more informative when profits are aggregated over longer forward windows. This suggests that bidders are not simply unable to identify valuable market conditions; rather, the main difficulty is pricing the realized value of a specific one-minute fast-lane interval ahead of time. This pattern is summarized in Table 2.

Consistent with this interpretation, auction bids exhibit strong and slowly decaying autocorrelation during the competition phase, as shown in Figure 2. This implies that bids reflect persistent market conditions rather than only round-specific realized opportunities. By contrast, per-round markout PnL is much less persistent. The all-round markout series exhibits little autocorrelation beyond the first two lags, while positive-markout rounds show only limited short-run clustering. This difference is consistent with the idea that bidders can track broader market conditions, but face substantial uncertainty about the realized value of any individual one-minute interval.

The intra-round timing distribution motivates a shorter-window analysis. As shown in Figure 3, time-boosted CEX–DEX transactions are not uniformly distributed within the round. Transaction activity initially declines and then exhibits a pronounced spike around second 10 of the Timeboost interval, which corresponds approximately to the beginning of the next wall-clock minute. Because we do not observe the searchers’ trading logic or off-chain opportunity signals, we cannot identify the cause of this spike directly.

■ **Table 2** Bid–PnL correlations at multiple horizons during the competition phase before the Kairos transition. Correlations are computed over forward windows of length W , where one auction round is approximately one minute. For each window, the bid variable is the mean top or paid bid over the window, and PnL is the total time-boosted PnL over the same window. The sample contains $n = 201,737$ round-level observations.


Bid type	5 min	10 min	15 min	30 min	1 hour
<i>Panel A: Spearman correlations</i>					
Top bid	+0.681	+0.743	+0.768	+0.794	+0.804
Paid bid	+0.650	+0.714	+0.742	+0.774	+0.792
<i>Panel B: Pearson correlations</i>					
Top bid	+0.195	+0.245	+0.278	+0.342	+0.401
Paid bid	+0.178	+0.215	+0.236	+0.316	+0.378



■ **Figure 2** Autocorrelation functions of auction bids and per-round markout PnL during the competition phase before the Kairos transition. *Top*: Autocorrelation of the top and paid bids up to 40 lags, where one lag corresponds approximately to one auction round, or one minute. *Bottom*: Autocorrelation of per-round time-boosted PnL for all rounds and for the positive-markout subsample. Dashed and dotted lines indicate 95% confidence bands.

One possible interpretation is clock-time scheduling. Hasbrouck and Saar [12] document periodicities in Nasdaq message activity associated with automated strategies that access the market at fixed wall-clock intervals. The alignment of the spike with the minute boundary is consistent with a similar mechanism. However, it could also reflect periodic updates on external venues, protocol-specific timing effects, or common reactions by the dominant searchers to the same market signals. We therefore interpret the timing pattern as suggestive rather than causal.

Because activity is concentrated later in the interval, the first seconds of each round contribute disproportionately little to total realized PnL, which is relevant when interpreting correlations measured over shorter markout windows. We therefore repeat the round-level correlation exercise using shorter markout windows. Instead of measuring PnL over the full



■ **Figure 3** Intra-round timing distribution of time-boosted CEX–DEX transactions by Wintermute and Selini during the competition phase. Each line represents the fraction of transactions submitted in a given second of the one-minute Timeboost round. Dashed line indicates the uniform distribution benchmark.

■ **Table 3** Correlations between auction bids and positive-PnL time-boosted markouts over shorter horizons during the competition phase before the Kairos transition. Markouts are computed over the first 15 and 30 seconds of the one-minute Timeboost interval. n denotes the number of auction rounds.

Winner	15 seconds		30 seconds		n
	Top bid	Paid bid	Top bid	Paid bid	
<i>Panel A: Pearson correlations</i>					
Wintermute	+0.240	+0.276	+0.282	+0.314	116,679
Selini	+0.087	+0.098	+0.101	+0.110	133,774
<i>Panel B: Spearman correlations</i>					
Wintermute	+0.385	+0.386	+0.421	+0.424	116,679
Selini	+0.292	+0.277	+0.315	+0.302	133,774

one-minute Timeboost interval, we compute markouts over the first 15 and 30 seconds of the interval. This test captures competing forces. On the one hand, bidders may be better informed about more immediate opportunities, which could increase the correlation between bids and realized PnL. On the other hand, shorter windows contain less aggregation, so realized profits are likely to be noisier than over the full minute. Moreover, because activity is concentrated later in the round, shorter windows also mechanically capture less of the realized opportunity, which works against finding strong correlations. The results in Table 3 support the latter interpretations: correlations remain positive but economically modest, and are generally weaker than for longer aggregation windows.

