

Chicago Booth School of Business Economics Research Papers

Smoothing the Cliff: Fair and Incentive-Compatible Priority Allocation via Randomized Mechanisms

Tingyi Lin
Central University of Finance and Economics

Weiyu Qi
The University of Chicago

Karl Yu
Boston College

Huanxi Zhang
University of Wisconsin-Madison

All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission, provided that full credit including © notice is given to the source.

This paper also can be downloaded without charge from the
Chicago Booth Microeconomics Working Paper Series
https://ssrn.com/abstract=6148512

Smoothing the Cliff: Fair and Incentive-Compatible Priority Allocation via Randomized Mechanisms*

Tingyi Lin† Weiyu Qi‡ Karl Yu§ Huanxi Zhang¶

Abstract

Can an allocation algorithm be designed to not only distribute resources efficiently but also generate reliable economic statistics? As AI and high-frequency data reshape financial infrastructure, traditional "winner-takes-all" rules are failing—they generate noisy, unstable data that undermines auditability and policy monitoring. To solve this, we propose FIRM-L (Fair and Incentive-compatible Randomized Mechanism). Think of it as a mechanism with a "volume knob" (temperature control) that turns chaotic cliff-edge competition into a smooth, predictable curve. We mathematically prove that FIRM-L ensures fairness (via Lipschitz stability) and honesty (truthfulness-in-expectation) without sacrificing speed. Furthermore, we design an AI-driven controller that automatically tunes this mechanism in real-time, adapting to market shocks to protect welfare while keeping the market stable and auditable.

JEL Classification: D44, D47, C73, D63

Keywords: Mechanism Design, Priority Allocation, Transaction Ordering, Economic Measurement, Algorithmic Game Theory, Online Learning

*We are deeply grateful to Dr. Chaofan Hao and Dr. Ke Dong (Tsinghua University PBC School of Finance), as well as other faculty members, for pre-reviewing this manuscript and providing constructive comments.

We extend our gratitude to Ruoran Lai (Sun Yat-sen University) for stimulating discussions and assistance during the writing process.

We also thank the International Conference on Machine Learning (ICML) reviewers and discussants for their valuable feedback. All remaining errors are our own.

Funding and Disclosure: This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. The author declares that has no relevant or material financial interests that relate to the research described in this paper.

†Central University of Finance and Economics. Email: hayashi@email.cufe.edu.cn**(Corresponding author)**

‡The University of Chicago. Email: weiyuq@uchicago.edu

§Boston College. Email: yulin@bc.edu

¶University of Wisconsin-Madison. Email: hzhang2488@wisc.edu

"Natura non facit saltus."

— Alfred Marshall, Principles of Economics (1890)

1 Introduction

Priority based allocation deciding who is served first, which transaction executes earlier, and whose request occupies scarce capacity has become a core design primitive in modern digital markets. In blockchain systems in particular, transaction ordering is not a neutral implementation detail: it directly affects execution outcomes, realized trading prices, and the distribution of rents associated with transaction-ordering dependencies (commonly discussed under the umbrella of miner/extractor value or MEV) (Daian et al., 2020). A salient empirical manifestation of this phenomenon is the emergence of fee-driven competition for early block position—often described as "priority gas auctions" in the context of decentralized exchanges (DEXs). In these settings, sophisticated actors pay high fees and optimize latency to obtain execution priority, and the resulting arms race can distort market quality and create broader system-level incentives and risks.

A striking and economically consequential feature of many deployed priority rules is their cliff-like discontinuity. Under strict deterministic ranking-by-bid, infinitesimal differences in a priority signal can induce macroscopic changes in rank, execution probability, and payments, yielding brittle outcomes that are hard to predict and difficult to justify to participants. This discontinuity is central to the priority gas auction dynamics documented in blockchain markets, where repeated bidding for marginal priority advantages can amplify volatility and strategic expenditure.

These discontinuities matter for at least three reasons. First, they create within-round volatility: tiny perturbations in bids or priority signals can flip allocation outcomes, producing sharp swings in realized utility for agents with nearly identical valuations. Second, they foster rent-seeking and wasteful escalation—costly bid shading and overbidding near cutoffs—consistent with the competitive priority dynamics observed in blockchain transaction ordering. Third, they undermine the stability and auditability of high-frequency

statistics derived from the allocation process itself such as measures of congestion pressure, effective execution costs, or "activity intensity", precisely when such measures are most valuable for monitoring and policy.

This paper proposes a principled approach to smoothing the cliff. We introduce **FIRM-L**, a reserve-and-temperature parameterized family of randomized priority allocation mechanisms that replaces hard sorting with a smooth probabilistic ranking rule. The mechanism is built on the Plackett–Luce (PL) model, a canonical distribution over rankings grounded in Luce's choice axiom and Plackett's classic permutation model (Luce, 1959; Plackett, 1975). Concretely, the designer chooses a reserve threshold $r \geq 0$ and a temperature parameter $\tau > 0$. Agents who bid below r are ineligible, while eligible agents receive scores that are increasing in their bids and scaled by τ. Priority positions are then assigned by sampling without replacement according to the PL rule: at each step, the next slot is allocated with probability proportional to the exponential of each remaining agent's score. The temperature τ acts as a continuous governance knob: as $\tau \to 0$, the mechanism approaches strict priority; as $\tau \to \infty$, it approaches a uniform lottery among eligible agents.

Randomization is only valuable if incentives remain well behaved. Our environment is single-parameter and quasi-linear, so implementability depends on monotonicity of each agent's interim allocation in her own bid, with truthful payments characterized by the Myerson identity. These results are foundational in mechanism design and have been formalized in classic work on optimal auctions and truthful mechanisms for one-parameter agents (Myerson, 1981; Archer and Tardos, 2001). A key theoretical contribution of this paper is to show that, despite sequential randomization, the PL top-K allocation induces an interim expected priority weight that is non-decreasing in an agent's bid. The argument uses the exponential-race representation of PL rankings: increasing an agent's bid increases her exponential rate, stochastically improving her arrival time and therefore her expected rank under a coupling.

Given monotonicity, expected transfers are characterized by the Myerson identity. For deployment, we distinguish the monotone base allocation rule from the transfer-

computation wrapper. Our stability results refer to the base FIRM-L allocation map. To implement truthful transfers with low online overhead, we use a single-call self-resampling reduction in the sense of Babaioff et al. (2015); Appendix B states the wrapper, the associated audit log, and the relation between allocation-side and payment-side objects.

Our main conceptual departure from standard priority-auction analysis is to elevate within-round stability to a first-class design objective, alongside welfare and revenue. Motivated by the "treat similar individuals similarly" paradigm in algorithmic fairness (Dwork et al., 2012), we formalize a notion of fairness tailored to priority markets: *within-agent opportunity stability*. Specifically, we define **Eligible-Region Stability** as a Lipschitz continuity requirement on each agent's interim expected allocation with respect to small changes in her own bid, conditional on being eligible (bidding above the reserve). This criterion directly targets the discontinuity of strict priority rules and provides an auditable metric of sensitivity. The temperature parameter τ provides explicit control over stability. We derive analytic sensitivity bounds showing that the Lipschitz constant scales inversely with τ: increasing τ smooths the allocation map and reduces marginal sensitivity, yielding a transparent stability–efficiency frontier. This complements the interpretability advantage of PL smoothing: unlike ad hoc random tie-breaking, the mechanism's sensitivity to bid perturbations is explicitly parameterized and can be reported, audited, and tuned.

These stability guarantees also support the construction of congestion-adjusted monitoring indices from interim expected allocations and expected payments. The relevant statistical objects are therefore expectation-level quantities computed from the announced policy and the logged bid profile, not one-shot realized outcomes. This distinction is important for both auditability and empirical validation.

A practical challenge is that the appropriate amount of smoothing depends on market conditions. We therefore study dynamic parameter selection under a pre-commitment model in which the platform publicly announces its parameters before observing contemporaneous bids. Our formal benchmark is a full-information online-learning problem over a discretized policy class, analyzed with multiplicative weights (Hedge) and evaluated

relative to the best fixed policy in hindsight. Extensions to policy-sensitive environments with partial feedback are discussed separately as bandit extensions rather than as part of the main theorem.

1.1 Summary of Contributions

This paper makes four primary contributions:

1. **Mechanism: Smoothed Priority Allocation.** We propose FIRM-L, a reserve-and-temperature parameterized randomized ranking mechanism based on Plackett–Luce top-K sampling, providing a continuous interpolation between strict priority and lottery.

2. **Incentives with a mechanism-design monotonicity result.** We show that reserve-screened, weighted top-K Plackett–Luce sampling induces an interim expected priority weight that is monotone in an agent's own bid. This is the mechanism-design object needed for Myerson implementation, and is stronger than a generic statement that higher scores improve expected rank. We then implement truthful transfers in expectation via a single-call wrapper around the monotone base rule.

3. **Auditable Fairness via Stability.** We formalize within-round fairness as eligible-region Lipschitz stability of the bid-to-allocation mapping, derive an explicit $O(1/\tau)$ certificate for the without-replacement PL process, and show how this certificate supports expectation-level auditing of derived monitoring indices.

4. **Learning-Augmented Governance.** We model parameter selection as a pre-committed online-learning problem and analyze a multiplicative-weights controller under a full-information benchmark, obtaining regret bounds relative to the best fixed policy in hindsight from a finite policy menu.

1.2 Roadmap

The remainder of the paper proceeds as follows. Sections 2 and 3 review the institutional context of blockchain priority markets and related literature. Section 4 formalizes the

model and planner objectives, while Section 5 defines the FIRM-L mechanism. Section 6 establishes incentive compatibility and details the single-call payment implementation. Section 7 derives the theoretical bounds for eligible-region stability. Section 8 introduces congestion-adjusted high-frequency indices derived from the mechanism's outputs to monitor market pressure. Section 9 presents the learning-augmented governance framework for parameter adaptation. Section 10 reports quantitative evaluations under bursty demand, followed by a discussion of policy implications in Section 11 and conclusions in Section 12.

2 Institutional Background: Blockchain Transaction Ordering & Priority Markets

2.1 Transaction Ordering as a Market-Design Problem

Blockchains transform what is usually a backend systems choice—how to order and include requests—into an explicit economic mechanism. In smart-contract platforms, transaction ordering determines payoffs whenever execution is order-sensitive (e.g., decentralized exchange trades, liquidations, and arbitrage). The resulting rents are not incidental but systemic: bots compete for earlier execution and pay higher fees to obtain priority, a phenomenon documented and formalized in the security literature as a form of high-frequency priority competition (Daian et al., 2020). From an economics/finance perspective, this environment can be interpreted as a high-frequency queue with endogenous priority prices. When demand spikes, the shadow price of early execution rises, and observed fees and priority bids become a revealed measure of scarcity and short-horizon opportunity. Large-scale evidence of pervasive frontrunning underscores that priority competition is continuous and strategic rather than occasional: attackers exploit the transaction pool by displacement, insertion, and suppression strategies .

2.2 From Mempool to Block: Implementing Priority in Practice

Most permissionless blockchains maintain a pool of pending transactions ("mempool") from which block producers select and order transactions. Historically, Ethereum's fee market closely resembled a pay-as-bid design in which users attach fees and miners/validators choose a subset and ordering to maximize revenue. This decentralized implementation introduces information heterogeneity and discretionary ordering, creating an adversarial arena in which ordering power can be systematically exploited (Daian et al., 2020; Ferreira Torres et al., 2021). A salient microstructure feature is the priority gas auction (PGA) dynamic: for fleeting opportunities, bots repeatedly revise fees to outbid competitors and obtain earlier block position, creating cliff-like discontinuities under strict priority sorting. These discontinuities imply that very small fee changes can produce discrete jumps in allocation (position), payments, and outcomes, which becomes a central design concern when the goal is to build stable and auditable measurement outputs from the mechanism itself Microstructure parallel. The priority gas auction dynamic closely parallels the "latency-arbitrage" mechanism studied in modern electronic markets: when a market design rewards infinitesimal priority advantages, competition can shift from economically meaningful price improvement to socially wasteful races for marginal speed and priority (Budish et al., 2015; Aquilina et al., 2022).. Market-design responses that reduce the value of tiny timing advantages have been proposed and empirically motivated in traditional markets; we pursue an analogous goal in priority allocation via a continuous smoothing parameter that explicitly bounds sensitivity and supports audit.

2.3 MEV/BEV: The Economic Content of Ordering Power

Maximal (miner/extractor) extractable value (MEV) refers to profits enabled by control over inclusion and ordering. Quantitative studies of blockchain extractable value document large aggregate profits and extreme tail events, with some episodes far exceeding normal block rewards (Qin et al., 2022). Such tail behavior makes strict priority mechanisms fragile and difficult to govern, especially in congested regimes where strategic entry and bidding escalation intensify . Private orderflow and private pools emerged partly to

mitigate negative externalities such as sandwich attacks. However, evidence suggests that these infrastructures can shift value extraction and may reinforce centralization, rather than eliminating ordering rents (Weintraub et al., 2022). This institutional reality matters for mechanism design: the "platform" is not a single designer but an ecosystem of intermediaries whose incentives shape the effective allocation rule faced by users .

2.4 EIP-1559 and the Modern Fee Market

Ethereum's EIP-1559 introduced a burned base fee and a separate priority fee (tip), alongside variable block sizes around a target. Formal work frames transaction fee mechanisms as mechanism-design objects with blockchain-specific incentive constraints and studies EIP-1559's guarantees and failure modes (Roughgarden, 2020, 2021). Empirical analysis suggests EIP-1559 improves certain aspects of user experience (e.g., intrablock fee dispersion and waiting time) but does not eliminate fee volatility or MEV-related distortions (Liu et al., 2022). For our purposes, the key implication is: even with improved fee predictability, strict priority ordering and order-sensitive strategies persist, and the market remains highly non-stationary with congestion bursts and regime changes. That motivates mechanisms that explicitly trade off efficiency against stability and provide parameters that can be tuned and audited under stress.

2.5 PBS / MEV-Boost: An Auction for Block Building

With proof of stake, Ethereum increasingly relies on proposer–builder separation (PBS), deployed in practice through MEV-Boost. PBS decouples the proposer (validator) from specialized builders and relays, creating a second-level market in which builders bid for the right to provide an optimized block (Bahrani et al., 2025). Measurement evidence highlights centralization among builders and relays and raises concerns about whether PBS fully achieves decentralization and censorship-resistance goals (Heimbach et al., 2023). This institutional context reinforces two design constraints: (i) allocation rules must be implementable with low overhead under tight latency constraints, and (ii) governance levers must be transparent and auditable even when the effective ordering is intermediated.

These constraints motivate parameterized randomized mechanisms that can be tuned and evaluated statistically.

2.6 Measurement and Governance Motivation

A central measurement insight is that priority markets produce rich high-frequency data—bids, realized inclusion/position, payments, and execution outcomes,yet strict priority rules can make derived statistics unstable because small input perturbations can flip discrete outcomes. The measurement agenda explicitly highlights transparency and auditability for algorithmically produced statistics, aligning with mechanisms whose sensitivity can be bounded and tuned. More broadly, official-statistics discussions note that generative and algorithmic systems can challenge reproducibility and traceability, motivating explicit monitoring and audit frameworks (UNECE, 2025). This perspective supports interpreting a parameterized allocation rule as a "statistical generator" whose sensitivity properties can be audited and used in policy and risk monitoring.

3 Related Literature

3.1 Ranking & PL Models: Luce/Plackett–Luce

Our mechanism uses probabilistic ranking as an allocation primitive. The Plackett–Luce (PL) model, derived from Luce's choice axiom and Plackett's distribution over permutations, represents rankings as a sequential choice process: each position is sampled from the remaining items with probability proportional to latent "worth" (Luce, 1959; Plackett, 1975). This structure makes PL a natural device for top-K priority assignment without replacement. On the algorithmic side, perturb-and-max and Gumbel-top-k representations yield exact sampling procedures and useful couplings (Kool et al., 2019, 2020); on the statistical side, scalable inference methods support large ranking problems (Guiver and Snelson, 2009; Maystre and Grossglauser, 2015).

What is novel here is not the classical PL model itself, nor the general observation that higher scores improve rank in a stochastic-order sense. The mechanism-design question

we need is narrower and, to our knowledge, not stated in this form in the prior literature: whether a reserve-screened, weighted top-K PL rule induces an *interim expected priority weight* that is monotone in an agent's own bid in a one-parameter environment, so that the Myerson envelope applies. Our coupling argument is tailored to that object. It must simultaneously handle (i) reserve eligibility, (ii) without-replacement top-K selection, and (iii) heterogeneous slot weights $w_1 \geq \cdots \geq w_K$, rather than merely rank comparison or score monotonicity in the abstract. This is the specific sense in which we import PL into mechanism design: as a tractable randomized allocation rule whose bid-monotonicity can be established for the weighted interim object that determines truthful implementation.

