Auctioning Time to Mitigate Latency Races: Theory and Evidence from Blockchains

Agostino Capponi* and Brian Zhu†

Abstract

High-frequency trading, in both traditional and decentralized markets, induces latency races and redundant order flow as traders spend resources to win time-sensitive opportunities. We show that auctioning artificial time priority can redirect resources away from wasteful speed races toward auction payments. While such waste is difficult to measure in traditional markets, blockchain transactions provide transparent records of these competitive costs through observable duplicate submissions. We study the introduction of Timeboost, a time-priority auction mechanism on Arbitrum, a blockchain that batches transactions before settlement on Ethereum, as a natural experiment. We find that redundant transactions decrease and platform revenue increases relative to comparable networks, consistent with our theoretical predictions.

The authors are grateful to Edward Felten and Akaki Mamageishvili for helpful comments and discussion, and gratefully acknowledge the support of Ripple's University Blockchain Research Initiative.

*Columbia University, Department of Industrial Engineering and Operations Research and Columbia Business School; ac3827@columbia.edu.

†Columbia University, Department of Industrial Engineering and Operations Research; bzz2101@columbia.edu.

1 Introduction

Speed is a central determinant of trading outcomes in high-frequency environments. When access to time-sensitive trading opportunities depends on being first to a platform, traders invest in costly technology, connectivity, and order management systems to reduce latency. These latency races create large expenditures with limited social upside as resources are focused toward speed acquisition and redundant order flow creates congestion externalities. Empirical evidence quantifying these races confirms their significance. Aquilina et al. (2022) find latency arbitrage races, where traders rapidly attempt to "snipe" profitable opportunities and cancel failed attempts, occur roughly once per minute per symbol in FTSE 100 stocks. However, the investment decisions underlying these races remain unobservable, because speed technology investment in traditional markets, such as co-location services and proprietary infrastructure, is private information and unable to be observed.

Many studies in the market-design literature have proposed altering the allocation of time priority, for example through frequent batch auctions (Budish et al. (2015)), speed bumps such as on IEX (Woodward (2018)), or Pigouvian taxes on investment technologies (Biais et al. (2015)). An alternative approach to mitigating latency races is to sell time priority explicitly.[1] Instead of rewarding whoever arrives first, a platform can auction an ordering advantage that determines execution precedence for a specified time interval. The rationale behind time-priority auctions is to convert competition over speed into competition over price: traders bid for priority rather than investing in latency reduction technologies. The idea is closely related to broader proposals in market design that replace technological races with price-based allocation of scarce trading privileges, including batch auctions and congestion pricing mechanisms. Whether this reallocation occurs in practice, however, is difficult to test in traditional markets because the underlying investments in speed technology are not directly observable.

We study this question in blockchains, where the outcomes of latency races are more directly observable and most notably, *costly*. Blockchain-based trading venues generate time-sensitive rent extraction opportunities known as maximal extractable value (MEV), which include centralized and decentralized exchanges (Capponi and Jia (2025)), cyclic arbitrage between exchanges (Solmaz et al. (2025)), frontruns (Capponi et al. (2025), Lehar and Parlour (2023)), backruns, and liquidations

[1] While the IEX introduces a 350 millisecond delay for all non-resting orders, this is not a winner-take-all approach and is mainly designed to protect resting orders from being picked off during periods of high volatility.

on decentralized lending protocols. Participants compete to be first to exploit profitable opportunities, but the nature of competition differs from traditional markets. Improvements to latency see diminished returns relative to traditional electronic markets due to how blockchain networks process transactions.[2] As a result, traders often submit multiple copies of the same transaction in rapid succession, particularly on so-called "Layer-2" blockchains, hoping one executes while the rest fail: this phenomenon where competitive resources are focused on redundant submissions is an equilibrium outcome microfounded by Mazorra et al. (2026).[3]

Unlike traditional electronic platforms where order submission and cancellation is costless, users incur infrastructural costs for failed attempts on blockchains. Moreover, while message data on traditional exchanges is proprietary and may require invoking financial market regulations to acquire, failed attempts are public and visible on blockchains. These duplicate submissions constitute an on-chain trail of evidence and provide a measurable analogue of otherwise unobservable costly efforts to achieve latency. We summarize these differences in Table 1.

We leverage the introduction of *Timeboost* on the Arbitrum blockchain, as a market-design intervention in this observable environment. Timeboost auctions an artificial time priority via a sealed-bid second-price auction. Rather than racing through duplicate submissions, traders can compete over priority directly via auction. We interpret this setting as a natural laboratory that introduces an explicit market for time priority into an environment where latency races are actively occurring. The experiment allows us to study a general question: what happens when a platform replaces speed races with price-based allocation of priority?

We develop a model in which traders expend costly effort, interpretable as latency investment in traditional markets or duplicate submissions in blockchain markets, to capture time-sensitive opportunities; these investments have convex costs. We show that introducing a time-priority auction reduces equilibrium effort, reallocating resources from wasteful speed competition into platform revenue. The blockchain environment represents a special case where effort costs are approximately linear, as they are gas fees paid from the reversion of non-winning duplicate transactions. Addition-

[2]See Section 2 for details on how blockchain transaction routing creates this competitive dynamic.

[3]Layer-2 blockchains execute transactions off the main Ethereum network and periodically settle batched results on-chain, enabling faster and cheaper execution. Most Layer-2 networks use a centralized sequencer to order transactions, creating a setting where time-priority can be explicitly auctioned. Due to different sequencing rules and block production methods, as well as private mempools, shorter block times, lower gas fees, the MEV ecosystem is different on Layer-2 blockchains (and EVM-compatible Layer-1 blockchains strictly using priority gas auctions to determine transaction ordering) relative to the Ethereum blockchain.

	Speed technology	Order submission
Centralized exchanges	Orders propagated via proprietary networks to matching engine → latency differences translate directly to priority → unobservable but costly competition	Traders cancel and resubmit orders via message data at no cost → observable but costless competition
Layer-2 blockchains	Orders propagated via public networks to sequencing node; latency advantages are noisy and bring diminished returns → low incentive for competition here	Traders submit multiple identical transactions; unsuccessful attempts revert and incur fees → observable and costly competition

Table 1: Dimensions of Latency Competition Across Trading Environments

ally, the user's revert costs are awarded to the blockchain platform, whereas effort costs in general may go to other service providers.[4] Having user costs being accrued by the platform as revenue also holds for traditional HFT settings, as many exchanges either own or tightly control the co-location facilities used by HFT entities. Under this setting, we show that despite accruing fewer fees from fewer transactions submitted, the auction revenue more than makes up for it, increasing overall platform revenue.

How much time priority should be allocated, and is allocating time priority to a single trader the most effective way to mitigate latency races? We extend the model to account for auctions that award multiple time advantage slots of the same amount. Along the intensive margin, i.e. the size of the time advantage, we show that broadly, total effort decreases and platform revenue increases in the amount of time priority given to the winner(s). Along the extensive margin i.e. how many traders receive a time advantage, we show that when the total number of users in the environment is large, increasing the number of winners may further reduce total effort and improve platform revenue.

We empirically test model predictions using transaction-level data across Layer-2 blockchains and an event-study design around Timeboost's deployment on Arbitrum, measuring latency effort via bursts of identical transactions and comparing outcomes to control networks. We find significant evidence that following adoption, redundant submissions decline and platform revenue increases relative to benchmark networks. We additionally find that non-priority traders experience higher failure rates in contested opportunities. These findings match the model's predictions.

[4]Revert costs typically go to the blockchain's sequencer node or associated decentralized autonomous organization (DAO). These entities may then choose to burn all or part of said fees for tokenomics purposes.

Our results have implications for the improved design of traditional markets. Blockchain markets make priority valuations observable through on-chain fees, with Timeboost introducing an explicit market for time priority. While institutional details differ between blockchain and traditional markets—execution speed, asset characteristics, and participant composition—the core economic question is common to both. Can price-based priority allocation substitute for costly speed competition? Our results suggest that explicit priority pricing can reallocate rents from technology investment to platform revenue without diminishing trading activity, providing evidence relevant to broader market design debates about latency races in high-frequency trading environments.

Our findings contribute to the broader debate on the market design of exchanges accommodating high-frequency trading. First, our model predicts that auctioning time priority reduces latency investment effort, consistent with theoretical arguments for batch auctions, and providing another parallel mechanism to reduce latency arms races. Second, we show that platform revenue increases under the auction. This speaks to concerns that exchanges resist efficiency-improving design because it takes away co-location and speed-technology revenue (Budish et al. (2020)). Our empirical analysis provides causal evidence supporting these predictions in a setting where latency competition is observable, and supports time-priority auctions as a parallel option alongside frequent batch auctions to mitigate latency races, with the added benefit that exchange platforms can capture this auction revenue.

Related Literature This paper contributes to three bodies of literature. First, we connect research on latency competition and market design in traditional HFT to decentralized trading environments through a common framework of costly time priority acquisition. A large market microstructure literature studies how HFT and latency competition shape trading strategies and market outcomes. Hasbrouck and Saar (2013) propose measures of low-latency activity on NASDAQ and document how millisecond-scale response strategies affect spreads, depth, and short-term volatility. O'Hara (2015) characterizes how technological change and HFT have transformed the organization of trading, emphasizing that speed and fragmentation alter liquidity provision and price discovery.

Similar interactions have been documented in decentralized trading environments. Daian et al. (2020) document priority gas auctions on Ethereum-based decentralized exchanges, in which

arbitrage bots bid up gas prices to obtain favorable transaction ordering, interpreting these races as a blockchain analogue of HFT latency competition. Subsequent work measures the scale and composition of MEV in DeFi: Weintraub et al. (2022) quantify MEV extraction in private relay pools, and Ferreira Torres et al. (2024) analyze MEV across multiple networks and show that MEV volume and composition on Layer-2s can rival Ethereum's. Solmaz et al. (2025) identify optimistic MEV as a major source of persistent spam-like probing on Layer-2s. Our paper connects these literatures by placing both traditional HFT and decentralized MEV extraction in a common model of costly time priority acquisition and by treating Layer-2 sequencers as matching engines whose design choices shape latency races.

Second, we contribute to market design literature in both traditional and blockchain financial markets by analyzing an explicit market for time priority and showing how it reallocates surplus from dissipation to revenue. In traditional settings, Budish et al. (2015) propose batch auctions to eliminate continuous-time races for queue priority, while Biais et al. (2015) study how slow markets and taxes on speed can implement socially optimal investment in trading technology. On blockchains, Flashbots introduced sealed-bid auctions over transaction bundles as a mechanism to internalize MEV and reduce observable priority gas auctions (Obadia (2020)). Li et al. (2023) and McLaughlin et al. (2025) propose sequencing rules to reduce MEV based on greedy and volatility-minimizing heursitics, respectively. Zhu et al. (2026) study how "revert protection," i.e. a feature where sequencers simulate transactions and exclude failed ones from the final block, saving users from paying gas fees on failed transactions, can make priority fee auctions more competitive and increase sequencer revenue. Mamageishvili et al. (2023) analyze a primitive version of the Timeboost mechanism that blends bidding and latency to allocate transaction priority. Mazorra et al. (2026) show that spamming arises as an equilibrium strategy to capture MEV opportunities such as backruns. Our paper contributes to this mechanism-design literature by analyzing single-winner time-priority auctions, as implemented by Timeboost, and extend it to the multi-winner case. We show that under certain regimes, multi-winner time-priority auctions may result in better efficiency outcomes.

Finally, we contribute to empirical work on decentralized exchange infrastructure by providing novel empirical evidence on how auctioning time priority affects order flow duplication and platform revenue. A growing DeFi microstructure literature studies automated market makers (AMMs),

arbitrage, and liquidity provision Lehar and Parlour (2025); Capponi and Jia (2025). Most closely related to our work is Messias and Ferreira Torres (2025), which provides a descriptive empirical analysis of Timeboost, documenting centralization of express-lane usage and high revert rates. By contrast, we use a difference-in-differences design to estimate the causal impact of Timeboost adoption on redundant orderflow and platform revenue.

2 Institutional Details

This section describes the institutional features that make blockchain-based trading venues a natural laboratory for studying time-priority auctions. We first contrast the infrastructure underlying traditional electronic markets with blockchain platforms, highlighting how differences in connectivity and network architecture determine the form of latency competition. We then explain the structure of Layer-2 blockchains, the types of profit opportunities that drive competition for transaction ordering, and the mechanics of the Timeboost mechanism that we study empirically.

Platform Infrastructure In modern electronic markets, trades occur via a centralized matching engine housed within a data center. Market participants connect directly to this engine using dedicated communication infrastructure, often by co-locating servers inside the exchange facility and routing orders through private cross-connects and fiber or microwave links. The exchange timestamps orders at when they enter its network interface, which determines priority. As connectivity is tightly engineered and shared routing uncertainty is minimized, latency differences translate directly into execution priority. As a result, competition centers on capital-intensive investment in communication technology designed to reduce physical transmission delay.

By contrast, transactions submitted to Layer-2 blockchains traverse heterogeneous paths before reaching the ordering intermediary, often called the *sequencer*. Traders typically transmit through remote procedure call (RPC) gateways, load balancers, and peer-to-peer forwarding layers, all of which introduce random propagation and queueing delays outside the user and sequencer's control. Although the sequencer ultimately processes transactions by timestamp, the arrival time it observes reflects both physical latency and network randomness. In this setting, marginal improvements in transmission may not translate directly into improved execution priority. Thus, as observed

empirically (Solmaz et al. (2025),Gogol et al. (2025)), users adapt by submitting multiple attempts of the same transaction.

Layer-2 Blockchains Ethereum's consensus and blockspace design prioritize decentralization and security, which constrains throughput and leads to congestion-induced fee spikes during active market periods. Layer-2 (L2) networks address this scalability bottleneck by batching and executing transactions off-chain before posting aggregated transaction data to Ethereum for settlement.[5] This division of labor allows L2s to offer lower fees, higher transaction rates, and faster execution while preserving the settlement assurances of the base layer. Users submit signed transactions to a sequencer, which aggregates and executes them in batches before committing the resulting state to the underlying "Layer-1" blockchain. The sequencing step determines which transactions obtain time priority and therefore which traders capture time-sensitive opportunities, such as arbitrage between automated market makers (AMMs) or liquidations in lending protocols.

Layer-2 blockchains implement their own transaction sequencing rules prior to settlement on Ethereum. Most use a *priority gas auction (PGA)* or a *first-come, first-served (FCFS)* sequencing rule. Under the PGA rule, transactions are ordered by a priority fee that users submit along with their transaction; ties in priority fee are commonly broken by the time at which the sequencer received the transaction. While top-of-block MEV (e.g. CEX-DEX arbitrage, liquidations) on PGA-using networks is contested over via first-price auctions, backruns and optimistic MEV does not require competition over priority fees, thus falling into the regime where ties are broken by timestamp. Under the FCFS rule, transactions are ordered directly by the time at which the sequencer received the transaction.

Sequencing rules on Layer-2 networks determine the relative ordering of transactions and therefore which traders capture time-sensitive opportunities. Most L2s use either priority gas auctions (PGAs) in which users bid higher gas fees to obtain earlier queue positions, creating a fee-based form of time priority or a first-come, first-served (FCFS) rule in which the sequencer orders transactions according to observed arrival time via public or private RPC endpoints. Under both FCFS and PGA sequencing, execution outcomes depend on latency races and/or fee bidding. In PGA settings, when two transactions offer the same priority fee, sequencing falls back to arrival timestamp as a

[5]Despite having a centralized sequencer node, L2 networks achieve consensus by either using a dispute period.

tie-breaking rule.

Maximal Extractable Value A significant share of economically meaningful trading activity on Layer-2 networks, also called *rollups* because they 'roll up' many transactions into compressed batches posted to Ethereum, is associated with maximal extractable value (MEV): rents that arise from the ability to influence transaction ordering relative to liquidity, oracle updates, or protocol state. Many common MEV opportunities consist of CEX-DEX arbitrage and DEX-DEX arbitrage. These opportunities are strongly time-sensitive: once a profitable mispricing is exploited, the payoff disappears for all subsequent transactions.

While MEV in the Ethereum mainnet is typically characterized by CEX-DEX arbitrage and sandwich attacks, the MEV landscape for L2s differs substantially: Gogol et al. (2026) finds that sandwich attacks are essentially non-existent on L2. Instead L2s exhibit a distinctive pattern of *optimistic MEV* and *backruns* (Solmaz et al. (2025), Gogol et al. (2025)), in which traders repeatedly submit transactions that would be profitable conditional on favorable sequencing or oracle updates, but that frequently revert when the required preconditions fail to materialize. Optimistic MEV generates large volumes of nearly-identical probing transactions whose result is binary: either capturing the opportunity or reverting at a cost. For the trader, reverts burn gas and sequencing fees, while for the venue they consume blockspace and network capacity, producing congestion externalities for other users. Because MEV payoffs are winner-take-all and highly sensitive to ordering, sequencing rules directly determine who captures MEV and who pays revert costs. Under FCFS or PGA sequencing, traders must race to obtain time priority through arrival latency or fee bidding, and in the absence of fine-grained latency tools, duplicate submissions serve as a costly substitute for speed.

The Timeboost Mechanism Timeboost modifies Arbitrum's default sequencing rule by introducing a sealed-bid, second-price auction for 200 milliseconds of transaction time priority. The mechanism operates at the beginning of each sequencing interval. Traders submit encrypted bids, denominated in Ether, to an off-chain autonomous auctioneer for access to the Express Lane, a short artificial time advantage that allows the winner's transactions to be placed ahead of the first-come, first-served (FCFS) queue. The auction has a parameterizable reserve price, set to 0.001 ETH at the

time of writing. The auction is resolved, with the winner submitted to Arbitrum's on-chain auction contract, where the highest bidder wins the express lane while paying the second-highest bid as an auction clearing price.[6] Auction payments are first transferred from the winner's address to the auction contract, and then to a parameterizable beneficiary address, currently set to the Arbitrum DAO treasury at the time of writing.

