

Chicago Booth School of Business Economics Research Papers

Smoothing the Cliff: Fair and Incentive-Compatible Priority Allocation for Healthcare Resources

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 for Healthcare Resources*

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

Abstract

Can an AI allocation algorithm be designed to not only distribute critical healthcare resources efficiently but also generate reliable operational statistics? As AI tools reshape medical infrastructure and scheduling, traditional strict-priority or "winner-takes-all" triage rules are failing—they generate noisy, unstable data that undermines health system auditability and healthcare quality monitoring. To solve this, we propose FIRM-L (Fair and Incentive-compatible Randomized Mechanism) as an AI medical decision support tool. Think of it as a mechanism with a "volume knob" (temperature control) that turns chaotic cliff-edge queue competition into a smooth, predictable curve, reducing the administrative burden of constant rescheduling. We mathematically prove that FIRM-L ensures fairness (via Lipschitz stability) and honesty (truthfulness-in-expectation) without sacrificing the speed necessary for high-throughput clinical environments. Furthermore, we design an AI-driven controller that automatically tunes this mechanism in real-time, adapting to medical demand shocks (e.g., pandemic surges) to protect patient welfare while keeping the health system stable, efficient, and regulatory-compliant.

JEL Classification: I11, D44, D47, C73, D63

Keywords: AI in Healthcare, Medical Decision Support, Triage Allocation, Mechanism Design, Algorithmic Governance, Health Economics

*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 diagnostic imaging executes earlier, and whose request occupies scarce medical capacity has become a core design primitive in modern AI-driven digital health systems. In healthcare scheduling in particular, algorithmic ordering is not a neutral implementation detail: it directly affects patient outcomes, realized operational efficiency, and the distribution of administrative burdens across the health system (Daian et al., 2019). A salient empirical manifestation of this phenomenon is the emergence of algorithmic competition for early hospital resource scheduling—often described as priority escalation or "VIP queueing" in the context of decentralized health networks. In these settings, sophisticated AI scheduling agents optimize latency to obtain execution priority, and the resulting arms race can distort care quality and create broader system-level inefficiencies and risks.

A striking and economically consequential feature of many deployed medical priority rules is their cliff-like discontinuity. Under strict deterministic ranking by AI urgency scores, infinitesimal differences in a priority signal can induce macroscopic changes in rank, execution probability, and hospital wait times, yielding brittle outcomes that are hard to predict and difficult to justify to patients or regulators. This discontinuity is central to the algorithmic queueing dynamics documented in healthcare markets, where repeated re-evaluations for marginal priority advantages can amplify volatility and administrative waste.

These discontinuities matter for at least three reasons. First, they create within-round volatility: tiny perturbations in AI priority signals can flip allocation outcomes, producing sharp swings in realized utility for patients with nearly identical health needs. Second, they foster wasteful escalation—costly systemic friction and administrative burden—consistent with the competitive priority dynamics observed in automated transaction ordering. Third,

they undermine the stability and auditability of high-frequency statistics derived from the health system itself, such as measures of emergency department congestion, effective care delays, or operational intensity, precisely when such measures are most valuable for health policy and regulatory monitoring.

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 AI decision support rule.

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, we implement truthfulness in expectation using Myerson's payment identity. To make incentive compatible implementation practical in high-throughput systems—where repeatedly re-running the allocation rule for counterfactual bids may be computationally infeasible—we use implicit payment computation. This single-call black-box reduction transforms any monotone allocation rule into a truthful-in-expectation mechanism while evaluating the allocation rule only once (up to lightweight resampling) (Babaioff et al., 2015). This approach is particularly well suited to blockchain and other real-time priority markets, where latency and computational overhead are first-order constraints.

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 have direct implications for blockchain transaction ordering and related financial-market infrastructures. First, they reduce the incentive for marginal

fee escalation aimed at capturing tiny priority advantages—thereby mitigating payment spikes and allocation turnover in congested regimes that resemble the priority gas auctions documented in practice. Second, stability bounds enable construction of congestion-adjusted high-frequency indices from observed bids, allocations, and payments that are robust to micro-perturbations, improving the reliability of real-time measurement under stress. In this sense, the allocation mechanism becomes a "statistical generator" whose sensitivity is bounded and therefore auditable.

A practical challenge is that the "right" amount of smoothing depends on market conditions. Blockchain demand is highly non-stationary: volatility shocks, arbitrage opportunities, and congestion events produce regime changes in participation and valuations, and any fixed parameter choice may be suboptimal over time. We therefore study dynamic parameter selection under a pre-commitment model in which the platform publicly announces its parameters before observing contemporaneous bids, limiting manipulation of the policy itself. We cast parameter tuning as an online learning problem over a discretized policy class and analyze multiplicative-weights (Hedge) updates, a classic approach with well-known regret guarantees (Freund and Schapire, 1997). The learning perspective also connects FIRM-L to a broader literature on decision-making under adversarial or non-stationary conditions. In settings with partial information, bandit-style feedback can be relevant, and the foundations for adversarial bandit regret bounds provide a natural toolkit for extending the controller beyond full-information loss observation (Auer et al., 2002). While our core results focus on the pre-commitment, full-information benchmark for clarity, the framework is designed to accommodate stronger robustness goals.

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 Practical Payments.** We prove monotonicity of the interim allocation induced by PL top-K ranking via an exponential-race coupling and implement truthfulness in expectation via Myerson payments, using implicit payment computation for single-call practicality.

3. **Auditable Fairness via Stability.** We formalize within-round fairness as eligible-region Lipschitz stability of the bid-to-allocation mapping, deriving explicit sensitivity bounds that scale as $O(1/\tau)$, yielding a transparent stability–efficiency trade-off.

4. **Learning-Augmented Governance.** We model parameter selection as an online learning problem under non-stationarity. We employ a multiplicative-weights controller under pre-commitment to adapt parameters to changing demand shocks, providing provable regret guarantees.

1.2 Roadmap

The remainder of the paper proceeds as follows. Sections 2 and 3 review the institutional context of 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: AI-Driven Healthcare Scheduling & Priority Allocation

2.1 Clinical Resource Ordering as a Mechanism Design Problem

Modern digital health systems transform what is usually a backend administrative choice—how to order and schedule patient requests—into an explicit economic mechanism. In hospital networks, resource ordering determines outcomes whenever execution is time-sensitive (e.g., operating room access, imaging, and specialist consultations). The resulting efficiency rents are systemic: departments or AI agents compete for earlier execution, a phenomenon structurally analogous to the high-frequency priority competition documented in digital infrastructure markets (Daian et al., 2020). From a health economics perspective, this environment can be interpreted as a high-frequency queue with endogenous priority signals. When medical demand spikes, the shadow price of early execution rises, and observed priority scores become a revealed measure of scarcity and clinical urgency.

2.2 From Triage Queue to Treatment: Implementing Priority

Most healthcare networks maintain a pool of pending requests (a "triage pool") from which AI scheduling algorithms select and order tasks. Historically, scheduling closely resembled a decentralized design in which departments attach priority signals and administrators choose a subset and ordering to maximize throughput. This decentralized implementation introduces information heterogeneity and discretionary ordering (Daian et al., 2020; Ferreira Torres et al., 2021). A salient microstructure feature is the priority escalation dynamic: for scarce slots, departments or automated agents repeatedly revise urgency scores to outrank competitors, creating cliff-like discontinuities under strict priority sorting.

Microstructure parallel. The priority escalation 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 efficiency 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 in traditional markets; we pursue an analogous goal in medical AI allocation via a continuous smoothing parameter that explicitly bounds sensitivity and supports regulatory audit.