3.2 Fairness & Online Learning: Stability and Governance

Priority markets invite objectives beyond welfare and revenue. We focus on within-round sensitivity: whether small changes in a bidder's own report can generate disproportionate changes in expected priority weight. This is related in spirit to individual-fairness-style Lipschitz ideas (Dwork et al., 2012), but adapted to the institutional problem created by strict priority sorting in high-frequency markets. Our contribution is not a general distributive theory of fairness across heterogeneous agents. Rather, it is a procedural stability notion that can be certified, logged, and audited in a mechanism with explicit policy parameters.

This emphasis connects naturally to measurement and auditability. In digital markets, allocation rules do not merely determine who wins; they also generate operational statistics used for monitoring congestion, execution risk, and tail behavior. Work on official statistics and algorithmic auditing stresses the importance of reproducibility, traceability, and sensitivity control for such systems (UNECE, 2025; Li and Goel, 2024). Our use of stability certificates sits in that tradition: the mechanism is designed so that certain expectation-level measurement outputs inherit formal sensitivity bounds.

A separate issue is policy adaptation. Blockchain demand and extractable-value opportunities vary sharply across time, which motivates a governance layer that chooses (r, τ) ex ante. We use Hedge/multiplicative weights as the canonical full-information

benchmark (Freund and Schapire, 1997; Cesa-Bianchi and Lugosi, 2006). The formal comparator is the best *fixed* policy in hindsight from a finite announced menu. This is an auditable benchmark because the menu, loss functional, and update rule can all be pre-committed and logged. It should not be confused with a dynamic-regret guarantee against a time-varying oracle. When participation or the type distribution is policy-sensitive, or when only realized losses are observed, the appropriate extensions are partial-information and switching-comparator problems, for which adversarial bandit methods provide the relevant toolkit (Auer et al., 2002; Bubeck and Cesa-Bianchi, 2012).

3.3 Economic and Financial-Market Foundations: Priority, Congestion, and Market Design

Our priority market can be interpreted as a capacity-constrained service system in which earlier execution is a scarce quality attribute. This view connects our setting to the economics and operations literature on congestion pricing and priority auctions in queueing systems (Afèche and Mendelson, 2004), which studies how admission rules and priority pricing trade off welfare, revenue, and delay externalities .

A complementary connection is to financial market microstructure (Budish et al., 2015). In modern electronic markets, deterministic priority rules and high-speed technology can create mechanical arbitrage opportunities and induce socially wasteful races for infinitesimal priority advantages. Market-design responses that reduce the value of tiny timing advantages have been proposed and empirically motivated as a way to curb arms-race incentives and improve market quality .

Finally, the microstructure perspective is increasingly relevant for crypto markets themselves. High-frequency, rule-generated data (fees, inclusion, ordering) can be interpreted through microstructure lenses, motivating the use of robust indices and audit-ready statistics (O'Hara, 2015).

4 The Model

4.1 Players, Information, and Timing

We consider a discrete-time priority market over rounds $t = 1, 2, \ldots$. Each round represents a discrete allocation window, such as a blockchain block or a batch auction interval. In each round t, a finite set of risk-neutral agents $N_t = \{1, \ldots, n_t\}$ competes for scarce priority execution. To preclude ex-post discrimination based on realized bids, we impose a *pre-commitment* timing structure: the platform announces policy parameters θ_t (formally defined in Section 5) prior to the submission of bids. This assumption aligns with standard mechanism design practice and reflects the institutional reality of protocol-governed ordering rules, where parameters are often encoded in open-source software or governance proposals (Roughgarden, 2021; Heimbach et al., 2023). Each agent i possesses a private type $v_{i,t} \in \mathcal{V} \subseteq \mathbb{R}_+$, representing their scalar willingness to pay for "priority-weighted execution." Agents submit bids $b_{i,t} \in \mathcal{B} \subseteq \mathbb{R}_+$ via a direct mechanism. While the platform may update θ_t across rounds using online learning techniques (Freund and Schapire, 1997; Auer et al., 2002), our incentive analysis focuses on the one-shot game within each round, targeting dominant strategies or truthfulness-in-expectation.

4.2 Scarce Slots and Priority Weights

The platform allocates up to K priority positions (slots), indexed by $p \in \{1, \ldots, K\}$. Each slot confers a deterministic, non-negative priority weight w_p, with strictly decreasing importance: $w_1 \geq w_2 \geq \cdots \geq w_K \geq 0$. This structure captures the diminishing marginal value of earlier execution; for instance, the first position in a block allows for the highest-value arbitrage (Daian et al., 2020; Qin et al., 2022). A feasible outcome is a partial matching between agents and slots. Let $\mathcal{Y}_t$ denote the set of feasible assignment matrices $y \in \{0, 1\}^{n_t \times K}$ satisfying the canonical constraints for position auctions:

$$\sum_{i \in N_t} y_{i,p} \leq 1 \quad \forall p, \quad \text{and} \quad \sum_{p=1}^{K} y_{i,p} \leq 1 \quad \forall i. \tag{1}$$

Agents unassigned to any slot $p \in \{1, \ldots, K\}$ are assumed to receive a reserve weight $w_{K+1} = 0$. This slot-weight formulation also corresponds to differentiated service in congestion/queueing models, where priority can be allocated and priced via priority auctions with admission controls (Afèche and Mendelson, 2004).

4.3 Preferences and Quasi-Linear Utility

Given a realized assignment y and a monetary transfer π_i, agent i's utility is quasi-linear:

$$u_i(v_{i,t}, y, \pi_i) = v_{i,t} \cdot \left(\sum_{p=1}^{K} w_p y_{i,p} \right) - \pi_i. \tag{2}$$

This formulation places us in a standard *single-parameter* environment (Myerson, 1981; Archer and Tardos, 2001). The agent's private type $v_{i,t}$ scales a one-dimensional allocation statistic, $X_i(y) := \sum_p w_p y_{i,p}$, which represents the total priority weight received. In the context of blockchain transaction ordering, $v_{i,t}$ proxies for the marginal surplus extractable from execution priority, such as arbitrage profits or liquidation rewards, which are often heavy-tailed and highly time-sensitive.

4.4 Mechanisms as Stochastic Kernels

A mechanism in round t is a tuple (A_t, P_t). The allocation rule $A_t : \mathcal{B}^{n_t} \to \Delta(\mathcal{Y}_t)$ is a stochastic kernel mapping bid profiles to a probability distribution over feasible assignments. The payment rule P_t determines the transfer $\pi_i(b, \omega) \in \mathbb{R}$, which may depend on the mechanism's internal randomness ω. The use of internal randomization is central to our design, as it allows for the smoothing of discrete priority cliffs while maintaining implementability (Babaioff et al., 2015; Dwork et al., 2012).

4.5 Interim Allocation and Implementability

Fix the bids of opponents $\mathbf{b}_{-i}$. We define the *interim expected priority weight* for agent i as:

$$x_i(b_i; \mathbf{b}_{-i}) := \mathbb{E}_{y \sim A_t(b_i, \mathbf{b}_{-i})}[X_i(y)] = \sum_{p=1}^{K} w_p \cdot q_{i,p}(b), \tag{3}$$

where $q_{i,p}(b) := \Pr(y_{i,p} = 1 \mid b)$ is the probability that agent i is assigned to slot p. In single-parameter domains, a necessary and sufficient condition for Incentive Compatibility in Expectation (IC-E) is the monotonicity of the interim allocation $x_i(b_i; \mathbf{b}_{-i})$ with respect to b_i. If this condition holds, the expected payments are uniquely determined (up to an additive constant) by the Myerson payment identity.

4.6 The Planner's Objective

We posit a social planner who evaluates outcomes based on a scalarized loss function $\mathcal{L}_t(\theta_t)$, incorporating efficiency, stability, and risk.

Welfare. The primary objective is the expected social welfare, defined as the sum of valuations weighted by assignment probabilities:

$$W_t(\theta) := \sum_{i \in N_t} v_{i,t} x_i(b_t; \theta). \tag{4}$$

Stability as Auditability. Motivated by the "individual fairness" paradigm of Dwork et al. (2012), we operationalize stability as a Lipschitz constraint on the allocation rule. Let r be an eligibility threshold. We define the *Eligible-Region Sensitivity* S_t as the Lipschitz constant of the interim allocation map for eligible bids:

$$S_t(\theta) := \sup_{i, \mathbf{b}_{-i}} \sup_{b, b' \geq r, b \neq b'} \frac{|x_i(b; \mathbf{b}_{-i}) - x_i(b'; \mathbf{b}_{-i})|}{|b - b'|}. \tag{5}$$

This metric provides an auditable bound on the "cliff" behavior of the mechanism.

Volatility. To account for the tail risks inherent in MEV markets such as payment spikes and allocation turnover, we include a volatility penalty $V_t(\theta)$, which may capture the variance of payments or allocation weights. Combining these terms, the planner seeks to minimize the per-round loss:

$$\ell_t(\theta) = -W_t(\theta) + \lambda S_t(\theta) + \mu V_t(\theta), \tag{6}$$

where $\lambda, \mu \geq 0$ are Lagrange multipliers governing the trade-off between efficiency, stability, and risk. This functional form allows us to treat the dynamic parameter selection problem in Section 9 as an instance of online regret minimization.

5 Mechanism: FIRM-L

We introduce **FIRM-L**, a randomized priority mechanism parameterized by a tuple $\theta = (r, \tau)$.

5.1 Parameters and Eligibility

The mechanism is governed by a reserve price $r \geq 0$ and a temperature parameter $\tau > 0$. Given a bid profile $\mathbf{b}$, the set of eligible agents is defined as $\mathcal{E}(\mathbf{b}) = \{i \in N_t : b_i \geq r\}$. Agents with bids strictly below r are effectively excluded from the priority queue. The reserve serves the standard role of congestion control, while the temperature governs the degree of randomization.

5.2 Scoring and the Plackett–Luce Allocation

For each eligible agent $i \in \mathcal{E}(\mathbf{b})$, we define a score s_i and an intensity α_i:

$$s_i := \frac{b_i - r}{\tau}, \quad \alpha_i := \exp(s_i). \tag{7}$$

Ineligible agents are assigned $\alpha_i = 0$. The exponential scoring rule connects our mechanism to the Plackett–Luce (PL) ranking model (Luce, 1959; Plackett, 1975). The allocation

rule proceeds by sequential sampling without replacement. The probability of observing a specific ordered tuple of winners $(i_1, \ldots, i_K)$ is given by the product of conditional probabilities:

$$\Pr(i_1, \ldots, i_K \mid \alpha) = \prod_{p=1}^{K} \frac{\alpha_{i_p}}{\sum_{j \in \mathcal{E} \setminus \{i_1, \ldots, i_{p-1}\}} \alpha_j}. \tag{8}$$

This aligns with the standard interpretation of the PL model as a multi-stage Luce choice process.

5.3 Theoretical Representations

To facilitate both the theoretical analysis of incentives and efficient implementation, we leverage two equivalent representations of the PL process.

Exponential Race (Arrival Times). The PL ranking is stochastically equivalent to a process where each agent i draws an independent arrival time $T_i \sim \text{Exp}(\alpha_i)$, and agents are ranked by increasing T_i. This representation is crucial for our proofs of incentive compatibility: increasing a bid b_i increases the rate α_i, which stochastically decreases the arrival time T_i, thereby improving the agent's expected position in a monotone fashion.

Gumbel-Top-K (Perturb-and-Select). Alternatively, the mechanism can be implemented by drawing i.i.d. Gumbel noise $g_i \sim \text{Gumbel}(0, 1)$ for each agent and selecting the K agents with the highest perturbed scores $\log \alpha_i + g_i$. This "perturb-and-max" view (Kool et al., 2019) is particularly computationally efficient and provides a direct link to the literature on differentiable sorting.

5.4 Limiting Behavior

The temperature τ acts as a continuous interpolation parameter:

- As $\tau \to 0$, the mechanism approaches a **Strict Priority Auction**. The agent with the highest bid is selected with probability approaching 1 (breaking ties arbitrarily).

- As $\tau \to \infty$, the scoring differences vanish, and the mechanism approaches a **Uniform Lottery** among eligible participants.

The temperature τ governs how sharply bid differences translate into allocation probabilities.

Figure 1 illustrates how the temperature parameter τ continuously smooths the bid-to-allocation mapping by flattening the top-K selection probability and expected priority weight as bids vary within the eligible region.

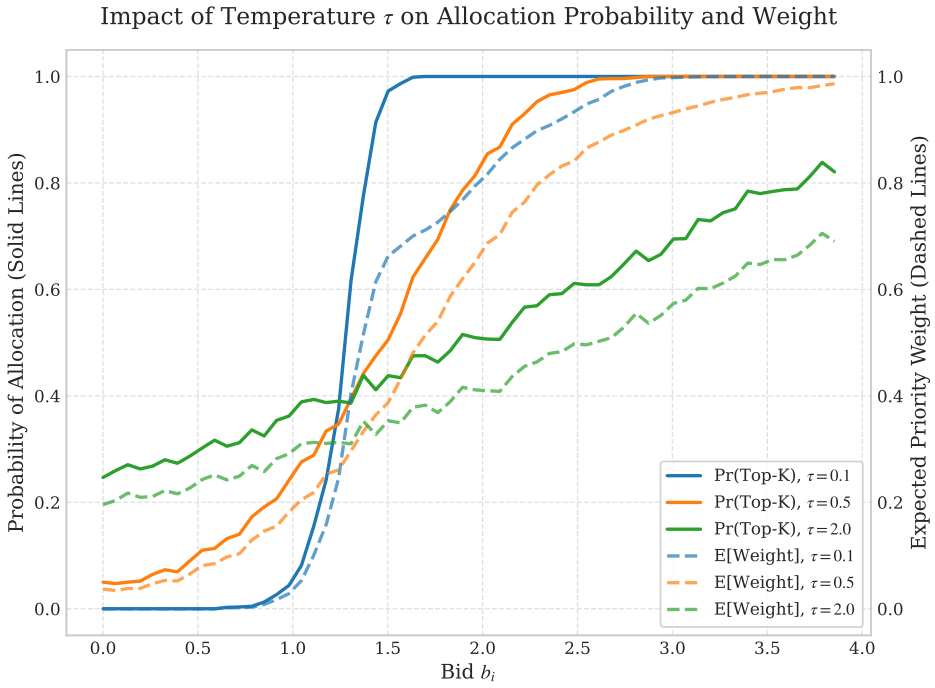


Figure 1: Temperature smooths priority allocation. Notes: The figure plots agent i's probability of receiving a top-K slot (left axis) and expected priority weight (right axis) as functions of i's bid b_i, holding a fixed opponent profile $\mathbf{b}_{-i}$ constant and conditioning on eligibility $b_i \geq r$. The illustrative configuration uses the common replication parameters from Appendix Table 2, namely $K = 6$, reserve $r = 0.85$, and slot weights $\mathbf{w} = (1.00, 0.88, 0.76, 0.64, 0.54, 0.46)$. The accompanying figure-generation script should archive the exact opponent profile, bid grid, and displayed temperature list.

This structural property underpins our stability analysis in Section 7, where we show that the Lipschitz constant of the allocation rule scales as $O(1/\tau)$.

5.5 Payment Implementation

The PL top-K rule is the *base allocation rule* analyzed throughout the paper. Transfers are then computed by a separate single-call wrapper around that base rule, described formally in Section 6 and Appendix B. This distinction matters: the monotonicity, stability, and measurement results in the main text refer to the interim allocation induced by the base FIRM-L rule on the reported bid profile $\mathbf{b}$.

Computing Myerson transfers by direct numerical integration would require repeated evaluation of the interim allocation over counterfactual bids and is therefore unsuitable for latency-sensitive deployment. We instead use the self-resampling reduction of Babaioff et al. (2015). In the implemented single-call version, the mechanism does *not* run the PL allocator once on the true bid vector and then again on a counterfactual profile. Rather, it makes one execution of the base allocator on a monotone resampled bid vector and computes a payment estimator from the realized allocation and the logged resampling weights. The wrapper is allocation-preserving at the level of truthful expected transfers, not pathwise on every realized sample. Appendix B states the precise construction and its audit log.