4.2 Bid Coordination and Just-in-Time Resale

In this section, we analyze the transition from direct competition in the primary Timeboost auction to Kairos-mediated fast-lane access. We focus on how this shift reshapes bidding behavior, primary auction revenue, and the distribution of surplus.

■ **Table 4** Share of auction rounds won by each entity per period.

Entity	Competition	Kairos	Res. Price Adapt.	Steady State	Overall
Wintermute	45.0% (7,682/17,065)	0.0% (1/8,593)	72.7% (141/194)	11.3% (2,478/21,837)	21.6% (10,302/47,689)
Selini	53.4% (9,109/17,065)	0.2% (18/8,593)	3.6% (7/194)	9.8% (2,134/21,837)	23.6% (11,268/47,689)
Kairos	1.4% (244/17,065)	99.6% (8,559/8,593)	23.7% (46/194)	78.4% (17,119/21,837)	54.5% (25,968/47,689)

4.2.1 Timeline

Within our study period, we identify three key on-chain events that delineate the observed market dynamics:

1. **2026-02-12 20:31:51 UTC** – Wintermute and Selini cease competitive bidding in the primary Timeboost auction and begin sourcing fast-lane access through Kairos.
2. **2026-02-18 20:01:51 UTC** – Arbitrum raises the Timeboost reserve price from 0.001 ETH to 0.0075 ETH.
3. **2026-02-25 19:49:51 UTC** – Arbitrum lowers the reserve price back to 0.001 ETH.

These events divide the February 1–March 12, 2026 transition-analysis window into four sub-periods: *Competition* (pre-transition February), *Kairos* (February 12–18), *Reserve Price Adaptation* (February 18–25), and *Steady State* (after February 25). The label *Competition* refers specifically to the February segment of the broader pre-transition competition phase. We use this shorter pre-transition window because auction dynamics in February are representative of the full pre-transition period: the August 2025–February 2026 data exhibit a stable structure (see Figure 1).

4.2.2 Bidding Behavior

During the coordination phase, we observe 47,689 Timeboost auctions, of which Kairos (54.5%), Selini (23.6%), and Wintermute (21.6%) together account for 99.7% of all wins (see Table 4). In total, 117,146 bids were placed, with an average of 2.46 bids per round across 14 distinct bidding addresses. Selini accounts for the highest share of bids (38.7%), followed by Wintermute (32.7%) and Kairos (25.0%); together these three entities are responsible for 96.4% of all bids. We summarize auction participation and active periods for each bidding address in the appendix of the full version [26]. We interpret these as distinct on-chain bidding accounts and do not assume that unattributed addresses necessarily correspond to distinct economic entities.


Looking only at the number of bids per round might suggest that the competitiveness of the auction was not significantly affected by the adoption of Kairos. Table 5 shows that the number of rounds in which all three entities placed a bid is highest after this transition. However, examining the gap between the winning bid (top bid) and the paid bid reveals a clear decline in auction competitiveness.

Figure 4 shows that during the *Competition* phase, both the top bid and the paid bid move closely with volatility in ETH prices⁹. This relationship is confirmed by the correlation measures reported in the appendix of the full version [26], which show strong and statistically significant linear relationships between volatility and both paid and top bids. The absolute bid gap also increases with volatility, indicating that while both bids respond to more valuable trading opportunities, the top bid tends to rise more strongly. During this period, the median daily relative bid gap is 37.3%.

⁹ We compute volatility as the standard deviation of log returns from consecutive one-second ETH/USDT mid-price observations.

■ **Table 5** Share of auction rounds by bidder participation combination in each period. WM = Wintermute, Sel = Selini, Kai = Kairos.

Combination	Competition	Kairos	Res. Price Adapt.	Steady State	Overall
WM + Sel + Kai	1.5% (262/17,065)	10.9% (935/8,593)	1.5% (3/194)	85.3% (18,621/21,837)	41.6% (19,821/47,689)
WM + Sel	88.8% (15,161/17,065)	0.0% (0/8,593)	4.1% (8/194)	5.1% (1,118/21,837)	34.2% (16,287/47,689)
WM + Kai	0.0% (4/17,065)	0.2% (14/8,593)	10.3% (20/194)	6.7% (1,471/21,837)	3.2% (1,509/47,689)
Sel + Kai	0.2% (41/17,065)	87.1% (7,485/8,593)	0.5% (1/194)	1.2% (257/21,837)	16.3% (7,784/47,689)
Only WM	1.2% (198/17,065)	0.0% (0/8,593)	67.5% (131/194)	1.4% (314/21,837)	1.3% (643/47,689)
Only Sel	8.2% (1,399/17,065)	0.0% (4/8,593)	3.1% (6/194)	0.2% (35/21,837)	3.0% (1,444/47,689)
Only Kai	0.0% (0/17,065)	1.8% (155/8,593)	12.9% (25/194)	0.1% (20/21,837)	0.4% (200/47,689)
None	0.0% (0/17,065)	0.0% (0/8,593)	0.0% (0/194)	0.0% (1/21,837)	0.0% (1/47,689)