Time priority lasts for 60 seconds per auction round. For the auction granting Express Lane access from time t to $t+60$ seconds, bidding starts at $t-60$ seconds and closes at $t-15$ seconds, with the remaining 15 seconds used to propagate auction results on-chain. Any transactions submitted via the Express Lane (which only the winner can submit to) are sequenced according to the timestamp at which the sequencer receives the transaction, which we henceforth call the *arrival timestamp*. All other transactions receive an artificial 200-millisecond delay to their arrival timestamp. Winners can also "resell" Express Lane access by allowing third parties to route transactions through their address for side payments on- or off-chain.

3 Model and Results

We start with a stylized, yet general model of a contest for a common-value time-sensitive trading opportunity of value $V > 0$. Transactions are ordered, e.g. by a matching engine or a sequencer, according to a first-come, first-served (FCFS) rule applied to the arrival timestamps. There are $N \geq 3$ traders indexed by $i \in \{1, \ldots, N\}$. Exactly one trader can capture the opportunity, and all other traders fail to capture the opportunity. Trader i can choose an effort level of $k_i \geq 0$ into latency at a cost of $C(k_i)$, which determines their probability of successfully winning the opportunity relative to the other traders' investments. Specifically, we assume that latency investments result in a Tullock contest over the opportunity, with trader i's probability of winning as proportionate to their investment relative to the total investment:

$$\mathbb{P}(\text{trader } i \text{ wins} \mid k_i, k_{-i}) = \frac{k_i}{k_i + \sum_{-i} k_{-i}}$$

Assumption 1. *We assume the following on the cost function $C : \mathbb{R}_+ \to \mathbb{R}_+$:*

[6] In the case where there is only one bidder, the price paid is zero.

- *C is twice-differentiable, strictly increasing, and convex;*
- *C has convex first derivative;*
- *if C is not linear, then $C'(0)$ is sufficiently small.*

Trader i's expected payoff is then

$$u_i(k_i; k_{-i}) = \frac{k_i}{k_i + \sum_{-i} k_{-i}} \cdot V - C(k_i)$$

We solve for symmetric pure-strategy Nash equilibria in latency technology investment level, which the following proposition characterizes.

Proposition 3.1. *There exists a unique symmetric pure-strategy Nash equilibrium where each trader chooses investment level $k^* > 0$ into latency technology.*

We now introduce an auction for artificial time priority T in the matching engine or sequencer taking place before the arrival of the trading opportunity. The auction format is a sealed-bid second-price auction where the trader with the highest bid pays the second-highest bid to the auctioneer. Traders now participate in a sequential game consisting of the following stages:

1. **Auction:** Trader i submits a bid b_i. The highest bidder wins the time priority for this opportunity and pays the second-highest bid.
2. **Latency Race:** Trader i chooses the level to invest into latency technology based on the results of the auction: $k_w^{(i)}$ if they win, and $k_l^{(i)}$ if they lose.

Instead of the value of a single profitable opportunity, the value V can now represent the total expected value of profitable opportunities occurring during the time window in which this time advantage is to be allocated. We assume that this advantage reduces the baseline contest win probabilities for those who did not win the auction by a factor of $\alpha(k_w, T)$, where $\alpha : \mathbb{R}_+^2 \to [0, 1]$ can be thought of as a disadvantage function on the losers as a result of time priority being awarded to the winner.

Assumption 2. *We assume the following on the disadvantage function $\alpha : \mathbb{R}_+^2 \to [0, 1]$:*

- *α is twice-differentiable, strictly decreasing, and convex in k_w;*

- $\alpha(0, T) = 1$.

Under this setup, a loser j's expected payoff in the continuation (latency race) game:

$$u_l^{(j)}(k_l^{(j)}; k_w, k_l^{(-j)}) = V\left(\frac{k_l^{(j)}}{k_w + k_l^{(j)} + \sum_{-j} k_l^{(-j)}} \cdot \alpha(k_w, T)\right) - C(k_l^{(j)}).$$

Adding up each auction loser's probability of winning the opportunities yields the auction winner's probability of winning and continuation payoff:

$$u_w(k_w; k_l^{(j)}) = V\left(1 - \frac{\sum_j k_l^{(j)}}{k_w + \sum_j k_l^{(j)}} \cdot \alpha(k_w, T)\right) - C(k_w).$$

The following result characterizes the latency race subgame.

Proposition 3.2. *There exists a unique symmetric pure-strategy Nash equilibrium (k_l^*, k_w^*) in the latency race subgame where k_l^* and k_w^* are the levels invested into latency technology by each trader conditional on losing and winning the auction, respectively, and satisfy $0 < k_l^* < k_w^*$. Moreover, letting u_l^* and u_w^* denote loser and winner equilibrium continuation payoffs, we have $0 < u_l^* < u_w^*$.*

Notably, the result $k_l^* < k_w^*$ follows from any α that is decreasing in k_w, revealing the economic mechanism behind this structural handicap. Once the success function of losers is multiplicatively scaled down by $\alpha(k_w, T) < 1$, any incremental effort by the advantaged trader further depresses the marginal return to investment for the rest of the field. In equilibrium, this induces a first-order downward shift in losers' best-response functions. With convex costs, the result is unambiguous: conditional on losing the auction, traders optimally scale back their latency expenditures relative to the symmetric baseline. Intuitively, the auction endogenously segments the field into an advantaged leader and disadvantaged followers, attenuating the competitive externality that sustains the symmetric arms race.

Both types retain positive continuation value in equilibrium. Even the loser obtains a strictly positive expected payoff, since its probability of winning the contest remains positive absent complete dominance by the advantaged trader. The winner, however, enjoys a strictly larger continuation surplus: at any fixed level of effort, the time advantage increases its contest success probability, and the resulting gain is not fully dissipated by equilibrium effort adjustments. This surplus gap,

$u_w^* - u_l^*$, determines willingness to pay for the time advantage in the auction stage. Hence, the mechanism can extract rents from the latency race while preserving participation incentives for individual traders.

Proposition 3.3. *Let $b^* = u_w^* - u_l^*$. Then (b^*, k_l^*, k_w^*) comprise the unique symmetric subgame-perfect Nash equilibrium in the full auction-submission sequential game.*

Proposition 3.3 ties everything together by showing that the unique equilibrium of the latency race subgame endogenously pins down a unique equilibrium of the preceding auction. In the sealed-bid second-price auction, bidding this value is the (unique symmetric) best response: any lower common bid can be profitably outbid to secure strictly positive surplus, while any higher common bid yields negative expected surplus relative to losing. As a result, $b^* = u_w^* - u_l^*$ is the unique symmetric equilibrium bid, and combining it with the submission strategies (k_l^*, k_w^*) yields a unique symmetric subgame-perfect Nash equilibrium of the full sequential game.

Our key result compares the total amount of investment level expended by traders in equilibrium without and with the time-priority auction, denoted as K^* and K_A^*, respectively. In the language of this section, we can express these as

$$K^* = Nk^*,$$
$$K_A^* = k_w^* + (N-1)k_l^*.$$

The following result shows that less total effort is spent to achieve latency under a time-priority auction.

Theorem 3.4. *Traders invest less into latency technology under a time-priority auction than without such an auction, i.e. $K_A^* < K^*$.*

Without an auction, priority is determined entirely by latency, so all traders must invest in speed to remain competitive. This produces a symmetric arms race: each trader's marginal gain from being slightly faster is offset by others doing the same, and in equilibrium the surplus associated with time priority are largely dissipated through latency investment. By contrast, when time priority is allocated via a single-winner auction, only the winner has strong incentives to invest further in latency, while the remaining traders face a systematic disadvantage and therefore scale back.

The reduction in additional latency investment by non-winners more than offsets the additional investment by the winner, lowering aggregate speed expenditure. In effect, the auction compresses a continuous latency race into a one-shot competition for priority, eliminating secondary races among disadvantaged traders and converting dissipative investment into auction payments. We explain the role of the convexity of α behind this salient result.

Whether aggregate investment declines depends on how strongly the winner's incentives expand in response to its advantage. When α is convex, the incremental benefit of reinforcing the time advantage exhibits diminishing sensitivity to the winner's own effort. Although the winner internalizes a higher marginal return than losers, convexity ensures that this is not self-reinforcing, as the marginal gain from additional latency declines sufficiently fast. The consequent reduction in investment by the disadvantaged traders dominates the expansion by the single winner, yielding lower total expenditure on speed under the auction, $K_A^* < K^*$. Economically, convexity prevents the time advantage from propagating into a secondary escalation dynamic where losers may try to collectively invest more to overcome their disadvantage, instead compressing a latency arms race into a discrete allocation of priority.

If convexity fails, i.e. when α is locally concave over some range of k_w values, the amplification mechanism can reverse. A concave disadvantage function implies that the winner's marginal return to latency may remain elevated, or even increase, as it invests more. In this case, the time advantage becomes self-reinforcing, and incremental speed investments disproportionately strengthen the winner's position. The resulting escalation by the advantaged trader can outweigh the retrenchment by losers, potentially driving aggregate latency investment above the symmetric benchmark. Thus, the curvature of the disadvantage function governs whether the auction internalizes and dampens the arms race or instead reallocates and intensifies it.[7]

4 The Blockchain Setting and Platform Revenue

By modeling transaction submissions and random latency delays, we microfound the blockchain setting as a reduced-form instance of the general model with linear effort costs, linking observable on-chain behavior to latency investment. Here, the competition is over an MEV opportunity of

[7] A counterexample is given by $C(k) = k^2/2$ and $\alpha(k_w, T) = \exp\big(-(k_w \lambda T)^2\big)$, $N = 2$, $V = 1$, and $\lambda T = 1$. Note that this choice of α is not globally convex in k_w.

common value $v > 0$. A sequencer orders transactions according to a first-come, first-served (FCFS) rule applied to their arrival timestamps. Exactly one trader can exploit the MEV opportunity by being sequenced first with the proper transaction; all other copies revert, failing to capture the opportunity. Traders can submit multiple copies of the same MEV-exploiting transaction, incurring a linear per-copy cost that represents the gas cost of failed attempts.[8] Mazorra et al. (2026) show how spamming can arises as an equilibrium strategy on blockchains, thus providing the microfoundations for our reduced-form model.

Each submitted transaction experiences a random delay before it is observed by the sequencer for inclusion. Formally, the timestamp is shifted by an i.i.d. delay $\varepsilon \sim F$ (with density f), capturing network latency. Suppose that trader i submits k_i copies of the transaction. Then the effective arrival timestamp relevant under FCFS is the minimum of k_i i.i.d. draws from F:

$$X_i(k_i) \;:=\; \min\{\varepsilon_{i1}, \ldots, \varepsilon_{ik_i}\}.$$

Under FCFS, trader i captures the opportunity iff $X_i(k_i) \leq \min_{j \neq i} X_j(k_j)$. Intuitively, submitting more copies reduces the minimum delay and increases the chance of being first. We assume that a successful transaction costs $g \in (0, v)$ in gas fees, and reverted transactions must pay a fraction $r \in (0, 1)$ of the gas fees for a successful transactions.

For tractability, we assume that latency is exponentially distributed: a single copy has a delay of $\varepsilon \sim \text{Exp}(\lambda)$ with rate $\lambda > 0$. Then we have that

$$X_i(k_i) \sim \text{Exp}(k_i\lambda).$$

In practice, k_i should be a nonnegative integer, but for tractability, we relax this and allow traders to choose fractional amounts for the number of copies to send, expanding their action space to $k_i \in \mathbb{R}_+$. While this breaks the interpretation that $X_i(k_i)$ is the minimum of several i.i.d. draws from F, we still have that $X_i(k_i) \sim \text{Exp}(k_i\lambda)$. Trader i's expected payoff is then

$$u_i(k_i; k_{-i}) = (v - (1-r)g) \cdot \Pr\left(X_i(k_i) \leq \min_{j \neq i} X_j(k_j)\right) - rgk_i$$

[8] We assume that traders submit transactions that are identical in recipient address and calldata, but have increasing nonces, meaning that failed attempts appear on-chain, such as the case for backruns and optimistic MEV.

$$:= V \cdot \Pr\left(X_i(k_i) \leq \min_{j \neq i} X_j(k_j) \right) - Ck_i$$

where we reparameterize with $V = v - (1 - r)g$ and $C = rg$.[9] Under independent exponentially-distributed latency delays, the win probability simplifies to

$$\Pr\left(X_i(k_i) \leq \min_{j \neq i} X_j(k_j) \right) = \frac{k_i}{k_i + \sum_{j \neq i} k_j},$$

resembling the Tullock contest-like win probability in the general model. Under linear costs, the equilibrium number of copies has a simple closed form:

$$k^* = \frac{N-1}{N^2} \cdot \frac{V}{C}.$$

Using this expression, we can compute important equilibrium quantities. Let K^* be the total number of transactions submitted in equilibrium. This represents the level of on-chain spam, since $K^* - 1$ of these transactions will revert on-chain, and has formula

$$K^* = Nk^* = \frac{N-1}{N} \cdot \frac{V}{C}.$$

We can also compute an individual trader's equilibrium utility u_i^*, and how the total value V is split between users, captured by the term W^*, and revenue R^* accrued by the sequencer/DAO from reverted gas fees:[10]

$$\begin{aligned} u_i^* &= \frac{V}{N} - Ck^* = \frac{V}{N^2}, \\ W^* &= Nu_i^* = \frac{V}{N}, \\ R^* &= V - W^* = \frac{N-1}{N} \cdot V. \end{aligned}$$

These corollaries highlight the core inefficiency motivating Timeboost-style mechanisms: competition to be first under FCFS induces redundant transaction copies. In later sections we show how

[9] Note that the cost accounting is done by adding the revert costs for all transactions plus an additional gas fee if the transaction is successful.

[10] To be precise in the case of Arbitrum, gas fees not spend on posting block data to the Ethereum mainnet go to the Arbitrum DAO as well.

auctioning a time advantage changes the competitive margin, reducing K^* and reallocating surplus away from revert-driven waste and toward auction payments.

We now consider the time-priority auction for a time advantage of T in the blockchain setting. As in the baseline model, each submitted copy experiences an exponential idiosyncratic latency delay before it is observed by the sequencer. Let $X_w(k_w) \sim \text{Exp}(k_w\lambda)$ denote the winner's earliest timestamp and $X_l^{(j)} \sim \text{Exp}(k_l\lambda)$ denote loser j's earliest timestamp, with being independent across arbitrageurs. Now, the winner's effective arrival time is $X_w - T$, while each loser's effective arrival time is still $X_l^{(j)}$. Therefore the winner succeeds if and only if

$$X_w - T \leq \min_j X_l^{(j)}.$$

Similarly, a particular loser j succeeds if and only if

$$X_l^{(j)} < \min_{-j} X_l^{(-j)} \text{ and } \; X_l^{(j)} \leq X_w - T.$$

Using standard probability calculations and exponential distribution properties, the continuation payoffs for the winner and loser are, respectively:

$$u_w(k_w; k_l^{(j)}) = V\left(1 - \frac{\sum_j k_l^{(j)}}{k_w + \sum_j k_l^{(j)}} \cdot \exp(-k_w\lambda T)\right) - Ck_w$$
$$u_l^{(j)}(k_l^{(j)}; k_w, k_l^{(-j)}) = V\left(\frac{k_l^{(j)}}{k_w + k_l^{(j)} + \sum_{-j} k_l^{(-j)}} \cdot \exp(-k_w\lambda T)\right) - Ck_l^{(j)}.$$

In the language of the general model, the blockchain setting corresponds to the disadvantage function $\alpha(k_w, T) = \exp(-k_w\lambda T)$.

Evaluating the same important quantities under Timeboost, we first have that the total number of transactions submitted in Timeboost equilibrium, denoted K_A^*, is

$$K_A^* = k_w^* + (N-1)k_l^*.$$

Since the winner pays $u_w^* - u_l^*$ in Timeboost equilibrium, all arbitrageurs receive an equilibrium payoff of u_l^*. The division of the surplus value V between arbitrageurs and sequencer, denoted by

components W_A^* and R_A^*, respectively, is

$$W_{K^*}^* = Nu_l^*,$$
$$R_{K^*}^* = V - W^* = V - Nu_l^*.$$

Proposition 4.1. *Arbitrageurs send fewer total transactions and the sequencer earns more revenue in equilibrium under Timeboost than in the baseline setting, i.e.* $K_{tb}^* < K^*$ *and* $R_{tb}^* > R^*$.

In the blockchain setting with linear costs, latency investment manifests as duplicate transaction submissions. The efficiency gain from the time-priority auction corresponds to a reduction in redundant submissions, freeing blockspace that would otherwise be consumed by failed race attempts. Interestingly, when considering sequencer/DAO revenue, one would expect there to be a tradeoff between auction revenue and less gas fees from reverted duplicate transactions, but Proposition 4.1 shows that equilibrium revenue unambiguously increases under a time-priority auction. Hence, this result indicates that time-priority auctions in blockchains, such as Timeboost, convert what would have been dissipated in revert-related costs for incremental latency improvement into direct payments for a substantial "latency improvement" via artificial time advantages.

5 Extensions and Discussion

5.1 Time-Priority Auctions with Multiple Winners

The analysis thus far has focused on a single-winner auction that allocates a single time advantage. We now generalize the mechanism to allow multiple winners, each receiving the same time advantage. This extension allows us to study how the allocation of priority along the extensive margin—the number of traders granted priority—interacts with the intensive margin—the size of the time advantage, in shaping equilibrium user effort and platform revenue. Additionally, holding the number of winners fixed, we examine how key equilibrium outcomes vary with the size of the time advantage. By varying the number of priority slots and the magnitude of each priority slot, we can characterize how competition, effort, and revenue respond to broader access to priority and more concentrated priority.