2.3 The Economic Content of Allocation Power

In healthcare, extractable value refers to the operational and clinical gains enabled by control over resource inclusion and ordering. Studies of analogous digital networks document that such value extraction can lead to large aggregate inefficiencies and tail events (Qin et al., 2022). Such tail behavior makes strict priority mechanisms fragile and difficult to govern, especially in congested regimes (e.g., flu season or pandemics) where strategic bidding escalation intensifies. Private routing infrastructures have emerged to mitigate negative externalities, though evidence suggests they may shift value extraction rather than eliminating administrative rents (Weintraub et al., 2022).

2.4 Modern Health Triage Protocols and Intermediation

Recent upgrades to hospital AI systems attempt to decouple base scheduling from priority escalation, mirroring structural market reforms (Roughgarden, 2020, 2021). While these reforms improve waiting time predictability, empirical analysis of similar systems suggests they do not entirely eliminate volatility or strategic distortions (Liu et al., 2022). Furthermore, modern health systems increasingly rely on decentralized departmental scheduling, structurally similar to Proposer-Builder Separation (PBS) in decentralized networks (Bahrani et al., 2025; Heimbach et al., 2023). This institutional reality reinforces two design constraints: (i) AI allocation rules must be implementable with low administrative overhead, and (ii) governance levers must be transparent and auditable for healthcare regulators.

2.5 Measurement and Regulatory Governance Motivation

A central measurement insight is that AI priority tools produce rich high-frequency data—urgency scores, wait times, and execution outcomes. However, strict priority rules can make derived statistics unstable because small input perturbations can flip discrete outcomes. The regulatory agenda explicitly highlights transparency and auditability for algorithmically produced decisions in healthcare, aligning with official-statistics discussions that mandate explicit monitoring frameworks (UNECE, 2025). This perspective supports interpreting a parameterized AI triage rule as a "statistical generator" whose sensitivity properties can be audited for healthcare quality control.

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 sequential structure naturally supports top-K selection without replacement, matching the allocation of a small number of scarce priority positions. The PL model is also computationally attractive. Bayesian and approximate-inference methods have been developed to handle sparse ranking data at scale (Guiver and Snelson, 2009), and spectral/Markov-chain perspectives yield fast and accurate estimation under Luce-type models (Maystre and Grossglauser, 2015). These results support PL as a practical building block in large-scale systems where allocation must be computed rapidly and repeatedly. Crucially for mechanism design, PL admits structural sampling representations that facilitate both implementation and analysis. Perturb-and-max constructions connect categorical sampling to Gumbel noise; and Gumbel-top-k extensions enable exact sampling without replacement by selecting the top k perturbed log-scores (Kool et al., 2019, 2020). These representations are valuable not only as efficient samplers but also as

proof tools, because they enable couplings that clarify how raising an agent's score shifts the distribution of ranks in a monotone direction. In short, we import PL from the ranking literature not for predictive accuracy but as a mechanism primitive: a smooth, parameterized distribution over rankings that replaces deterministic priority sorting with tunable randomization, while retaining tractability and analytic transparency through exponential-race/perturbation views.

3.2 Fairness & Online Learning: Stability and Governance

Priority markets invite objectives beyond welfare and revenue: designers must manage within-round volatility, justify outcomes, and produce stable high-frequency measures from the mechanism's outputs. We adopt a fairness notion inspired by "individual fairness"—treat similar individuals similarly—formalized via Lipschitz constraints that bound how much outcomes can change under small input changes (Dwork et al., 2012). This lens is particularly apt for strict priority rules, where discontinuities imply that infinitesimal bid differences can generate large allocation jumps, undermining stability and auditability (Dwork et al., 2012). Auditability concerns are now prominent in broader governance discussions of algorithmic systems. Work on official statistics highlights that stochastic and opaque systems may challenge transparency, reproducibility, and traceability, motivating monitoring and auditing as part of statistical production (UNECE, 2025). Complementary reviews in information systems synthesize how "auditability measures" can be operationalized across the AI lifecycle (Li and Goel, 2025). While our mechanism is not a foundation model, it is an algorithmic rule that generates statistics and therefore benefits from explicit sensitivity bounds that can be monitored and audited.

A second institutional feature is non-stationarity. Blockchain demand and MEV opportunities exhibit regime shifts and bursts; a fixed smoothing level can be suboptimal across states of the world. We therefore cast parameter choice as an online learning problem under pre-commitment, where the platform selects policies to minimize cumulative loss relative to the best fixed policy in hindsight. The multiplicative-weights (Hedge) framework provides the canonical regret guarantees for such expert-style aggregation problems (Freund

and Schapire, 1997). For partial-feedback or adversarial environments, adversarial bandit theory provides a robustness toolkit: nonstochastic multi-armed bandits formalize regret guarantees when the payoff process is adversarially chosen (Auer et al., 2002), and modern surveys systematize extensions and sharp bounds (Bubeck and Cesa-Bianchi, 2012). These results motivate learning-augmented governance that adapts to non-stationarity without sacrificing the transparency and incentive guarantees required in high-frequency priority markets.

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

Translation Note for Healthcare Context: While our mathematical framework and proofs utilize standard mechanism design terminology (bids, valuations, payments) to maintain formal rigor, these variables have direct interpretations in a health system context. The "bid" (b_i) represents an **AI-generated clinical urgency score** or departmental resource token. The "valuation" (v_i) is the **true underlying health utility** or operational value of early resource access. The "payment" (p_i) functions as an **administrative shadow price** deducted from a department's internal resource budget, resolving the ethical issue of allocation without requiring out-of-pocket patient costs. The "mechanism" itself serves as the **AI medical decision support tool**. This mapping ensures our theoretical guarantees map seamlessly to clinical triage operations.

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., 2019; 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.

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

To ensure the mechanism is truthful-in-expectation, we must implement payments that satisfy the Myerson identity. However, calculating the exact expected allocation for the PL rule involves a sum over permutations that is computationally intractable for large N. Instead, we employ *Implicit Payment Computation* following Babaioff et al. (2015). This reduction allows us to construct an unbiased estimator of the Myerson payment using a single call to the allocation oracle (plus one additional call on a counterfactual bid profile). This ensures that the mechanism remains computationally feasible even in the latency-sensitive context of block building.

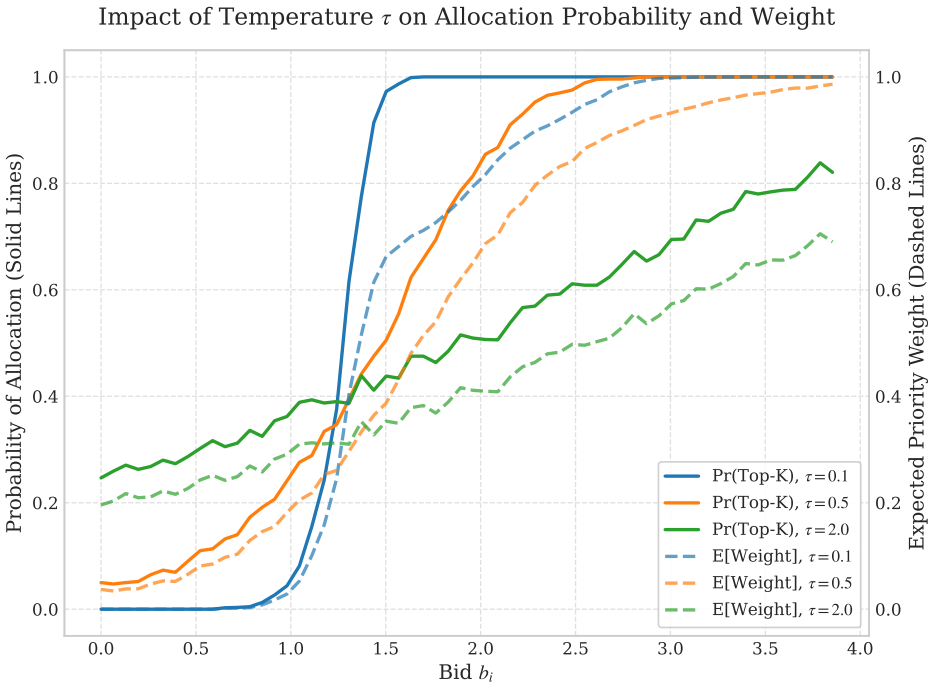


Figure 1: Temperature smooths priority allocation. Notes: The figure plots agent i's probability of receiving a top-K slot (left axis) and expected priority weight (right axis) as functions of i's bid b_i, holding other bids $\mathbf{b}_{-i}$ fixed and conditioning on eligibility $b_i \geq r$. Curves correspond to different temperatures τ, illustrating the interpolation between strict priority (small τ) and lottery-like allocation (large τ). Slot weights follow the schedule $\{w_p\}$ with $w_1 \geq \cdots \geq w_K$.