■ **Figure 4** Timeboost auction bidding dynamics during the transition-analysis window. *Top*: Top bid and paid bid per auction round, measured in ETH. The dotted line shows daily ETH volatility on the right axis. *Bottom*: Relative bid gap between the top bid and the paid bid for each auction round, shown as points, together with the daily median, shown as a line. Breaks in the daily-median line indicate periods with no auction rounds receiving bids.

Once Wintermute and Selini retreat from the auction and begin sourcing fast-lane access through Kairos (the *Kairos* period), the median daily relative bid gap increases to 57.5%. At the same time, the paid bid only weakly tracks price volatility, while the top bid shows no statistically significant relationship. During this period, Kairos wins 99.6% of the rounds and predominantly submits a fixed bid of 0.004 ETH, while the paid price typically remains only marginally above the reserve price, usually submitted by Selini (see Table 5). As a result, the bid gap becomes negatively correlated with volatility, which can be explained by the paid price adjusting slightly with volatility while the top bid remains largely fixed. Consequently, auction pricing in this regime remains largely compressed despite changing market conditions.

One notable deviation during the *Kairos* period occurs in the final 24 hours before the reserve-price increase. At 20:00:51 UTC on February 17, the previously inactive and unattributed address `0x3141592601bf441700f150bef82db6040f4fe373` begins submitting several different bid amounts. After approximately two hours, it switches to a fixed bid

of 0.003970 ETH, only 0.75% below Kairos’s predominant fixed bid of 0.004 ETH, and submits this amount in every round from 23:00 on February 17 until approximately 06:00 on February 18. Because Timeboost is a second-price auction, the address becomes the price-setting second-highest bidder in most of these rounds, sharply reducing the relative bid gap and increasing the price paid by Kairos.


From 06:00 until the reserve-price update approximately 14 hours later, the address instead submits a fixed bid of 0.002921 ETH, producing the roughly 27% relative-bid-gap plateau visible in Figure 4. The address wins only 13 rounds, approximately 1% of the rounds in which it participates, but sets the paid bid in 1,238 rounds, or 96% of its active rounds. In the remaining rounds, another address submits the second-highest bid. The address is listed in the appendix of the full version [26], but we find no public attribution for it.

During the *Reserve Price Adaptation* period, the reserve price increases from 0.001 ETH to 0.0075 ETH and bidding activity nearly disappears. Only 194 rounds receive at least one bid, the majority of which (131 rounds) involve Wintermute as the sole bidder. Among these active rounds, bidding activity is again positively correlated with price volatility, consistent with Wintermute deriving value from external price movements as a CEX–DEX searcher. The apparent increase in the daily median relative bid gap in Figure 4 should therefore not be interpreted as a gradual change in bidding behavior, as it is based on a small number of active rounds. In rounds with only one bidder, the paid price is determined by the reserve, so the relative bid gap reflects the difference between Wintermute’s bid and the elevated reserve price. Bidding resumes more consistently only after the reserve price returns to 0.001 ETH.


Finally, the market enters what we refer to as the *Steady State* once Arbitrum lowers the reserve price back to its original value of 0.001 ETH. Despite having the same reserve price as the earlier *Kairos* period, bidding behavior differs substantially. All three entities submit bids in 85.3% of rounds, while Kairos wins 78.4% of rounds. Both the top bid and the paid bid again exhibit strong positive correlations with volatility, as does the absolute bid gap. However, the paid bid remains effectively capped in most rounds, as reflected in Figure 4. As a result, the top bid rises more strongly than the paid bid when trading opportunities become more valuable, and the median relative bid gap stabilizes at 85.2%. Thus, the paid bid amounts to only approximately 14.8% of the top bid, indicating that the primary auction captures only a limited share of the value reflected in submitted bids.