Let $n_w, n_l \in \mathbb{N}$ denote the number of winners and losers, respectively, with $n_w + n_l = N$. We

now assume that traders participate in a (sealed-bid) *generalized second-price auction* when there are multiple winners. In such an auction, the n_w highest bids receive the time advantage, and the i-th highest bidder pays the value of the $(i+1)$-th highest bid. While this auction format is not dominant-strategy incentive compatible, it is the natural extension of the sealed-bid second-price auction used in the single-winner model and admits weak Nash equilibria where each winner pays the continuation value of having priority, similarly to the single-winner case.

Letting $K_w = k_w^{(i)} + \sum_{-i} k_w^{(-i)}$, $K_l = k_l^{(j)} + \sum_{-j} k_l^{(-j)}$, and $K = K_w + K_l$. Under exponential random delays, continuation payoffs are given by

$$u_w^{(i)}(k_w^{(i)}; k_w^{(-i)}, k_l^{(j)}) = V \cdot \frac{k_w^{(i)}}{K_w}\left(1 - \frac{K_l}{K} \cdot \exp(-\lambda T K_w)\right) - C k_w^{(i)},$$
$$u_l^{(j)}(k_l^{(j)}; k_w^{(i)}, k_l^{(-j)}) = V\left(\frac{k_l^{(j)}}{K} \cdot \exp(-\lambda T K_w)\right) - C k_l^{(j)}.$$

Similarly to the previous section, we search for interior type-symmetric equilibria $k_l^*, k_w^* > 0$ where all auction losers choose k_l^* and all auction winners choose k_w^* in the resulting continuation game. We show that if such an interior type-symmetric equilibrium exists, then $0 < k_l^* < k_w^*$ and $0 < u_l^* < u_w^*$, extending Proposition 3.2 to the multi-winner case. This places the "value" of the time-priority slot at $u_w^* - u_l^*$, and a standard auction theory argument shows that the symmetric bid profile $b^* = u_w^* - u_l^*$ constitutes a Nash equilibrium in the generalized second-price auction. The combined profile (b^*, k_l^*, k_w^*) is thus *type-symmetric* a subgame-perfect Nash equilibrium.

Our key results in the section concern the analysis of the multi-winner setting along the intensive and extensive margins. Proposition 5.1 extends the effort and revenue results from the single-winner environment to the multi-winner setting. The same economic forces apply here: time-priority becomes a more reliable determinant of winning the trading opportunity, lowering the marginal value of additional submission intensity. We expand on this result by examining total effort and platform revenue as the size of the time advantage itself varies.

The comparative statics on effort and revenue are globally monotonic in advantage size in the single-winner case, as more time priority further increases the value of the priority slot, redirecting more financial resources into the auction rather than effort. With multiple winners, we show that there exists an interval $[0, \bar{T})$ such that if $T \in [0, \bar{T})$, a unique type-symmetric equilibrium exists.

Along this interval, equilibrium total effort decreases and platform revenue increases in T, indicating that some advantage is always better than none in terms of mitigating latency races and capturing revenue.

Proposition 5.1. *Suppose that user costs are linear and contribute to platform revenue. For any number of winners, there exists a range of time advantage sizes for which a unique interior type-symmetric equilibrium (k_l^*, k_w^*) exists. In this range, increasing the priority advantage reduces total effort and raises platform revenue.* [11]

Proposition 5.2 characterizes the extensive margin of time-priority allocation. Increasing the number of winners spreads priority across more traders, presenting an interesting tradeoff of incentives. On one hand, competition between the loser and winner types is further dampened as auction losers are facing multiple users with time advantages, rather than just one. Yet, on the other hand, allowing for multiple winners now introduces competition between winners themselves. Although priority becomes less exclusive, the aggregate demand for priority remains high because more traders can benefit from obtaining it. In equilibrium, over a range $n_w \in [1, \bar{n}_w)$ where the symmetric equilibrium is well behaved, the former force dominates the latter, resulting in a decline in total effort and increase platform revenue.

Proposition 5.2. *Suppose that user costs are linear and contribute to platform revenue. When the total number of users in the environment is large, for any time advantage size, there exists a range of winner counts for which a unique type-symmetric equilibrium exists. In this range, increasing the number of winners reduces total effort and raises platform revenue.*

Together, the propositions highlight distinct roles for the intensive and extensive margins. Increasing the size of the time advantage primarily weakens incentives to race conditional on priority allocation, while increasing the number of winners redistributes competitive pressure across traders. Both channels reduce wasteful competition and convert rent dissipation into platform revenue, but they operate through different economic mechanisms: the intensive margin changes the *saliency* of priority, whereas the extensive margin changes its *scope*.

[11] In the single-winner case, a unique type-symmetric equilibrium exists for any $T \geq 0$. Equilibrium total effort decreases and platform revenue increases monotonically in T.

5.2 Implications for Centralized Exchanges

An important distinction between the blockchain setting and centralized exchanges is the frequency with which priority shoule be reassigned. Current blockchain-based implementations of time-priority auctions provide the time priority for 60-seconds, but in traditional markets the relevant margins of competition (co-location, communication links, and routing technology) are chosen over much longer horizons and cannot be reoptimized from minute to minute. For that reason, the natural analogue in centralized markets would be to auction priority over longer intervals, rather than through short-duration auctions.

These longer intervals can capture natural improvements in technologies or the renewal of co-location agreements, so agents' investment choices (e.g. k, k_w, or k_l) can thus represent decisions on whether to update or not. However, the underlying economics are unchanged: explicit pricing of time priority can shift competition away from socially costly speed investment and toward auction payments, allowing exchanges to capture rents that would otherwise be dissipated in the latency arms race. Our results therefore provide evidence for a general design principle, namely that priority should be sold on a horizon aligned with the durability of the underlying technology investment.

For centralized exchanges, another important feature is that they can absorb part of the latency race through co-location, connectivity, and related access fees. Proposition 4.1 shows that auction payments coupled with residual investment payments can exceed investment payments in the absence of time-priority auctions. The relevant implication for these venues here is that these auctions can both increase increase and make more transparent the venue's capture of latency rents by shifting revenue sources away from various private technology expenditures and toward direct payments for execution priority.

6 Empirical Analysis

This section evaluates the model's predictions using the introduction of Timeboost as an empirical implementation of a time-priority auction. As a market-design intervention that converts latency competition into explicit price competition for execution priority, the mechanism therefore provides a setting in which the allocation of time priority changes while trading opportunities and participant incentives remain otherwise comparable, allowing us to study how equilibrium behavior responds

when priority is priced rather than obtained through racing.

Blockchain markets make this exercise feasible because latency competition is observable. In traditional HFT environments, investments are privately-held information and unobservable to researchers. By contrast, on blockchain platforms, attempts to win latency races manifest as bursts of identical transactions submitted within short intervals. We interpret these duplicate submissions as the on-chain analogue of speed investment and use them as a proxy for effort. Our model predicts that a time-priority auction should therefore reduce redundant submissions and redirect competition toward payments for priority.

Timeboost implements precisely such a change. Prior to adoption, transaction ordering on Arbitrum followed a first-come, first-served rule, under which traders competed by submitting multiple copies of the same transaction to increase the probability of early inclusion. Timeboost introduces a sealed-bid auction that grants a short execution advantage to winning bidders, replacing pure arrival-time competition with an explicit price-based allocation of priority. In the context of the model, this corresponds to shifting from a race for execution probability to an auction for execution rights.

We study this transition using transaction-level data across multiple Layer-2 blockchains as control groups and an event-study design around Timeboost's deployment. Despite other Layer-2 blockchains ordering transactions via priority gas fee, ties among priority fees are still broken by timestamp, and many MEV opportunities on Layer-2 networks rely on being the first to exploit it in the zero priority fee regime. This results in a latency race in those settings as well for these MEV opportunities. By comparing outcomes on Arbitrum to other Layer-2 networks that did not introduce time-priority auctions, we isolate how explicitly pricing execution priority affects equilibrium behavior. There are three testable implications of our model:

1. Timeboost decreases transaction spamming. (Theorem 3.4)

2. Timeboost increases sequencer/DAO revenue. (Theorem 4.1)

To test these implications, we conduct a panel data analysis using data across Layer 2 blockchains and a within-Arbitrum analysis.

6.1 Data Collection

We assemble a panel of on-chain activity for Arbitrum and a set of major EVM L2 comparison networks using Dune Analytics. We collect data from the Arbitrum, Optimism, Polygon, Base, and Avalanche (C-Chain) blockchains between February 17 and June 17 of 2025, a two-month window before and after the deployment of Timeboost, which happened on April 17, 2025.[12] Our objective is to measure spamming and wasteful competition among bots and searchers. This phenomenon is relevant both in FCFS sequencer settings, where arrival timing drives inclusion, and PGA environments, where actors repeatedly submit or adjust transactions to win priority.

To measure spamming in a comparable way across chains, we look for bursts of near-identical transactions within short second-level windows. Setting a window length of 2 seconds, we track the transaction's sender, recipient, value, function selector, and calldata. If transactions with these same characteristics appear multiple times in the same burst window, we label all transactions beyond the first as repeated (redundant) attempts, interpreting them as a conservative proxy for automated resubmission and competitive duplication. This is motivated by common bot execution logic: when inclusion and ordering are uncertain, whether due to FCFS latency dispersion, PGA-style priority competition, or optimistic MEV transactions Solmaz et al. (2025).

A key challenge is that the objects of interest, i.e. attempts to win ordering priority, are not easily identifiable in blockchain data. Unlike traditional exchanges, on-chain transactions do not have a well-labeled field identifying whether an action corresponds to arbitrage, liquidation, and so on. In particular, economically identical actions can be implemented through arbitrary contracts. Users frequently deploy custom contracts, proxy contracts, and externally-owned accounts (EOAs) that encode trading logic internally. Thus, many MEV-related transactions can be executed with minimal or even empty calldata (i.e. a call that triggers internal state-dependent execution). Since attempting to manually classify MEV-related transactions using contract allowlists or application binary interface (ABI) decoding is difficult and could exclude economically relevant activity, biasing measurement toward more well-established protocols, we adopt a protocol-agnostic approach: rather

[12]While the Avalanche C-Chain is a Layer-1 blockchain, it is compatible with the Ethereum Virtual Machine, with many applications deployed on Ethereum and its L2s also being deployed on Avalanche (e.g. Uniswap), and used a PGA sequencing rule until June 17, 2025. On June 17, 2025, a protocol called *MEV Zone* was deployed on the Avalanche C-Chain that provides proposer-builder separation. In the appendix, we display results when the Avalanche C-Chain is excluded from the control groups: our key empirical results are robust to this specification.

than focusing on a narrow set of known MEV contracts or function calls, we identify behavior consistent with a latency race, i.e. multiple submissions of the same economic action within a short interval.

This broader classification introduces potential false positives, such as automated resubmissions unrelated to strategic activity. Importantly, however, this noise is not specific to the treated chain: since the same data collection procedure is applied uniformly across all Layer-2 blockchains in the panel, any automation or wallet retry logic enters both treatment and control groups. In the difference-in-differences design, this noise will be absorbed by chain and time fixed effects. The key empirical identification strategy relies on differential changes in outcome variables of interest *after* the introduction of Timeboost, rather than on the level of repeated activity. Remaining misclassification would mainly add noise, in fact attenuating estimates, rather than creating a false treatment effect.

Recent evidence supports interpreting repeated failed transactions on rollups as strategic competition. Solmaz et al. (2025) find that in L2s, searchers submit transactions that repeatedly probe liquidity pools for cyclic arbitrage opportunities, executing the trade if profitable and reverting the transaction otherwise; these types of transactions are responsible for up to 50% of all gas fees on some networks (e.g. Base and Optimism). Gogol et al. (2025) find similar empirical evidence of spamming in L2s as an arbitrage strategy, and further document that searchers often prefer submitting duplicate transactions over bidding priority fees, even on networks using PGAs, indicating competition over arrival time rather than price. Their findings support that bursts of near-identical transactions within short time intervals—the behavioral signature that our data collection process leverages—is indicative of MEV-related competition on L2s.

Another important nuance in the MEV landscape for L2s is how certain types of MEV differ depending on whether a trader can replace a pending transaction or must submit independent attempts. In "top-of-block" MEV (e.g. arbitrage or frontruns), traders repeatedly update a single transaction using the same nonce while adjusting timing or fees; each submission replaces the previous one, so only one execution can occur and losing attempts leave little on-chain footprint. In "rest-of-block" MEV (e.g. backruns, optimistic arbitrage), outcomes depend on realized block state, so replacement is infeasible and traders submit multiple transactions with increasing nonces. Thus, unsuccessful attempts revert and remain recorded, resulting in costly duplication. Because

this distinction holds across both PGA and FCFS sequencing rules, our data measure a consistent object across networks: observable contest effort rather than protocol-specific behavior. These differences are summarized in Table 2.

	Top-of-block MEV (e.g. arbitrage, frontruns)	**Rest-of-block MEV** (e.g. backruns, optimistic MEV)
PGA L2s	Contested via fee bidding wars; failed attempts are costless and leave no on-chain record	Contested via multi-nonce spamming with zero priority fee; failed attempts are costly and leave an on-chain record
FCFS L2s	Contested via same-nonce spamming; failed attempts are costless and leave no on-chain record	Contested via multi-nonce spamming; failed attempts are costly and leave an on-chain record

Table 2: MEV Opportunity Types Across Execution Environments

Our empirical strategy focuses on this observable component of latency competition. While this does not account for all MEV activity, it measures the subset for which competition requires multiple independent submissions and therefore appears on chain. Crucially, this subset exists both on Arbitrum and on comparison Layer-2 networks, where ordering among these transactions is likewise determined by arrival time. The identification therefore compares how traders adjust observable effort when priority is priced instead of obtained through racing.

We first record the number of repeated transactions and the gas consumed by failed repeated transactions, denoted by $\mathsf{RepTXs}_{c,t}$ and $\mathsf{RepGas}_{c,t}$, where c indexes a blockchain and t indexes the time period. On Arbitrum after the implementation of Timeboost, the auction proceeds also contribute to the revenue. We thus collect data for each Timeboost auction round within our sample window, focusing on the payment, denoted by $\mathsf{AuctionPayment}_t$, of the winner (the second-highest bid). The $\mathsf{Revenue}_{c,t}$ variable then integrates revenue from gas fees and auctions across Layer 2s:

$$\mathsf{Revenue}_{c,t} = \mathsf{RepGas}_{c,t} + \mathsf{AuctionPayment}_t \cdot 1_{\{c=\mathsf{Arbitrum}\}}.$$

We collect a set of controls, including base gas fees and DEX volume on each blockchain, as well as returns and volatility of Ether (relative to US dollars). While Arbitrum, Optimism, and Base use ETH for gas payments, Polygon and Avalanche use native tokens; to account for this, we convert

values into ETH. All variables are aggregated at daily frequencies.

	RepTXs				Revenue			
Chain	**Mean**	SD	Min	Max	**Mean**	SD	Min	Max
Arbitrum	68,416	92,305	2,850	643,983	57.14	78.64	0.322	592.3
Avalanche	2,774	4,625	304	39,127	0.026	0.046	0.001	0.262
Base	216,987	142,519	62,578	857,769	0.503	0.879	0.108	8.215
Optimism	43,092	56,968	2,701	516,531	0.012	0.017	0.002	0.156
Polygon	11,290	7,686	2,577	36,718	0.001	0.002	0.000	0.009

Table 3: Summary statistics by chain for total repeated transactions and sequencer revenue.

Summary statistics in Table 3 show substantial heterogeneity across chains in all three metrics of interest. Base and Arbitrum exhibit the highest average repeated transaction counts, indicating more activity and frequent competition over MEV, while Polygon and Avalanche are far lower in levels. Failure rates also differ markedly: Arbitrum and Polygon have relatively high mean failure rates while Optimism stands out with a much lower mean failure share. Revenue varies by orders of magnitude as well, as Arbitrum's mean is far larger than the other chains, likely due to higher base gas fees and the magnitude of transactions processed.

6.2 Cross-Network Analysis

We estimate the causal effect of Timeboost on transaction spamming and sequencer revenue using a two-way fixed-effects (TWFE) design that compares changes on Arbitrum (the treated chain) before and after the Timeboost launch to contemporaneous changes on a set of comparison L2 networks with no major changes to their sequencer ordering rules. The identifying variation is thus cross-sectional (treated vs. control chains) interacted with time-series variation (post- vs. pre-implementation). We supplement the baseline specification with controls for fee conditions and market intensity, respectively: an L2 base fee measure in Ether (BaseGasL2) and within-chain DEX volume in log USD as a DeFi activity proxy (VolumeDEX).

All outcome variables and time-varying covariates (except for the binary treatment indicator) are standardized (z-scored) prior to estimation. This choice is practical and methodological: chains differ by orders of magnitude in baseline transaction counts, gas usage, and revenue, and raw-unit levels can cause the regression to be dominated by the largest networks, conflating scale with

treatment response. Standardization produces estimates in standard-deviation units, improving comparability across chains and outcomes and making effect sizes interpretable on a common scale. Importantly, standardization does not change the sign or statistical significance patterns relative to monotone transformations.

Empirical Specification and Identification. Our specification is given by the following model:

$$Y_{c,t} = \beta \cdot \mathsf{Post}{\times}\mathsf{Treated}_{c,t} + \gamma^\top X_{c,t} + \alpha_c + \delta_t + \varepsilon_{c,t}$$

where the dependent variable of interest (repeated transactions or revenue) is

$$Y \in \{\log \mathsf{RepTXs}, \log \mathsf{Revenue}\},$$

$\mathsf{Post}{\times}\mathsf{Treated}_{c,t} = 1_{\{c=\mathsf{Arbitrum},\ t\geq \text{2025-04-17}\}}$ denotes the treatment effect (i.e. the differential change on Arbitrum after Timeboost relative to the change on control chains over the same period),

$$X \subseteq \{\mathsf{BaseGasL2}, \mathsf{VolumeDEX}\}$$

is a set of controls, and $\varepsilon_{c,t}$ is idiosyncratic noise. We include chain fixed effects α_c that absorb time-invariant differences across networks (e.g., average activity levels, application mix, sequencer infrastructure) and day fixed effects δ_t that absorb common shocks shared across chains (e.g., market-wide volatility, macro crypto news, broad changes in user demand).