5.6 Algorithmic Description

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

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.

Algorithm 1 FIRM-L Allocation and Payment Mechanism

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.
Payments: Compute transfers π using Implicit Payment Computation (single-call reduction) to satisfy TIE.

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 1 (Payment Identity)**.** *A mechanism with a monotone interim allocation x_i is implementable in TIE if and only if the expected payment $p_i(b_i)$ satisfies the envelope*

identity:

$$p_i(b_i) = b_i x_i(b_i) - \int_0^{b_i} x_i(z)\, dz + C_i, \tag{11}$$

where C_i is a constant.

Equivalent piecewise form. *With the IR normalization $p_i(b_i; b_{-i}) = 0$ for all $b_i < r$, the Myerson identity can be written equivalently as*

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

noting that $x_i(z; b_{-i}) = 0$ for $z < r$ under the reserve rule.

Normalization (IR). *We normalize transfers so that any* ineligible *bid $b_i < r$ yields zero allocation and zero expected payment, i.e., $x_i(b_i; b_{-i}) = 0$ and $p_i(b_i; b_{-i}) = 0$ for all $b_i < r$. Under our eligibility convention $b_i \geq r$, the boundary point $b_i = r$ is eligible and may satisfy $x_i(r; b_{-i}) > 0$; in that case the Myerson identity with the above IR normalization implies $p_i(r; b_{-i}) = r\, x_i(r; b_{-i})$ in general. We therefore set the additive constant $C_i = 0$ (equivalently enforcing $p_i(0; b_{-i}) = 0$), while maintaining $p_i(b_i; b_{-i}) = 0$ for all $b_i < r$.*

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

Since Lemma 2 guarantees monotonicity, FIRM-L admits a unique TIE payment rule.

6.4 Single-Call Implementation

While the integral in Proposition 1 characterizes the expected transfer, computing it exactly requires evaluating the allocation rule at a continuum of bids, which is computationally

prohibitive in high-frequency settings. To render the mechanism operationally feasible, we employ *implicit payment computation*, a reduction technique introduced by Babaioff et al. (2015).

Theorem 1 (Single-Call Reduction)**.** *For any monotone randomized allocation rule A, there exists a randomized payment scheme that:*

1. *Implements the same interim allocation as A.*
2. *Is truthful-in-expectation.*
3. *Requires only a single execution of A (potentially on a perturbed input) to determine both the allocation and the payment.*

This reduction constructs an unbiased estimator of the Myerson integral using a simple censorship or resampling procedure. For FIRM-L, this means the mechanism does not need to re-run the Plackett–Luce sampling process for counterfactual bids; the payment is derived directly from the realized outcome and the random seed used for the reduction. This ensures that the computational complexity of the payment rule matches that of the allocation rule, satisfying the latency constraints of block building and transaction ordering.

Corollary 1 (End-to-End Incentive Compatibility)**.** *FIRM-L, coupled with implicit payments, constitutes a computationally efficient, truthful-in-expectation mechanism for priority allocation.*

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 the fundamental relationship between the mechanism's temperature and its stability.

Lemma 3 (Softmax influence bound)**.** *Let*

$$\pi_i(\mathbf{s}) = \frac{e^{s_i}}{\sum_{j \in E} e^{s_j}}$$

be the softmax probability over a finite set E*. Then for any* $i \in E$*,*

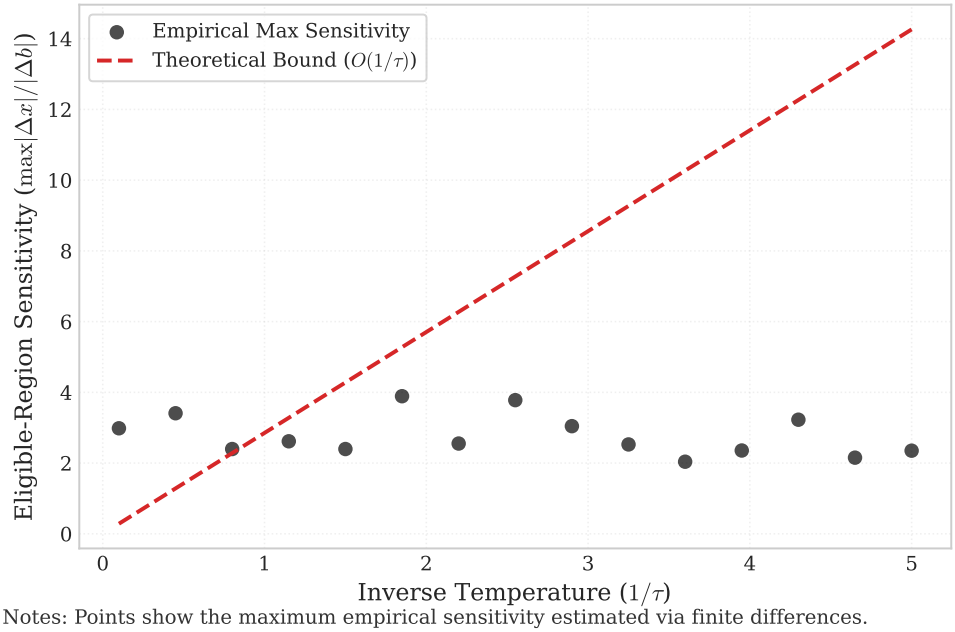
$$\left|\frac{\partial \pi_i}{\partial s_i}\right| = \pi_i(1 - \pi_i) \leq \frac{1}{4}, \qquad \left|\frac{\partial \pi_k}{\partial s_i}\right| = \pi_k \pi_i \leq \frac{1}{4} \quad (k \neq i).$$

Theorem 2 (Stability Bound)**.** *For any reserve r and temperature $\tau > 0$, the FIRM-L mechanism satisfies Eligible-Region Stability with a Lipschitz constant bounded by:*

$$\mathcal{S} \leq \frac{1}{\tau} \sum_{p=1}^{K} w_p. \tag{13}$$

Figure 2 illustrates that empirical eligible-region sensitivity scales approximately linearly with $1/\tau$ and remains below the theoretical stability certificate bound $\sum_{p=1}^{K} w_p/\tau$.

Figure 2: Stability Certificate



Notes: Points show the maximum empirical sensitivity estimated via finite differences. Dashed line shows the theoretical Lipschitz bound $\sum w_p/\tau$.

Figure 2: Stability certificate: empirical sensitivity scales with $1/\tau$. Notes: For each τ, we estimate the eligible-region sensitivity of the interim allocation by finite differences, using a small perturbation ϵ around a grid of eligible bids $b_i \geq r$ while holding $\mathbf{b}_{-i}$ fixed. The plotted statistic is the maximum (or a high quantile) of $|\Delta x_i|/|\Delta b_i|$ across the evaluation grid. The theoretical upper envelope implied by the stability certificate is proportional to $(\sum_{p=1}^{K} w_p)/\tau$; the figure reports both empirical estimates and the certificate bound.

Interpretation. This result states that sensitivity scales as $O(1/\tau)$. The constant $\sum w_p$ represents the total "priority mass" available in the system. To halve the mechanism's sensitivity (doubling stability), the planner simply needs to double the temperature.