Unlike in the preceding *Kairos* period, Kairos no longer submits a predominantly fixed bid or wins nearly every auction. Across the 21,837 rounds with at least one bid, Kairos participates in 93.3% and wins 84.0%. The median relative gap between Kairos’s winning bid and the second-highest bid is approximately 86.0%, whereas it falls to 32.0% in rounds that Kairos loses. The difference is even more pronounced in absolute terms: the median gap is approximately 0.015 ETH when Kairos wins, compared with 0.001 ETH when another bidder wins.

This asymmetry is also reflected in the bid distributions shown in Figure 5. During the *Competition* period, median bids are relatively close across Wintermute, Selini, and Kairos (0.0075, 0.0062, and 0.0069 ETH, respectively). In the *Steady State*, they diverge substantially: Kairos’s median bid rises to 0.0182 ETH, whereas the corresponding medians for Wintermute and Selini fall to 0.0027 and 0.0021 ETH. The same pattern appears in the upper tail of the distributions (p99: Wintermute 0.0096 ETH, Selini 0.0101 ETH, and Kairos 0.0703 ETH). A more detailed distribution is reported in the appendix of the full version [26]. These patterns show that the renewed participation of Wintermute and Selini does not constitute a return to the competitive bidding observed before the adoption of Kairos. Kairos bids substantially more aggressively, while the other two searchers generally submit bids close to the reserve.



■ **Figure 5** Bid distributions for Wintermute, Selini, and Kairos. The left half of each violin corresponds to the *Competition* period, while the right half corresponds to the *Steady State*. The dashed horizontal line indicates the reserve price (0.001 ETH).



■ **Figure 6** Kairos loss rate and median bids by hour of day in the *Steady State* period. The shaded region indicates UTC hour bins overlapping the regular NYSE session (14:30–21:00 UTC during EST).

To better understand when Kairos loses in the *Steady State*, we examine the temporal distribution of its losses. As shown in Figure 6, Kairos loses more frequently during U.S. market hours, particularly around the NYSE opening: its loss rate is 19.0% during NYSE hours, compared with 14.5% outside those hours. Selini, in particular, submits somewhat higher bids during these periods. Nevertheless, Kairos’s median bid in rounds it loses remains substantially below its median winning bid. Heightened market activity therefore explains only part of the variation in Kairos’s losses. More generally, losses tend to occur when Kairos submits a relatively low bid or does not participate, rather than when Wintermute or Selini submit bids comparable to Kairos’s usual winning bids.

One possible interpretation is that the temporary reserve-price increase altered participants’ expectations about future protocol responses, making a complete withdrawal from the primary auction less attractive after the reserve returned to 0.001 ETH. This could help explain why all three entities maintain visible bidding activity in the *Steady State*, while Kairos submits higher and more variable bids to preserve its access to the fast-lane. Because participants’ expectations and off-chain arrangements are unobserved, we cannot identify the cause of this shift directly. We therefore interpret the *Steady State* as a modified coordination regime with limited residual competition, rather than a return to the earlier competitive environment.

4.2.3 Searcher Activity

We next examine how the changing auction dynamics affect searcher activity and the use of the Timeboost fast-lane versus the regular-lane. The top-left and top-right panels in Figure 7 show the number of time-boosted and regular CEX–DEX transactions submitted over time by Wintermute and Selini, respectively. The bottom-left panel reports daily PnL, while the bottom-right panel shows the average PnL per trade per day for time-boosted and regular transactions. Consistent with our baseline PnL specification, we restrict attention to trades with positive estimated PnL. Results are summarized in the appendix of the full version [26].


Before the adoption of Kairos, in the *Competition* phase, searchers submit CEX–DEX transactions through both the Timeboost fast-lane and the regular-lane when they do not control the fast-lane. Their usage patterns differ meaningfully: Wintermute earns more total PnL from time-boosted trades and submits relatively more of them, whereas Selini earns more total PnL from regular-lane trades and also achieves higher average profit per trade there. More broadly, although Wintermute trades more frequently and earns higher total profits overall, Selini tends to realize higher average profit per trade. These differences remain largely stable across auction regimes and suggest that the two searchers follow somewhat different trading models. At the same time, there are periods in which average profit per trade is similar across lanes, and sometimes even higher in the regular-lane for both searchers, indicating that Timeboost is not reserved only for the most profitable opportunities but also helps secure smaller trades that might otherwise be lost to latency competition. For both searchers, PnL from both time-boosted and regular-lane trades remains positively correlated with volatility (see the appendix of the full version [26]), consistent with the idea that higher volatility creates more opportunities to exploit stale DEX prices relative to rapidly moving CEX prices.