A key identification assumption is parallel trends conditional on fixed effects and controls: absent Timeboost, would Arbitrum's outcomes have behaved similarly to the control L2s, after conditioning on fixed effects and control variables? This assumption is plausible because all networks share the same EVM execution environment and no other contemporaneous changes to ordering rules occurred on the control chains during the sample period. The outcomes we study arise from latency-sensitive MEV opportunities exist on all rollups, with ordering ultimately depending on arrival time even under priority-fee mechanisms. Hence, traders face the same competitive environment across chains. In the Appendix, we empirically validate this assumption by estimating the model under a placebo treatment date and performing lead tests.

	(1) log RepTXs	(2) log RepTXs	(3) log Revenue	(4) log Revenue
Post×Treated	-0.699**	-0.698***	0.726***	0.711***
	(0.281)	(0.254)	(0.123)	(0.112)
BaseGasL2		0.096		0.150**
		(0.076)		(0.071)
VolumeDEX		0.159		0.078
		(0.125)		(0.150)
Observations	605	605	605	605
N. of groups	5	5	5	5
R^2	0.025	0.041	0.001	0.032

Includes chain and day fixed effects. Standard errors are clustered by chain.
Note: *p<0.1; **p<0.05; ***p<0.01

Table 4: This table reports two-way fixed-effects estimates of the impact of Timeboost adoption on Arbitrum relative to control Layer-2 blockchains. The sample period spans February 17 to June 17, 2025, with Timeboost implemented on April 17, 2025. The dependent variables are the natural logarithm of repeated transactions (columns 1-2) and the natural logarithm of platform revenue in ETH (columns 3-4). Repeated transactions are defined as bursts of identical transactions (same sender, recipient, value, function selector, and calldata) within 2-second windows. Platform revenue includes gas fees from reverted transactions and, for Arbitrum post-Timeboost, auction payments from express lane winners. Post×Treated is an indicator equal to one for Arbitrum observations after April 17, 2025. BaseGasL2 is the Layer-2 base gas fee in ETH. VolumeDEX is the natural logarithm of DEX trading volume in USD.

Discussion of Results. Table 4 reports two-way fixed-effects estimates of the effect of Timeboost adoption on repeated transaction activity and sequencer revenue. The coefficient on Post×Treated is negative and statistically significant (about -0.70 standard deviations), indicating that after Timeboost implementation, Arbitrum experienced a sizable decline in repeated transactions relative to the comparison L2s. This estimate is nearly unchanged when adding controls, suggesting that this reduction is not merely driven by coincident changes in transaction costs or market-wide trading activity.

The PostTreated coefficient for revenue is positive and significant, implying that after Timeboost, revenue from reverted transactions and gas fees increased overall. In the appendix, we provide additional robustness checks: we vary the window of measuring repeated transactions from 2 seconds to 5 seconds, and change the duration of the window before and after the Timeboost adoption date. These results carry through, with the decline in repeated transactions and gain in revenue after Timeboost adoption being significant across these changes. Additionally, as clustered standard errors may introduce bias when the number of groups is small, we also quantify the uncertainty of our estimated coefficients using heteroskedasticity-robust standard errors ("HC3"), find that the empirical results hold under this as well.

6.3 Within-Arbitrum Analysis

Our cross-chain empirical design controlled variation across chains and time by introducing fixed effects. Now, we complement the analysis with a within-Arbitrum analysis by directly incorporating variables that would potentially influence transaction volumes. This exercise does not replace the TWFE identification strategy, as it lacks an external control group, but provides a useful additional robustness check: if Timeboost meaningfully altered submission behavior, we should observe significant changes on Arbitrum itself, particularly after conditioning on proxies for activity (DEX volume) and transaction costs (L1 and L2 base fees).

We estimate a model of the form

$$Y_t = \alpha + \beta \cdot \mathsf{Post}_t + \gamma^\top X_t + \varepsilon_t$$

where the dependent variable of interest is

$$Y \in \{\log \mathsf{RepTXs}, \log \mathsf{Revenue}\},$$

$\mathsf{Post}_t = 1_{\{t \geq \text{2025-04-17}\}}$ is a post-implementation indicator, and

$$X \subseteq \{\mathsf{VolumeDEX}, \mathsf{BaseGasMain}, \mathsf{BaseGasArb}, \mathsf{VolatilityETH}\}$$

is a vector of time-varying controls. In all specifications we include a measure of DEX volume and base gas fees on Ethereum; we additionally include Arbitrum base gas fee and ETH volatility to capture additional shifts in arbitrage intensity. As a robustness check (reported in the appendix), we also verify that conclusions are not sensitive to alternative bandwidths around the implementation date and alternative time intervals for short-window repetition. Here, we do *not* standardize variables before estimation.

Results. Table 5 reports the within-Arbitrum estimates. The post-Timeboost indicator is negative and statistically significant for RepTXs, implying a sizable decrease in repeated submissions after the implementation date even after controlling for DEX volume and fee conditions; the estimated coefficients remains similar when adding mainnet base fees and market controls. Similarly

to the cross-chain analysis, there is a positive significant relation between the post-Timeboost period and revenue, and the post-Timeboost period is associated with higher failure rates among non-timeboosted transactions.

	(1)	(2)	(3)	(4)
	log RepTXs	log RepTXs	log Revenue	log Revenue
Intercept	-9.423	-8.955	2.621	-4.377
	(5.991)	(6.110)	(7.327)	(7.449)
Post	-0.578**	-0.453*	0.858**	0.594*
	(0.238)	(0.240)	(0.354)	(0.360)
VolumeDEX	1.004***	0.882***	-0.026	0.427
	(0.299)	(0.312)	(0.364)	(0.382)
BaseGasMain	-0.050	-0.146	0.631***	0.690***
	(0.158)	(0.155)	(0.189)	(0.198)
BaseGasArb		187.350***		-94.886*
		(54.911)		(50.138)
VolatilityETH		0.262		-3.216***
		(0.768)		(0.846)
Observations	121	121	121	121
R^2	0.173	0.257	0.068	0.141
Adjusted R^2	0.152	0.218	0.044	0.096

Standard errors are computed using the Newey-West estimator.
Note: *p<0.1; **p<0.05; ***p<0.01

Table 5: This table reports estimates from a within-Arbitrum analysis of the impact of Timeboost adoption on redundant transactions and platform revenue. The sample consists of daily observations on Arbitrum from February 17 to June 17, 2025, with Timeboost implemented on April 17, 2025. The dependent variables are the natural logarithm of repeated transactions (columns 1-2) and the natural logarithm of platform revenue in ETH (columns 3-4). Repeated transactions are defined as bursts of identical transactions (same sender, recipient, value, function selector, and calldata) within 2-second windows. Platform revenue includes gas fees from reverted transactions and, post-Timeboost, auction payments from express lane winners. Post is an indicator equal to one for observations after April 17, 2025. VolumeDEX is the natural logarithm of DEX trading volume in USD on Arbitrum. BaseGasMain is the Ethereum mainnet base gas fee in ETH. BaseGasArb is the Arbitrum base gas fee in ETH. VolatilityETH is the daily return volatility of ETH/USD.

DEX volume is positively related to repeated transactions, consistent with higher trading activity increasing the number of contested opportunities and associated transaction flow. The coefficient on BaseGasArb is large and positive in the repeated transactions regressions, which is expected as Arbitrum gas fees are determined by congestion and would be higher under more spammed transactions. The coefficients on BaseGasMain and BaseGasArb for the revenue regressions tell a

more interesting story: revenue from reverted gas fees and auction proceeds appears to be driven by base gas on Ethereum, with Arbitrum gas fees and volatility having a second-order effect. Overall, the within-chain estimates provide complementary evidence that Timeboost's implementation is associated with a reduction in repeated submissions on Arbitrum.

7 Conclusion

This paper studies how explicit markets for time priority reshape competition for time-sensitive trading opportunities. Our theoretical framework isolates a simple mechanism: when access to a time advantage is allocated via a market—rather than via a latency race in a public queue—equilibrium incentives to invest in speed shift from wasteful effort (duplicate submissions or latency technology) to auction bids. In our general model, a single-winner time-priority mechanism reduces aggregate latency investment and reallocates surplus from dissipation into payments for priority. In the blockchain specialization, where latency effort manifests as duplicate on-chain transactions, this translates into less spam, lower revert-driven resource waste, and higher platform-level revenue from time-priority sales.

We bring this framework to data by exploiting the introduction of Timeboost, an express-lane time-priority auction on Arbitrum, and combining cross-chain and within-chain empirical designs. The cross-chain difference-in-differences evidence shows a large and robust post-adoption decline in repeated transaction activity on Arbitrum relative to comparable Layer-2 networks, consistent with reduced redundant competition in the public queue. At the same time, we document an increase in protocol revenue from gas fees and auction proceeds, indicating that a larger share of MEV-related surplus is captured through explicit payments rather than revert-driven dissipation. A within-Arbitrum analysis yields qualitatively similar patterns, reinforcing the interpretation that the introduction of an explicit time-priority mechanism causally reshapes the structure of competition for MEV. Overall, our theoretical and empirical findings provide novel insights into the design of exchanges in the world of high-frequency trading, presenting another approach to mitigating externalities that arise from speed technology competition while improving revenue implications for the platforms themselves.

References

Aquilina, Matteo, Eric Budish, and Peter O'Neill, "Quantifying the High-Frequency Trading "Arms Race"," *The Quarterly Journal of Economics*, February 2022, *137* (1), 493–564.

Biais, Bruno, Thierry Foucault, and Sophie Moinas, "Equilibrium Fast Trading," *Journal of Financial Economics*, 2015, *116* (2), 292–313.

Budish, Eric, Peter Cramton, and John Shim, "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response," *Quarterly Journal of Economics*, 2015, *130* (4), 1547–1621.

_ , Robin S. Lee, and John J. Shim, "Will the Market Fix the Market? A Theory of Stock Exchange Competition and Innovation," *American Economic Review*, 2020, *110* (6), 1607–1635.

Capponi, Agostino and Ruizhe Jia, "Liquidity Provision on Blockchain-Based Decentralized Exchanges," *Review of Financial Studies*, 2025, *38* (10), 3040–3085.

_ , _ , and Kanye Ye Wang, "Maximal Extractable Value and Allocative Inefficiencies in Public Blockchains," *Journal of Financial Economics*, 2025, *172*, 104132.

Daian, Philip, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels, "Flash Boys 2.0: Frontrunning, Transaction Reordering, and Consensus Instability in Decentralized Exchanges," in "2020 IEEE Symposium on Security and Privacy (SP)" 2020, pp. 910–927.

Gogol, Krzysztof M., Manvir Schneider, and Claudio J. Tessone, "First-Spammed, First-Served: MEV Extraction on Fast-Finality Blockchains," 2025. arXiv:2506.01462.

Gogol, Krzysztof, Manvir Schneider, Jan Gorzny, and Claudio Tessone, "How to Serve Your Sandwich? MEV Attacks in Private L2 Mempools," 2026. arXiv:2601.19570.

Hasbrouck, Joel and Gideon Saar, "Low-Latency Trading," *Journal of Financial Markets*, 2013, *16* (4), 646–679.

Lehar, Alfred and Christine Parlour, "Decentralized Exchange: The Uniswap Automated Market Maker," *Journal of Finance*, 2025, *80* (1), 321–374.

Li, Yuhao, Mengqian Zhang, Jichen Li, Elynn Chen, Xi Chen, and Xiaotie Deng, "MEV Makes Everyone Happy under Greedy Sequencing Rule," 2023. arXiv:2309.12640.

Mamageishvili, Akaki, Mahimna Kelkar, Jan Christoph Schlegel, and Edward W. Felten, "Buying Time: Latency Racing vs. Bidding in Transaction Ordering," in "5th Conference on Advances in Financial Technologies (AFT 2023)," Vol. 282 of Leibniz International Proceedings in Informatics (LIPIcs) Schloss Dagstuhl – Leibniz-Zentrum für Informatik 2023. Preprint available as arXiv:2306.02179.

Mazorra, Bruno, Christoph Schlegel, and Akaki Mamageishvili, "Timing Games: Probabilistic backrunning and spam," *arXiv preprint*, 2026, *arXiv:2602.22032*.

McLaughlin, Robert, Nir Chemaya, Dingyue Liu, and Dahlia Malkhi, "CLVR Ordering of Transactions on AMMs," 2025. arXiv:2408.02634.

Messias, Johnnatan and Christof Ferreira Torres, "The Express Lane to Spam and Centralization: An Empirical Analysis of Arbitrum's Timeboost," 2025. arXiv:2509.22143.

Obadia, Alex, "Flashbots: Frontrunning the MEV Crisis," `https://medium.com/flashbots/frontrunning-the-mev-crisis-40629a613752` 2020. Accessed January 26, 2026.

O'Hara, Maureen, "High Frequency Market Microstructure," *Journal of Financial Economics*, 2015, *116* (2), 257–270.

Solmaz, Ozan, Lioba Heimbach, Yann Vonlanthen, and Roger Wattenhofer, "Optimistic MEV in Ethereum Layer 2s: Why Blockspace Is Always in Demand," in "7th Conference on Advances in Financial Technologies (AFT 2025)" 2025.

Torres, Christof Ferreira, Albin Mamuti, Ben Weintraub, Cristina Nita-Rotaru, and Shweta Shinde, "Rolling in the Shadows: Analyzing the Extraction of MEV Across Layer-2 Rollups," in "Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security (CCS '24)" Association for Computing Machinery 2024, pp. 2591–2605.

Weintraub, Ben, Christof Ferreira Torres, Cristina Nita-Rotaru, and Radu State, "A Flash(bot) in the Pan: Measuring Maximal Extractable Value in Private Pools," in "Proceed-

ings of the 22nd ACM Internet Measurement Conference (IMC '22)" Association for Computing Machinery Nice, France 2022.

Woodward, Megan, "Bumping Up the Competition: The Influence of IEX's Speed Bump in US Financial Markets," *SSRN Electronic Journal*, 2018. Available at SSRN: `https://ssrn.com/abstract=3202843`.

Zhu, Brian Z., Xin Wan, Ciamac C. Moallemi, Dan Robinson, and Brad Bachu, "Quantifying the Value of Revert Protection," in Christina Garman and Pedro Moreno-Sanchez, eds., *Financial Cryptography and Data Security (FC 2025)*, Vol. 15752 of Lecture Notes in Computer Science Springer, Cham 2026, pp. 72–88.

A Proofs of Results

A.1 Proof of Proposition 3.1

Let $K_{-i} = \sum_{-i} k_{-i}$. Differentiating, we have, for $k_i > 0$

$$\frac{\partial u_i}{\partial k_i}(k_i; k_{-i}) = V \frac{K_{-i}}{(k_i + K_{-i})^2} - C'(k_i),$$

$$\frac{\partial^2 u_i}{\partial k_i^2}(k_i; k_{-i}) = -2V \frac{K_{-i}}{(k_i + K_{-i})^3} - C''(k_i) < 0$$

since $C'' > 0$ and $K_{-i} \geq 0$. Thus for each fixed $K_{-i} \geq 0$, $u_i(\cdot; K)$ is strictly concave on $(0, \infty)$, so player i has at most one interior best response. A strategy profile $k_i = k$ for all $i \in [N]$ is an interior symmetric equilibrium iff

$$\frac{\partial u_i}{\partial k_i}(k; (N-1)k) = 0,$$

which is exactly

$$\begin{aligned} f(k) &:= V \cdot \frac{(N-1)k}{(NK)^2} - C'(k) \\ &= V \cdot \frac{N-1}{N^2} \cdot \frac{1}{k} - C'(k) = 0. \end{aligned}$$

We now show f has a unique positive root. Differentiating,

$$f'(k) = -V \cdot \frac{N-1}{N^2} \cdot \frac{1}{k^2} - C''(k) < 0 \quad \text{for all } k > 0,$$

so f is strictly decreasing on $(0, \infty)$. As $k \downarrow 0$, the term $(1/k)$ dominates while $C'(k)$ remains finite (by continuity of C'), so $\lim_{k \downarrow 0} f(k) = +\infty$. As $k \to \infty$, strict convexity and $C'(k)$ increasing imply $C'(k) \to \infty$, while the first term $O(1/k) \to 0$, hence $\lim_{k \to \infty} f(k) = -\infty$. By continuity and the intermediate value theorem, there exists at least one $k^* > 0$ with $f(k^*) = 0$. By strict monotonicity, this root is unique. At that k^*, strict concavity of $u_i(\cdot; K^*)$ where $K^* = Nk^*$ implies that k^* is the unique best response to k^*, so $(k^*, \ldots, k^*)$ is the unique symmetric equilibrium.