5.6 Algorithmic Description

The complete procedure for a single round is formalized in Algorithm 1.

Algorithm 1 Base FIRM-L Allocation Rule

Parameters: Reserve $r \geq 0$, Temperature $\tau > 0$, Weights $\mathbf{w} = (w_1, \ldots, w_K)$.
Input: Bid profile $\mathbf{b} = (b_1, \ldots, b_n)$.

Eligibility: Let $\mathcal{E} \leftarrow \{i \in N : b_i \geq r\}$.
Scoring: For each $i \in \mathcal{E}$, set intensity $\alpha_i \leftarrow \exp((b_i - r)/\tau)$.
Allocation (PL Top-K):
for $p = 1$ to $\min(K, |\mathcal{E}|)$ **do**
    Sample winner i_p from $\mathcal{E}$ with probability $\frac{\alpha_{i_p}}{\sum_{j \in \mathcal{E}} \alpha_j}$.
    Assign slot p to agent i_p: $y_{i_p,p} \leftarrow 1$.
    Remove winner: $\mathcal{E} \leftarrow \mathcal{E} \setminus \{i_p\}$.
end for
Output: Assignment y and realized priority weights $X(y)$.

Transfers are computed by the single-call wrapper in Section 6; the monotonicity and

stability results in Sections 6–7 refer to the base allocation rule above.

Proposition 1 (Operational complexity)**.** *Let n denote the number of submitted bids and let $K \leq n$ be the slot count. One round of the base FIRM-L allocation admits two standard implementations:*

1. *A direct sequential PL sampler with time complexity $O(nK)$ and linear memory;*
2. *A perturb-and-select implementation via Gumbel-top-K sampling with time complexity $O(n \log K)$ using a size-K heap, plus $O(K \log K)$ if the selected keys are sorted to recover ordered slots.*

The single-call payment wrapper adds only $O(n)$ auxiliary resampling and arithmetic on top of one execution of the base allocator.

Proof. For the direct sampler, each of the K draws requires scanning the current eligible set and sampling from the associated categorical distribution, which yields $O(nK)$ work in the worst case. For the perturb-and-select implementation, draw one Gumbel perturbation per bidder, compute the perturbed key $\log \alpha_i + g_i$, and maintain the top K keys in a size-K heap while scanning the n bidders; this takes $O(n \log K)$ time. If the ordered top-K list is required, sorting the retained K keys adds $O(K \log K)$. The payment wrapper samples one auxiliary resampling object per bidder and computes the transfer estimator from the realized allocation and the logged resampling weights, which is linear in n once the base allocation has been executed. □

6 Incentives and Payments

This section establishes the incentive properties of FIRM-L. We proceed in three steps: first, we prove the monotonicity of the interim allocation using a coupling argument; second, we characterize the transfer rule required for Truthfulness-in-Expectation (TIE) via the Myerson identity; and third, we discuss the practical implementation of these transfers using single-call implicit payment computation.

6.1 Incentive Target and Objects

We operate in a single-parameter quasi-linear environment where each agent's private type v_i is a scalar representing their value per unit of priority weight. The relevant allocation statistic for incentives is the expected weight received. Let $X_i(y) := \sum_{p=1}^{K} w_p y_{i,p}$ denote the realized priority weight for agent i given assignment y. Fixing the bids of opponents $\mathbf{b}_{-i}$, the interim expected allocation is:

$$x_i(b_i; \mathbf{b}_{-i}) := \mathbb{E}[X_i(y) \mid b_i, \mathbf{b}_{-i}]. \tag{9}$$

A randomized mechanism is said to be **Truthful-in-Expectation (TIE)** if for every agent i, every opponent profile $\mathbf{b}_{-i}$, and every misreport b_i', truth-telling maximizes expected utility:

$$\mathbb{E}[u_i(v_i, b_i = v_i; \mathbf{b}_{-i})] \geq \mathbb{E}[u_i(v_i, b_i'; \mathbf{b}_{-i})], \tag{10}$$

where the expectation is taken over the mechanism's internal randomness.

6.2 Monotonicity of FIRM-L

The Plackett–Luce allocation rule is inherently stochastic. To establish incentive compatibility, we must show that the probability of obtaining higher-weighted slots strictly increases with one's bid. We utilize the "exponential race" representation of the Plackett–Luce process to construct a coupling proof. Recall that for an eligible agent with bid $b_i \geq r$, the intensity is $\alpha_i(b_i) = \exp((b_i - r)/\tau)$.

Lemma 1 (Rank Monotonicity via Coupling)**.** *Fix $\mathbf{b}_{-i}$ and consider two bids $b_i \leq b_i'$ with both $b_i, b_i' \geq r$. Under the FIRM-L mechanism, the random rank of agent i among eligible agents is stochastically improved (in the sense of first-order stochastic dominance) when b_i increases to b_i'.*

Proof Sketch. The Plackett–Luce ranking can be generated by drawing independent arrival times $T_j \sim \text{Exp}(\alpha_j)$ for all eligible agents and sorting them in increasing order. We construct a coupling using common unit-rate exponential variables $E_j \sim \text{Exp}(1)$ such that

$T_j = E_j/\alpha_j$. When agent i increases their bid from b_i to b_i', their intensity increases to $\alpha_i' > \alpha_i$, which implies a strictly smaller arrival time $T_i' = E_i/\alpha_i' < T_i$ for every realization of E_i. The arrival times of all other agents $j \neq i$ remain unchanged. Consequently, in every coupled sample path, the set of agents arriving before i can only shrink, ensuring that agent i's rank weakly improves. □

Lemma 2 (Monotonicity of Interim Priority Weight)**.** *For any fixed $\mathbf{b}_{-i}$, the interim allocation $x_i(b_i; \mathbf{b}_{-i})$ is non-decreasing in b_i on the domain $[r, \infty)$, and is identically zero for $b_i < r$.*

Proof. The realized weight $X_i(y)$ is a non-increasing function of rank (since $w_1 \geq \cdots \geq w_K$). By Lemma 1, an increase in b_i leads to a stochastic improvement in rank, which directly translates to a non-decrease in the expected priority weight. For bids below the reserve r, the agent is ineligible, resulting in a deterministic weight of 0. □

6.3 Truthfulness via Myerson Payments

With the monotonicity of the interim allocation established, we invoke the standard characterization of incentive compatibility for single-parameter domains (Myerson, 1981; Archer and Tardos, 2001).

Proposition 2 (Payment Identity and Reserve Normalization)**.** *Fix $\mathbf{b}_{-i}$. A monotone interim allocation $x_i(\cdot; \mathbf{b}_{-i})$ is implementable in truthfulness in expectation if and only if the expected payment satisfies*

$$p_i(b_i; \mathbf{b}_{-i}) = b_i x_i(b_i; \mathbf{b}_{-i}) - \int_0^{b_i} x_i(z; \mathbf{b}_{-i})\, dz + C_i(\mathbf{b}_{-i}). \tag{11}$$

Under the reserve rule, $x_i(z; \mathbf{b}_{-i}) = 0$ for all $z < r$. Choosing the standard IR normalization $C_i(\mathbf{b}_{-i}) = 0$ therefore yields the piecewise form

$$p_i(b_i; \mathbf{b}_{-i}) = \begin{cases} 0, & b_i < r, \\ b_i x_i(b_i; \mathbf{b}_{-i}) - \displaystyle\int_r^{b_i} x_i(z; \mathbf{b}_{-i})\, dz, & b_i \geq r. \end{cases}$$

In particular, strictly below the reserve the agent receives zero allocation and pays zero, whereas at the boundary $b_i = r$ we generally have

$$p_i(r; \mathbf{b}_{-i}) = r\, x_i(r; \mathbf{b}_{-i}).$$

Corollary 1 (Interim Individual Rationality)**.** *Under truthful bidding, the mechanism is interim individually rational. If $v_i < r$, truthful utility is zero. If $v_i \geq r$,*

$$U_i(v_i; v_i, \mathbf{b}_{-i}) = v_i x_i(v_i; \mathbf{b}_{-i}) - p_i(v_i; \mathbf{b}_{-i}) = \int_r^{v_i} x_i(z; \mathbf{b}_{-i})\, dz \geq 0.$$

Remark 1 (Reserve boundary convention)**.** *Our baseline definition treats bids with $b_i \geq r$ as eligible. If a designer prefers the stronger normalization $p_i(r; \mathbf{b}_{-i}) = 0$, one may equivalently adopt the strict convention "eligible iff $b_i > r$" (or replace the hard reserve with a soft gate). This affects only a measure-zero boundary event and does not alter the monotonicity or stability arguments.*

6.4 Single-Call Implementation

We distinguish the *base allocation rule* from the *payment wrapper*. Let A^θ denote the base FIRM-L allocation rule under policy $\theta = (r, \tau)$, and let $x_i^\theta(b_i; \mathbf{b}_{-i})$ be its interim expected priority weight. Sections 7–10 analyze this base allocation object. Because A^θ is monotone in the sense established above, the expected transfers required for truthful implementation are characterized by Proposition 2.

To compute these transfers without numerically integrating $x_i^\theta(z; \mathbf{b}_{-i})$ over a dense grid of counterfactual bids, we invoke the single-parameter self-resampling reduction of Babaioff et al. (2015) as a *wrapper* around A^θ. Concretely, after agents report $\mathbf{b}$, the wrapper draws auxiliary bids $\widetilde{\mathbf{b}}$ from a monotone self-resampling scheme supported on $[0, b_i]$ coordinate-wise, executes the base allocation rule once on $\widetilde{\mathbf{b}}$, and computes a randomized payment estimator from the realized allocation and the resampling weights. By the theorem of Babaioff et al. (2015), every monotone single-parameter base rule admits such a truthful-in-expectation implementation.

Three clarifications are important. First, the monotonicity proof and the stability certificate in Sections 6 and 7 refer to the base FIRM-L allocation rule A^θ. Second, the single-call reduction is an implementation device for transfers; it uses one call to the base allocation routine and auxiliary randomness, but it need not preserve the exact realized allocation kernel pathwise. Third, whenever we report welfare, exposure, sensitivity, turnover, or other allocation-side statistics, we report them for the base rule A^θ; payment-side statistics are reported under the induced TIE transfer rule. This separation prevents conflating the stability certificate for x_i^θ with the auxiliary randomness introduced by the payment wrapper.

Corollary 2 (Deployable truthful implementation)**.** *For every policy $\theta = (r, \tau)$, the monotone base allocation rule A^θ together with its Myerson payment rule is truthful in expectation. The Babaioff–Kleinberg–Slivkins single-call wrapper provides an operational implementation that computes these transfers with one execution of the PL allocation routine plus auxiliary resampling.*

7 Stability, Fairness, and Auditability

This section formalizes within-round stability as a primary design objective for priority markets. We derive explicit sensitivity bounds for FIRM-L and demonstrate how these bounds translate into auditable certificates of fairness. This focus is motivated by the empirical reality of blockchain transaction ordering, where strict priority mechanisms generate "cliffs"—discontinuities that amplify micro-perturbations into large allocation and payment jumps (Daian et al., 2020).

7.1 Stability as a Market Design Objective

Priority markets differ from standard auctions in a crucial structural respect: allocation discontinuities are inherent when priority is implemented by deterministic sorting. In a strict priority regime, an infinitesimal increase in a bid can flip an agent's rank across a critical cutoff, resulting in a discrete swing in execution position and realized utility. This

fragility creates specific market failures:

1. **Procedural Fairness.** When similarly situated agents (those with nearly identical bids) face drastically different outcomes due to noise or timing granularity, the mechanism violates the "treat similar individuals similarly" principle that underpins individual fairness (Dwork et al., 2012).

2. **Auditability.** If a mechanism's output is hypersensitive to input changes, ex-post auditing becomes fragile. It becomes computationally difficult to distinguish strategic manipulation from system noise, or to certify that the mechanism behaved as intended under network perturbations (Heimbach et al., 2023).

3. **Measurement Reliability.** Mechanism outputs—such as realized fees and inclusion probabilities—are often used as high-frequency indicators of market congestion. Discontinuities render these indices unstable precisely during periods of stress, undermining their utility for policy monitoring.

FIRM-L is designed so that the temperature parameter τ serves as an explicit, reportable "sensitivity knob," enabling the social planner to target a specific point on the efficiency–stability frontier based on theoretical bounds rather than ex-post heuristics.

7.2 Eligible-Region Lipschitz Stability

We adapt the principle of individual fairness to the context of priority allocation. We focus on the sensitivity of the interim allocation for eligible agents.

7.2.1 Definitions and Objects

Recall that $x_i(b_i; \mathbf{b}_{-i})$ denotes the interim expected priority weight for agent i. Let r be the reserve price. The interim weight is derived from the exposure probabilities $q_{i,p}(b) := \Pr(y_{i,p} = 1 \mid b)$ via $x_i(b) = \sum_{p=1}^{K} w_p q_{i,p}(b)$. We restrict our attention to the domain of eligible bids, $b_i \geq r$, as the eligibility decision itself induces a necessary discontinuity at the boundary.

Definition 1 (Eligible-Region Sensitivity)**.** *Fix the opponents' profile* $\mathbf{b}_{-i}$*. The* eligible-region sensitivity *for agent* i *is the Lipschitz constant of the interim allocation map on the domain* $[r, \infty)$*:*

$$Lip_i(\mathbf{b}_{-i}) := \sup_{b,b' \geq r, b \neq b'} \frac{|x_i(b; \mathbf{b}_{-i}) - x_i(b'; \mathbf{b}_{-i})|}{|b - b'|}. \tag{12}$$

The mechanism's ***Eligible-Region Stability*** $\mathcal{S}$ *is the supremum of this sensitivity over all agents and opponent profiles.*

This definition provides a scalar summary of the mechanism's "smoothness." A finite $\mathcal{S}$ certifies that small changes in a bid result in proportionally small changes in the expected opportunity.

7.3 Main Result: Sensitivity Bounds

We now establish an explicit eligible-region sensitivity bound. The key technical point is that, conditional on the opponents' exponential-race arrival times, agent i's probability of occupying slot p is the probability that an exponential clock with rate $\lambda_i = \exp((b_i - r)/\tau)$ falls into an interval between consecutive order statistics. This yields a derivative bound for the full without-replacement PL process itself, rather than for a one-shot softmax approximation.

Theorem 1 (Stability Bound)**.** *For any reserve* r *and temperature* $\tau > 0$*, the FIRM-L mechanism satisfies Eligible-Region Stability with*

$$\mathcal{S} \leq \frac{1}{\tau} \sum_{p=1}^{K} w_p.$$

Equivalently, for every agent i*, every opponent profile* $\mathbf{b}_{-i}$*, and every eligible bids* $b_i, b_i' \geq r$*,*

$$|x_i(b_i; \mathbf{b}_{-i}) - x_i(b_i'; \mathbf{b}_{-i})| \leq \frac{\sum_{p=1}^{K} w_p}{\tau} |b_i - b_i'|.$$

Proof. Fix $\mathbf{b}_{-i}$ and write $s_i = (b_i - r)/\tau$ and $\lambda_i = e^{s_i}$. Use the exponential-race rep-

resentation of the PL ranking. For each eligible opponent $j \neq i$, draw $T_j \sim \text{Exp}(\lambda_j)$. Let

$$O_{(1)} \leq \cdots \leq O_{(m)}$$

denote the order statistics of the opponents' arrival times, where m is the number of eligible opponents, and set $O_{(0)} := 0$ and $O_{(m+1)} := \infty$.