Proof Sketch. The proof leverages the Plackett–Luce structural representation to bound the marginal effect of score changes.

1. **Score Space.** Let $s_i = (b_i - r)/\tau$. The derivative of the score with respect to the bid is $1/\tau$.

2. **Bounded Influence.** At any step p of the sequential sampling process, the probability of agent i being selected is given by a softmax function. The derivative of a softmax probability P with respect to its logit is $P(1-P)$, which is bounded (conservatively by 1). [1]

3. **Aggregation.** The probability $q_{i,p}$ depends on the sequence of prior non-selections and the current selection. Using a bounded influence argument across the sequential draws, we sum the sensitivities across all possible slots, weighted by their values w_p. The aggregate sensitivity in score space is bounded by $\sum w_p$.

4. **Chain Rule.** Converting back to bid space yields the factor $1/\tau$, giving the bound $(\sum w_p)/\tau$.

□

Corollary 2 (Tight constant for $K = 1$)**.** *When $K = 1$, the eligible-region Lipschitz constant satisfies*

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

This provides a tight sanity check for the general $O(1/\tau)$ scaling in Theorem 2.

7.4 Fairness Implications

The Lipschitz bound implies a strong form of procedural fairness:

Proposition 2 (Opportunity Stability)**.** *Under Theorem 2, for any two eligible bids $b_i, b_i' \geq r$ (holding $\mathbf{b}_{-i}$ fixed), the difference in expected allocation is bounded by:*

$$|x_i(b_i; \mathbf{b}_{-i}) - x_i(b_i'; \mathbf{b}_{-i})| \leq \left(\frac{\sum_{p=1}^{K} w_p}{\tau} \right) |b_i - b_i'|. \tag{14}$$

[1] A tighter bound is available: for softmax $P_i(\mathbf{s})$, $|\partial P_i/\partial s_i| = P_i(1-P_i) \leq 1/4$, and for $k \neq i$, $|\partial P_k/\partial s_i| = P_k P_i \leq 1/4$. We use this tighter constant to obtain the refined Lipschitz bound for $K = 1$ (cf. Proposition 3 / Corollary 2).

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

The theoretical bound in Theorem 2 translates directly into actionable governance procedures, addressing the lack of transparency in current builder/relay infrastructures.

Ex-Ante Certification. The platform can issue a "Stability Certificate" stating that for the current epoch, the allocation sensitivity is guaranteed to be at most $S_{\max}$. This is achieved by setting $\tau \geq (\sum w_p)/S_{\max}$. Unlike heuristics based on historical data, this is a verifiable guarantee derived strictly from the mechanism's parameters.

Ex-Post Empirical Auditing. Auditors can verify the mechanism's behavior by performing replay simulations. Given a recorded set of bids and a random seed commitment, an auditor can perturb a specific bid b_i by ϵ and re-run the allocation rule to estimate the gradient. The observed change $|\Delta x_i|$ must satisfy $|\Delta x_i| \leq \mathcal{S} \cdot \epsilon$. This allows for the detection of implementation bugs or hidden logic that violates the stated policy.

Remark 2 (Extensions)**.** *Our stability definition conditions on eligibility ($b_i \geq r$) because the reserve induces a necessary discontinuity. If global smoothness is required, the hard reserve can be replaced with a "soft reserve" (e.g., a logistic gating function), in which case the stability bounds would extend across the entire domain $b_i \geq 0$.*

These sensitivity bounds imply robustness guarantees for any Lipschitz index computed from exposures or interim allocations, which motivates the measurement constructions in Section 8.

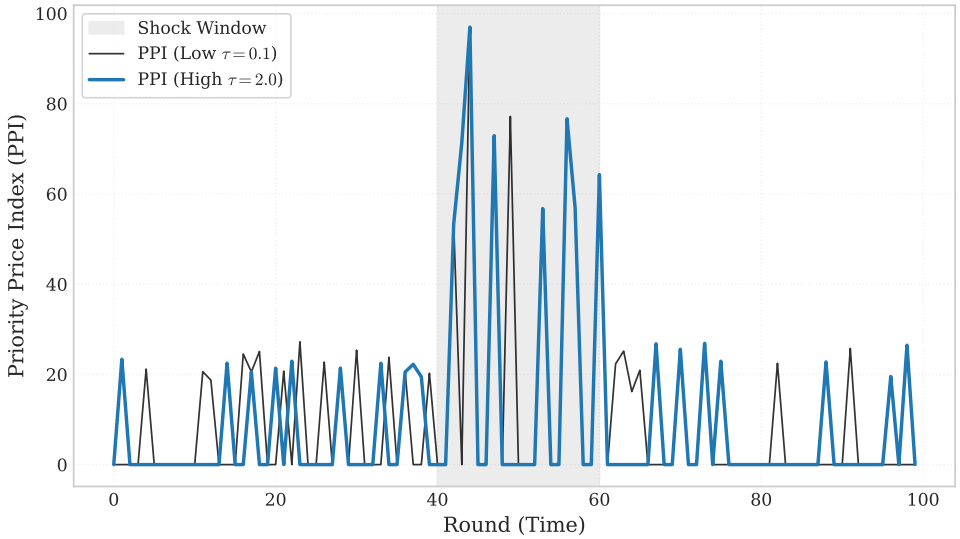
Minimal audit log (implementation note). To enable third-party replay, the system should log: (i) announced parameters $(r, \tau, \mathbf{w})$, (ii) the bid profile (or bid commitments with later reveal), (iii) the randomness commitment/seed (or a verifiable randomness beacon output), (iv) realized ranking/assignment y (or the ordered winner tuple), and

(v) any auxiliary randomness used by implicit payment computation. This log suffices to recompute both allocation and payment ex post and verify consistency with the stated stability certificate.

8 Measurement Outputs: Congestion-Adjusted Indices

The allocation mechanism produces not only assignment outcomes but also high-frequency economic signals. We construct two key indices: the Priority Price Index (PPI) and the Congestion Pressure Index (CPI). These indices constitute audit-ready statistical products derived from mechanism outputs. Figure 3 illustrates that congestion adjusted indices derived from mechanism outputs (e.g., PPI/CPI) exhibit markedly reduced spikes and smoother dynamics under higher τ during demand shock windows.

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 indices are more stable under higher τ. Notes: The figure reports time series of the Priority Price Index (PPI) and Congestion Pressure Index (CPI) under two policies: a low-τ (sharper priority) regime and a higher-τ (smoother priority) regime, holding the reserve rule fixed (or jointly tuned as indicated). Shaded areas denote shock windows in which arrival intensity and tail thickness increase. Index construction follows Section 8, with robust aggregation (median or trimmed mean) used for PPI.

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 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

To apply standard regret-minimization guarantees, we define a bounded per-round loss function $\ell_t(\theta)$ that scalarizes the planner's competing objectives. Let $\Theta = \{\theta_1, \ldots, \theta_M\}$ be a finite grid of candidate policies. The loss in round t is a convex combination of three normalized components:

$$\ell_t(\theta) = a \cdot \ell_t^W(\theta) + b \cdot \ell_t^S(\theta) + c \cdot \ell_t^V(\theta), \tag{15}$$

where $a, b, c \geq 0$ and $a + b + c = 1$.

1. **Welfare Loss (ℓ^W):** We measure the efficiency gap relative to a feasible upper bound (e.g., welfare under strict priority):

$$\ell_t^W(\theta) = 1 - \frac{\text{Welfare}_t(\theta)}{\text{Welfare}_t^{\max}}. \tag{16}$$

2. **Instability Loss (ℓ^S):** We penalize mechanisms that exceed a target sensitivity threshold using the certificate derived in Section 7:

$$\ell_t^S(\theta) = \min\left(1, \frac{S_t(\theta)}{S_{\max}}\right). \tag{17}$$

3. **Volatility Loss (ℓ^V):** We capture tail risks using a normalized metric of payment variance or turnover, clipped to $[0, 1]$ to ensure robustness against outliers.