A.2 Proof of Proposition 3.2

Existence of Equilibrium Write winner and loser payoffs are

$$u_w(k_w; k_{-w}) = V\left(1 - \alpha(k_w, T)\frac{s}{k_w + s}\right) - C(k_w),$$

$$u_j(k_j; k_{-j}) = V\,\alpha(k_w, T)\frac{k_j}{k_j + S_j} - C(k_j),$$

where $S_j := k_w + \sum_{m \neq w,j} k_m$. Fix $j \neq w$ and fix opponents' efforts $(k_i)_{i \neq j}$, so that k_w and $S_j => 0$ are constants and $\alpha(k_w, T)$ is constant with respect to k_j. Define

$$f(k) := \frac{k}{k + S_j}.$$

Taking derivatives yields

$$f'(k) = \frac{S_j}{(k + S_j)^2} > 0,$$

$$f''(k) = -\frac{2S_j}{(k + S_j)^3} < 0.$$

Then, the benefit term $V \cdot \alpha(k_w, T) \cdot f(k)$ is concave in k (since $V \cdot \alpha(k_w, T) \geq 0$) and $-C(k)$ is concave because C is convex, so

$$\frac{\partial^2 u_j}{\partial k_j^2}(k_j; k_{-j}) = V\alpha(k_w, T)f''(k_j) - C''(k_j) = -V\alpha(k_w, T)\frac{2S_j}{(k_j + S_j)^3} - C''(k_j) < 0.$$

Fix the losers' effort profile k_{-w}, so that $s :== \sum_{j \neq w} k_j, > 0$ is constant. Define

$$g(k) := \alpha(k, T) \cdot \frac{s}{k + s}.$$

Then we can write the winner's payoff as

$$u_w(k_w; k_{-w}) = V(1 - g(k_w)) - C(k_w).$$

We show g is convex in k. Write $d := k + s > 0$, and denote $\alpha = \alpha(k, T)$, $\alpha_k = \partial\alpha/\partial k$, $\alpha_{kk} = \partial^2\alpha/\partial k^2$. Differentiating yields

$$g'(k) = s\left(\frac{\alpha_k}{d} - \frac{\alpha}{d^2}\right) = s\frac{\alpha_k d - \alpha}{d^2}.$$

Differentiating again yields

$$\begin{aligned} g''(k) &= s\frac{d}{dk}\left(\frac{\alpha_k d - \alpha}{d^2}\right) = s\frac{(\alpha_{kk}d + \alpha_k - \alpha_k)d^2 - (\alpha_k d - \alpha)\cdot 2d}{d^4} \\ &= s\frac{\alpha_{kk}d^3 - 2\alpha_k d^2 + 2\alpha d}{d^4} = s\frac{\alpha_{kk}d^2 - 2\alpha_k d + 2\alpha}{d^3}. \end{aligned}$$

Under $\alpha_{kk} \geq 0$, $\alpha_k \leq 0$, and $\alpha \geq 0$, each term in the numerator $\alpha_{kk}d^2$, $-2\alpha_k d$, and 2α is nonnegative; hence $g''(k) \geq 0$ for all $k \geq 0$. Thus g is convex and $-Vg$ is concave. Since $C''(k) > 0$, we have

$$\frac{\partial^2 u_w}{\partial k_w^2}(k_w; k_{-w}) = -Vg''(k_w) - C''(k_w) < 0.$$

Because each player's payoff is strictly concave in own action and continuously differentiable, each player's best response is single-valued and characterized by the first-order condition whenever it is interior. Consider the symmetric-loser class: all losers choose the same effort k_l. Then $s = (N-1)k_l$ and $K = k_w + (N-1)k_l$, so a representative loser's payoff (given k_w and other losers at k_l) is

$$u_l(k_l; k_w) = V\,\alpha(k_w, T)\frac{k_l}{k_w + (N-1)k_l} - C(k_l).$$

and the winner's payoff (given k_l) is

$$u_w(k_w; k_l) = V\left(1 - \alpha(k_w, T)\frac{(N-1)k_l}{k_w + (N-1)k_l}\right) - C(k_w).$$

The interior first-order conditions are then:

$$C'(k_l) = V \cdot \alpha(k_w, T) \cdot \frac{k_w + (N-2)k_l}{(k_w + (N-1)k_l)^2}, \tag{L}$$

$$C'(k_w) = V\left[-\alpha_k(k_w, T) \cdot \frac{(N-1)k_l}{k_w + (N-1)k_l} + \alpha(k_w, T) \cdot \frac{(N-1)k_l}{(k_w + (N-1)k_l)^2}\right]. \tag{W}$$

Define the best responses in the symmetric-loser class:

$$BR_l(k_w) \in \arg\max_{k_l \geq 0} u_l(k_l; k_w), \qquad BR_w(k_l) \in \arg\max_{k_w \geq 0} u_w(k_w; k_l).$$

Strict concavity in own action implies each best response is unique. Because u_l and u_w are C^2 and strictly concave in own choice, the maximizers depend continuously on the parameter. Boundedness follows from $\lim_{k\to\infty} C'(k) = \infty$ and bounded marginal benefits. Indeed, for the loser,

$$\frac{\partial}{\partial k_l}\left(V \cdot \alpha(k_w, T) \cdot \frac{k_l}{k_w + (N-1)k_l}\right) = V \cdot \alpha(k_w, T) \cdot \frac{k_w}{(k_w + (N-1)k_l)^2} \leq \frac{V}{k_w},$$

which tends to 0 as $k_l \to \infty$. For the winner, note that the winner's probability is in $[0, 1]$ and thus $u_w(k_w; k_l) \leq V - C(k_w)$, which implies maximizers cannot diverge because $C(k_w) \to \infty$ (since C' is unbounded and C is increasing). Hence there exists $M > 0$ such that $BR_l(k_w) \leq M$ for $k_w \in [0, M]$ and $BR_w(k_l) \leq M$ for $k_l \in [0, M]$.

A sufficient condition to ensure interiority is that each player's marginal payoff at zero is strictly positive at the relevant parameter values. Fix $k_l > 0$ (so $s = (N-1)k_l > 0$). The winner's marginal benefit at $k_w = 0$ is

$$\frac{\partial}{\partial k_w}\left[V\left(1 - \alpha(k_w, T)\frac{s}{k_w + s}\right)\right]\Bigg|_{k_w=0} = V\left(-\alpha_k(0, T) + \frac{\alpha(0, T)}{s}\right) \geq V\big(-\alpha_k(0, T)\big).$$

Thus if $C'(0) < V\big(-\alpha_k(0, T)\big)$ and $-\alpha_k(0, T) > 0$, then the winner's marginal payoff at 0 exceeds marginal cost for every $s > 0$, so $BR_w(k_l) > 0$ for all $k_l > 0$. Fix $k_w > 0$. The loser's marginal benefit at $k_l = 0$ is

$$\frac{\partial}{\partial k_l}\left[V\alpha(k_w, T)\frac{k_l}{k_w + (N-1)k_l}\right]\Bigg|_{k_l=0} = V\alpha(k_w, T)\frac{1}{k_w}.$$

Hence if $C'(0)$ is finite and $\alpha(k_w, T) > 0$, then for k_w in any bounded interval away from infinity, the loser has positive marginal payoff at 0 and thus $BR_l(k_w) > 0$.

Define the continuous map on $[0, M]^2$:

$$\Psi(k_w, k_l) := \big(BR_w(k_l),\, BR_l(k_w)\big).$$

By boundedness, for sufficiently large M we have $\Psi([0,M]^2) \subseteq [0,M]^2$. By Brouwer's Fixed Point Theorem, there exists $(k_w^*, k_l^*) \in [0,M]^2$ such that

$$(k_w^*, k_l^*) = \Psi(k_w^*, k_l^*).$$

Under the condition for interiority and $\alpha(k,T) > 0$ on the relevant region, this fixed point is interior. Because payoffs are strictly concave in own actions, both best responses are single-valued. To obtain uniqueness of the fixed point, it suffices to show that the best responses are monotone in opposite directions and that they are continuous. Then the graphs can intersect at most once. The loser's objective is

$$u_l(k_l; k_w) = V \cdot \alpha(k_w, T) \cdot \frac{k_l}{k_w + (N-1)k_l} - C(k_l).$$

For fixed $k_l > 0$, the function $k_w \mapsto \alpha(k_w, T)\frac{k_l}{k_w+(N-1)k_l}$ is decreasing because both $\alpha(k_w, T)$ and $\frac{1}{k_w+(N-1)k_l}$ decrease in k_w. Thus the loser's payoff has decreasing differences in (k_l, k_w), implying $BR_l(k_w)$ is weakly decreasing in k_w. The winner's payoff can be written as

$$u_w(k_w; k_l) = V - V \cdot \alpha(k_w, T) \cdot \frac{(N-1)k_l}{k_w + (N-1)k_l} - C(k_w).$$

For fixed k_w, the term $\frac{(N-1)k_l}{k_w+(N-1)k_l}$ is increasing in k_l, hence the "penalty" $V\alpha(k_w, T)\frac{(N-1)k_l}{k_w+(N-1)k_l}$ increases in k_l, making higher k_l increase the marginal benefit of k_w. This yields increasing differences in (k_w, k_l) and therefore $BR_w(k_l)$ is weakly increasing in k_l. Since BR_w is weakly increasing and BR_l is weakly decreasing, continuity implies a unique intersection. Thus the interior symmetric equilibrium is unique.

We now show that $0 < k_l^* < k_w^*$. Suppose, for contradiction, that $k_w^* = k_l^* = k > 0$. Then $K = Nk$ and $s = (N-1)k$. The RHS of the loser FOC becomes

$$V \cdot \alpha(k) \cdot \frac{k + (N-2)k}{(Nk)^2} = V \cdot \alpha(k) \cdot \frac{N-1}{N^2} \cdot \frac{1}{k},$$

and the RHS of the winner FOC becomes

$$V\left[-\alpha'(k)\frac{(N-1)k}{Nk} + \alpha(k)\frac{(N-1)k}{(Nk)^2}\right] = V \cdot \frac{N-1}{N} \cdot -\alpha'(k) + V \cdot \alpha(k) \cdot \frac{N-1}{N^2} \cdot \frac{1}{k}.$$

Therefore

$$\text{RHS(W)} - \text{RHS(L)} = V \cdot \frac{N-1}{N} \cdot -\alpha'(k).$$

Since $\alpha'(k) < 0$, we have $-\alpha'(k) > 0$, hence RHS(W) > RHS(L). However, as both FOCs require equality with the same $C'(k)$, which is a contradiction. Thus, we must have that $0 < k_l^* < k_w^*$..

We finally show that $0 < u_l^* <_w^*$. Define

$$\pi_l(k_w) = \alpha(k_w) \cdot \frac{k_l}{k_w + (N-1)k_l}.$$

Differentiating yields

$$\pi_l'(k_w) = k_l \cdot \frac{\alpha'(k_w) \cdot (k_w + (N-1)k_l) - \alpha(k_w)}{(k_w + (N-1)k_l)^2}.$$

Since $\alpha'(k_w) \leq 0$ and $\alpha(k_w) \geq 0$, the numerator is strictly negative, hence $\pi_l'(k_w) < 0$. Thus $u_l(k_l; k_w)$ is strictly decreasing in k_w. Since $k_w^* > k_l^*$,

$$u_l(k_l^*; k_w^*) < u_l(k_l^*; k_l^*). \tag{A}$$

At $k_w = k_l = k_l^*$, we have $K = Nk_l^*$. Then

$$P_l = \frac{\alpha(k_l^*)}{N}$$

$$P_w = 1 - \alpha(k_l^*) \cdot \frac{N-1}{N}.$$

Thus $u_w(k_l^*; k_l^*) - u_l(k_l^*; k_l^*) = V\big(1 - \alpha(k_l^*)\big)$. Since $k_l^* > 0$ and α is strictly decreasing, we have

$$\alpha(k_l^*) < 1.$$

so we have

$$u_l(k_l^*; k_l^*) < u_w(k_l^*; k_l^*) < u_w(k_w^*; k_l^*) = u_w^*. \tag{B}$$

Combining (A) and (B) yields the result.

A.3 Proof of Proposition 3.3

Let $N \geq 3$ bidders participate in a sealed-bid second-price auction and

$$v_A \;:=\; u_w^* - u_l^*$$

denote each bidder's (common, known) incremental value of winning the Express Lane right, i.e. the difference between being the winner and being a loser in the continuation game. Conditional on the submission-subgame equilibrium being played after the auction, a bidder who wins and pays price p obtains utility $u_w^* - p$, while a bidder who loses obtains utility u_l^*. Thus the incremental gain from winning at price p is $v_A - p$.

Consider a symmetric pure bid profile in which all bidders submit the same bid $b \geq 0$. Under uniform tie-breaking, each bidder wins with probability $1/N$ and pays the (second-highest) price b whenever they win. Hence each bidder's expected utility is

$$U(b) = \frac{1}{N}(u_w^* - b) + \frac{N-1}{N}u_l^* = u_l^* + \frac{1}{N}(v_A - b).$$

Fix such a symmetric profile and consider a unilateral deviation by bidder i:

1. If $b < v_A$, bidder i can deviate to $b + \varepsilon$ for any sufficiently small $\varepsilon > 0$. Then bidder i wins with probability 1 and pays the second-highest bid, which remains b. The deviation payoff is therefore $u_w^* - b = u_l^* + (v_A - b)$, which strictly exceeds $U(b) = u_l^* + (v_A - b)/N$ because $v_A - b > 0$ and $N \geq 3$. Hence no symmetric profile with $b < v_A$ can be a Nash equilibrium.

2. If $b > v_A$, bidder i can deviate to a sufficiently low bid (e.g. 0), guaranteeing that they lose and obtain u_l^*. Under the symmetric profile, their payoff is $U(b) = u_l^* + (v_A - b)/N < u_l^*$. Thus any symmetric profile with $b > v_A$ admits a profitable deviation and cannot be a Nash equilibrium.

3. If $b = v_A$, then $U(v_A) = u_l^*$. Deviating to any bid $b_i < v_A$ guarantees losing and yields u_l^*, which equals $U(v_A)$. Deviating to any bid $b_i > v_A$ guarantees winning and paying v_A, yielding $u_w^* - v_A = u_l^*$, again equal to $U(v_A)$. Therefore no deviation yields a strictly higher payoff, so $b = v_A$ is a best response to the symmetric profile.

Combining (1)–(3), it follows that the only symmetric pure bid profile that can be a Nash equilibrium is $b^* = v_A$, and it is indeed a Nash equilibrium. We proceed to verify sequential rationality in every subgame.

1. From Proposition 3.2, after any realized auction winner, the continuation strategies form a Nash equilibrium of the corresponding submission subgame. Hence no player can profitably deviate after observing the auction outcome.

2. Given that the continuation equilibrium is played after the auction, each bidder's value for winning the Express Lane right is $v_A = u_w^* - u_l^*$. By the previous proposition, bidding $b^* = v_A$ is a best response (indeed, the unique symmetric pure-strategy equilibrium bid) to the others bidding b^*. Therefore no bidder can profitably deviate in the auction stage.

Thus, the overall strategy profile is subgame-perfect.

A.4 Proof of Theorem 3.4

Recall that the symmetric-loser FOCs are:

$$C'(k_l^*) = V \cdot \alpha(k_w^*, T) \cdot \frac{k_w^* + (N-2)k_l^*}{(K_A^*)^2}, \tag{L}$$

$$C'(k_w^*) = V\left[-\alpha_k(k_w^*, T) \cdot \frac{(N-1)k_l^*}{K_A^*} + \alpha(k_w^*, T) \cdot \frac{(N-1)k_l^*}{(K_A^*)^2}\right]. \tag{W}$$

Multiply (L) by $(N-1)$ and add to (W). Simiplifying yields

$$(N-1) \cdot C'(k_l^*) + C'(k_w^*) = \frac{V(N-1)}{K_A^*}\Big(\alpha(k_w^*, T) - \alpha_k(k_w^*, T) \cdot k_l^*\Big).$$

Since C' is convex and

$$\frac{K_A^*}{N} = \frac{1}{N} \cdot k_w^* + \frac{N-1}{N} \cdot k_l^*,$$

Jensen's inequality implies

$$C'\left(\frac{K_A^*}{N}\right) \leq \frac{1}{N}\Big(C'(k_w^*) + (N-1)C'(k_l^*)\Big) \leq \frac{V(N-1)}{NK_A^*}\Big(\alpha(k_w^*, T) - \alpha_k(k_w^*, T)k_l^*\Big).$$

Convexity of α implies that for every $k > 0$,

$$1 = \alpha(0, T) \geq \alpha(k, T) - k \cdot \alpha_k(k, T).$$

Since $0 < k_l^* < k_w^*$ and $\alpha_k < 0$,

$$-\alpha_k(k_w^*, T) \cdot k_l^* < -\alpha_k(k_w^*, T) \cdot k_w^* \implies \alpha(k_w^*, T) - \alpha_k(k_w^*, T) k_l^* < 1.$$

Now, the following inequality holds:

$$C'\left(\frac{K_A^*}{N}\right) < \frac{V(N-1)}{NK_A^*}. \tag{5}$$

Define

$$F(x) := C'\left(\frac{x}{N}\right) - \frac{V(N-1)}{Nx}.$$

Because C is convex, C' is increasing; and $x \mapsto -1/x$ is increasing on $(0, \infty)$, so F is strictly increasing. By construction, we have $F(K^*) = 0$. However, as we have just shown, $F(K_A^*) < 0$. Since F is strictly increasing, it follows that $F(K_A^*) < F(K^*)$, so $K_A^* < K^*$ as well.