Conditional on T_{-i}, agent i receives slot p exactly when

$$O_{(p-1)} < T_i < O_{(p)}, \qquad T_i \sim \text{Exp}(\lambda_i).$$

Hence the conditional slot probability is

$$q_{i,p}(b_i; \mathbf{b}_{-i} \mid T_{-i}) = e^{-\lambda_i O_{(p-1)}} - e^{-\lambda_i O_{(p)}}, \qquad p = 1, \ldots, K,$$

with the convention $e^{-\lambda_i \infty} = 0$. Differentiating with respect to the score s_i gives

$$\frac{\partial q_{i,p}}{\partial s_i}(b_i; \mathbf{b}_{-i} \mid T_{-i}) = \lambda_i \left[O_{(p)} e^{-\lambda_i O_{(p)}} - O_{(p-1)} e^{-\lambda_i O_{(p-1)}} \right].$$

For every $x \geq 0$, $xe^{-x} \leq 1/e$. Therefore

$$\left| \frac{\partial q_{i,p}}{\partial s_i}(b_i; \mathbf{b}_{-i} \mid T_{-i}) \right| \leq \frac{2}{e} < 1.$$

Taking expectations over T_{-i} and using dominated convergence yields

$$\left| \frac{\partial q_{i,p}}{\partial s_i}(b_i; \mathbf{b}_{-i}) \right| \leq 1 \qquad \text{for each } p \leq K.$$

Since $x_i(b_i; \mathbf{b}_{-i}) = \sum_{p=1}^{K} w_p q_{i,p}(b_i; \mathbf{b}_{-i})$,

$$\left| \frac{\partial x_i}{\partial s_i}(b_i; \mathbf{b}_{-i}) \right| \leq \sum_{p=1}^{K} w_p \left| \frac{\partial q_{i,p}}{\partial s_i}(b_i; \mathbf{b}_{-i}) \right| \leq \sum_{p=1}^{K} w_p.$$

Finally, $\partial s_i/\partial b_i = 1/\tau$, so

$$\left|\frac{\partial x_i}{\partial b_i}(b_i; \mathbf{b}_{-i})\right| \leq \frac{1}{\tau}\sum_{p=1}^{K} w_p.$$

The mean-value theorem then implies the stated Lipschitz bound. □

Figure 2 reports Monte Carlo estimates of the interim-allocation slope using common random numbers across perturbed bid profiles. The figure is therefore an empirical check of the expectation-level bound in Theorem 1, rather than a finite-difference calculation on a single realized PL draw.

Figure 2: Stability Certificate

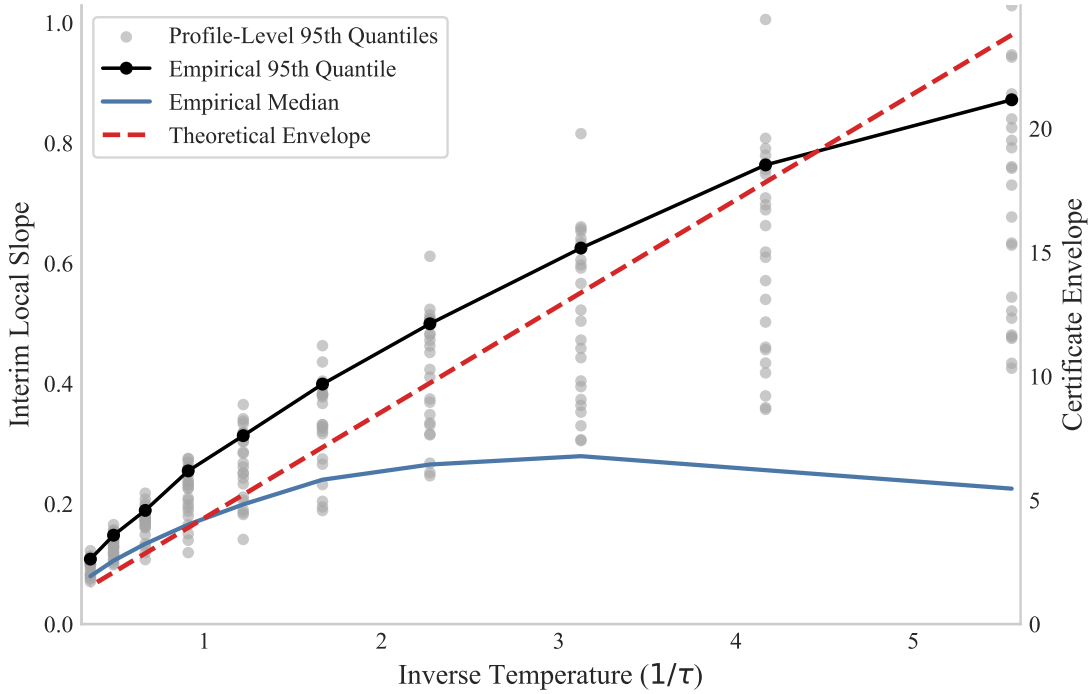


Figure 2: Stability certificate: empirical sensitivity of the *interim* allocation scales with $1/\tau$. Notes: Empirical slope summaries are plotted on the left axis and the analytical certificate on the right axis. For each τ, we evaluate 24 stored opponent profiles (12 calm and 12 storm), use 11 eligible bid-grid points per profile, and estimate the interim expected allocation $x_i(b_i; \mathbf{b}_{-i})$ by Monte Carlo integration of the PL sampler with 1,536 common-random-number draws per evaluation. Gray points show profile-level 95th-percentile slopes, the black line reports the cross-profile 95th percentile, the blue line reports the median, and the dashed red line reports the certificate $(\sum_{p=1}^{K} w_p)/\tau$.

Corollary 3 (Exact constant for $K = 1$)**.** *When $K = 1$, the eligible-region Lipschitz*

constant satisfies

$$\mathrm{Lip}_i \le \frac{w_1}{4\tau}.$$

Proof. When $K = 1$,

$$q_{i,1}(b_i; \mathbf{b}_{-i}) = \frac{\alpha_i}{\alpha_i + \sum_{j \neq i} \alpha_j}, \qquad \alpha_i = e^{(b_i - r)/\tau}.$$

Differentiating with respect to $s_i = \log \alpha_i$ gives

$$\frac{\partial q_{i,1}}{\partial s_i} = q_{i,1}(1 - q_{i,1}) \le \frac{1}{4}.$$

Hence

$$\left| \frac{\partial x_i}{\partial b_i} \right| = \frac{w_1}{\tau} \left| \frac{\partial q_{i,1}}{\partial s_i} \right| \le \frac{w_1}{4\tau}.$$

□

Remark 2 (Conservative versus tight constants)**.** *Theorem 1 is a uniform certificate for general top-K without-replacement PL sampling. Corollary 3 shows that the exact constant can be smaller in special cases. We therefore interpret Theorem 1 as a conservative but fully proved $O(1/\tau)$ bound, not as a claim of tightness for all K.*

7.4 Fairness Implications

The Lipschitz bound implies a strong form of procedural fairness:

Proposition 3 (Opportunity Stability)**.** *Under Theorem 1, for any two eligible bids $b_i, b_i' \ge r$ (holding $\mathbf{b}_{-i}$ fixed), the difference in expected allocation is bounded by:*

$$|x_i(b_i; \mathbf{b}_{-i}) - x_i(b_i'; \mathbf{b}_{-i})| \le \left(\frac{\sum_{p=1}^{K} w_p}{\tau} \right) |b_i - b_i'|. \tag{13}$$

This guarantees that similarly situated agents receive similar expected opportunities, avoiding the arbitrary discrimination inherent in tie-breaking or granular sorting noise.

7.5 Auditability and Governance Procedures

Auditability in our setting has three distinct layers.

1. Replayability of the announced rule. Given the announced parameters $(r, \tau, \mathbf{w})$, the reported bid vector, the randomness seed for the PL sampler, and the auxiliary randomness used by the payment wrapper, a third party can replay the exact executed procedure. This is a software-verification claim about whether the implemented mechanism matches the announced rule.

2. Expectation-level stability auditing. Theorem 1 concerns the interim map $b_i \mapsto x_i(b_i; \mathbf{b}_{-i})$, not a single realized PL draw. Accordingly, ex-post auditing of sensitivity must be based on analytical evaluation or repeated simulation of the expected allocation using common random numbers across perturbed bid profiles. A single-seed replay of the realized allocation can detect implementation bugs, but it is not itself a valid estimate of the Lipschitz constant.

3. Input integrity. Any audit is conditional on complete and authenticated visibility of the eligible bid set. In proposer–builder–relay environments, this means that the relevant bid pool must itself be committed to or attested. Otherwise one can certify only *conditional* correctness given the reported inputs, not unconditional fairness of the market interface.

Ex-ante certification. The platform can issue a stability certificate by publishing a target sensitivity level $S_{\max}$ and choosing $\tau \geq (\sum_{p=1}^{K} w_p)/S_{\max}$. This certificate is analytical and can be checked directly from the announced parameter values.

Minimal audit log. To enable all three layers of audit, the log must include: (i) announced parameters $(r, \tau, \mathbf{w})$; (ii) the bid vector or cryptographic commitments with later reveal; (iii) the full randomness used by the PL sampler and by the payment wrapper; (iv) the realized ranking or assignment; and (v) the realized transfers. For expectation-

level auditing, the code version and Monte Carlo configuration used to estimate interim quantities should also be archived.

8 Measurement Outputs: Congestion-Adjusted Indices

The mechanism generates two economically meaningful interim objects: the expected priority-weight vector

$$x^{\theta}(\mathbf{b}) = \big(x_1^{\theta}(\mathbf{b}), \dots, x_n^{\theta}(\mathbf{b})\big)$$

and the expected truthful-payment vector

$$p^{\theta}(\mathbf{b}) = \big(p_1^{\theta}(\mathbf{b}), \dots, p_n^{\theta}(\mathbf{b})\big).$$

The stability theorem in Section 7 pertains to the mapping $\mathbf{b} \mapsto x^{\theta}(\mathbf{b})$. To construct monitoring statistics that inherit this regularity, we work with clipped transformations of expected allocations and expected payments on a bounded monitoring domain.

Fix clipping constants $\underline{x} > 0$, $\overline{u} < \infty$, $\underline{u} > 0$, and $\overline{c} < \infty$. Define the clipped expected unit priority price for agent i by

$$u_i^{\theta}(\mathbf{b}) := \operatorname{clip}\left(\frac{p_i^{\theta}(\mathbf{b})}{\max\{x_i^{\theta}(\mathbf{b}), \underline{x}\}}, 0, \overline{u}\right).$$

Let $W := \sum_{p=1}^{K} w_p$ denote total available priority mass and let

$$\mathcal{A}^{\theta}(\mathbf{b}) := \{i : \; x_i^{\theta}(\mathbf{b}) \geq \underline{x}\}$$

be the active set for monitoring.

8.1 Price and Congestion Indices

Definition 2 (Priority Price Index (PPI))**.** *For trimming proportion $\rho \in [0, 1/2)$, define*

$$\mathrm{PPI}^{\theta}(\mathbf{b}) := \mathrm{TrimMean}_{\rho}\Big(\{u_i^{\theta}(\mathbf{b}) : i \in \mathcal{A}^{\theta}(\mathbf{b})\}\Big).$$

This is the robust central tendency of expected unit priority prices among materially exposed agents.

Definition 3 (Congestion Pressure Index (CPI))**.** *Define*

$$\mathrm{CPI}^{\theta}(\mathbf{b}) := \mathrm{clip}\left(\frac{1}{W}\sum_{i=1}^{n} p_i^{\theta}(\mathbf{b}), 0, \overline{c}\right).$$

This is the average expected payment pressure per unit of available priority mass.

8.2 Tail and Concentration Indices

Definition 4 (Payment Spike Index (PSI))**.** *Let $Q_q(\cdot)$ denote the empirical q-quantile on the active set. Define*

$$\mathrm{PSI}^{\theta}(\mathbf{b}) := \frac{Q_{0.99}\big(\{u_i^{\theta}(\mathbf{b}) : i \in \mathcal{A}^{\theta}(\mathbf{b})\}\big)}{\max\Big\{Q_{0.50}\big(\{u_i^{\theta}(\mathbf{b}) : i \in \mathcal{A}^{\theta}(\mathbf{b})\}\big), \underline{u}\Big\}}.$$

The denominator floor $\underline{u}$ prevents the ratio from becoming mechanically explosive when the median price is near zero.

Definition 5 (Ordering Rent Concentration Index (ORCI))**.** *Suppose bids are routed through intermediaries $m \in \{1, \ldots, B\}$, and let $g(i)$ denote the intermediary associated with agent i. Define intermediary-level expected rent*

$$R_m^{\theta}(\mathbf{b}) := \sum_{i:g(i)=m} p_i^{\theta}(\mathbf{b}),$$

shares

$$s_m^{\theta}(\mathbf{b}) := \frac{R_m^{\theta}(\mathbf{b})}{\max\Big\{\sum_{\ell=1}^{B} R_{\ell}^{\theta}(\mathbf{b}), \underline{u}\Big\}},$$

and

$$\mathrm{ORCI}^{\theta}(\mathbf{b}) := \sum_{m=1}^{B} \left(s_m^{\theta}(\mathbf{b})\right)^2.$$

This is a Herfindahl-style concentration index for expected ordering rents under intermediated execution.

8.3 Why these indices inherit stability

The key design choice is that all four indices are built from clipped versions of x^{θ} and p^{θ} through operations that are Lipschitz on compact domains: clipping, trimmed means, empirical quantiles, and quadratic concentration scores. In particular, if $x_i^{\theta}(\cdot; \mathbf{b}_{-i})$ is S_x-Lipschitz on the eligible region and bids are monitored on a bounded domain $[r, \bar{b}]$, then the Myerson payment identity implies that $p_i^{\theta}(\cdot; \mathbf{b}_{-i})$ is also Lipschitz on that domain, with derivative

$$\frac{d}{db_i} p_i^{\theta}(b_i; \mathbf{b}_{-i}) = b_i \frac{d}{db_i} x_i^{\theta}(b_i; \mathbf{b}_{-i})$$

at differentiability points. Appendix A.4 gives the general transfer lemma.

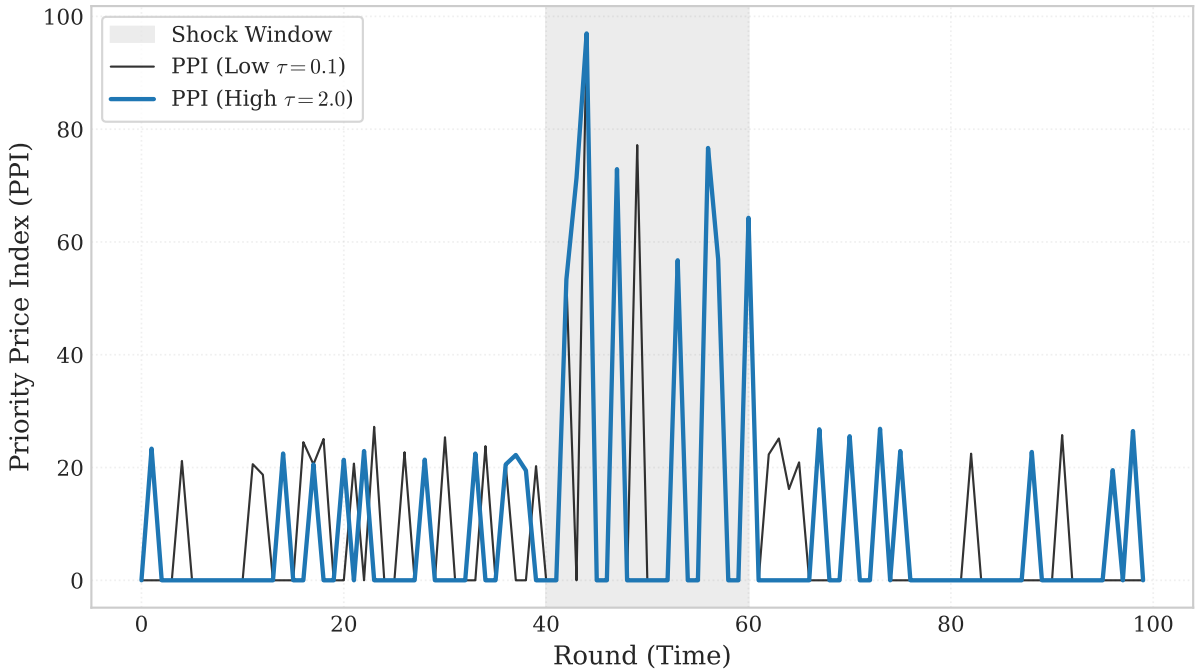
Operationally, the platform computes these indices from logged bids and the announced policy by either analytical evaluation or Monte Carlo estimation of the interim objects x^{θ} and p^{θ} using common random numbers. A single realized allocation draw is sufficient for replay, but not for certifying the sensitivity of the derived indices.

Figure 3 should therefore be interpreted as a time series of *estimated interim indices*, not as a plot of raw single-draw realized outcomes.

9 Dynamic Policy: Learning to Tune (r, τ)

In this section, we formalize the governance problem as an online learning process. We interpret the platform not merely as a mechanism designer, but as a *statistical regulator* that must repeatedly choose parameters $\theta_t = (r_t, \tau_t)$ to balance allocative efficiency with measurement integrity. This perspective is motivated by the operational reality of blockchain infrastructure: ordering rules are part of the protocol, observed by strategic

Figure 3: Congestion-Adjusted Indices Stability



Notes: Shaded area indicates a demand shock (increased bid mean/variance).
Low τ index exhibits sharp spikes; High τ index is smoother.