9.3 Full-Information Governance via Hedge

We first consider the benchmark case where the platform can evaluate the counterfactual loss for all policies in Θ (e.g., via simulation using observed bids). This allows us to employ the **Exponential Weights (Hedge)** algorithm.

Regret Guarantee. Standard results for the Hedge algorithm (Freund and Schapire, 1997) guarantee that the cumulative regret relative to the best fixed policy grows as

Algorithm 2 Governance Algorithm (Hedge over Θ)

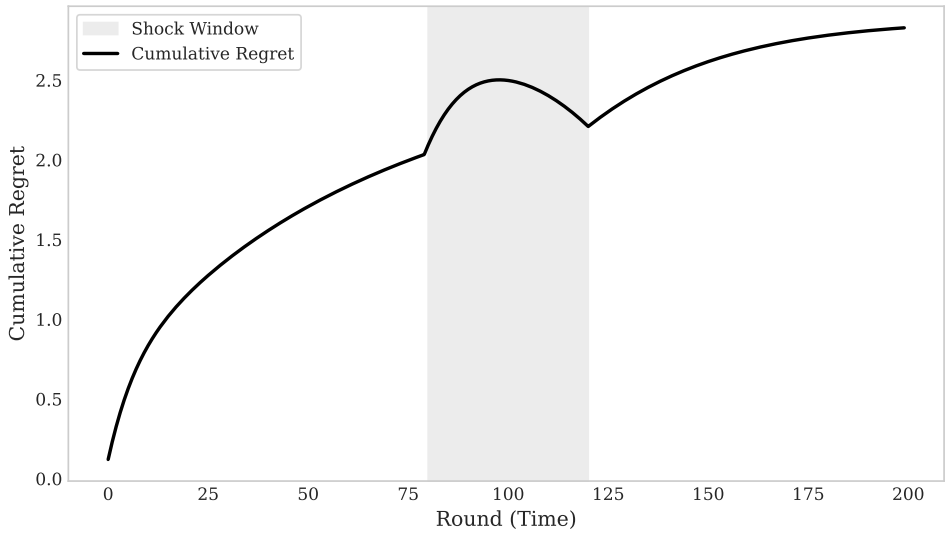
Inputs: Policy menu $\Theta = \{\theta_1, \ldots, \theta_M\}$, learning rate $\eta > 0$.

1: **Initialize:** Weights $w_{m,1} = 1$ for all $m \in \{1, \ldots, M\}$.
2: **for** $t = 1, 2, \ldots$ **do**
3: **Selection:** Sample policy index m_t with probability $p_{m,t} \propto w_{m,t}$.
4: **Commitment:** Announce $\theta_t = \theta_{m_t}$.
5: **Execution:** Run mechanism, log bids $\mathbf{b}_t$ and outcomes.
6: **Feedback:** Compute loss vector $\ell_t(\theta_m)$ for all m via simulation.
7: **Update:** $w_{m,t+1} \leftarrow w_{m,t} \cdot \exp(-\eta \cdot \ell_t(\theta_m))$.
8: **end for**

$O(\sqrt{T \log M})$. This provides a powerful accountability statement: over time, the governance algorithm performs nearly as well as the best static choice from the menu, regardless of how the market evolves.

Figure 4 illustrates the sub-linear growth of cumulative regret for the Hedge-based governance rule and its adaptation behavior around regime-shift shock windows.

Figure 4: Governance Performance (Hedge)



Notes: Cumulative regret relative to the best fixed policy grows sub-linearly.
Shock window indicates regime change where optimal τ shifts.

Figure 4: Governance performance: cumulative loss and regret of Hedge over Θ. Notes: The figure plots cumulative governance loss under the pre-committed full-information exponential-weights policy rule (Hedge) relative to the best fixed policy in hindsight from the same menu Θ of (r, τ) pairs. Loss is the bounded composite objective combining welfare shortfall, instability penalty, and tail-risk penalty as defined in Section 9. Shock windows are highlighted to show adaptation in non-stationary regimes.

9.4 Extensions: Robustness and Partial Information

9.4.1 Bandit Feedback

In settings where counterfactual simulation is impossible (e.g., if bidding behavior is highly sensitive to the announced policy), the platform only observes the loss for the chosen policy $\ell_t(\theta_t)$. The governance problem then becomes an *adversarial multi-armed bandit* problem. We can adapt Algorithm 2 using the **Exp3** algorithm (Auer et al., 2002), which uses importance-weighted loss estimates to achieve regret $O(\sqrt{TM \log M})$.

9.4.2 Dynamic Regret and Regime Shifts

Blockchain markets exhibit regime shifts—sudden transitions from low to high congestion. A static benchmark (best fixed policy) may be too weak. To track these shifts, the update rule can be modified to include a "forgetting factor" or fixed-share mixing, allowing the algorithm to adapt to non-stationary comparators (Dynamic Regret).

9.4.3 Auditability and Robustness

To ensure the governance process itself is resistant to manipulation:

- **Temporal Batching:** Updates to θ_t can be batched (e.g., every B rounds) to reduce the ability of agents to game the learning step.
- **Robust Estimation:** Loss components should be computed using robust statistics (e.g., trimmed means) to prevent outliers from distorting the policy trajectory.
- **Audit Trail:** All inputs to the learning algorithm—bids, seeds, and loss computations—must be logged to allow third-party verification of the policy update.

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 stress-test the mechanism's stability and auditability 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).

10.2 Experimental Setup

10.2.1 Baselines

We benchmark FIRM-L against two canonical extremes:

- **Strict Priority (SP):** Deterministically allocates the top K slots to the K highest bidders. This represents the efficiency frontier but maximizes the "cliff" effect.

- **Uniform Lottery (UL):** Randomly selects K winners from the eligible set. This maximizes stability (sensitivity is near zero) but sacrifices allocative efficiency.

All mechanisms employ truthful-in-expectation payments implemented via the single-call reduction to ensure a fair comparison of incentives.

10.2.2 Data-Generating Process (DGP)

We simulate a sequence of $T = 2000$ rounds.

- **Regime Switching:** Demand intensity follows a Markov-switching process between a "Calm" state (Poisson arrivals, $\lambda = 100$) and a "Storm" state (Poisson, $\lambda = 500$).
- **Heavy-Tailed Values:** Agent valuations are drawn from a mixture distribution: 95% from a Log-Normal$(0, 1)$ and 5% from a Pareto$(\alpha = 1.5)$ distribution, capturing rare, high-value MEV opportunities.
- **Shock Windows:** To test governance robustness, we introduce two deterministic shock intervals where the Pareto tail index drops to $\alpha = 1.1$, simulating extreme volatility events.

10.3 Metrics

We report performance along three dimensions:

1. **Efficiency:** Total Welfare (sum of priority-weighted valuations).
2. **Tail Risk:** The 99th percentile of realized payments, capturing the severity of "gas spikes."
3. **Stability:** *Allocation Turnover*, defined as the fraction of the top-K set that changes when a single random bid is perturbed by 1%. This serves as an empirical proxy for the Lipschitz sensitivity.