A.5 Preliminaries for 2-Type Allocations

Suppose that $n_l, n_w \geq 1$. Let $K_w = k_w^{(i)} + \sum_{-i} k_w^{(-i)}$, $K_l = k_l^{(j)} + \sum_{-j} k_l^{(-j)}$, and $K = K_w + K_l$. Payoffs are then

$$u_w^{(i)}(k_w^{(i)}; k_w^{(-i)}, k_l^{(j)}) = V \cdot \frac{k_w^{(i)}}{K_w}\left(1 - \frac{K_l}{K} \cdot \exp(-\lambda T K_w)\right) - C k_w^{(i)},$$

$$u_l^{(j)}(k_l^{(j)}; k_w^{(i)}, k_l^{(-j)}) = V\left(\frac{k_l^{(j)}}{K} \cdot \exp(-\lambda T K_w)\right) - C k_l^{(j)}.$$

For fixed winner i and loser j, let $K_{-w} = \sum_{-i} k_w^{(-i)}$ and $K_{-l} = \sum_{-j} k_l^{(-j)}$. Differentiating, we have

$$\frac{du_w^{(i)}}{dk_w^{(i)}} = V\left[\frac{K_{-w}}{K_w^2}\left(1 - \frac{K_l}{K} e^{-\lambda T K_w}\right) + \frac{k_w^{(i)}}{K_w}\left(\frac{K_l}{K^2} + \lambda T \frac{K_l}{K}\right) e^{-\lambda T K_w}\right] - C$$

$$\frac{du_l^{(j)}}{dk_l^{(j)}} = V \cdot \frac{K_w + K_{-l}}{K^2} \cdot e^{-\lambda T K_w} - C$$

Differentiating the loser's condition again, noting that K_w is constant in $k_l^{(j)}$ and $\frac{dK}{dk_l^{(j)}} = 1$, we have

$$\frac{d^2 u_l^{(j)}}{d(k_l^{(j)})^2} = V\left(K_w + K_{-l}\right) e^{-\lambda T K_w} \frac{d}{dk_l^{(j)}}(K^{-2}) = -\frac{2V\left(K_w + K_{-l}\right)}{K^3} e^{-\lambda T K_w} \leq 0.$$

For the winner's condition, we have

$$\frac{d^2 u_w^{(i)}}{d(k_w^{(i)})^2} < 0.$$

The proof is given in the Additional Proofs section.

Imposing symmetry, i.e. $K_l = n_l k_l$, $K_{-l} = (n_l - 1)k_l$, $K_w = n_w k_w$, and $K_{-w} = (n_w - 1)k_w$, the FOCs are then

$$C = V\left[\frac{n_w - 1}{n_w^2\, k_w}\left(1 - \frac{n_l k_l}{n_w k_w + n_l k_l} e^{-\lambda T n_w k_w}\right) + \frac{1}{n_w}\left(\frac{n_l k_l}{(n_w k_w + n_l k_l)^2} + \lambda T \frac{n_l k_l}{n_w k_w + n_l k_l}\right) e^{-\lambda T n_w k_w}\right],$$

$$C = V \frac{n_w k_w + (n_l - 1)k_l}{(n_w k_w + n_l k_l)^2} e^{-\lambda T n_w k_w}.$$

Equilibrium effort is $K^* = n_w k_w^* + n_l k_l^*$ and equilibrium continuation payoffs are then

$$u_w^* = V \cdot \frac{1}{n_w}\left(1 - \frac{n_l k_l^*}{n_l k_l^* + n_w k_w^*} \cdot \exp(-\lambda T n_w k_w^*)\right) - C k_w^*$$

$$u_l^* = V \cdot \frac{k_l^*}{n_l k_l^* + n_w k_w^*} \exp(-\lambda T n_w k_w^*) - C k_l^*$$

Let $W = n_w k_w$, $L = n_l k_l$, $K = W + L$, $\alpha_l = 1 - n_l^{-1}$, $\alpha_w = 1 - n_w^{-1}$, and $\tau = \lambda T$. Assume that $V, C, \lambda, T > 0$ and $\alpha_l, \alpha_w \in [0, 1)$. Rewrite the FOCs as

$$F_w(W, L, T) := \alpha_w\Big(\frac{1}{W} \cdot \frac{W + L(1 - e^{-\tau W})}{W + L}\Big) + (1 - \alpha_w)\Big(\frac{L}{(W + L)^2} + \frac{\tau L}{W + L}\Big) \exp(-\tau W) - \rho$$

$$F_l(W, L, T) := \frac{W + \alpha_l L}{(W + L)^2} e^{-\tau W} - \rho,$$

$$\alpha_l = \frac{(N - 1)(1 - \alpha_w) - 1}{N(1 - \alpha_w) - 1}$$

$$u_w^* = V\left[(1-\alpha_w)\left(1-\frac{L}{W+L}e^{-\tau W}\right)-\rho(1-\alpha_w)W\right]$$

$$u_l^* = V\left[(1-\alpha_l)\cdot\frac{L}{W+L}e^{-\tau W}-\rho(1-\alpha_l)L\right]$$

A.6 Proof of Proposition 5.1

Less Redundant Submissions. Let $K(T) = n_w k_w(T) + n_l k_l(T)$ denote total equilibrium effort as a function of the advantage difference $T = T$, where $(k_w(T), k_l(T))$ solves the symmetric equilibrium first-order conditions. At $T = 0$, the game is symmetric across all $N = n_w + n_l$ players, and the unique interior equilibrium k_0 is given by the baseline case.

Differentiating the symmetric equilibrium first-order conditions with respect to T and evaluating at $T = 0$ yields

$$k_w'(0) = \lambda n_l(N-1)k_0^2,$$

$$k_l'(0) = -\lambda n_w(N-1)k_0^2.$$

so $K'(0) = n_w k_w'(0) + n_l k_l'(0) = 0$. Differentiating again and solving for $k_w''(0)$ and $k_l''(0)$ yields

$$k_w''(0) = -\frac{V^3\lambda^2}{C^3}\frac{n_l\,(N-1)^2}{N^6}\Big(2n_l^3+5n_l^2n_w-3n_l^2+4n_ln_w^2-3n_ln_w+n_l+n_w^3\Big),$$

$$k_l''(0) = -\frac{V^3\lambda^2}{C^3}\frac{n_w\,(N-1)^2}{N^6}\Big(n_l^2n_w+n_l^2+2n_ln_w^2-n_ln_w+n_w^3-2n_w^2+n_w\Big),$$

$$K''(0) = n_w k_w''(0) + n_l k_l''(0) = -\frac{V^3\lambda^2}{C^3}\frac{n_wn_l(N-1)^2}{N^5}\big(2N(N-1)+1\big) < 0$$

since $n_w, n_l \geq 1$ and $\lambda > 0$. By Taylor's theorem,

$$K'(T) = K'(0) + K''(0)T + o(T) = K''(0)T + o(T).$$

Since $K''(0) < 0$, there exists$\delta > 0$ such that $K'(T) < 0$ for all $T \in (0, \delta)$. Thus, total equilibrium effort is locally decreasing in T to the right of 0.

More Sequencer Revenue. Define $L(T) = n_l k_l^*(T)$, $W(T) = n_w k_w^*(T)$, $K(T) = W(T) + L(T)$, and $E(T) = e^{-\lambda T L(T)}$. From earlier, we have

$$W'(0) = -\lambda n_w n_l (N-1) k_0^2, \quad L'(0) = \lambda n_w n_l (N-1) k_0^2, \quad E'(0) = -\lambda B(0) = -\lambda n_w k_0$$

Define $f(T) = \frac{W(T)}{K(T)} E(T)$; then for the winner's payoff:

$$u_w(T) = \frac{V}{n_w}\big(1 - f(T)\big) - C k_w^*(T),$$

$$u_w'(T) = -\frac{V}{n_w} f'(T) - C k_w'(T).$$

Using $K'(0) = 0$ and $E(0) = 1$, we have

$$f'(0) = \frac{W'(0)}{K(0)} + \frac{W(0)}{K(0)} \cdot E'(0) = -\lambda n_w n_l k_0.$$

so, substituting the expression for k_0, it follows that

$$u_w'(0) = \lambda n_l \big(V k_0 - C(N-1)k_0^2\big) \propto \frac{V^2}{C} \frac{(N-1)(2N-1)}{N^4} > 0.$$

For the loser's payoff, define $g(T) = \frac{k_l^*(T)}{K(T)} E(T)$; then

$$u_l(T) = V g(T) - C k_l^*(T),$$

$$u_l'(T) = V g'(T) - C k_l'(T)$$

Again using $K'(0) = 0$,

$$g'(0) = \frac{k_l'(0)}{K(0)} + \frac{k_l(0)}{K(0)} E'(0) = -\lambda n_w k_0$$

Thus $u_l'(0) = \lambda n_w \big(-V k_0 + C(N-1)k_0^2\big) = -\lambda n_w \big(V k_0 - C(N-1)k_0^2\big) < 0$. Auction revenue $n_w(u_w - u_l)$ and sequencer revenue $V - N u_l$ are thus increasing in T around a right-neighborhood of $T = 0$ for any $n_w, n_l \geq 1$.

A.7 Proof of Proposition 5.2

Less Redundant Submissions. Fix $\tau \coloneqq \lambda T > 0$ (slight abuse of notation) and $\rho \coloneqq C/V > 0$. For each $N \geq 2$, let $(W(\alpha_w; N), L(\alpha_w; N))$ be an equilibrium branch satisfying

$$F_w(W(\alpha_w; N), L(\alpha_w; N), \alpha_w) = 0,$$

$$F_l(W(\alpha_w; N), L(\alpha_w; N), \alpha_w) = 0,$$

Let J denote the Jacobian of the equilibrium system, i.e.

$$J = \begin{pmatrix} F_{wW} & F_{wL} \\ F_{lW} & F_{lL} \end{pmatrix}.$$

From the baseline model with time-priority auction, there exists $W, L > 0$ satisfying both equations at $\alpha_w = 0$ for any $N \geq 2$. Since the system has at most one interior solution, it follows that J is nonsingular at $\alpha_w = 0$ for any $N \geq 2$ as well. We claim that

$$\lim_{N\to\infty} \big(W'(0) + L'(0)\big) < 0,$$

where $W'(0) = \frac{dW}{d\alpha_w}\big|_{\alpha_w=0}$ and $L'(0) = \frac{dL}{d\alpha_w}\big|_{\alpha_w=0}$. Let $K \coloneqq W + L$ and $E \coloneqq e^{-\tau W}$. Define

$$A(W, L) \coloneqq \frac{1}{W} \cdot \frac{W + L(1 - E)}{K}$$

$$B(W, L) \coloneqq \frac{L}{K^2} + \frac{\tau L}{K}$$

Then FOCs can be written as

$$F_w(W, L, \alpha_w) = \alpha_w \cdot A(W, L) + (1 - \alpha_w) \cdot B(W, L) \cdot E - \rho$$

$$F_l(W, L, \alpha_w) = \frac{W + \alpha_l L}{K^2} E - \rho$$

where $\alpha_l = \frac{(N-1)(1-\alpha_w)-1}{N(1-\alpha_w)-1}$. At $\alpha_w = 0$, we denote evaluation by a superscript 0.

Differentiate the equilibrium conditions with respect to α_w:

$$F_{wW}W' + F_{wL}L' + F_{w\alpha_w} = 0,$$

$$F_{lW}W' + F_{lL}L' + F_{l\alpha_w} = 0.$$

Then at $\alpha_w = 0$, by Cramer's rule,

$$W'(0) = \frac{F^0_{wL}F^0_{l\alpha_w} - F^0_{lL}F^0_{w\alpha_w}}{\det J^0},$$
$$L'(0) = \frac{F^0_{lW}F^0_{w\alpha_w} - F^0_{wW}F^0_{l\alpha_w}}{\det J^0}.$$

Summing yields

$$K'(0) = \frac{\left(F^0_{lW} - F^0_{lL}\right)F^0_{w\alpha_w} + \left(F^0_{wL} - F^0_{wW}\right)F^0_{l\alpha_w}}{\det J^0}.$$

From $F_w = \alpha_w A + (1-\alpha_w)BE - \rho$, differentiating and evaluating at $\alpha_w = 0$ yields $F^0_{w\alpha_w} = A - BE$. At $\alpha_w = 0$, the equilibrium condition $F_w = 0$ gives $BE = \rho$, hence

$$F^0_{w\alpha_w} = A - \rho.$$

Next, α_w enters F_l only through α_l, so

$$F_{l\alpha_w} = \frac{\partial F_l}{\partial \alpha_l} \cdot \frac{d\alpha_l}{d\alpha_w} = -\frac{L}{K^2} \cdot E \cdot \frac{1}{(N(1-\alpha_w)-1)^2}.$$

At $\alpha_w = 0$,

$$F^0_{l\alpha_w} = -\frac{L}{K^2}\frac{E}{(N-1)^2} = O\left(\frac{1}{(N-1)^2}\right).$$

At $\alpha_w = 0$ we have $\alpha_l = \frac{N-2}{N-1} \to 1$ as $N \to \infty$. For fixed (W, L), the formulas for the (W, L)-partials of F_l are

$$F^0_{lW} = E\left[\frac{-W + (1-2\alpha_l)L}{K^3} - \tau\frac{W + \alpha_l L}{K^2}\right],$$

$$F^0_{lL} = E\left[\frac{(\alpha_l - 2)W - \alpha_l L}{K^3}\right].$$

Letting $\alpha_l \to 1$ yields

$$F^0_{lW} \to E\left[\frac{-W-L}{K^3} - \tau\frac{W+L}{K^2}\right] = -E\left(\frac{1}{K^2} + \frac{\tau}{K}\right), \qquad F^0_{lL} \to E\left[\frac{-W-L}{K^3}\right] = -\frac{E}{K^2}.$$

Along $\alpha_w = 0$, we may pass to the limit $(W, L) \to (W_\infty, L_\infty)$, so

$$F^0_{lW} - F^0_{lL} \longrightarrow -\frac{E\tau}{K} \quad \text{as } N \to \infty.$$

Define $B(W, L) := \frac{L}{K^2} + \tau\frac{L}{K}$. Then the winner's FOC at $\alpha_w = 0$ is

$$F_w(W, L, 0) = B(W, L) \cdot E - \rho$$

Hence, using the product rule and $E_W = -\tau E$, $E_L = 0$,

$$F^0_{wW} = \frac{\partial}{\partial W}\big(BE\big) = B_W\, E + B\, E_W = E\big(B_W - \tau B\big), \tag{1}$$

$$F^0_{wL} = \frac{\partial}{\partial L}\big(BE\big) = B_L\, E + B\, E_L = E\, B_L. \tag{2}$$

Writing $B = LK^{-2} + \tau LK^{-1}$ and noting that $K_W, K_L = 1$ we have

$$B_W = L \cdot (-2)K^{-3} \cdot K_W + \tau L \cdot (-1)K^{-2} \cdot K_W - \frac{2L}{K^3} - \tau\frac{L}{K^2}, \tag{3}$$

$$B_L = \Big(K^{-2} + L \cdot (-2)K^{-3} \cdot K_L\Big) + \tau\Big(K^{-1} + L \cdot (-1)K^{-2} \cdot K_L\Big) = \frac{W-L}{K^3} + \tau\frac{W}{K^2}. \tag{4}$$

Substituting (3)–(4) into (1)–(2) gives the explicit Jacobian entries:

$$F^0_{wW} = E\left(-\frac{2L}{K^3} - \tau\frac{L}{K^2} - \tau\left(\frac{L}{K^2} + \tau\frac{L}{K}\right)\right) = E\left(-\frac{2L}{K^3} - 2\tau\frac{L}{K^2} - \tau^2\frac{L}{K}\right),$$

$$F^0_{wL} = E\left(\frac{W-L}{K^3} + \tau\frac{W}{K^2}\right).$$

Since $F^0_{l\alpha_w} < 0$ and

$$F^0_{wL} - F^0_{wW} = \left(\frac{1}{K^2} + \tau \frac{W + 2L}{K^2} + \tau^2 \frac{L}{K} \right) E > 0,$$

it follows that $(F^0_{wL} - F^0_{wW}) F^0_{l\alpha_w} < 0$. Since equilibrium is unique when $\alpha_w = 0$ for any N, $\det J^0$ stays bounded away from 0 for large N:

$$\exists\, c > 0,\ \exists\, N_0 \text{ such that } \quad |\det J^0| \geq c \quad \forall\, N \geq N_0. \tag{5}$$

Next, F^0_{wW} and F^0_{wL} are given explicitly by (A.7)–(A.7). In particular, for fixed (W, L) these are finite, so there exists $M < \infty$ such that for all large N,

$$|F^0_{wL} - F^0_{wW}| \leq M. \tag{6}$$

Combining the expressions derived , (5), and (6) yields the explicit estimate

$$\left| \frac{(F^0_{wL} - F^0_{wW})\, F^0_{l\alpha_w}}{\det J^0} \right| \leq \frac{M}{c} \left| F^0_{l\alpha_w} \right| = \frac{M}{c} \cdot \frac{L}{K^2} \cdot \frac{E}{(N-1)^2}.$$

Because L/K^2 and E are bounded along the convergent equilibrium sequence, the right-hand side tends to zero as $N \to \infty$. Therefore,

$$\frac{(F^0_{wL} - F^0_{wW})\, F^0_{l\alpha_w}}{\det J^0} \to 0 \quad \text{as } N \to \infty.$$

and thus,

$$\lim_{N\to\infty} \big(W'(0) + L'(0) \big) = \lim_{N\to\infty} \frac{(F^0_{lW} - F^0_{lL})\, F^0_{w\alpha_w}}{\det J^0}. \tag{7}$$

It remains to find the sign of $\det J^0$ and $F^0_{w\alpha_w}$. Let

$$\det J^0 = E^2\, \widetilde{D}_N,$$

where $\widetilde{D}_N$ is the determinant of the E-free matrix. Let $N \to \infty$, equivalently $\alpha_l \to 1$, and define

$\widetilde{D}_\infty := \lim_{N\to\infty} \widetilde{D}_N$. A direct algebraic simplification yields

$$\widetilde{D}_\infty = \frac{1}{K^4} + \frac{\tau(L+2W)}{K^4} + \frac{\tau^2}{K^2} > 0.$$

Consequently, for all sufficiently large N,

$$\widetilde{D}_N > 0 \implies \det J^0 = E^2 \widetilde{D}_N > 0.$$

Let $N \to \infty$ so $\alpha_l \to 1$. Then the loser's FOC at $\alpha_w = 0$ is

$$\frac{W+L}{K^2} E = \rho \implies E = \rho K.$$

Combining with the winner's first-order condition gives $B\rho K = K$, i.e. $BK = 1$. Since $B = \frac{L}{K^2} + \frac{\tau L}{K}$, this implies

$$\frac{L}{K} = \frac{1}{1+\tau K},$$
$$W = K - L = \frac{\tau K^2}{1+\tau K}.$$

By (A.7), $F^0_{w\alpha_w} = A - \rho$, so dividing by $E > 0$ and using $\rho/E = 1/K$ (because $E = \rho K$),

$$\frac{F^0_{w\alpha_w}}{E} = \frac{A}{E} - \frac{1}{K}.$$

We also have

$$A = \frac{1}{W} \cdot \frac{W + L(1-E)}{K} = \frac{1}{W}\left(1 - \frac{L}{K}E\right) \implies \frac{A}{E} = \frac{1}{W}\left(\frac{1}{E} - \frac{L}{K}\right).$$

Using $E = \rho K$, $\frac{L}{K} = \frac{1}{1+\tau K}$, and $\frac{1}{W} = \frac{1+\tau K}{\tau K^2}$ yields

$$\frac{F^0_{w\alpha_w}}{E} = \frac{1+\tau K}{\tau K^2}\left(\frac{1}{\rho K} - 1\right).$$

Because $W > 0$ and $T > 0$, we have $E = e^{-\tau W} \in (0,1)$, hence $\rho K = E \in (0,1)$ and thus $\frac{1}{\rho K} - 1 > 0$.