Figure 3: Congestion-adjusted interim indices under alternative smoothing policies. Notes: The figure reports time series of the Priority Price Index (PPI) and Congestion Pressure Index (CPI) defined in Section 8. Each series is computed from clipped interim expected allocations and expected payments under the announced policy. Shaded regions denote shock windows. Higher τ dampens transient spikes by smoothing the underlying bid-to-allocation map.

agents, and subject to intense scrutiny. A discretionary, data-contingent parameter change within a round would be difficult to justify and even harder to audit. Instead, we propose a **Regulatory Rule**—a governance algorithm that adapts parameters based on realized market conditions while maintaining strict ex-ante commitment.

9.1 The Governance Problem

9.1.1 Policy as a Rule

We view the sequence of parameter choices not as isolated decisions, but as the output of a deterministic algorithm operating on historical data. This aligns with the measurement governance agenda, which emphasizes transparency, explainability, and auditability for algorithmically produced statistics (UNECE, 2025).

9.1.2 Pre-Commitment Timing

To prevent manipulation and ensure auditability, we impose a strict timing constraint:

1. **Announcement:** At the start of round t, the platform computes θ_t based on history H_{t-1} and publicly announces it.
2. **Participation:** Agents observe θ_t and submit bids $\mathbf{b}_t$.
3. **Execution:** The mechanism executes allocation and payments according to θ_t.

This pre-commitment prevents "policy endogeneity," where the rules of the game are tweaked ex-post to favor specific outcomes, a key concern in decentralized mechanism design.

9.2 A Bounded, Auditable Loss Function

Our formal guarantee is stated for a pre-committed, full-information benchmark. In each round t, before bids are submitted, the platform chooses $\theta_t \in \Theta$ using only the public history H_{t-1}. After the round, it observes a loss vector $\ell_t(\theta) \in [0,1]$ for every $\theta \in \Theta$. This benchmark is appropriate when the platform has access to a policy-invariant replay environment or simulator in which counterfactual evaluation does not itself change participation or the type process. The comparator in the theorem below is the best *fixed* policy in hindsight. We do not claim a dynamic-regret guarantee in the main theorem.

Let $\Theta = \{\theta_1, \dots, \theta_M\}$ and let $a, b, c \geq 0$ with $a + b + c = 1$. We define the bounded per-round loss by

$$\ell_t(\theta) = a\,\ell_t^W(\theta) + b\,\ell_t^S(\theta) + c\,\ell_t^V(\theta),$$

where

$$\ell_t^W(\theta) = 1 - \frac{\mathrm{Welfare}_t(\theta)}{\mathrm{Welfare}_t^{\max}},$$

$$\ell_t^S(\theta) = \min\left\{1, \frac{\widehat{S}_t(\theta)}{S_{\max}}\right\},$$

$$\ell_t^V(\theta) = \min\left\{1, \omega_P \frac{\mathrm{P99}_t(\theta)}{\overline{P}_{99}} + (1-\omega_P)\frac{\mathrm{PSI}_t(\theta)}{\overline{\Psi}}\right\},$$

with $\omega_P \in [0,1]$. Here $\widehat{S}_t(\theta)$ is the empirical slope estimate of the interim allocation map computed with common random numbers, $\mathrm{P99}_t(\theta)$ is the round-t 99th percentile of clipped unit priority prices, and $\mathrm{PSI}_t(\theta)$ is the Payment Spike Index from Section 8. The scale parameters $(S_{\max}, \overline{P}_{99}, \overline{\Psi})$ are fixed ex ante and made public as part of the governance rule. All loss components are therefore auditable and uniformly bounded in $[0,1]$.

9.3 Full-Information Governance via Hedge

The boundedness claim used by Hedge is immediate from the definitions above. By construction,

$$0 \le \ell_t^W(\theta) \le 1$$

because $\mathrm{Welfare}_t(\theta) \le \mathrm{Welfare}_t^{\max}$ by definition of the benchmark. Likewise,

$$0 \le \ell_t^S(\theta) \le 1$$

because it is a clipped ratio, and

$$0 \le \ell_t^V(\theta) \le 1$$

because the convex combination inside the outer clipping operator is non-negative and the operator truncates it at 1. Therefore, with non-negative weights a, b, c summing to one, the composite loss $\ell_t(\theta) = a\ell_t^W(\theta) + b\ell_t^S(\theta) + c\ell_t^V(\theta)$ also belongs to $[0,1]$.

We allow governance updates to be batched. Fix a batch size $B \ge 1$ and, for simplicity, assume the horizon satisfies $T = BU$ for some integer number of updates U. For update index $u \in \{1, \dots, U\}$, define the batch-average loss

$$\bar{\ell}_u(\theta) := \frac{1}{B} \sum_{t=(u-1)B+1}^{uB} \ell_t(\theta).$$

Because each per-round loss lies in $[0,1]$, each batch-average loss $\bar{\ell}_u(\theta)$ also lies in $[0,1]$.

When U is known, we set

$$\eta = \sqrt{\frac{2 \ln M}{U}},$$

where ln denotes the natural logarithm. If U is not known in advance, the same guarantee can be obtained by a standard doubling trick. Let $h_{m,1} = 1$ for all $m \leq M$, and at update u define policy probabilities

$$p_{m,u} = \frac{h_{m,u}}{\sum_{\ell=1}^{M} h_{\ell,u}}.$$

The governance rule samples a policy index $m_u \sim p_u$, announces $\theta_u = \theta_{m_u}$ before the batch begins, applies that policy for the next B rounds, observes the full batch-loss vector $\bar{\ell}_u(\theta_m)$ after the batch, and updates

$$h_{m,u+1} = h_{m,u} \exp(-\eta \, \bar{\ell}_u(\theta_m)).$$

Algorithm 2 Pre-Committed Hedge Governance over Θ

Inputs: Policy menu $\Theta = \{\theta_1, \ldots, \theta_M\}$, learning rate $\eta > 0$.

1: Initialize $h_{m,1} = 1$ for all $m \in \{1, \ldots, M\}$.
2: **for** $u = 1, 2, \ldots, U$ **do**
3: Form probabilities $p_{m,u} \propto h_{m,u}$.
4: Sample $m_u \sim p_u$ and announce $\theta_u = \theta_{m_u}$ for the next batch of B rounds.
5: Execute FIRM-L under θ_u across rounds $(u-1)B+1, \ldots, uB$ and log bids, outcomes, and randomness.
6: Compute the full batch-loss vector $\bar{\ell}_u(\theta_m)$ for all m from the replay environment.
7: Update $h_{m,u+1} = h_{m,u} \exp(-\eta \, \bar{\ell}_u(\theta_m))$ for all m.
8: **end for**

Theorem 2 (Static-regret guarantee)**.** *Suppose that for every update u and every $m \leq M$, the counterfactual batch-loss $\bar{\ell}_u(\theta_m)$ is observed after the batch and belongs to $[0, 1]$. Then the batched Hedge rule above, with learning rate $\eta = \sqrt{2 \ln M / U}$, satisfies*

$$\mathbb{E}\left[\sum_{u=1}^{U} \bar{\ell}_u(\theta_u)\right] - \min_{m \leq M} \sum_{u=1}^{U} \bar{\ell}_u(\theta_m) \leq \sqrt{2U \ln M}.$$

The benchmark is the best fixed policy in hindsight from the same finite menu Θ.

Proof. The boundedness verification above shows that each per-round loss belongs to $[0, 1]$,

so each batch average $\bar{\ell}_u(\theta)$ also belongs to $[0, 1]$. Applying the standard multiplicative-weights guarantee to the update sequence $\{\bar{\ell}_u\}_{u=1}^{U}$ therefore yields the stated regret bound; see Freund and Schapire (1997). The only modeling requirements are pre-commitment, full-information observation of the batch-loss vector, and the policy-invariant replay benchmark used to construct that vector. □

Figure 4 should therefore be interpreted as a *static-regret* figure: it compares the cumulative loss of the announced policy sequence with the best fixed policy from the same menu, not with a time-varying oracle. Figure 4 plots the resulting cumulative static regret under the released replay configuration.

Figure 4: Governance Performance under Pre-Committed Hedge

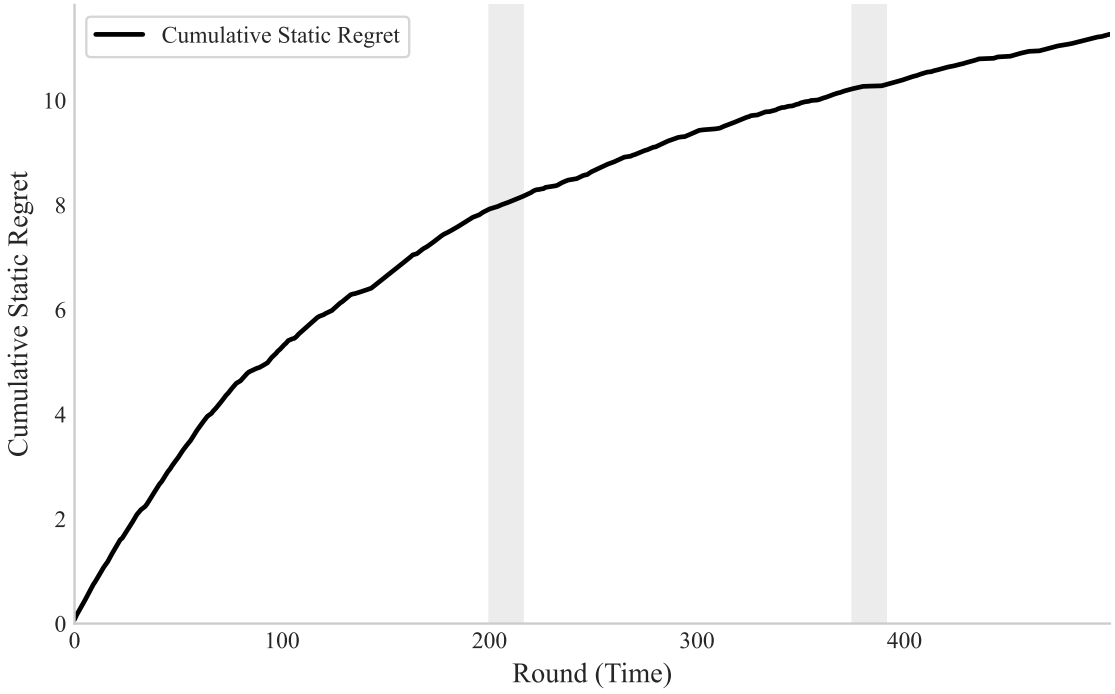


Figure 4: Governance performance under pre-committed Hedge. Notes: The figure reports cumulative *static regret* of the announced policy sequence relative to the best fixed policy in hindsight from the same finite menu $\Theta = \{0.18, 0.32, 0.60, 1.10, 2.05\}$. Hedge updates are batched every four rounds, so the effective update horizon is $U = T/B = 500$ and the learning rate is $\eta = \sqrt{2 \ln |\Theta|/500} \approx 0.0802$. Shaded regions denote update batches intersecting the exogenous shock windows.

9.4 Policy Sensitivity, Bandit Feedback, and Robustness

The full-information theorem above relies on a policy-invariant replay assumption: bids observed under one announced policy are taken to be valid inputs for evaluating other policies in the menu. In many real deployments this assumption can fail because participation, entry, or the type distribution may itself be policy-sensitive. In that case the platform no longer observes a full counterfactual loss vector; it observes only $\ell_t(\theta_t)$ for the policy that was actually announced.

This motivates, but does not form part of, our main formal result. Under policy-sensitive environments the relevant benchmark is a partial-information problem, and adversarial bandit algorithms such as Exp3 become the appropriate formal tool. Likewise, dynamic-regret guarantees against a time-varying comparator require additional structure such as variation budgets or fixed-share updates. We view these as natural extensions rather than as claims proved in the present paper.

Robust governance still requires three practical safeguards. First, policy updates should be temporally batched so that agents cannot target single-round update rules. Second, loss components should be winsorized or otherwise robustified to limit the influence of manipulated tails. Third, the full update log—including the announced policy, realized bids, replay code version, random seeds, and computed loss vector—must be archived for third-party verification.

10 Quantitative Evaluation

This section evaluates FIRM-L through the lens of *statistical governance*. Our goal is not merely to maximize average welfare, but to characterize how a truthful and certifiable smoothing rule behaves under the extreme conditions characteristic of blockchain priority markets.

10.1 Empirical Motivation and Evaluation Principles

Our experimental design is grounded in three stylized facts of transaction ordering markets:

1. **Bursty Congestion:** Demand arrives in waves, creating sudden spikes in the shadow price of block space (Ferreira Torres et al., 2021).

2. **Heavy-Tailed Rents:** The distribution of extractable value (MEV) is extremely fat-tailed; rare arbitrage opportunities drive the bulk of strategic bidding behavior (Qin et al., 2022).

3. **Intermediated Concentration:** Execution is increasingly routed through a small set of sophisticated builders, concentrating value extraction and amplifying the impact of priority rules.

Accordingly, we evaluate the mechanism as a *statistical production process*, placing equal emphasis on tail behavior (e.g., payment spikes) and index stability as on standard efficiency metrics.

Microstructure-style evaluation. Our focus on spikes and turnover mirrors standard microstructure concerns about race-driven volatility and trading costs in electronic markets (Aquilina et al., 2022; O'Hara, 2015).

Not every smoothed priority rule is equally informative. Any parametric randomization family can mechanically trace some efficiency–stability frontier. The non-trivial question is whether one can obtain, within a single mechanism class, all three properties simultaneously: truthful implementation in a one-parameter environment, an explicit analytical stability certificate with transparent dependence on a governance parameter, and derived measurement objects that inherit that certificate. Our quantitative exercise should therefore be read as characterizing the *certified frontier* generated by FIRM-L, rather than as claiming that every smoothing heuristic would perform identically.

10.2 Experimental Setup

10.2.1 Baselines

We benchmark FIRM-L against two canonical extremes:

- **Strict Priority (SP):** deterministically assigns the top K slots to the highest bidders.

- **Uniform Lottery (UL):** randomly allocates the K slots among eligible bidders.

All policies use the same slot-weight vector $\mathbf{w}$ and the same reserve convention. Allocation-side statistics are computed from the base allocation rule. Payment-side statistics are computed from the induced truthful-in-expectation transfer rule.

10.2.2 Data-Generating Process (DGP)

We simulate a horizon of $T = 2000$ rounds. Demand evolves through a two-state latent regime process $s_t \in \{\text{Calm}, \text{Storm}\}$. Conditional on the regime, the number of arriving bids is Poisson with mean $\lambda_{\text{Calm}} = 100$ and $\lambda_{\text{Storm}} = 500$, respectively. Valuations are drawn from a heavy-tailed mixture distribution: with probability 0.95, a bid is drawn from a log-normal component, and with probability 0.05, from a Pareto tail. During designated shock windows, the Pareto tail index is lowered to create rare but severe tail events. All figures and tables are generated from a fully specified simulation manifest that records the exact transition matrix, tail parameters, shock windows, slot weights, reserve rule, τ-grid, replication count, and random seeds. Appendix Table 2 reports the released parameter settings used by the current replication code.

10.3 Metrics

We report performance along three dimensions.

1. **Efficiency.** Total priority-weighted welfare.

2. **Tail risk.** The 99th percentile of clipped unit priority prices and the Payment Spike Index (PSI) from Section 8.

3. **Frontier instability.** Ordered top-K turnover near the cutoff. Let $L_t(\theta) = (i_{1,t}, \ldots, i_{K,t})$ denote the ordered winner list under policy θ, and let $\mathcal{C}_t$ be the set of eligible bids within a pre-specified neighborhood of the cutoff bid. For each $i \in \mathcal{C}_t$, we perturb the bid upward by 1%, recompute the ordered top-K list using common

random numbers, and define

$$\mathrm{Turn}_t(\theta) = \frac{1}{|\mathcal{C}_t|} \sum_{i \in \mathcal{C}_t} \frac{1}{K} \sum_{p=1}^{K} \mathbf{1}\{L_{t,p}(\theta) \neq L_{t,p}^{(i,+)}(\theta)\}.$$

This is intentionally a *local cutoff-instability* metric rather than a whole-population random-perturbation metric.