Once Wintermute and Selini begin using Kairos, both almost completely stop submitting transactions through the regular-lane, as they can now access the fast-lane via Kairos. This indicates that Kairos effectively substitutes for direct participation in the Timeboost auction, allowing searchers to retain fast-lane access without bidding competitively in the primary auction. During this *Kairos* phase, the PnL from time-boosted trades remains strongly correlated with volatility for both searchers, whereas the relationship is considerably weaker for regular-lane trades, suggesting a clear shift in how the two submission lanes are used.

During the *Reserve Price Adaptation* phase, when bidding activity was minimal and resumed only toward the end, both searchers mainly continued operating through the regular transaction lane. Their average PnL remains similar to the *Kairos* period, but the volatility relationships appear to reverse: regular-lane transactions now show a strong correlation with volatility for both searchers, while the much less frequent time-boosted trades exhibit a weaker relationship.

Finally, once the reserve price returns to its original level, both searchers move back to the fast-lane while submitting relatively few regular-lane transactions. In this *Steady State* period, both execute significantly more time-boosted trades relative to regular-lane trades than in the *Competition* period, reflecting improved access to the fast-lane through Kairos.

Overall, the adoption of Kairos allows searchers to access the fast-lane almost continuously, largely eliminating the need to rely on the slower regular-lane while time-boosted activity continues to track price movements closely. However, consistent with the weaker bidding competition documented above, this increased reliance on the fast-lane does not translate into a proportional increase in primary-auction revenue.



■ **Figure 7** CEX-DEX arbitrage activity by Wintermute and Selini during the transition-analysis window. *Top-left*: time-boosted and regular transaction count by Wintermute. *Top-right*: time-boosted and regular transaction count by Selini. *Bottom-left*: daily PnL from time-boosted and regular transactions for Wintermute and Selini. The dotted line shows daily ETH volatility (right axis). *Bottom-right*: average PnL per trade for time-boosted and regular transactions for Wintermute and Selini.

4.2.4 Surplus

Lastly, we examine how the changing auction dynamics affect the total surplus generated and its distribution across the involved parties. We define net surplus as

$$\text{Net Surplus} := \sum_{\text{Wintermute, Selini}} \text{PnL} - \text{transaction fees} - \text{auction bids paid},$$

where transaction fees and auction bids paid together represent revenue captured by Arbitrum, while total searcher PnL captures the value generated by CEX-DEX arbitrage opportunities and ultimately shared among participants in the Timeboost ecosystem.

Figure 8 presents the evolution of total surplus over time (top panel) and the share of searcher PnL captured by Arbitrum through auction bids and transaction fees (bottom panel). Throughout the study period, searcher PnL is strongly positively correlated with daily ETH volatility ($r = +0.96$, $p < 0.001$), with visible spikes during the *Competition* phase when volatility was highest. Importantly, total searcher PnL remains of similar magnitude across regimes, suggesting that the observed changes in primary-auction competition mainly reflect a redistribution of surplus rather than an increase in total value generated.

We find that bids paid correspond on average to 7.4% of total searcher PnL during the *Competition* phase¹⁰, but this declines following the adoption of the Kairos secondary market (*Kairos*: 5.8%, *Reserve Price Adaptation*: 0.4%, *Steady State*: 6.0%). This decline is even

¹⁰This share is likely understated, as we restrict attention to positive-PnL trades. Including loss-making trades would reduce net searcher PnL and thus mechanically increase the captured share.




Figure 8 Surplus dynamics during the February 1–March 12, 2026 transition-analysis window. *Top*: daily surplus decomposition into searcher PnL, transaction fees, and Timeboost auction bids paid. The black solid line shows the net surplus, while dotted line shows daily ETH volatility (right axis). *Bottom*: share of searcher PnL retained as net surplus versus captured by Arbitrum through transaction fees and paid auction bids. The hatched area shows the difference between the share of surplus reflected in the top bid and the share captured through the paid bid.

more pronounced when focusing only on time-boosted surplus and excluding regular-lane transactions (*Competition*: 12.5%, *Steady State*: 6.7%). This pattern is consistent with increasing reliance on the fast-lane via Kairos together with weaker direct competition in the primary auction.