All remaining factors are positive, so as $N \to \infty$

$$\frac{F^0_{w\alpha_w}}{E} > 0 \implies F^0_{w\alpha_w} > 0$$

Combining everything meas that for large N, the quotient in (7) converges to a strictly negative limit, and hence $\lim_{N\to\infty}\big(W'(0) + L'(0)\big) < 0$.

More Sequencer Revenue. Using our reparameterizations, we can write a loser's equilibrium continuation payoff, which we now denote by U, as

$$U := V\Big(\frac{L}{K}\Big)(1-\alpha_l)E - C\,L(1-\alpha_l) = V(1-\alpha_l)\Big(\frac{L}{K}E - \rho L\Big),$$

and assume the equilibrium path in the (α_w, N)-space, given by $(W(\alpha_w; N), L(\alpha_w; N))$ satisfies

$$F_w(W, L, \alpha_w) = 0,$$

$$F_l(W, L, \alpha_w) = 0.$$

Let $H := \frac{L}{K}E - \rho L$. Then

$$U = V(1-\alpha_l)H$$

$$U' = V\big[(1-\alpha_l)H' - \alpha_l' H\big].$$

Since

$$\left(\frac{L}{K}\right)' = \frac{WL' - LW'}{K^2}, \qquad E' = -TW'E,$$

we obtain

$$H' = E\left(\frac{WL' - LW'}{K^2} - \tau\frac{L}{K}W'\right) - \rho L'.$$

Substituting, we then have

$$U' = V\left\{(1-\alpha_l)\left[E\left(\frac{WL' - LW'}{K^2} - \tau\frac{L}{K}W'\right) - \rho L'\right] - \alpha_l'\left(\frac{L}{K}E - \rho L\right)\right\}.$$

In equilibrium, $F_l = 0$ implies

$$\frac{W + \alpha_l L}{K^2}E = R \implies E = \frac{\rho K^2}{W + \alpha_l L}.$$

Thus, we have

$$\frac{L}{K}E - \rho L = \rho L\left(\frac{K - (W + \alpha_l L)}{W + \alpha_l L}\right) = \rho L\,\frac{(1-\alpha_l)L}{W + \alpha_l L}.$$

We can write the loser-type's equilibrium utility as

$$U = V(1-\alpha_l)\,\rho L\,\frac{(1-\alpha_l)L}{W + \alpha_l L} = C\,\frac{(1-\alpha_l)^2 L^2}{W + \alpha_l L}.$$

Let $\varepsilon := 1 - \alpha_l$ and $D := W + \alpha_l L$. Then

$$U = C\,\frac{\varepsilon^2 L^2}{D}.$$

Differentiating with respect to α_w gives

$$U' = C\left(\frac{2\varepsilon\varepsilon' L^2 + 2\varepsilon^2 LL'}{D} - \frac{\varepsilon^2 L^2 D'}{D^2}\right),$$

$$D' = W' + \alpha_l' L + \alpha_l L'.$$

At $\alpha_w = 0$,

$$\varepsilon_N := 1 - \alpha_l(0) = \frac{1}{N-1} \implies \varepsilon_N' = \varepsilon_N^2,\ \alpha_l'(0) = -\varepsilon_N^2.$$

Substituting into U' yields

$$U'(0) = C\,\varepsilon_N^2 \left(\frac{2\varepsilon_N L^2 + 2L \cdot L'(0)}{D_0} - \frac{L^2\big(W'(0) + (1-\varepsilon_N)L'(0)\big)}{D_0^2} \right) + O(\varepsilon_N^4),$$

where $D_0 = W + (1-\varepsilon_N)L \to S$ as $N \to \infty$. Dividing by $C\varepsilon_N^2$ and letting $N \to \infty$ gives

$$\lim_{N\to\infty} \frac{U'(0)}{C\varepsilon_N^2} = \frac{2L_\infty L'_\infty}{S_\infty} - \frac{L_\infty^2(W'_\infty + L'_\infty)}{S_\infty^2}. \tag{7}$$

As $N \to \infty$, $\alpha_l \to 1$ and $F_{l\alpha_w}(0) \to 0$, so by the Implicit Function Theorem,

$$\begin{pmatrix} F_{wW}^\infty & F_{wL}^\infty \\ F_{lW}^\infty & F_{lL}^\infty \end{pmatrix} \begin{pmatrix} W'_\infty \\ L'_\infty \end{pmatrix} = -\begin{pmatrix} F_{w\alpha_w}^\infty \\ 0 \end{pmatrix}.$$

The second row implies

$$F_{lW}^\infty W'_\infty + F_{lL}^\infty L'_\infty = 0.$$

At $\alpha_l = 1$,

$$F_{lW}^\infty = -E\left(\frac{1}{K^2} + \frac{T}{K}\right),$$

$$F_{lL}^\infty = -\frac{E}{K^2},$$

so $L'_\infty = -(1+\tau K_\infty)W'_\infty$. From the limiting $\alpha_w = 0$ equilibrium we also have

$$\frac{L_\infty}{S_\infty} = \frac{1}{1+\tau K_\infty}.$$

Moreover, since $\lim_{N\to\infty}(W'(0) + L'(0)) < 0$, it follows that $W'_\infty > 0$. Substituting yields

$$\frac{2L_\infty L'_\infty}{K_\infty} - \frac{L_\infty^2(W'_\infty + L'_\infty)}{K_\infty^2} = W'_\infty \left[-2(1+x)\frac{1}{1+x} + x\frac{1}{(1+x)^2} \right],$$

where we denote $x := \tau K_\infty > 0$. Simplifying,

$$= W'_\infty \left(-2 + \frac{x}{(1+x)^2} \right) = -\,W'_\infty \frac{2 + 3x + 2x^2}{(1+x)^2} < 0,$$

since $W'_\infty > 0$ and $x > 0$. Therefore, $\lim_{N\to\infty} \frac{U'(0)}{C\varepsilon_N^2} < 0$ so it follows that $\lim_{N\to\infty} U'(0) = 0^-$. In particular, $U'(0) < 0$ for all sufficiently large N.

B Additional Proofs

B.1 Multiple Winners: Concavity of Winner's Payoff

Write the winner's payoff as

$$u_w^{(i)}(x; a, L) = V \cdot \frac{x}{K_w}\left(1 - \frac{L}{K}\exp(-\tau K_w)\right) - Cx.$$

Since $(-Cx)'' = 0$ and $V > 0$, it suffices to prove concavity of

$$s(x) := \frac{x}{K_w}\left(1 - \frac{K_l}{K}e^{-\tau K_w}\right)$$

where $K_w = K_{-w} + x$ and $K = K_l + K_{-w} + x$. A direct differentiation yields the identity

$$s''(x) = -\frac{e^{-\tau K_w}}{K_w^3 K^3}\, P(x; K_l, K_{-w}, \tau), \tag{8}$$

where $P(x; K_{-w}, K_l, \tau) = e^{\tau K_w} A(x; K_{-w}, K_l) + B(x; K_{-w}, K_l, \tau)$ with

$$\begin{aligned}
A(x; K_{-w}, K_l) \coloneqq\ & 2K_l^3K_{-w} + 6K_l^2K_{-w}^2 + 6K_l^2K_{-w}x + 6K_lK_{-w}^3 + 12K_lK_{-w}^2x + 6K_lK_{-w}x^2 \\
& + 2K_{-w}^4 + 6K_{-w}^3x + 6K_{-w}^2x^2 + 2K_{-w}x^3, \\
B(x; K_{-w}, K_l, \tau) \coloneqq\ & K_l^3K_{-w}^2\tau^2x - 2K_l^3K_{-w}^2\tau + 2K_l^3K_{-w}\tau^2x^2 - 2K_l^3K_{-w}\tau x - 2K_l^3K_{-w} + K_l^3\tau^2x^3 \\
& + 2K_l^2K_{-w}^3\tau^2x - 4K_l^2K_{-w}^3\tau + 6K_l^2K_{-w}^2\tau^2x^2 - 6K_l^2K_{-w}^2\tau x - 6K_l^2K_{-w}^2 \\
& + 6K_l^2K_{-w}\tau^2x^3 - 6K_l^2K_{-w}x + 2K_l^2\tau^2x^4 + 2K_l^2\tau x^3 + K_lK_{-w}^4\tau^2x \\
& - 2K_lK_{-w}^4\tau + 4K_lK_{-w}^3\tau^2x^2 - 4K_lK_{-w}^3\tau x - 4K_lK_{-w}^3 \\
& + 6K_lK_{-w}^2\tau^2x^3 - 6K_lK_{-w}^2x + 4K_lK_{-w}\tau^2x^4 + 4K_lK_{-w}\tau x^3 \\
& + K_l\tau^2x^5 + 2K_l\tau x^4 + 2K_lx^3.
\end{aligned}$$

Thus $s''(x) \leq 0$ if $P(x; K_{-w}, K_l, \tau) \geq 0$. Each term in has a nonnegative coefficient, so

$$A(x; K_{-w}, K_l) \geq 0.$$

Moreover, for all $t \geq 0$, $e^t \geq 1 + t$ (e.g. by convexity of e^t or by the monotonicity of $e^t - 1 - t$). Applying this with $t = \tau K_w \geq 0$ gives

$$e^{\tau K_w} \geq 1 + \tau K_w.$$

so it follows that

$$P(x; K_{-w}, K_l, \tau) \geq (1 + \tau K_w) A(x; K_{-w}, K_l) + B(x; K_{-w}, K_l, \tau).$$

Define the right-hand side of (**??**) as

$$R(x; K_{-w}, K_l, \tau) := (1 + \tau K_w) \cdot A(x; K_{-w}, K_l) + B(x; K_{-w}, K_l, \tau).$$

We can write $R(x; K_{-w}, K_l, \tau)$ as

$$R(x; K_{-w}, K_l, \tau) = (K_{-w} + x)^2 \, S(x; K_{-w}, K_l, \tau),$$

where

$$\begin{aligned} S(x; K_{-w}, K_l, \tau) := {} & K_l^3 \tau^2 x + 2K_l^2 K_{-w} \tau^2 x + 2K_l^2 K_{-w} \tau + 2K_l^2 \tau^2 x^2 + 2K_l^2 \tau x \\ & + K_l K_{-w}^2 \tau^2 x + 4K_l K_{-w}^2 \tau + 2K_l K_{-w} \tau^2 x^2 + 6K_l K_{-w} \tau x + 2K_l K_{-w} \\ & + K_l \tau^2 x^3 + 2K_l \tau x^2 + 2K_l x + 2K_{-w}^3 \tau + 4K_{-w}^2 \tau x + 2K_{-w}^2 \\ & + 2K_{-w} \tau x^2 + 2K_{-w} x. \end{aligned}$$

Every term in S has a nonnegative coefficient, so $S(x; K_{-w}, K_l, \tau)$ and thus $R(x; K_{-w}, K_l, \tau) \geq 0$ as well as $P(x; K_{-w}, K_l, \tau)$ are nonnegative. Moreover, if $(K_{-w}, K_l) \neq (0, 0)$ and $K_w = K_{-w} + x > 0$, then the inequality is strict, so $u_w^{(i)}$ is concave in $k_w^{(i)}$.

B.2 Global Monotonicity of $K(T)$ when $n_w = 1$

Fix $T \geq 0$. Define $E(T) = \exp\big(-\lambda T\, k_w\big)$ and $K(T) = k_w(T) + n_l \cdot k_l(T)$. The symmetric equilibrium FOCs are then

$$\rho = \left(\frac{n_l \cdot k_l(T)}{K(T)^2} + \lambda T\, \frac{n_l k_l(T)}{K(T)} \right) E(T),$$
$$\rho = \frac{k_w(T) + (n_l - 1) \cdot k_l(T)}{K(T)^2}\, E(T).$$

Suppressing dependence on T, equating the right-hand sides and cancelling $E(T) > 0$ yields

$$\frac{n_l k_l}{K^2} + \lambda T\, \frac{n_l k_l}{K} = \frac{k_w + (n_l - 1) k_l}{K^2} \implies K = (n_l + 1) k_l + \lambda T\, n_l k_l K.$$

Defining $\theta = \lambda T K$ and $D(\theta) = (1 + \theta) n_l + 1$, we obtain

$$k_l = \frac{K}{D(\theta)},$$
$$k_w = \frac{(1 + n_l \theta) K}{D(\theta)}.$$

Substituting the above expression for k_w into the loser's FOC, we have

$$K = k_w + (n_l - 1) k_l = \frac{(1 + n_l \theta) K}{D(\theta)} + \frac{(n_l - 1) K}{D(\theta)} = \frac{n_l (1 + \theta) K}{D(\theta)}.$$

so we obtain a single equation in K and T:

$$H(K, T) \coloneqq \frac{n_l (1 + \theta)}{K \cdot D(\theta)} \exp\left(-\lambda T\, \frac{(1 + n_l \theta) K}{D(\theta)} \right) = \rho,$$

Define $G(K, T) \coloneqq H(K, T) - \rho$. Along the equilibrium path $G(K(T), T) = 0$, the implicit-function theorem yields

$$K'(T) = -\frac{G_T}{G_K} = -\frac{H_T}{H_K}.$$

Thus it suffices to show $H_T < 0$ and $H_K < 0$ for all $T > 0$.

We first show that $H_T < 0$. Note that

$$\log H(K,T) = \log n_l + \log(1+\theta) - \log K - \log D(\theta) - \lambda T k_w.$$

Holding K fixed, we have $\partial_T \theta = \lambda K$ and $\partial_T D(\theta) = n_l \partial_T \theta = \lambda n_l K$. Moreover,

$$\frac{\partial k_w}{\partial \theta} = \frac{\partial}{\partial \theta}\left(\frac{(1+n_l\theta)K}{D(\theta)}\right) = K \cdot \frac{n_l D(\theta) - (1+n_l\theta)n_l}{D(\theta)^2} = K \cdot \frac{n_l^2}{D(\theta)^2},$$

so we have that

$$\frac{\partial k_w}{\partial T} = \frac{\partial k_w}{\partial \theta} \cdot \frac{\partial \theta}{\partial T} = \lambda n_l^2 \frac{K^2}{D(\theta)^2}.$$

Therefore,

$$\begin{aligned}\frac{\partial}{\partial T}\log H(K,T) &= \frac{\lambda K}{1+\theta} - \frac{\lambda n_l K}{D(\theta)} - \lambda k_w - \lambda T \frac{\partial k_w}{\partial T} \\ &= \frac{\lambda K}{1+\theta} - \frac{\lambda n_l K}{D(\theta)} - \lambda \frac{(1+n_l\theta)K}{D(\theta)} - \lambda T\left(\lambda n_l^2 \frac{K^2}{D(\theta)^2}\right).\end{aligned}$$

Dividing by $\lambda K > 0$, we have

$$\frac{1}{\lambda K}\frac{\partial}{\partial T}\log H(K,T) = \frac{1}{1+\theta} - \frac{n_l}{D(\theta)} - \frac{1+n_l\theta}{D(\theta)} - \frac{\theta n_l^2}{D(\theta)^2} = -\frac{\theta}{1+\theta} - \frac{\theta n_l^2}{D(\theta)^2} < 0$$

Hence $\partial_T \log H(K,T) < 0$, so $H_T(K,T) < 0$ for all $T, K > 0$.

We then show that $H_K < 0$, holding T fixed. Note that $\partial_K \theta = \lambda T = \theta/K$ and $\partial_K D(\theta) = n_l \partial_K \theta = n_l \theta / K$. Then

$$\begin{aligned}\frac{\partial}{\partial K}\log H(K,T) &= \frac{1}{1+\theta}\frac{\partial \theta}{\partial K} - \frac{1}{K} - \frac{1}{D(\theta)}\frac{\partial D}{\partial K} - \lambda T \frac{\partial k_w}{\partial K} \\ &= \frac{a/K}{1+\theta} - \frac{1}{K} - \frac{n_l\theta/K}{D(\theta)} - \frac{a}{K}\left(\frac{1+n_l\theta}{D(\theta)} + \frac{a n_l^2}{D(\theta)^2}\right).\end{aligned}$$

Multiplying by $K > 0$, we have

$$K\frac{\partial}{\partial K}\log H(K,T) = \frac{\theta}{1+\theta} - 1 - \frac{n_l\theta}{D(\theta)} - \frac{\theta(1+n_l\theta)}{D(\theta)} - \frac{\theta^2 n_l^2}{D(\theta)^2} = -\frac{1}{1+\theta} - \theta - \frac{\theta^2 n_l^2}{D(\theta)^2} < 0$$

so $H_K(K,T) < 0$.