10.4 Main Results

10.4.1 The Efficiency–Stability Frontier

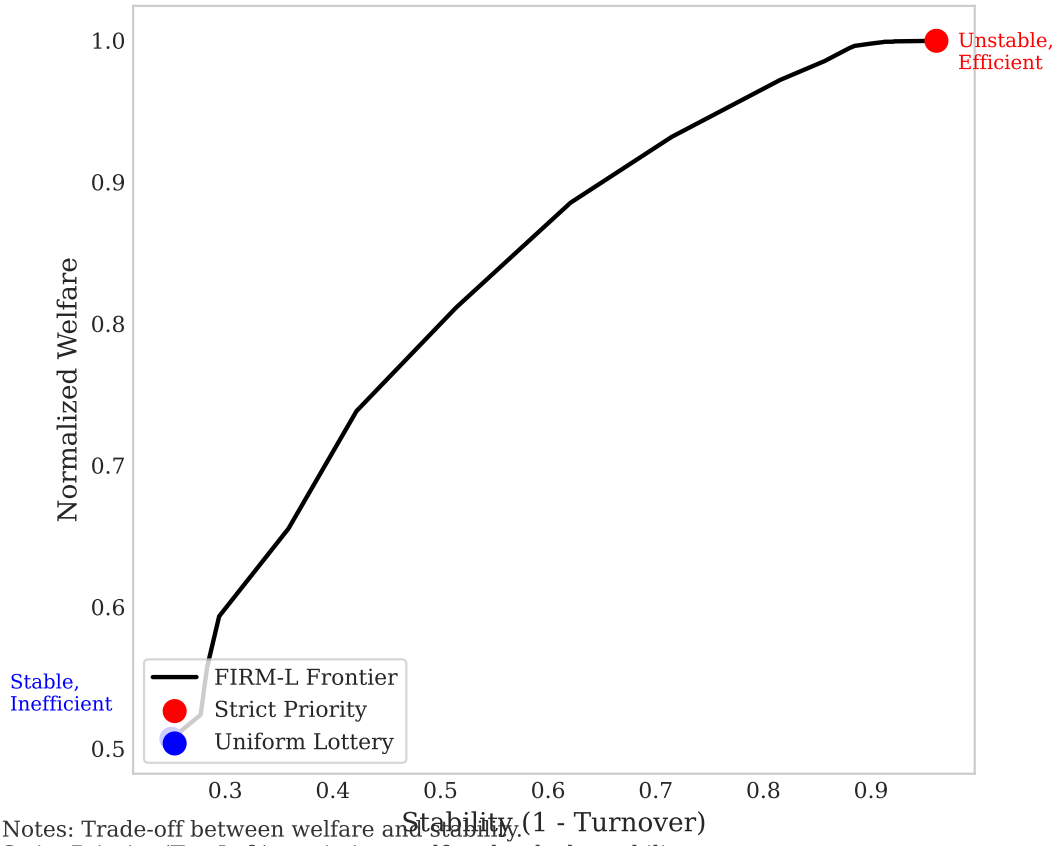
Table 1 now reports means and cross-replication standard errors from $R = 12$ independent macro-replications of the unified simulation pipeline. Under the released manifest, the low- and medium-temperature FIRM-L policies preserve essentially all strict-priority welfare while lowering ordered cutoff turnover from 0.0313 to 0.0289 and 0.0244, respectively. The higher-temperature calibration lowers turnover further to 0.0180 at a visible welfare cost. We therefore use Table 1 as the headline welfare–stability comparison and treat tail behavior separately in Figure 7, because the monitored tail-payment statistic is conditional on the activity threshold $\underline{x} = 0.30$ and is not a monotone unconditional risk measure across policies.

Figure 5 should be read in the same way: each point corresponds to a policy on this frontier, and movement to the right reflects a deliberate reduction in local cutoff instability or tail risk rather than a coding artifact.

10.4.2 Validating the Theoretical Bound

The empirical validation is conducted at the level of interim expected allocations, not single realized draws. For each stored opponent profile $\mathbf{b}_{-i}$, each eligible bid b_i, and each temperature τ, we estimate $x_i(b_i; \mathbf{b}_{-i})$ by repeated PL simulations and compute local slopes by symmetric finite differences using common random numbers. We then regress the resulting slope estimates on $1/\tau$. The appropriate empirical claim is therefore that the fitted slope is increasing in $1/\tau$ and that the cloud of estimates remains

Figure 5: Welfare-Stability Pareto Frontier



Notes: Trade-off between welfare and stability.
Strict Priority (Top-Left) maximizes welfare but lacks stability.
FIRM-L offers a convex interpolation.

Figure 5: Welfare–stability (and tail-risk) Pareto frontier across (r, τ). Notes: Each point corresponds to a fixed policy (r, τ) evaluated over the full simulation horizon (and optionally separately by regime). The vertical axis reports normalized welfare; the horizontal axis reports a stability proxy (allocation turnover or empirical sensitivity) and, in the alternative panel, a tail-risk metric (99th percentile payment spike). Baselines (strict priority and uniform lottery) are marked for comparison.

below the analytical certificate from Theorem 1. Appendix Table 3 reports the released slope-regression diagnostics together with the companion governance and tail-summary diagnostics. Figure 6 visualizes the resulting slope cloud, the fitted line, and the analytical certificate on the same inverse-temperature grid.

10.5 Dynamic Governance Experiment

We evaluate the pre-committed Hedge governance rule from Section 9 on a discrete policy menu ranging from low-τ (aggressive) to high-τ (conservative) settings. The formal comparison is against the best *fixed* policy in hindsight from the same menu. During

Table 1: Comparative Welfare and Stability under MEV Storm Conditions

Mechanism	**Welfare (% of Opt)**	**Turnover (Ordered Stability)**
Strict Priority ($\tau \to 0$)	100.00 (0.00)	0.0313 (0.0008)
Uniform Lottery ($\tau \to \infty$)	20.69 (0.07)	**0.0000 (0.0000)**
FIRM-L (Low $\tau = 0.18$)	99.90 (0.00)	0.0289 (0.0003)
FIRM-L (Med $\tau = 0.44$)	99.32 (0.01)	0.0244 (0.0002)
FIRM-L (High $\tau = 0.82$)	96.71 (0.04)	0.0180 (0.0001)

Notes: Entries are means with cross-replication standard errors in parentheses from $R = 12$ independent macro-replications of the same released simulation manifest used for Figures 2, 4, 6, and 7. **Welfare** is normalized relative to strict priority within each macro-replication's storm rounds. **Turnover** is the ordered top-K cutoff-instability metric defined in Section 10: it averages, over bids in a $\pm 6\%$ neighborhood of the strict cutoff bid, the fraction of ordered slots that change after a 1% upward perturbation under common random numbers. Tail-payment diagnostics are reported separately in Figure 7 because the monitored tail statistic in our released pipeline is conditional on the activity threshold $\underline{x} = 0.30$ and is therefore not an unconditional risk measure across policies.

calm periods the algorithm tends to place more weight on low-τ policies; during storm and shock windows it shifts weight toward smoother policies, thereby reducing tail-risk exposure. We therefore interpret Figure 4 as evidence of adaptation under a static-regret benchmark, not as a dynamic-regret theorem. Figure 7 complements the regret plot by isolating the first exogenous shock episode and comparing tail-payment dynamics across strict priority and two nearby FIRM-L calibrations.

10.6 Robustness

We track concentration risk via ORCI under intermediated execution. Figure 8 illustrates the time-series behavior of the Ordering Rent Concentration Index (ORCI) under intermediated execution, comparing strict priority with FIRM-L to highlight concentration-risk implications.

10.7 Reproducibility and Audit Statement

Every table and figure should be generated from a machine-readable manifest that records: the regime transition matrix; all bid-distribution parameters; the slot count K and slot weights $\mathbf{w}$; the reserve rule; the full τ-grid and policy menu Θ; the replication count; the learning-rate rule; the shock-window calendar; and all random seeds. Allocation-side

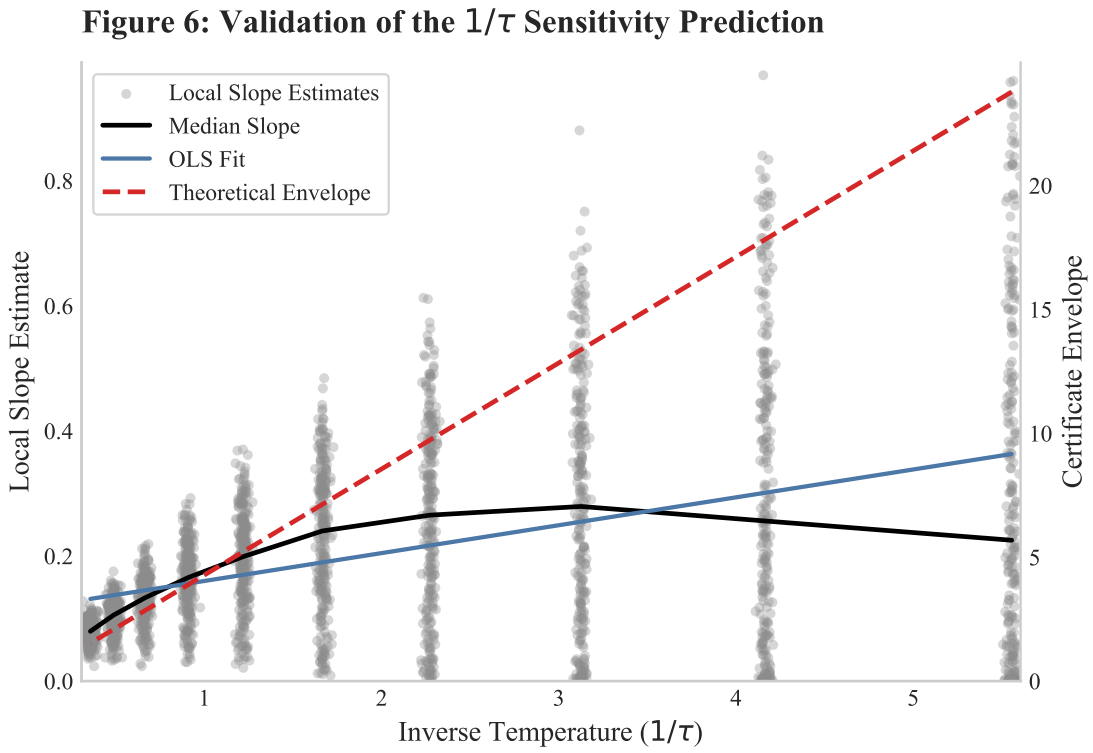


Figure 6: Validation of the $1/\tau$ sensitivity prediction. Notes: Local slope estimates are plotted on the left axis and the analytical envelope from Theorem 1 on the right axis. The black line reports the median local slope at each τ, the blue line reports the OLS fit, and the dashed red line reports the certificate $(\sum_{p=1}^{K} w_p)/\tau$. Under the released manifest, the fitted slope coefficient is 0.0446 with 95% confidence interval [0.0409, 0.0482], $R^2 = 0.177$, and 2,640 local-slope observations.

sensitivity calculations should use common random numbers across perturbed profiles. In the replication package, each reported scalar statistic should be accompanied by either a standard error or a bootstrap confidence interval. The analytical stability certificate remains a theorem; the simulations are a falsification and calibration exercise built around fully logged code and seeds.

11 Discussion and Implications

11.1 The Mechanism as a Statistical Generator

A central conceptual contribution of this work is to treat a priority rule as more than an allocation device. In high-frequency digital markets, the mechanism maps bids, eligibility

Figure 7: Tail Behavior around Shock Windows

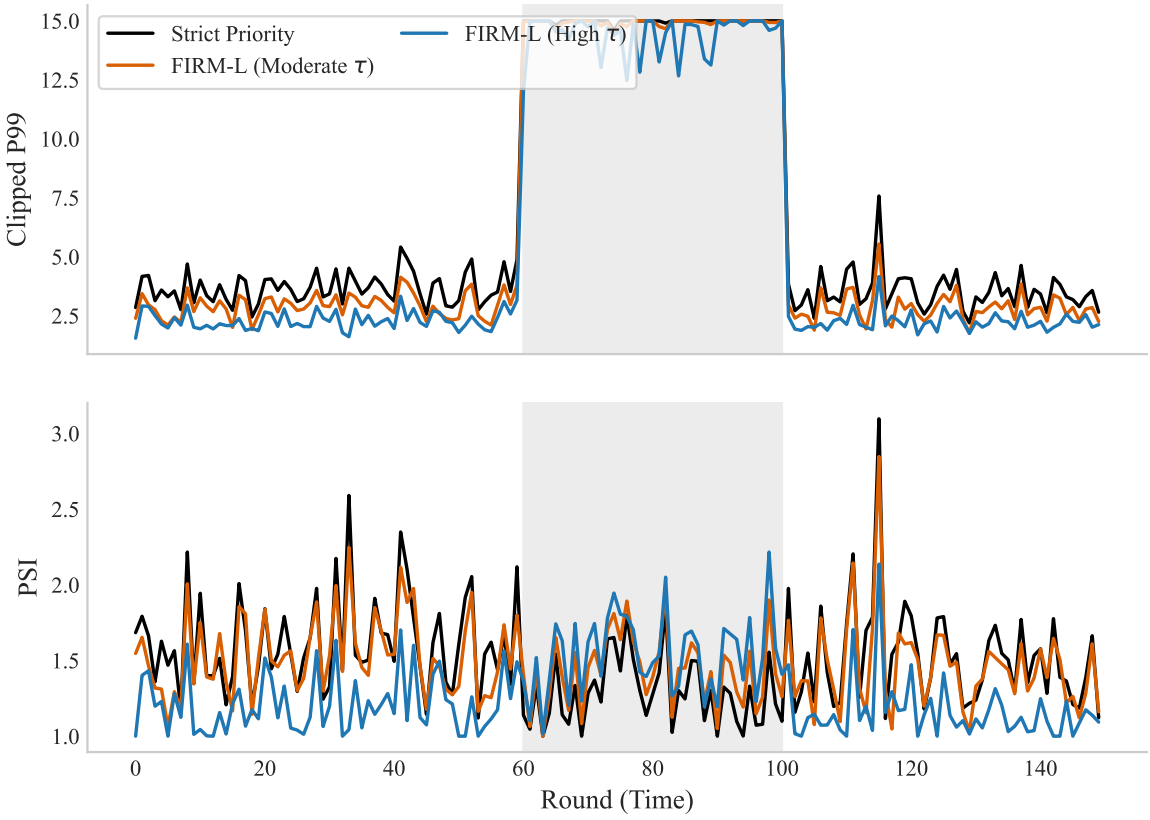
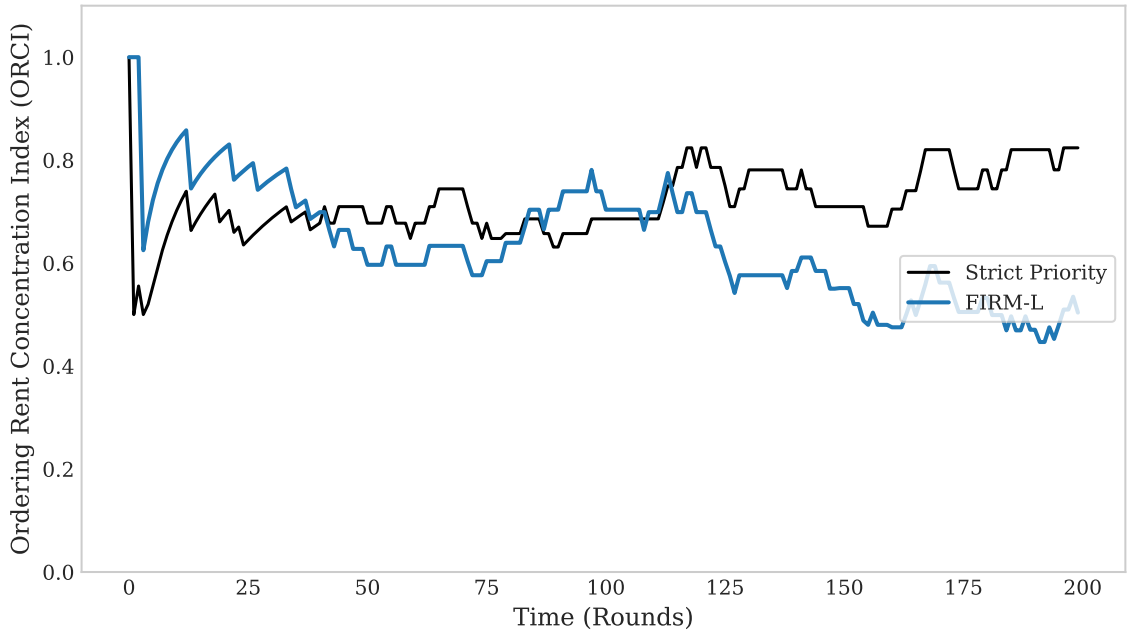


Figure 7: Tail behavior around shock windows. Notes: The figure plots clipped unit priority-price tails over the first shock episode, with 45 rounds of padding on each side, reporting the round-level 99th percentile and the Payment Spike Index (PSI) defined in Section 8. The monitoring set is restricted to materially exposed agents with $x_i^\theta \geq \underline{x} = 0.30$, and unit prices are clipped at $\overline{u} = 15$. We compare strict priority with FIRM-L at $\tau = 0.14$ and $\tau = 0.24$. Shaded regions denote the exogenous shock window.

criteria, and randomness into operational statistics that policymakers and market operators will inevitably inspect: expected execution priority, congestion measures, payment-pressure indices, and tail-risk indicators. The design of the mechanism therefore influences not only allocative outcomes but also the stability and interpretability of the monitoring objects built from its outputs.