Although bids in the Timeboost auction need not directly reveal true valuations because of the common-value component and strategic shading, stronger competition in the *Competition* period likely allowed a larger share of surplus to be captured through the auction than in the *Steady State*, where Kairos is often the only competitive participant with a high value for winning the lane. Consistent with this interpretation, the top bid corresponds to a substantially larger share of surplus in the final regime (*Competition*: 13.6%, *Steady State*: 36.3%), while the gap between the surplus reflected in the top bid and the surplus captured through the paid bid also widens markedly (bottom panel of Figure 8). The key change, therefore, is not that bidding ceases to respond to valuable trading opportunities, but that weaker competition causes a smaller fraction of the underlying value to be captured through the paid bid. Whether this reflects more aggressive shading, coordination, or lower standalone value from controlling the lane for bidders other than Kairos, the implication is the same: weaker competition increases the disconnect between value and revenue.

We now examine payments to Kairos to better understand how surplus may be distributed outside the primary auction. Figure 9 compares the total daily bids paid by Kairos in the Timeboost auction with the on-chain payments it receives from transactions using its service. While Kairos’s auction payments increase substantially once it becomes the dominant auction participant, the observed on-chain payments it receives do not increase proportionally.




Figure 9 Kairos primary auction payments and resale revenue during the February 1–March 12, 2026 transition-analysis window. The figure compares the daily Timeboost auction bids paid by Kairos to the on-chain payments Kairos receives from transactions using its service.

Over the study period, Kairos pays a total of 151,302 USD in auction bids, while the on-chain payments we observe sum to only 8,001 USD. The resulting gap of more than 143,000 USD suggests that a substantial portion of the payments to Kairos likely occurs through alternative mechanisms that are not visible on-chain, such as its subscription model for CEX–DEX arbitrage searchers.

During periods when Kairos controls the fast-lane, Wintermute and Selini submit a combined total of 1,004,876 time-boosted CEX–DEX transactions with positive PnL, generating approximately 2,541,351 USD. However, none of the identified on-chain payments to Kairos originate from their transactions, suggesting that these entities primarily rely on Kairos’s off-chain payment structure.

Taken together, these findings are consistent with the evidence above that, once Kairos becomes dominant, a smaller share of the value generated around Timeboost is captured through the primary auction. However, because payments to Kairos are only partially observable on-chain, the ultimate division of the remaining surplus between the intermediary and the searchers remains unclear. Hence, we cannot establish the extent using Kairos benefits Wintermute and Selini as the final surplus allocation is unobservable on-chain.

5 Concluding Discussion

We conclude by discussing what our empirical results imply for value extraction in AOT auctions. We first consider the effectiveness of AOT auctions when access is allocated through the primary auction alone, then examine how JIT resale changes revenue capture and surplus allocation. Finally, we discuss broader implications for Ethereum.

5.1 Value Extraction Through Ahead-of-Time Auctions

Analyzing Timeboost auction data in the competition phase shows that estimating immediate future arbitrage value is difficult. Bids are reasonably informative about longer-horizon average arbitrage profits, but they are noisy predictors of one-minute or shorter-horizon profits. This finding casts doubt on whether AOT auctions are effective MEV extraction mechanisms for rollups or other blockchains, even when no JIT resale market forms.

In a pure common-value auction with risk-neutral bidders, noisy ex-ante signals do not necessarily imply low expected seller revenue: bidders should condition on the winner’s curse and shade their bids relative to their private signals, but competition can still transfer a substantial share of expected value to the seller. In practice, however, the risk-neutrality

assumption might not hold perfectly. Bidders may be risk averse, face capital or inventory constraints, or differ in their private ability to exploit opportunities. Such private-value components weaken the link between common arbitrage value and bids. They also imply that selling ordering rights AOT may fail to allocate the fast-lane to the bidder with the highest realized value for it.

5.2 Value Extraction if There is a Just-in-Time Resale Market

The adoption of the Kairos resale market effectively dampens competition in Arbitrum’s Timeboost auction, consistent with the weakened-competition channel being the dominant mechanism – though we find some limited supporting evidence for un-bundling and information effects as well. As a result, the protocol appears to capture a substantially smaller share of the value generated around Timeboost, while a larger share remains outside the primary auction. However, the exact distribution of this additional surplus between the searchers and Kairos remains unclear, as the payments associated with the resale mechanism are only partially observable on-chain.

The success of the Kairos resale market is also reflected in the increasing volume of time-boosted transactions once Kairos wins most auction rounds. However, this effect is not clearly visible in PnL, which is a noisier signal. As a result, the unbundling effect appears to be second-order relative to the impact of weaker competition on auction outcomes. This channel may become more prominent if Kairos attracts additional searchers in the future, particularly those employing more diverse strategies such as atomic arbitrage.


To improve revenue capture, Arbitrum could pursue several directions. Beyond adopting an entirely different mechanism or market structure, the current auction format itself could be modified to better align payments with the value reflected in bids.