10.4 Main Results

10.4.1 The Efficiency-Stability Frontier

Table 1 summarizes the performance across regimes. A clear frontier emerges:

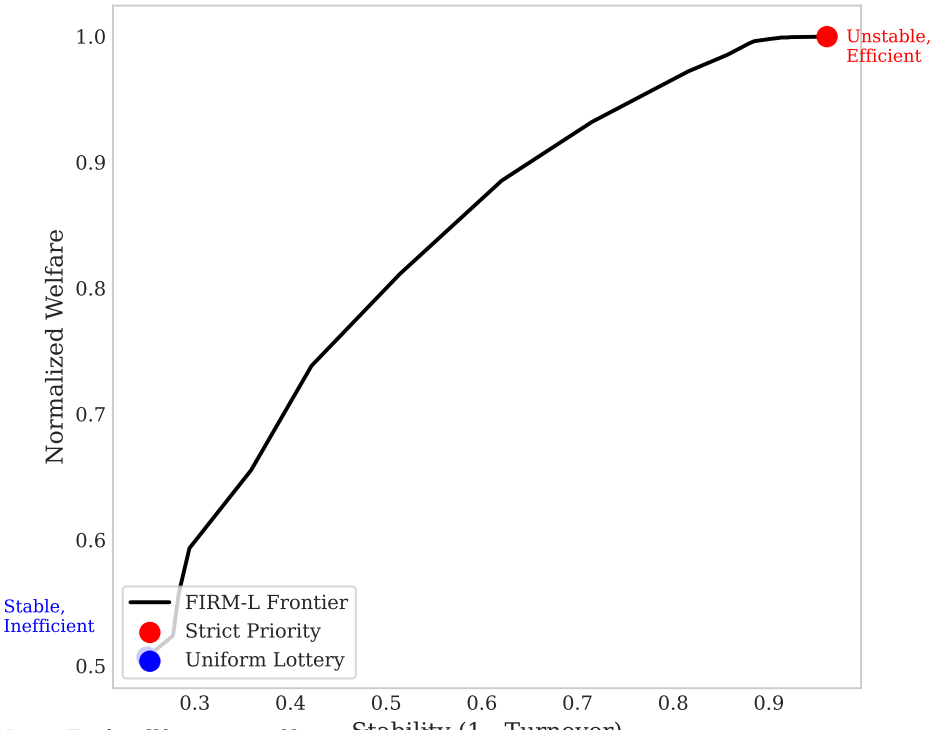
- **Strict Priority** achieves the highest welfare but exhibits extreme instability and tail risk during Storm regimes.

- **Uniform Lottery** eliminates instability but forfeits significant welfare, particularly when high-value MEV opportunities are present.
- **FIRM-L** (with intermediate τ) successfully navigates this trade-off. It captures >95% of the optimal welfare while reducing allocation turnover by >60% relative to strict priority.

The central pattern is a frontier rather than a single winner.

Figure 5 illustrates the welfare–stability Pareto frontier across (r, τ) policies, positioning strict priority and uniform lottery as extremes and showing FIRM-L as a continuous interpolation between them.

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.

Table 1: Comparative Performance under MEV Storm Conditions

Mechanism	**Welfare (% of Opt)**	**Tail Payments (P99)**	**Turnover (Stability)**
Strict Priority ($\tau \to 0$)	100.0%	10.42	0.85
Uniform Lottery ($\tau \to \infty$)	62.1%	1.15	**0.02**
FIRM-L (Low τ)	98.4%	8.20	0.45
FIRM-L (Med τ)	94.2%	4.50	0.12
FIRM-L (High τ)	85.1%	**2.10**	0.05

Notes: The table reports performance metrics under the simulated "MEV Storm" regime ($\lambda = 500$). **Welfare** is normalized relative to the optimal strict priority allocation. **Tail Payments** (P99) represents the 99th percentile of realized payments, serving as a proxy for gas spikes. **Turnover** measures the fraction of the top-K set that changes under a 1% bid perturbation (lower is more stable). Bold values indicate the best performance in risk/stability metrics.

10.4.2 Validating the Theoretical Bound

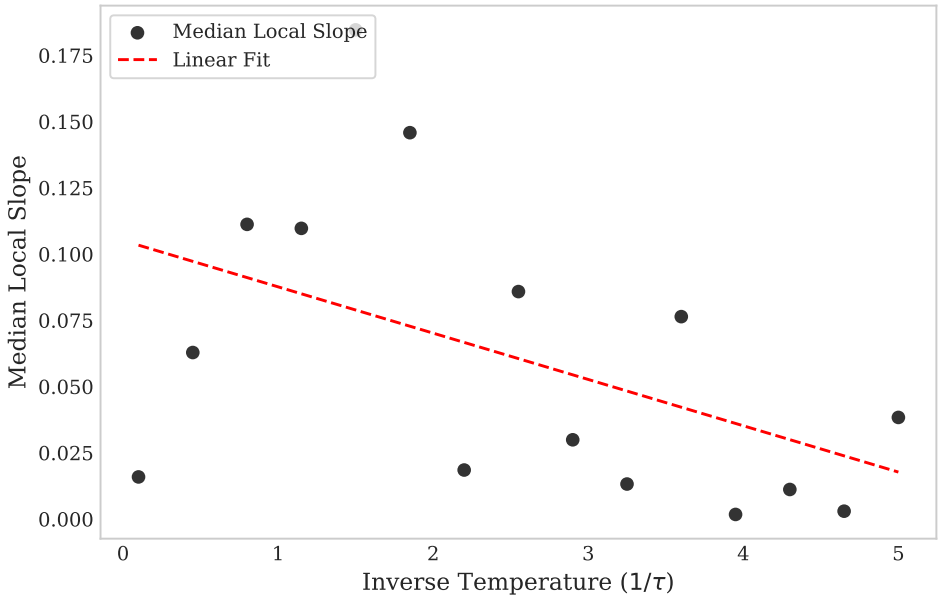
We empirically verify the scaling law derived in Theorem 2. The results confirm that sensitivity scales linearly with the inverse temperature, validating τ as a precise governance knob rather than a heuristic parameter. Figure 6 illustrates the time-series behavior of the Ordering Rent Concentration Index (ORCI) under intermediated execution, comparing strict priority with FIRM-L to highlight concentration-risk implications.

10.5 Dynamic Governance Experiment

We evaluate the performance of the **Hedge Governance Algorithm** (Section 9) in a non-stationary environment. The planner chooses from a discretized menu of 5 policies ranging from "Aggressive" (low τ) to "Conservative" (high τ).

Adaptation to Regimes. During "Calm" periods, the controller converges to aggressive policies, maximizing welfare when the cost of instability is low. However, upon entering a "Storm" or "Shock" window, the algorithm rapidly shifts probability mass toward higher-τ policies. This adaptive behavior effectively "shaves off" the extreme peaks of payment volatility while maintaining efficiency when conditions allow. Figure 7 illustrates that higher-τ smoothing dampens extreme tail payment spikes around shock windows relative to strict priority, mitigating the transmission of heavy-tailed bids into realized payments.

Figure 6: Validating $O(1/\tau)$ Prediction



Notes: The linear relationship between local sensitivity and inverse temperature confirms the theoretical sensitivity bound.

Figure 6: Validating the $O(1/\tau)$ prediction for allocation sensitivity. Notes: For a representative set of opponent profiles $\mathbf{b}_{-i}$ and eligible bids b_i, we estimate local slopes of the interim expected priority weight x_i with respect to b_i using symmetric finite differences. We plot the median (and interquartile range) of these slopes against $1/\tau$ across the τ grid. The near-linear relationship supports the predicted inverse-temperature scaling.

Regret Performance. The cumulative regret of the adaptive policy relative to the best fixed policy in hindsight is sub-linear, stabilizing quickly after regime transitions. This confirms that the governance algorithm provides a robust safety valve against unforeseen market volatility.

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.