Our claim is not that FIRM-L by itself solves the general problem of official-statistics production. The narrower and defensible point is that a mechanism with explicit sensitivity certificates is a better generator of policy-relevant monitoring statistics than a discontinuous strict-priority rule, because the mapping from reported conditions to expectation-level

Figure 8: Concentration Risk under Intermediated Execution



Notes: The figure plots the Ordering Rent Concentration Index (ORCI) over time for strict priority and FIRM-L under a simulated Proposer-Builder Separation (PBS) environment. Higher values indicate greater concentration.

Figure 8: Concentration risk under intermediated execution. Notes: The figure plots the Ordering Rent Concentration Index (ORCI) over time for strict priority and FIRM-L under a simulated Proposer-Builder Separation (PBS) environment. Higher values indicate greater concentration of value extraction among a few dominant builders.

measurement objects can be bounded, audited, and compared across policy choices. In that sense, the measurement contribution of FIRM-L is institutional and operational: it makes the production of governance-relevant statistics less fragile to microscopic perturbations in bids or market conditions.

11.2 Within-Agent Stability vs. Cross-Agent Fairness

We have operationalized fairness as *within-agent opportunity stability*: conditional on eligibility, small changes in an agent's bid should not result in disproportionate jumps in expected priority weight. This definition allows us to obtain an auditable, parameter-explicit certificate tied to τ. It is important to clarify the scope of this claim. We do not assert that FIRM-L resolves all normative fairness concerns, particularly those related to *cross-agent* distributional parity, which requires defining contentious similarity

metrics across heterogeneous agents. Instead, our contribution is a minimal, procedural fairness notion that addresses the specific externality of deterministic sorting—the "cliff edge"—while remaining compatible with the incentives of a high-frequency financial market.

11.3 Policy Realism: Intermediated Execution

The institutional reality of transaction ordering—characterized by Proposer–Builder Separation (PBS) and relay intermediation—motivates our focus on low-overhead implementation. Efficient PL sampling and single-call transfer computation keep the online procedure compatible with tight latency budgets. At the same time, our auditability claim is deliberately conditional rather than absolute. Replayability of the announced rule is achievable given logged randomness and code; expectation-level verification of the stability certificate is achievable given repeated simulation or analytical evaluation; but both exercises remain conditional on complete and authenticated visibility of the eligible bid set. In intermediated environments, data availability at the relay/builder interface is therefore part of the governance problem rather than something solved automatically by the mechanism.

11.4 Robust Governance via Pre-Commitment

Pre-commitment is not merely institutional window dressing; it is the mechanism-design constraint that separates rule-based governance from discretionary ex-post intervention. The platform announces the policy menu, loss functional, update frequency, and learning rule before bids arrive. This removes the most direct channel through which the platform could manipulate contemporaneous allocations in response to particular bids.

A different concern is strategic manipulation of the *governance signal*. Sophisticated agents may attempt to influence future parameter updates by inducing extreme realized losses—for example through burst bidding, order splitting, spammy near-cutoff submissions, or the strategic timing of high-variance orderflow during rounds that receive large weight in the controller update. Our main theorem does not rule out such behavior. The Hedge

guarantee is a performance statement conditional on the logged loss process, not a proof that the loss process itself is strategy-proof.

The correct way to read the governance result is therefore layered. At the *allocation layer*, FIRM-L is truthful in expectation for a fixed announced policy. At the *governance layer*, the announced update rule is auditable and admits static-regret guarantees under the policy-invariant full-information benchmark of Section 9. Strategic robustness of that governance layer requires additional design choices. In practice, three measures are especially important: (i) low-frequency or batched parameter updates, so that no single round can materially shift the next policy; (ii) clipped and winsorized loss components, especially for tail-payment and instability terms; and (iii) public commitment to a finite policy menu together with minimum dwell times or smoothing of the controller weights. These design choices do not make manipulation impossible, but they raise its cost and keep the update rule transparent enough for ex-post audit.

A full game-theoretic analysis of strategic interaction between bidders and the policy-update rule is an important extension. It would require modelling how agents trade off current-round allocation rents against the future effect of induced policy changes. We view this as a governance-layer mechanism-design problem in its own right, not as something silently resolved by the present static-regret theorem.

11.5 Limitations and Future Work

Our analysis is intentionally built in a single-parameter environment because both the monotonicity result and the Myerson payment characterization exploit a scalar private type. This choice is not innocuous. In transaction-ordering environments, sophisticated agents may have multi-dimensional private information: not only willingness to pay for priority, but also cross-transaction complementarities, latency technology, bundle value, and state-contingent liquidation or arbitrage opportunities. In such environments, monotonicity of a one-dimensional interim allocation is no longer sufficient for truthful implementation. The correct implementability object is cyclical monotonicity, and the simple envelope formula used here generally fails without stronger structure (Rochet,

1987).

What is likely to extend is the *allocation-side* insight that score-based randomization can smooth deterministic cliffs and generate explicit sensitivity certificates. What does *not* extend automatically is the clean payment characterization and the single-call truthful wrapper. A multi-parameter analogue would require either a restricted type space, a different transfer architecture, or an approximation-based notion of incentive compatibility.

A second limitation is the hard reserve. Our stability notion is conditional on eligibility, so the reserve threshold r creates an unavoidable discontinuity at entry. Soft-reserve or probabilistic-gating mechanisms may deliver global continuity, but doing so while preserving implementability and interpretability would require a distinct analysis.

Finally, our governance theorem uses a full-information, policy-invariant replay benchmark. This is the right benchmark for auditable ex-post policy comparison in settings with a trusted simulator or replay environment, but it is not the last word on adaptive governance. When participation or the type distribution is policy-sensitive, the platform faces a partial-information problem, and formal guarantees should be derived from bandit or switching-comparator frameworks rather than imported from the full-information case.

12 Conclusion

This paper develops **FIRM-L**, a reserve-and-temperature family of randomized priority mechanisms for environments in which strict rank-by-bid rules create discontinuous and difficult-to-audit outcomes. In a single-parameter setting, we show that the reserve-screened weighted top-K Plackett–Luce allocation is monotone in an agent's own bid, which yields truthful implementation in expectation, and we derive an explicit eligible-region sensitivity certificate of order $O(1/\tau)$. These results convert the temperature parameter from a heuristic tuning device into a formal governance instrument: it controls the amount of smoothing, the magnitude of the stability certificate, and the behavior of derived expectation-level monitoring indices.

Operationally, the base allocation rule admits a low-overhead implementation, and

truthful transfers can be computed through a single-call payment wrapper. On the governance side, we analyze pre-committed policy choice over (r, τ) with a multiplicative-weights controller under a full-information benchmark and obtain regret relative to the best fixed policy in hindsight from a finite announced menu. Quantitative exercises illustrate the resulting efficiency–stability trade-off and the policy value of explicit stability certificates.

Taken together, these results show that one can replace cliff-edge priority competition with a mechanism family that is incentive-compatible, explicitly certifiable, and suitable for policy-relevant measurement in blockchain transaction ordering and related priority markets.

References

Afèche, P., and H. Mendelson. 2004. "Pricing and Priority Auctions in Queueing Systems with a Generalized Delay Cost Structure." *Management Science* 50 (7): 869–882.

Aquilina, M., E. Budish, and P. O'Neill. 2022. "Quantifying the High-Frequency Trading 'Arms Race'." *The Quarterly Journal of Economics* 137 (1): 493–564.

Archer, A., and É. Tardos. 2001. "Truthful Mechanisms for One-Parameter Agents." In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science (FOCS)*, 482–491. IEEE.

Auer, P., N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. 2002. "The Nonstochastic Multiarmed Bandit Problem." *SIAM Journal on Computing* 32 (1): 48–77.

Babaioff, M., R. D. Kleinberg, and A. Slivkins. 2015. "Truthful Mechanisms with Implicit Payment Computation." *Journal of the ACM* 62 (2): Article 10.

Bahrani, M., P. Garimidi, and T. Roughgarden. 2025. "Centralization in Block-Building and Proposer-Builder Separation." In *Financial Cryptography and Data Security (FC 2024)*, Lecture Notes in Computer Science 14744, 331–349. Springer.

Bubeck, S., and N. Cesa-Bianchi. 2012. "Regret Analysis of Stochastic and Nonstochastic Multi-Armed Bandit Problems." *Foundations and Trends in Machine Learning* 5 (1): 1–122.

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

Cesa-Bianchi, N., and G. Lugosi. 2006. *Prediction, Learning, and Games*. Cambridge: Cambridge University Press.

Daian, P., S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, and A. Juels. 2020. "Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable

Value, and Consensus Instability." In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE.

Dwork, C., M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. 2012. "Fairness Through Awareness." In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS)*, 214–226.

Ferreira Torres, C., R. Camino, and R. State. 2021. "Frontrunner Jones and the Raiders of the Dark Forest: An Empirical Study of Frontrunning on the Ethereum Blockchain." In *30th USENIX Security Symposium*, 1343–1359.

Flashbots. 2023. "MEV-Boost: Transparency and Ecosystem Metrics." *Flashbots Transparency Reports*. Available at https://docs.flashbots.net.

Freund, Y., and R. E. Schapire. 1997. "A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting." *Journal of Computer and System Sciences* 55 (1): 119–139.

Guiver, J., and E. Snelson. 2009. "Bayesian Inference for Plackett–Luce Ranking Models." In *Proceedings of the 26th International Conference on Machine Learning (ICML)*.

Heimbach, L., L. Kiffer, C. F. Torres, and R. Wattenhofer. 2023. "Ethereum's Proposer-Builder Separation: Promises and Realities." arXiv preprint arXiv:2305.19037; also ACM IMC '23.

Kool, W., H. van Hoof, and M. Welling. 2019. "Stochastic Beams and Where to Find Them: The Gumbel-Top-k Trick for Sampling Sequences Without Replacement." In *Proceedings of ICML 2019*, 3499–3508.

Kool, W., H. van Hoof, and M. Welling. 2020. "Ancestral Gumbel-Top-k Sampling for Sampling Without Replacement." *Journal of Machine Learning Research* 21: 1–36.

Li, Y., and S. Goel. 2024. "Making It Possible for the Auditing of AI: A Systematic Review of AI Audits and AI Auditability." *Information Systems Frontiers* 27: 1121–1151.

Liu, Y., et al. 2022. "Empirical Analysis of EIP-1559: Transaction Fees, Waiting Time, and Consensus Security." arXiv preprint arXiv:2201.05574; also ACM CCS 2022.

Luce, R. D. 1959. *Individual Choice Behavior: A Theoretical Analysis*. New York: Wiley.

Maystre, L., and M. Grossglauser. 2015. "Fast and Accurate Inference of Plackett–Luce Models." In *Advances in Neural Information Processing Systems (NeurIPS)*.

Myerson, R. B. 1981. "Optimal Auction Design." *Mathematics of Operations Research* 6 (1): 58–73.

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

Plackett, R. L. 1975. "The Analysis of Permutations." *Applied Statistics* 24 (2): 193–202.

Qin, K., L. Zhou, and A. Gervais. 2022. "Quantifying Blockchain Extractable Value: How Dark Is the Forest?" In *2022 IEEE Symposium on Security and Privacy (SP)*.

Rochet, J.-C. 1987. "A Necessary and Sufficient Condition for Rationalizability in a Quasi-Linear Context." *Journal of Mathematical Economics* 16 (2): 191–200.

Roughgarden, T. 2020. "Transaction Fee Mechanism Design for the Ethereum Blockchain: An Economic Analysis of EIP-1559." Report, December 1, 2020.

Roughgarden, T. 2021. "Transaction Fee Mechanism Design." arXiv preprint arXiv:2106.01340.

UNECE. 2025. *Generative Artificial Intelligence in Official Statistics: Mitigating and Monitoring*. United Nations Economic Commission for Europe.

Weintraub, B., C. F. Torres, C. Nita-Rotaru, and R. State. 2022. "A Flash (bot) in the Pan: Measuring Maximal Extractable Value in Private Pools." arXiv preprint arXiv:2206.04185; also ACM IMC 2022.

A Proofs and Technical Lemmas

This appendix formalizes the stochastic objects used in the paper and provides complete proofs for (i) the monotonicity of FIRM-L's interim allocation, (ii) the eligible-region Lipschitz stability bound, and (iii) a general measurement-robustness transfer lemma connecting mechanism stability to derived indices. The exposition follows standard algorithmic mechanism design conventions, distinguishing between mechanism randomness and exogenous "nature" randomness (Babaioff et al., 2015).

A.1 Notation and Measurability Preliminaries

A.1.1 Measurable Spaces and Stochastic Kernels

Fix a round t and suppress the index. Let $N = \{1, \dots, n\}$ be the set of agents, K the number of slots, and $\mathcal{Y}$ the finite set of feasible assignments (partial matchings). We equip $\mathcal{Y}$ with the discrete σ-algebra $2^{\mathcal{Y}}$. The bid space is $\mathcal{B}^n$, where $\mathcal{B} \subseteq \mathbb{R}_+$ is Borel measurable. An allocation rule is a stochastic kernel A from $(\mathcal{B}^n, \mathfrak{B})$ to $(\mathcal{Y}, 2^{\mathcal{Y}})$. For each bid profile $\mathbf{b} \in \mathcal{B}^n$, $A(\cdot \mid \mathbf{b})$ is a probability measure on $\mathcal{Y}$. Since $\mathcal{Y}$ is finite, measurability follows immediately provided the probabilities are explicit functions of bids.

A.1.2 Exposure Probabilities and Interim Allocation

For agent i and slot $p \in \{1, \dots, K\}$, let $y_{i,p}(y) \in \{0, 1\}$ be the indicator that i is assigned slot p in outcome y. Define the *exposure probability*:

$$q_{i,p}(\mathbf{b}) := \mathbb{E}_{y \sim A(\cdot|\mathbf{b})}[y_{i,p}(y)]. \tag{14}$$

Given priority weights $w_1 \geq \cdots \geq w_K \geq 0$, the realized priority weight is $X_i(y) := \sum_{p=1}^{K} w_p y_{i,p}(y)$. The *interim allocation* is:

$$x_i(\mathbf{b}) := \mathbb{E}_{y \sim A(\cdot|\mathbf{b})}[X_i(y)] = \sum_{p=1}^{K} w_p q_{i,p}(\mathbf{b}). \tag{15}$$

This is the canonical single-parameter allocation statistic used to establish implementability.

A.1.3 Randomness: Mechanism vs. Nature

We explicitly distinguish two sources of randomness:

- **Mechanism Randomness (U):** Internal random seeds used for sampling (e.g., Plackett–Luce sampling, payment resampling).
- **Nature Randomness (Z):** Exogenous stochasticity (e.g., click realizations), which we condition on for incentive analysis.

When we write $\mathbb{E}[\cdot \mid \mathbf{b}]$, the expectation is taken over U. Incentive properties are typically required to hold in expectation over U for every realization of Z (an ex-post perspective w.r.t. nature).

A.2 Monotonicity of Interim Allocation

We prove that the interim allocation induced by FIRM-L is monotone, satisfying the necessary condition for truthful implementation.

Theorem 3 (Eligible-Region Monotonicity)**.** *Fix an opponent profile $\mathbf{b}_{-i}$. Under FIRM-L with reserve r and temperature τ, the function $b_i \mapsto x_i(b_i; \mathbf{b}_{-i})$ is non-decreasing on $[r, \infty)$ and equals 0 on $(-\infty, r)$.*

Proof. Let $\alpha_j(b_j) = \exp((b_j - r)/\tau)$ for $b_j \geq r$, and 0 otherwise. The Plackett–Luce (PL) top-K ranking can be generated by an "exponential race": draw independent arrival times $T_j \sim \text{Exp}(\alpha_j)$ and rank agents by increasing T_j.