First, switching to a first-price auction could make anti-competitive bidding more difficult [17] and may reduce the gap between value and payments. Second, the auction could respond more effectively to changing demand by adjusting the reserve price dynamically rather than relying on a fixed reserve. In particular, incorporating external price volatility – which appears to be a key driver of searcher activity and auction values – into the reserve-price design could push bids closer to bidders’ willingness to pay.

We illustrate this with a simple volatility-based reserve-price rule. Specifically, we set the reserve price equal to $c\sigma_t^2$, where c is a constant calibrated using bid data from the *Competition* period and σ_t denotes the standard deviation of log returns, $\log(p_{i+1}/p_i)$, computed over a window of length t before the auction ends. Using this rule, we recover at least 80% of the revenue benchmarked against the sum of top bids during the period from August 2025 to January 2026. Figure 10 shows the performance of this dynamic reserve-price rule across different choices of t and c .

5.3 Lessons for Ethereum L1

The lessons from Timeboost’s AOT auction may also be relevant for Ethereum L1, subject to two caveats. First, Timeboost is mainly used for CEX–DEX arbitrage, which captures only part of Ethereum MEV, albeit an important and highly latency-sensitive part. A rough analogue on Ethereum would be an AOT auction for top-of-block rights over several consecutive blocks. Second, while revenue capture through Timeboost is a primary concern for Arbitrum, protocol-level MEV capture on Ethereum – for example through MEV-burn-type designs [6] – is arguably less central than preserving broader system properties.



■ **Figure 10** Performance of the dynamic reserve-price rule for different choices of t and c , measured relative to the sum of top bids.

With these caveats in mind, two lessons appear particularly relevant. First, on Timeboost the population of dominant searchers is narrow, and the value they extract from the fast-lane is largely concentrated in the same type of opportunity: CEX–DEX arbitrage on a few liquid pairs. On Ethereum L1, MEV opportunities are more heterogeneous and accrue to different searcher classes at different times. An AOT auction allocates ordering rights before this realization is known and is therefore more likely to misallocate them relative to a JIT mechanism. We would expect this allocative inefficiency to be larger on L1, even before any secondary market emerges.

Second, limited competition among dominant searchers can create favorable conditions for secondary markets that move value outside the primary auction. Ethereum’s searcher and builder markets remain more competitive than Arbitrum’s and are not currently as concentrated as the duopolistic structure we observe here. Still, the searcher market on Ethereum has become more concentrated over time, and similar resale dynamics could plausibly emerge if competition weakens substantially.

More broadly, the Arbitrum case suggests that AOT allocation mechanisms may be particularly vulnerable to secondary-market intermediation once competition among dominant participants becomes weak. This vulnerability may itself be amplified by AOT design. In a JIT auction, valuations are likely to be more dispersed across auctions, allowing participants who are usually not competitive to occasionally outbid dominant players when they identify especially valuable opportunities. Such episodic entry can help preserve some competitive pressure. When ordering rights are sold ahead of time, this source of occasional competitiveness is weaker, potentially making uncompetitive equilibria even easier to sustain.

References

- 1 Arbitrum Foundation. How to use timeboost: How to view historical bid data. <https://docs.arbitrum.io/how-arbitrum-works/timeboost/how-to-use-timeboost#how-to-view-historical-bid-data>, 2026. Accessed: 2026-03-11.