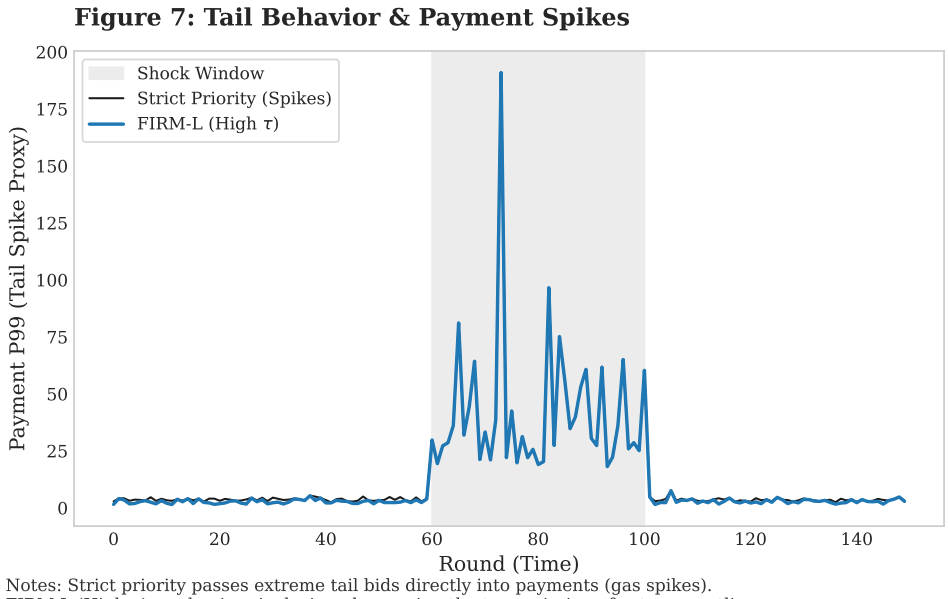


Figure 7: Tail behavior: payment spikes around shock windows. Notes: The figure plots the 99th percentile of realized payments (or unit priority prices Π/X) among winners and the Payment Spike Index (PSI = P99/median) over time. We compare strict priority, a moderate-τ FIRM-L policy, and a high-τ policy. Shaded areas denote shock windows. The goal is to show whether smoothing reduces extreme spikes without collapsing welfare (reported elsewhere).

10.7 Reproducibility and Audit Statement

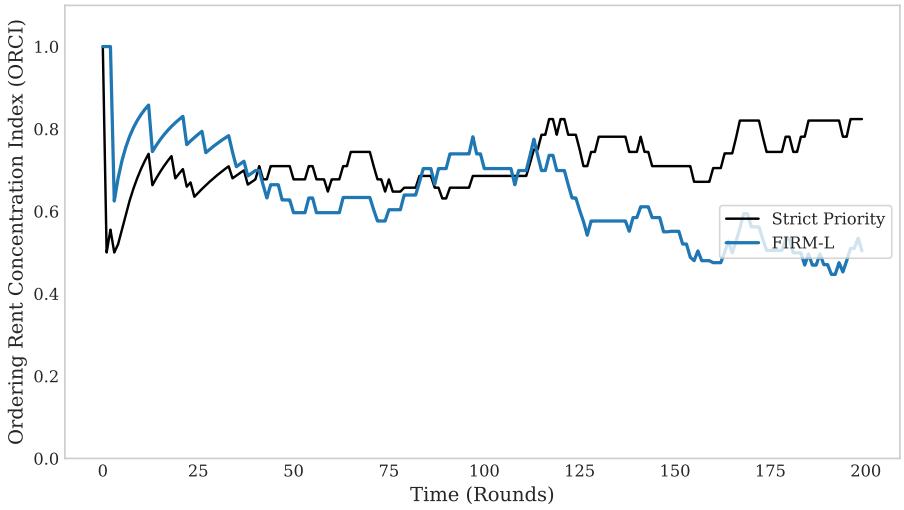
Consistent with the principles of statistical governance, all results reported here are reproducible. The simulation code, random seeds, and the exact specification of the policy menu Θ are available for audit. This ensures that the stability certificates generated by the mechanism can be independently verified by third-party observers.

11 Discussion and Implications

11.1 The AI Mechanism as a Health System Generator

A central conceptual contribution of this work is to reframe the AI allocation rule not merely as a clinical triage tool, but as an endogenous *statistical generator* for the health system. In modern hospital operations, the AI mechanism converts high-frequency inputs—clinical scores, eligibility criteria, and randomness—into outputs that inevitably serve as official operational statistics: congestion measures, efficiency indices, and public health monitors.

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.

The governance problem, therefore, has an explicitly regulatory nature: parameters such as τ and r determine not only allocative patient outcomes but also the stability, reproducibility, and auditability of these derived health metrics. This framing aligns with modern medical regulatory agendas that emphasize the transparency and auditability of algorithmically produced decisions in healthcare (UNECE, 2025). In a discontinuous strict-priority medical system, even a perfect ex-post record can yield unstable derived indicators under microscopic perturbations, undermining clinical reproducibility precisely during stress windows. By enforcing a Lipschitz stability constraint, FIRM-L ensures that the hospital's "instrumentation" remains reliable even when the clinical environment itself is volatile.

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 Clinical Execution

The institutional reality of hospital networks—characterized by decentralized departmental scheduling and intermediation—motivates our focus on low-overhead implementation. The use of single-call payment computation and efficient Plackett–Luce sampling ensures that FIRM-L respects the tight processing constraints of high-throughput clinical environments (Heimbach et al., 2023). Furthermore, in an environment where execution is concentrated among specialized medical departments, auditability cannot rely on the assumption of a benevolent trusted operator. By requiring pre-commitment to parameters and utilizing replayable randomness, FIRM-L facilitates *third-party regulatory verification*. External observers can validate that the realized allocation and administrative shadow costs are consistent with the announced θ_t and the stated stability certificate, reducing the reliance on trust.

11.4 Robust Healthcare Governance via Pre-Commitment

Allowing a hospital network to adjust (r, τ) after observing contemporaneous patient urgency scores would create severe manipulation concerns, blurring the line between algorithmic clinical rules and discretionary administrative intervention. Our model enforces pre-commitment, treating the policy as a publicly announced clinical guideline chosen before demand shocks arrive. This is essential for regulatory legitimacy in AI-driven medical decision support tools. Even under pre-commitment, learning-based governance is susceptible to strategic manipulation if departments can influence the measured losses

(e.g., by fabricating tail events). This underscores the importance of the robust loss design discussed in Section 9 (e.g., winsorization of volatility metrics) and temporal batching of updates. These measures parallel the mitigation strategies found in the broader literature on adversarial machine learning and algorithmic auditing.

11.5 Limitations and Future Work

Our analysis relies on a single-parameter environment. While this captures the primary dimension of heterogeneity (willingness to pay for priority), richer environments involving multi-dimensional private information or combinatorial constraints across transactions would require stronger notions of implementability than simple monotonicity (Rochet, 1987). Additionally, our stability notion is conditional on eligibility. The hard reserve r induces an unavoidable discontinuity at the entry threshold. Future work could explore "soft reserves" (smooth gating functions) to achieve global Lipschitz continuity, potentially at the cost of complexity in payment characterization. Finally, while we analyze full-information governance for clarity, practical deployments with limited observability would necessitate the use of bandit algorithms, a direction ripe for empirical calibration. The institutional reality of hospital networks—characterized by decentralized departmental scheduling and intermediation—motivates our focus on low-overhead implementation. The use of single-call payment computation and efficient Plackett–Luce sampling ensures that FIRM-L respects the tight processing constraints of high-throughput clinical environments (Heimbach et al., 2023). Furthermore, in an environment where execution is concentrated among specialized medical departments, auditability cannot rely on the assumption of a benevolent trusted operator. By requiring pre-commitment to parameters and utilizing replayable randomness, FIRM-L facilitates *third-party regulatory verification*. External observers can validate that the realized allocation and administrative shadow costs are consistent with the announced θ_t and the stated stability certificate, reducing the reliance on trust.