B.3 Multiple Winners: Winner Copies Greater Than Loser Copies

Write $W = n_w k_w^*$, $L = n_l k_l^*$, and $E = e^{-\tau W}$. Dividing both equations by V and rewriting them in terms of (W, L) yields

$$\frac{C}{V} = \frac{n_w - 1}{n_w W}\left(1 - \frac{L}{W+L}E\right) + \frac{1}{n_w}\left(\frac{L}{(W+L)^2} + \tau\frac{L}{W+L}\right)E,$$

$$\frac{C}{V} = \frac{W + \frac{n_l - 1}{n_l}L}{(W+L)^2}E.$$

Subtracting the loser's FOC from the winner's FOC and multiplying by $n_l n_w W(W+L)^2 > 0$ yields

$$n_l(n_w - 1)(W+L)^2 + E \cdot \Xi(W, L) = 0, \tag{9}$$

where $\Xi(W, L) = -n_l n_w(W+L)^2 + \tau n_l W L(W+L) + 2n_l W L + n_l L^2 + n_w W L$.

Step 3. Use of $E < 1$. For $n_l > 0$, $n_w > 1$, and $W + L > 0$, the first term in above is strictly positive, so $\Xi(W, L) < 0$ and

$$E = \frac{n_l(n_w - 1)(W+L)^2}{-\Xi(W, L)}.$$

Since $\tau > 0$ and $W > 0$, we have $E = e^{-\tau W} \in (0, 1)$, and therefore $E < 1$ implies

$$n_l(n_w - 1)(W+L)^2 < -\Xi(W, L).$$

Expanding Ξ and collecting terms yields

$$0 < W\Big(n_l W - n_w L - \tau n_l L(W+L)\Big).$$

As $W > 0$, this implies $n_l W > n_w L + \tau n_l L(W+L) > n_w L$. Diving by $n_l n_w$ yields $k_l^* < k_w^*$.

B.4 Multiple Winners: Winner Utility Greater Than Loser Utility

Let $K = W + L$ and $E = \exp^{-\tau W}$. Using the equilibrium condition for ρ, we have

$$\begin{aligned} u_w^* - u_l^* &= (1-\alpha_w) - [(1-\alpha_l) - (1-\alpha_w)]\frac{L}{K}E - [(1-\alpha_w)W + (1-\alpha_l)L]\frac{W+\alpha_l L}{K^2}E \\ &= \frac{1}{n_w} - \left(\frac{1}{n_w} + \frac{1}{n_l}\right)\frac{n_l k_l}{K}E - (k_w - k_l)\cdot\frac{K-k_l}{K^2}E \end{aligned}$$

Multiplying by n_w, it suffices to show that

$$1 > \left[\left(1 + \frac{n_w}{n_l}\right)\frac{n_l k_l}{n_l k_l + n_w k_w} - n_w(k_w - k_l)\frac{K-k_l}{K^2}\right]E.$$

Since $E < 1$, it remains to show that the bracketed expression is less than one as well. Note that

$$\left(1 + \frac{n_w}{n_l}\right)\frac{n_l k_l}{n_l k_l + n_w k_w} - n_w(k_w - k_l)\frac{K-k_l}{K^2} = \frac{1}{K^2}[(n_l k_l - n_w k_w + 2n_w k_l)K + n_w k_l(k_w - k_l)]$$

We claim that $(n_l k_l - n_w k_w + 2n_w k_l)K + n_w k_l(k_w - k_l) < K^2$. Dividing by K yields

$$n_l k_l + 2n_w k_l - n_w k_w + n_w k_l(k_w - k_l) < n_l k_l + n_w k_w$$

Collecting terms and simplifying yields

$$\frac{k_l}{K}(k_w - k_l) < 2(k_w - k_l)$$

which holds by construction.

C Robustness Checks for the Empirical Analysis

C.1 Narrower Study Window

We narrow the study window by 15 days on both ends, so that the new sample period is from March 2 to June 2, 2025.

	(1) log RepTXs	(2) log RepTXs	(3) log Revenue	(4) log Revenue
Post×Treated	-0.564**	-0.587**	0.964***	0.924***
	(0.242)	(0.238)	(0.044)	(0.052)
BaseGasL2		0.059		0.104
		(0.050)		(0.067)
VolumeDEX		0.058		0.132
		(0.091)		(0.149)
Observations	465	465	465	465
N. of groups	5	5	5	5
R^2	0.016	0.020	0.060	0.081

Includes chain and day fixed effects. Standard errors are clustered by chain.
Note: *p<0.1; **p<0.05; ***p<0.01

Table A.1: This table reports two-way fixed-effects estimates using a narrower sample window that excludes the first and last 15 days of the original sample period. The sample spans March 2 to June 2, 2025, with Timeboost implemented on April 17, 2025. The dependent variables are the natural logarithm of repeated transactions (columns 1-2) and the natural logarithm of platform revenue in ETH (columns 3-4). Repeated transactions are defined as bursts of identical transactions (same sender, recipient, value, function selector, and calldata) within 2-second windows. Platform revenue includes gas fees from reverted transactions and, for Arbitrum post-Timeboost, auction payments from express lane winners. Post×Treated is an indicator equal to one for Arbitrum observations after April 17, 2025. BaseGasL2 is the Layer-2 base gas fee in ETH. VolumeDEX is the natural logarithm of DEX trading volume in USD. All variables except the treatment indicator are standardized (z-scored) prior to estimation. All specifications include chain and day fixed effects. Standard errors are clustered by chain. This narrower window tests whether results are sensitive to the choice of pre- and post-treatment periods.

	(1) log RepTXs	(2) log RepTXs	(3) log Revenue	(4) log Revenue
Intercept	-11.000*	-11.939*	-3.732	-8.377
	(5.681)	(6.153)	(6.958)	(7.223)
Post	-0.462*	-0.390	1.307***	1.155***
	(0.259)	(0.263)	(0.355)	(0.370)
VolumeDEX	1.090***	1.053***	0.286	0.563
	(0.283)	(0.315)	(0.345)	(0.366)
BaseGasMain	-0.200	-0.264*	0.355**	0.376**
	(0.149)	(0.145)	(0.169)	(0.175)
BaseGasArb		169.113***		-19.931
		(52.415)		(34.822)
VolatilityETH		-0.114		-1.926**
		(0.839)		(0.775)
Observations	123	123	123	123
R^2	0.224	0.255	0.313	0.342
Adjusted R^2	0.205	0.224	0.296	0.314

Standard errors are computed using the Newey-West estimator.
Note: *p<0.1; **p<0.05; ***p<0.01

Table A.2: This table reports estimates from a within-Arbitrum analysis using a wider 5-second window to identify repeated transactions. The sample consists of daily observations on Arbitrum from March 2 to June 2, 2025, with Timeboost implemented on April 17, 2025. The dependent variables are the natural logarithm of repeated transactions (columns 1-2) and the natural logarithm of platform revenue in ETH (columns 3-4). Repeated transactions are defined as bursts of identical transactions (same sender, recipient, value, function selector, and calldata) within 5-second windows, compared to the 2-second baseline. Platform revenue includes gas fees from reverted transactions and, post-Timeboost, auction payments from express lane winners. Post is an indicator equal to one for observations after April 17, 2025. VolumeDEX is the natural logarithm of DEX trading volume in USD on Arbitrum. BaseGasMain is the Ethereum mainnet base gas fee in ETH. BaseGasArb is the Arbitrum base gas fee in ETH. VolatilityETH is the daily return volatility of ETH/USD. Variables are not standardized in this specification. Standard errors are computed using the Newey-West estimator to account for heteroskedasticity and autocorrelation. This specification tests whether results are sensitive to the choice of time window for identifying repeated transactions.

C.2 Wider Interval for Repeated Transactions

We widen the interval under which we search for repeated transactions from 2 seconds to 5 seconds.

	(1) log RepTXs	(2) log RepTXs	(3) log Revenue	(4) log Revenue
Post×Treated	-0.671***	-0.673***	0.772***	0.724***
	(0.191)	(0.159)	(0.195)	(0.135)
BaseGasL2		0.145***		0.403***
		(0.044)		(0.115)
VolumeDEX		0.207*		0.139
		(0.108)		(0.128)
Observations	605	605	605	605
N. of groups	5	5	5	5
R^2	0.022	0.056	0.033	0.204

Includes chain and day fixed effects. Standard errors are clustered by chain.
Note: *p<0.1; **p<0.05; ***p<0.01

Table A.3: This table reports two-way fixed-effects estimates using a wider time interval to detect repeated identical transactions. The sample period spans February 17 to June 17, 2025, with Timeboost implemented on Arbitrum on April 17, 2025. Control Layer-2 blockchains include Optimism, Base, Polygon, and Avalanche. The dependent variables are the natural logarithm of repeated transactions (columns 1-2) and the natural logarithm of platform revenue in ETH (columns 3-4). Repeated transactions are defined as bursts of identical transactions (same sender, recipient, value, function selector, and calldata) within 2-second windows. Platform revenue includes gas fees from reverted transactions and, for Arbitrum post-Timeboost, auction payments from express lane winners. Post×Treated is an indicator equal to one for Arbitrum observations after April 17, 2025. BaseGasL2 is the Layer-2 base gas fee in ETH. VolumeDEX is the natural logarithm of DEX trading volume in USD. All variables except the treatment indicator are standardized (z-scored) prior to estimation. All specifications include chain and day fixed effects. This specification addresses concerns that cluster-robust standard errors may be biased when the number of clusters is small (five chains).

	(1) log RepTXs	(2) log RepTXs	(3) log Revenue	(4) log Revenue
Intercept	-9.453	-9.028	0.566	-6.362
	(5.784)	(5.859)	(7.011)	(7.060)
Post	-0.542**	-0.416*	0.783**	0.527
	(0.228)	(0.228)	(0.344)	(0.350)
VolumeDEX	1.015***	0.894***	0.084	0.529
	(0.288)	(0.299)	(0.348)	(0.360)
BaseGasMain	-0.062	-0.159	0.608***	0.663***
	(0.152)	(0.149)	(0.182)	(0.191)
BaseGasArb		190.117***		-86.322*
		(52.041)		(50.565)
VolatilityETH		0.243		-3.182***
		(0.723)		(0.830)
Observations	121	121	121	121
R^2	0.262	0.315	0.229	0.313
Adjusted R^2	0.243	0.285	0.209	0.283

Standard errors are computed using the Newey-West estimator.
Note: *p<0.1; **p<0.05; ***p<0.01

Table A.4: This table reports two-way fixed-effects estimates of the impact of Timeboost adoption on Arbitrum relative to control Layer-2 blockchains. The sample period spans February 17 to June 17, 2025, with Timeboost implemented on April 17, 2025. The dependent variables are the natural logarithm of repeated transactions (columns 1-2) and the natural logarithm of platform revenue in ETH (columns 3-4). Repeated transactions are defined as bursts of identical transactions (same sender, recipient, value, function selector, and calldata) within 2-second windows. Platform revenue includes gas fees from reverted transactions and, for Arbitrum post-Timeboost, auction payments from express lane winners. Post×Treated is an indicator equal to one for Arbitrum observations after April 17, 2025. BaseGasL2 is the Layer-2 base gas fee in ETH. VolumeDEX is the natural logarithm of DEX trading volume in USD.

C.3 Heteroskedasticity and Autocorrelation Robust Standard Errors

	(1)	(2)	(3)	(4)
	log RepTXs	log RepTXs	log Revenue	log Revenue
Post×Treated	-0.699***	-0.698***	0.760***	0.745***
	(0.172)	(0.170)	(0.174)	(0.179)
BaseGasL2		0.096**		0.142***
		(0.041)		(0.043)
VolumeDEX		0.159**		0.074
		(0.074)		(0.061)
Observations	605	605	605	605
N. of groups	5	5	5	5
R^2	0.025	0.041	0.037	0.063

Includes chain and day fixed effects. Standard errors are HC3 robust.
Note: *p<0.1; **p<0.05; ***p<0.01

Table A.5: This table reports two-way fixed-effects estimates using heteroskedasticity-robust (HC3) standard errors instead of cluster-robust standard errors. The treatment group is Arbitrum after April 17, 2025, and the control group consists of other Layer-2 blockchains. The dependent variables are the natural logarithm of repeated transactions (columns 1–2) and the natural logarithm of platform revenue in ETH (columns 3–4). Repeated transactions are defined as bursts of identical transactions within 2-second windows. Platform revenue includes gas fees from reverted transactions and auction payments from express lane winners. BaseGasL2 is the Layer-2 base gas fee in ETH and VolumeDEX is the natural logarithm of DEX trading volume in USD. All specifications include chain and day fixed effects and use HC3 robust standard errors.

C.4 Excluding Avalanche (C-Chain) Transactions

	(1)	(2)	(3)	(4)
	log RepTXs	log RepTXs	log Revenue	log Revenue
Post×Treated	-0.60018*	-0.63922*	0.87571***	0.83390***
	(0.36374)	(0.34658)	(0.12181)	(0.11920)
BaseGasL2		0.14103**		0.10714*
		(0.06402)		(0.06054)
VolumeDEX		-0.00456		-0.10126
		(0.04273)		(0.10770)
Observations	484	484	484	484
N. of groups	4	4	4	4
R^2	0.02295	0.04007	0.06205	0.07805

Includes chain and day fixed effects. Standard errors are clustered by chain.
Note: *p<0.1; **p<0.05; ***p<0.01

Table A.6: This table reports two-way fixed-effects estimates using cluster-robust standard errors at the chain level. The dependent variables are the natural logarithm of repeated transactions (columns 1–2) and the natural logarithm of platform revenue (columns 3–4). BaseGasL2 denotes the standardized Layer-2 base gas fee and VolumeDEX denotes standardized DEX trading volume. All specifications include chain and day fixed effects.

	(1) log RepTXs	(2) log RepTXs	(3) log Revenue	(4) log Revenue
Post×Treated	-0.60018***	-0.63922***	0.87571***	0.83390***
	(0.19891)	(0.19871)	(0.19065)	(0.19351)
BaseGasL2		0.14103***		0.10714**
		(0.04668)		(0.05129)
VolumeDEX		-0.00456		-0.10126*
		(0.07507)		(0.05948)
Observations	484	484	484	484
N. of groups	4	4	4	4
R^2	0.02295	0.04007	0.06205	0.07805

Includes chain and day fixed effects. Standard errors are HC3-robust.
Note: *p<0.1; **p<0.05; ***p<0.01

Table A.7: This table reports two-way fixed-effects estimates using heteroskedasticity-robust (HC3) standard errors. The dependent variables are the natural logarithm of repeated transactions (columns 1–2) and the natural logarithm of platform revenue (columns 3–4). BaseGasL2 denotes the standardized Layer-2 base gas fee and VolumeDEX denotes standardized DEX trading volume. All specifications include chain and day fixed effects.

C.5 Prior Treatment Date as Placebo

For a properly identified setup, treatment effects should attenuate or disappear when designating a date before the actual treatment date. We move the "treatment date" 15 days prior to April 2, 2025.

	(1) log RepTXs	(2) log RepTXs	(3) log Revenue	(4) log Revenue
Post×Treated	-0.318	-0.347	0.213	0.158
	(0.314)	(0.294)	(0.144)	(0.131)
BaseGasL2		0.097		0.153**
		(0.071)		(0.069)
VolumeDEX		0.167		0.069
		(0.121)		(0.160)
Observations	605	605	605	605
N. of groups	5	5	5	5
R^2	0.005	0.023	0.003	0.033

Includes chain and day fixed effects. Standard errors are clustered by chain.
Note: *p<0.1; **p<0.05; ***p<0.01

Table A.8: This table reports two-way fixed-effects estimates using a placebo treatment date of April 2, 2025—15 days prior to the actual Timeboost implementation on April 17, 2025. The sample period spans February 17 to June 17, 2025. The dependent variables are the natural logarithm of repeated transactions (columns 1-2) and the natural logarithm of platform revenue in ETH (columns 3-4). Repeated transactions are defined as bursts of identical transactions (same sender, recipient, value, function selector, and calldata) within 2-second windows. Platform revenue includes gas fees from reverted transactions and, for Arbitrum after the placebo date, hypothetical auction payments. Post×Treated is an indicator equal to one for Arbitrum observations after April 2, 2025. BaseGasL2 is the Layer-2 base gas fee in ETH. VolumeDEX is the natural logarithm of DEX trading volume in USD. All variables except the treatment indicator are standardized (z-scored) prior to estimation. All specifications include chain and day fixed effects. Standard errors are clustered by chain. Under the parallel trends assumption, we should not observe significant treatment effects when using a pre-treatment placebo date.

Event Time	Coefficient	Std. Error	t statistic	p-value
D_{-5}	-0.0259	0.2538	-0.1021	0.9187
D_{-4}	-0.3175	0.4190	-0.7579	0.4489
D_{-3}	0.0150	0.2695	0.0556	0.9557
D_{-2}	-0.1095	0.6163	-0.1777	0.8591
D_{-1}	-0.2743	0.3229	-0.8493	0.3961
D_{+1}	0.0390	0.2210	0.1765	0.8600
D_{+2}	0.1589	0.1998	0.7955	0.4268
D_{+3}	0.1165	0.2157	0.5403	0.5892
D_{+4}	0.9690	0.2478	3.9100	0.0001
D_{+5}	0.0676	0.1465	0.4616	0.6446

Table A.9: Lead Test Results — Two-Way Fixed Effects. This table reports coefficients from a lead test of $\log$ RepTXs of event indicators D_k in a two-way fixed effects regression. Each D_k is a dummy variable equal to 1 for observations k periods away from the event (negative values indicating leads and positive values indicating lags), and 0 otherwise. Coefficient estimates, standard errors, t-statistics, and p-values are shown.