- 2 Arbitrum Foundation. Timeboost: A gentle introduction. <https://docs.arbitrum.io/how-arbitrum-works/timeboost/gentle-introduction>, 2026. Accessed: 2026-03-11.
- 3 Binance. Binance market data. <https://data.binance.vision>, 2026. Accessed: 2026-03-11.
- 4 Agostino Capponi and Brian Zhu. Does timeboost reduce mev-related spam? theory and evidence from layer-2 transactions. *CoRR*, abs/2512.10094, 2025. doi:10.48550/arXiv.2512.10094.
- 5 Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *2020 IEEE Symposium on Security and Privacy, SP 2020, San Francisco, CA, USA, May 18-21, 2020*, pages 910–927. IEEE, 2020. doi:10.1109/SP40000.2020.00040.
- 6 Justin Drake. Mev burn—a simple design. <https://ethresear.ch/t/mev-burn-a-simple-design/15590>, 2023.
- 7 Dune Analytics. Dune: Blockchain data analytics platform. <https://dune.com/>, 2026. Accessed: 2026-03-11.
- 8 Flashbots. What is mev-boost? <https://docs.flashbots.net/flashbots-mev-boost/introduction>, 2026. Accessed: 2026-03-18.
- 9 Robin Fritsch, Maria Inês Silva, Akaki Mamageishvili, Benjamin Livshits, and Edward W. Felten. MEV capture through time-advantaged arbitrage. *CoRR*, abs/2410.10797, 2024. doi:10.48550/arXiv.2410.10797.
- 10 Gattaca. Kairos timeboost documentation. <https://docs.kairos-timeboost.xyz/>, 2026. Accessed: 2026-03-11.
- 11 Gattaca. Parallel kairos lanes. <https://docs.kairos-timeboost.xyz/parallel-kairos-lanes>, 2026. Accessed: 2026-03-12.
- 12 Joel Hasbrouck and Gideon Saar. Low-latency trading. doi:10.2139/ssrn.1695460.
- 13 Boyi Li and Weixuan Xia. Crypto inverse-power options and fractional stochastic volatility. *Quantitative Finance*, 25(7):1073–1099, 2025. doi:10.1080/14697688.2025.2518252.
- 14 Julian Ma. Block vs. Slot Auction PBS. *Paragraph*, 2022. Available at <https://paragraph.com/@julianma/block-vs-slot-auction-pbs>.
- 15 Julian Ma. Trusted advantage in slot auction epbs. <https://ethresear.ch/t/trusted-advantage-in-slot-auction-epbs/20456>, 2024.
- 16 Akaki Mamageishvili, Mahimna Kelkar, Jan Christoph Schlegel, and Edward W. Felten. Buying Time: Latency Racing vs. Bidding for Transaction Ordering. In Joseph Bonneau and S. Matthew Weinberg, editors, *5th Conference on Advances in Financial Technologies (AFT 2023)*, volume 282 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 23:1–23:22, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.AFT.2023.23.
- 17 Robert C. Marshall and Leslie M. Marx. Bidder collusion. *Journal of Economic Theory*, 133(1):374–402, March 2007. URL: <https://ideas.repec.org/a/eee/jetheo/v133y2007i1p374-402.html>.
- 18 Johnatan Messias and Christof Ferreira Torres. The express lane to spam and centralization: An empirical analysis of arbitrum’s timeboost. *CoRR*, 2025. doi:10.48550/arXiv.2509.22143.
- 19 Paul R. Milgrom and Robert J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50, 1982.
- 20 Jason Milionis, Ciamac C Moallemi, Tim Roughgarden, and Anthony Lee Zhang. Automated market making and loss-versus-rebalancing. *arXiv preprint*, 2022. arXiv:2208.06046.
- 21 Ciamac C. Moallemi, Mallesh M. Pai, and Dan Robinson. Latency advantages in common-value auctions. *CoRR*, abs/2504.02077, 2025. arXiv:2504.02077.
- 22 Mallesh Pai and Max Resnick. Centralization in attester-proposer separation. <https://arxiv.org/abs/2408.03116>, 2024. doi:10.48550/arXiv.2408.03116.

- 23 Christoph Schlegel. Transaction ordering auctions. In *International Conference on Financial Cryptography and Data Security*, pages 91–104. Springer, 2024. doi:10.1007/978-3-031-69231-4_7.
- 24 Caspar Schwarz-Schilling, Fahad Saleh, Thomas Thiery, Jennifer Pan, Nihar Shah, and Barnabé Monnot. Time is money: Strategic timing games in proof-of-stake protocols. In Joseph Bonneau and S. Matthew Weinberg, editors, *5th Conference on Advances in Financial Technologies, AFT 2023, October 23-25, 2023, Princeton, NJ, USA*, volume 282 of *LIPIcs*, pages 30:1–30:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.AFT.2023.30.
- 25 Fei Wu, Danning Sui, Thomas Thiery, and Mallesh Pai. Measuring CEX-DEX extracted value and searcher profitability: The darkest of the MEV dark forest. *CoRR*, abs/2507.13023, 2025. doi:10.48550/arXiv.2507.13023.
- 26 Burak Öz, Akaki Mamageishvili, Christoph Schlegel, and Ali Taslimi. Just-in-time resale in an ahead-of-time auction for faster execution, 2026. arXiv:2603.20175.
- 27 Burak Öz, Christof Ferreira Torres, Christoph Schlegel, Bruno Mazorra, Jonas Gebele, Filip Rezabek, and Florian Matthes. Cross-chain arbitrage: The next frontier of mev in decentralized finance. <https://arxiv.org/abs/2501.17335>, 2025. arXiv:2501.17335.