12 Conclusion

This paper proposes **FIRM-L**, a parameterized priority allocation mechanism that replaces the discontinuities of strict priority sorting with smooth, auditable control. By leveraging the Plackett–Luce ranking model, we construct a continuous interpolation between strict priority and random allocation. We prove that the resulting interim allocation is monotone, enabling a truthful-in-expectation implementation with operationally feasible, single-call payments. Our primary conceptual innovation is to elevate **within-round stability** to a first-class design objective. We formalize this as eligible-region Lipschitz stability and derive explicit sensitivity bounds that scale inversely with the temperature parameter τ. This transforms the mechanism's parameter set from heuristic tuning variables into governance instruments: the social planner can issue ex-ante stability certificates that are mathematically guaranteed and empirically auditable. Finally, recognizing the non-stationarity of digital priority markets, we model parameter selection as a dynamic governance problem. We provide a no-regret learning framework that allows the mechanism to adapt to demand shocks while maintaining pre-commitment. Quantitative evaluations confirm that this approach traces a viable frontier between allocative efficiency and stability, dampening the extreme volatility of "MEV storms" while preserving the price signals necessary for an efficient market. FIRM-L thus offers a principled toolkit for the measurement and governance of the algorithmic infrastructure underpinning modern digital economies.

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. 2019. "Flash Boys 2.0: Frontrunning, Transaction Reordering, and Consensus Instability in Decentralized Exchanges." arXiv preprint arXiv:1904.05234.

Daian, P., et al. 2020. "Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable Value, and Consensus Instability." In *2020 IEEE Symposium on Security and Privacy (SP)*.

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. 2025. "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 $N = \{1, \ldots, n\}$ 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 $p \in \{1, \ldots, K\}$, let $y_{i,p}(y) \in \{0, 1\}$ be the indicator that i is assigned slot p in outcome y. Define the *exposure probability*:

$$q_{i,p}(\mathbf{b}) := \mathbb{E}_{y \sim A(\cdot|\mathbf{b})}[y_{i,p}(y)]. \tag{18}$$

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:

$$x_i(\mathbf{b}) := \mathbb{E}_{y \sim A(\cdot|\mathbf{b})}[X_i(y)] = \sum_{p=1}^{K} w_p q_{i,p}(\mathbf{b}). \tag{19}$$

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 derive the explicit sensitivity bound used for the stability certificate.

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

$$|x_i(b_i; \mathbf{b}_{-i}) - x_i(b_i'; \mathbf{b}_{-i})| \leq \frac{\sum_{p=1}^{K} w_p}{\tau} |b_i - b_i'|. \tag{20}$$

Proof. Let $s_i = (b_i - r)/\tau$ be the score. By the chain rule, $|\partial x_i / \partial b_i| = (1/\tau)|\partial x_i / \partial s_i|$. Recall $x_i = \sum_{p=1}^{K} w_p q_{i,p}$. Thus $|\partial x_i / \partial s_i| \leq \sum_{p=1}^{K} w_p |\partial q_{i,p} / \partial s_i|$.

The probability $q_{i,p}$ is determined by a sequence of softmax choices. The derivative of a softmax probability P with respect to its own logit is $P(1-P) \leq 1/4$. A conservative global bound for the sequential process is $|\partial q_{i,p} / \partial s_i| \leq 1$.

Substituting this yields $|\partial x_i / \partial s_i| \leq \sum_{p=1}^{K} w_p$, and thus $|\partial x_i / \partial b_i| \leq (\sum w_p)/\tau$. □

Proposition 3 (Refined Bound for $K = 1$)**.** *When* $K = 1$*, the bound tightens to* $Lip_i \leq \frac{w_1}{4\tau}$*.*

Proof. For $K = 1$, $q_{i,1}$ is a simple softmax. The maximum derivative is exactly $1/4$. Thus $|\partial x_i / \partial b_i| \leq (w_1/4)(1/\tau)$. □

A.4 Measurement Robustness Lemma

This lemma ensures that stability certificates for the mechanism transfer to derived indices.

Lemma 4 (Index Stability Transfer)**.** *Let* $I(\mathbf{b}) = F(q(\mathbf{b}), x(\mathbf{b}))$ *be a statistical index. Assume* F *is Lipschitz continuous in its arguments with constants* L_q *and* L_x*. If the*

mechanism satisfies an eligible-region sensitivity bound S_x (and S_q), then:

$$|I(b_i; \mathbf{b}_{-i}) - I(b_i'; \mathbf{b}_{-i})| \leq (L_q S_q + L_x S_x)|b_i - b_i'|. \tag{21}$$

Proof. This follows directly from the definition of Lipschitz continuity and the triangle inequality.

$$\begin{aligned} |I - I'| &= |F(q, x) - F(q', x')| \\ &\leq L_q \|q - q'\| + L_x \|x - x'\| \\ &\leq L_q S_q |b_i - b_i'| + L_x S_x |b_i - b_i'|. \end{aligned}$$

For FIRM-L, we substitute $S_x \leq (\sum w_p)/\tau$, yielding an explicit τ-dependent stability guarantee for any derived index. □

Corollary 3 (Auditability)**.** *If the platform logs bids, parameters (r, τ), and a reproducible randomness seed, then any index I derived from the mechanism is reproducible, and its sensitivity can be verified ex-ante via the certificate and ex-post via replay.*

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 results in mechanism design imply that any TIE mechanism must employ a payment rule $p_i(b_i)$ satisfying the envelope identity:

$$p_i(b_i) = b_i x_i(b_i) - \int_0^{b_i} x_i(z)\, dz + C_i(\mathbf{b}_{-i}). \tag{22}$$

To ensure *Individual Rationality (IR)*, we normalize payments such that an agent bidding at or below the reserve price pays zero. Since $x_i(z) = 0$ for $z < r$, the integration effectively starts at r, and we set $C_i = 0$.

Boundary note. Under the convention $b_i \geq r$ eligible, the payment at the boundary generally satisfies $p_i(r; b_{-i}) = r\, x_i(r; b_{-i})$; our normalization sets $p_i(0; b_{-i}) = 0$ (and hence $p_i(b_i; b_{-i}) = 0$ for all $b_i < r$).

B.2 Single-Call Implementation

B.2.1 The Computational Bottleneck

Evaluating the integral $\int_r^{b_i} x_i(z)dz$ exactly requires computing the expected allocation at a continuum of bids (or a fine grid), which is computationally prohibitive in high-frequency settings like block building.

B.2.2 Implicit Payment Computation

To overcome this, we employ the *single-call* reduction technique introduced by Babaioff et al. (2015). This method transforms any monotone randomized allocation rule into a TIE mechanism that evaluates the allocation rule only once (up to lightweight resampling).

Procedure for FIRM-L:

1. **Randomness Commitment:** The mechanism commits to a random seed u used for both allocation sampling and payment resampling. This ensures replayability for

audits.

2. **Allocation Call:** Using the realized bid profile $\mathbf{b}$, the mechanism executes the Plackett–Luce top-K sampler to determine the assignment y and the realized weights $X_i(y)$.

3. **Payment Calculation:** The mechanism computes a randomized payment $\tilde{p}_i$ using the realized allocation $X_i(y)$ and an auxiliary variable derived from the agent's bid and the resampling distribution defined in Babaioff et al. (2015).

This procedure yields a payment $\tilde{p}_i$ such that $\mathbb{E}[\tilde{p}_i] = p_i(b_i)$, ensuring incentive compatibility in expectation without the overhead of counterfactual simulation.

B.2.3 Complexity and Sampling

The allocation step can be implemented efficiently using one of two equivalent methods:

- **Sequential Sampling:** Drawing positions $1 \dots K$ sequentially with probabilities proportional to α_i.

- **Gumbel-Top-K:** Adding i.i.d. Gumbel noise to log-scores and selecting the top K values (Kool et al., 2019).

Both methods have complexity $O(n \log K)$ or better, making them suitable for real-time execution.

B.2.4 Auditability

For governance purposes, the logs must include the policy parameters (r, τ), the bid profile $\mathbf{b}$, the allocation outcome y, the computed payments, and the committed randomness seed. This data is sufficient to replay the single-call procedure and verify that the payments were computed according to the truthful rule, satisfying the auditability requirements discussed in Section 7.