Fix $\mathbf{b}_{-i}$ and consider two eligible bids $b_i \leq b_i'$ (both $\geq r$). Let $\alpha_i = \exp((b_i - r)/\tau)$ and $\alpha_i' = \exp((b_i' - r)/\tau)$, so $\alpha_i' \geq \alpha_i$. We construct a pathwise coupling using i.i.d. variables $E_j \sim \text{Exp}(1)$ for all j. Define arrival times:

$$T_j = E_j/\alpha_j \quad \text{under } b_i,$$
$$T_j' = E_j/\alpha_j \quad \text{for } j \neq i, \quad \text{and} \quad T_i' = E_i/\alpha_i' \quad \text{under } b_i'.$$

Since $\alpha_i' \geq \alpha_i$, we have $T_i' \leq T_i$ for every realization of E_i, while all other T_j remain unchanged. Thus, in every sample path, agent i's rank (arrival position) under $\mathbf{b}'$ is numerically less than or equal to their rank under $\mathbf{b}$.

Let $g(k) = w_k$ for $k \leq K$ and 0 otherwise. Since weights are non-increasing, $g(\text{rank})$ is non-increasing in numerical rank. Therefore, the realized weight $X_i' = g(\text{rank}_i') \geq g(\text{rank}_i) = X_i$ pointwise. Taking expectations yields $x_i(b_i'; \mathbf{b}_{-i}) \geq x_i(b_i; \mathbf{b}_{-i})$. □

A.3 Eligible-Region Lipschitz Bound

We now give a complete proof of the eligible-region stability certificate stated in Theorem 1.

Theorem 4 (Lipschitz Stability)**.** *Under FIRM-L, for any fixed* $\mathbf{b}_{-i}$ *and eligible bids* $b_i, b_i' \geq r$*,*

$$|x_i(b_i; \mathbf{b}_{-i}) - x_i(b_i'; \mathbf{b}_{-i})| \leq \frac{\sum_{p=1}^{K} w_p}{\tau} |b_i - b_i'|.$$

Proof. Fix $\mathbf{b}_{-i}$ and write $s_i = (b_i - r)/\tau$ and $\lambda_i = e^{s_i}$. In the exponential-race representation of PL, for each eligible opponent $j \neq i$ let $T_j \sim \text{Exp}(\lambda_j)$, and denote the order statistics of these opponent arrival times by

$$O_{(1)} \leq \cdots \leq O_{(m)},$$

where m is the number of eligible opponents. Set $O_{(0)} := 0$ and $O_{(m+1)} := \infty$.

Conditional on T_{-i}, agent i receives slot p exactly when $O_{(p-1)} < T_i < O_{(p)}$, where $T_i \sim \text{Exp}(\lambda_i)$. Therefore

$$q_{i,p}(b_i; \mathbf{b}_{-i} \mid T_{-i}) = \Pr\left(O_{(p-1)} < T_i < O_{(p)} \mid T_{-i}\right) = e^{-\lambda_i O_{(p-1)}} - e^{-\lambda_i O_{(p)}}.$$

Differentiating with respect to s_i yields

$$\frac{\partial q_{i,p}}{\partial s_i}(b_i; \mathbf{b}_{-i} \mid T_{-i}) = \lambda_i \left[O_{(p)} e^{-\lambda_i O_{(p)}} - O_{(p-1)} e^{-\lambda_i O_{(p-1)}}\right].$$

Because $xe^{-x} \leq 1/e$ for all $x \geq 0$,

$$\left| \frac{\partial q_{i,p}}{\partial s_i}(b_i; \mathbf{b}_{-i} \mid T_{-i}) \right| \leq \frac{2}{e} < 1.$$

The bound is integrable and does not depend on the realization of T_{-i}, so dominated convergence implies

$$\left| \frac{\partial q_{i,p}}{\partial s_i}(b_i; \mathbf{b}_{-i}) \right| \leq 1.$$

Summing across slots,

$$\left| \frac{\partial x_i}{\partial s_i}(b_i; \mathbf{b}_{-i}) \right| = \left| \sum_{p=1}^{K} w_p \frac{\partial q_{i,p}}{\partial s_i}(b_i; \mathbf{b}_{-i}) \right| \leq \sum_{p=1}^{K} w_p.$$

Since $\partial s_i / \partial b_i = 1/\tau$,

$$\left| \frac{\partial x_i}{\partial b_i}(b_i; \mathbf{b}_{-i}) \right| \leq \frac{1}{\tau} \sum_{p=1}^{K} w_p.$$

Applying the mean-value theorem on the interval between b_i and b_i' proves the theorem. □

Proposition 4 (Exact constant for $K = 1$)**.** *When $K = 1$, the bound tightens to*

$$\mathrm{Lip}_i \leq \frac{w_1}{4\tau}.$$

Proof. For $K = 1$, the PL rule reduces to a softmax over eligible agents:

$$q_{i,1}(b_i; \mathbf{b}_{-i}) = \frac{\alpha_i}{\alpha_i + \sum_{j \neq i} \alpha_j}, \qquad \alpha_i = e^{(b_i - r)/\tau}.$$

Differentiating with respect to $s_i = \log \alpha_i$ gives

$$\frac{\partial q_{i,1}}{\partial s_i} = q_{i,1}(1 - q_{i,1}) \leq \frac{1}{4}.$$

Therefore

$$\left| \frac{\partial x_i}{\partial b_i} \right| = \frac{w_1}{\tau} \left| \frac{\partial q_{i,1}}{\partial s_i} \right| \leq \frac{w_1}{4\tau}.$$

□

A.4 Measurement Robustness Lemma

This subsection formalizes how stability of the mechanism transfers to clipped statistical indices constructed from interim expected allocations and expected payments.

Lemma 3 (Index Stability Transfer)**.** *Let $I(\mathbf{b}) = F\big(x(\mathbf{b}), p(\mathbf{b})\big)$, where $x(\mathbf{b}) \in \mathbb{R}^n_+$ is the vector of interim expected priority weights and $p(\mathbf{b}) \in \mathbb{R}^n_+$ is the vector of expected truthful payments. Suppose that on the monitoring domain $[r, \bar{b}]$,*

$$\|x(b_i; \mathbf{b}_{-i}) - x(b'_i; \mathbf{b}_{-i})\|_\infty \leq S_x |b_i - b'_i|$$

and

$$\|p(b_i; \mathbf{b}_{-i}) - p(b'_i; \mathbf{b}_{-i})\|_\infty \leq S_p |b_i - b'_i|.$$

If F is Lipschitz on the clipped domain with constants L_x and L_p, then

$$|I(b_i; \mathbf{b}_{-i}) - I(b'_i; \mathbf{b}_{-i})| \leq (L_x S_x + L_p S_p) |b_i - b'_i|.$$

Proof. By Lipschitz continuity of F,

$$|F(x, p) - F(x', p')| \leq L_x \|x - x'\|_\infty + L_p \|p - p'\|_\infty.$$

Applying the assumed bounds for x and p gives the claim immediately. □

Corollary 4 (Expectation audit versus replay audit)**.** *A one-seed replay audit verifies that the implemented code reproduces the announced realized allocation and payment given the logged inputs and randomness. An expectation-level stability audit verifies the sensitivity of $x_i(\cdot; \mathbf{b}_{-i})$ or of a derived index by analytical evaluation or repeated simulation with common random numbers. These are distinct exercises, and Theorem 1 pertains to the second one.*

B Payments and Implementation Details

This appendix provides the formal characterization of the payment rule and details its efficient implementation via implicit payment computation.

B.1 Payment Characterization

B.1.1 Setup and Monotonicity

We consider a single-parameter environment where agent i's private type v_i is their value per unit of priority weight. Let $x_i(b_i; \mathbf{b}_{-i})$ denote the interim expected priority weight. In Appendix A.2, we established that FIRM-L's allocation rule is monotone: $x_i(b_i; \mathbf{b}_{-i})$ is non-decreasing in b_i on the eligible domain $[r, \infty)$. This monotonicity is the necessary and sufficient condition for *Truthfulness-in-Expectation (TIE)* implementation (Myerson, 1981).

B.1.2 The Envelope Identity

Standard single-parameter mechanism design implies that any truthful-in-expectation implementation of a monotone interim allocation must satisfy

$$p_i(b_i; \mathbf{b}_{-i}) = b_i x_i(b_i; \mathbf{b}_{-i}) - \int_0^{b_i} x_i(z; \mathbf{b}_{-i})\, dz + C_i(\mathbf{b}_{-i}).$$

Because FIRM-L uses a reserve rule, $x_i(z; \mathbf{b}_{-i}) = 0$ for all $z < r$. With the standard IR normalization $C_i(\mathbf{b}_{-i}) = 0$, the expected payment becomes

$$p_i(b_i; \mathbf{b}_{-i}) = \begin{cases} 0, & b_i < r, \\ b_i x_i(b_i; \mathbf{b}_{-i}) - \displaystyle\int_r^{b_i} x_i(z; \mathbf{b}_{-i})\, dz, & b_i \geq r. \end{cases}$$

Thus strictly below the reserve an agent is ineligible and pays zero. At the boundary $b_i = r$, however, eligibility is active under our convention $b_i \geq r$, so in general

$$p_i(r; \mathbf{b}_{-i}) = r\, x_i(r; \mathbf{b}_{-i}).$$

This is the correct boundary statement.

Truthful interim utility is therefore

$$U_i(v_i; v_i, \mathbf{b}_{-i}) = \begin{cases} 0, & v_i < r, \\ \displaystyle\int_r^{v_i} x_i(z; \mathbf{b}_{-i})\, dz, & v_i \geq r, \end{cases}$$

which is non-negative because $x_i(z; \mathbf{b}_{-i}) \geq 0$.

B.2 Single-Call Wrapper for Deployable Payments

Let A^θ denote the monotone base allocation rule under policy $\theta = (r, \tau)$, and let $X_i(A^\theta(\mathbf{b}))$ be the realized priority weight assigned to agent i by the base rule. We use the self-resampling framework of Babaioff et al. (2015).

Fix a resampling parameter $\mu \in (0, 1)$. For each agent i, given reported bid b_i, independently draw an auxiliary pair (χ_i, η_i) from a self-resampling procedure satisfying:

1. $0 \leq \eta_i \leq \chi_i \leq b_i$ almost surely;
2. with probability $1 - \mu$, $\chi_i = \eta_i = b_i$;
3. conditional on $\eta_i < b_i$, the variable η_i has a continuously differentiable cdf $F_i(\cdot \mid b_i)$ on $[0, b_i)$ with strictly positive density $f_i(\cdot \mid b_i)$;
4. the procedure is monotone in the reported bid, in the sense required by Babaioff et al. (2015).

Run the base allocation rule once on the modified bid vector $\chi = (\chi_1, \ldots, \chi_n)$, producing realized weights $X_i(A^\theta(\chi))$. The payment estimator takes the generic form

$$\widetilde{p}_i = b_i\, X_i(A^\theta(\chi)) - \mathbf{1}\{\eta_i < b_i\} \frac{X_i(A^\theta(\chi))}{\mu\, f_i(\eta_i \mid b_i)}.$$

Under the regularity conditions of Babaioff et al. (2015, Theorem 4), this wrapper is truthful in expectation and implements the Myerson payment of the base monotone rule

in expectation:

$$\mathbb{E}[\widetilde{p}_i \mid \mathbf{b}] = p_i^{\theta}(\mathbf{b}).$$

The paper's theoretical object is therefore the pair (A^{θ}, p^{θ}): the base FIRM-L allocation rule together with its Myerson payment rule. The single-call resampling layer is the online implementation device used to avoid repeated counterfactual calls to the allocator.

Audit log for the payment wrapper. A third-party verifier needs the following objects in the log: (i) the announced policy parameters $(r, \tau, \mathbf{w})$; (ii) the reported bid vector $\mathbf{b}$; (iii) the self-resampling parameter μ and the committed randomness used to generate $(\chi_i, \eta_i)_{i=1}^n$; (iv) the realized allocation produced by the base PL sampler; and (v) the computed transfers. Given these records, the verifier can replay the wrapper exactly and check that the transfer computation matches the announced truthful implementation.

C Supplementary Tables

Table 2: Baseline simulation parameters

Parameter	Value used in the replication code
Horizon T	2000 rounds
Number of slots K	6
Slot weights $\mathbf{w}$	$(1.00, 0.88, 0.76, 0.64, 0.54, 0.46)$
Reserve rule r	0.85 with eligibility convention $b_i \geq r$
Temperature sets	Slope grid $(0.18, 0.24, 0.32, 0.44, 0.60, 0.82, 1.10, 1.50, 2.05, 2.80)$; Hedge menu $\Theta = (0.18, 0.32, 0.60, 1.10, 2.05)$; tail comparison $(0.14, 0.24)$
Regime transition matrix	Calm $\to (0.965, 0.035)$; Storm $\to (0.090, 0.910)$
Arrival rates $(\lambda_{\text{Calm}}, \lambda_{\text{Storm}})$	$(100, 500)$
Log-normal parameters	Calm: $(\mu, \sigma) = (0.00, 0.55)$; Storm: $(0.18, 0.78)$
Pareto parameters	Calm: $(x_m, \alpha) = (3.2, 3.4)$; Storm: $(3.8, 2.7)$; shock override $\alpha = 1.75$
Shock windows	$[800, 860]$ and $[1500, 1560]$
Monitoring thresholds	$\underline{x} = 0.30$, $\underline{u} = 0.30$, $\overline{u} = 15$
Replication count	Table 1: 12 independent macro-replications; slope validation: 24 profiles, 11 bid-grid points, 1,536 PL draws per evaluation; governance replay: $(56, 72, 96)$ draws for allocation, payment, and slope terms; tail summaries: $(96, 128)$ draws for allocation and payment terms
Learning-rate rule	Hedge uses batch size $B = 4$, update horizon $U = T/B = 500$, and $\eta = \sqrt{2 \ln \lvert\Theta\rvert / U} \approx 0.0802$
Random seeds	Environment 20260306; slopes 314159; governance 271828; tails 161803

Notes: This table reports the machine-readable manifest currently used by the released replication code in `firm_l_real_simulation.py`. The common slot weights, reserve rule, and monitoring thresholds are shared across Figures 2, 4, and 7; figure-specific Monte Carlo budgets are listed explicitly because the slope, governance, and tail diagnostics use different replay granularities.

Table 3: Released-manifest diagnostics for the simulation figures

Object	Estimate	Diagnostic details
Figure 6: OLS slope on $1/\tau$	0.0446	95% CI [0.0409, 0.0482]; intercept 0.1155; $R^2 = 0.1765$; $n = 2{,}640$ local-slope estimates
Figure 6: profile bank	24 profiles	12 calm and 12 storm profiles; 11 eligible bid-grid points per profile; 95th-percentile and median summaries plotted in Figure 2
Figure 4: final static regret	11.2633	Hedge menu $\Theta = (0.18, 0.32, 0.60, 1.10, 2.05)$; final policy probabilities $(0.4919, 0.3881, 0.1129, 0.0068, 0.0002)$
Figure 4: theoretical regret bound	40.1178	Computed as $\sqrt{2U \ln \lvert\Theta\rvert}$ with $U = 500$ batched updates and $\lvert\Theta\rvert = 5$
Figure 4: learning rate	0.0802	Computed as $\eta = \sqrt{2 \ln \lvert\Theta\rvert / U}$ with $U = 500$ under four-round batching over a 2,000-round horizon
Figure 7: strict priority mean tail level	P99 = 4.6515	Mean PSI = 1.1602 over the displayed first-shock window with the materially exposed monitoring set
Figure 7: FIRM-L at $\tau = 0.14$	P99 = 4.6212	Mean PSI = 1.1456 over the displayed first-shock window
Figure 7: FIRM-L at $\tau = 0.24$	P99 = 4.6433	Mean PSI = 1.1449 over the displayed first-shock window

Notes: This appendix table reports auxiliary diagnostics emitted by the same unified pipeline that also generates Table 1. The slope-regression entries come directly from `figure2_6_slope_diagnostics.json`; the governance entry comes from `figure4_governance_diagnostics.json`; the tail entries come from `figure7_tail_diagnostics.json`; and the cross-replication frontier summary comes from `table1_macro_replications.json`. Main-table uncertainty is now reported directly in Table 1, while this appendix table records the complementary figure-specific diagnostics.