

A Framework for Combined Transaction Posting and Pricing for Layer 2 Blockchains

Shouqiao Wang¹, Davide Crapis², and Ciamac C. Moallemi¹

¹Columbia University

²Ethereum Foundation

Abstract. This paper presents a comprehensive framework for transaction posting and pricing in Layer 2 (L2) blockchain systems, focusing on challenges stemming from fluctuating Layer 1 (L1) gas fees and the congestion issues within L2 networks. Existing methods have focused on the problem of optimal posting strategies to L1 in isolation, without simultaneously considering the L2 fee mechanism. In contrast, our work offers a unified approach that addresses the complex interplay between transaction queue dynamics, L1 cost variability, and user responses to L2 fees. We contribute by (1) formulating a dynamic model that integrates both posting and pricing strategies, capturing the interplay between L1 gas price fluctuations and L2 queue management, (2) deriving an optimal threshold-based posting policy that guides L2 sequencers in managing transactions based on queue length and current L1 conditions, and (3) establishing theoretical foundations for a dynamic L2 fee mechanism that balances cost recovery with congestion control. We validate our framework through simulations.

1 Introduction

The scalability roadmap for Ethereum is rapidly coming to fruition, marked by the deployment of numerous Layer 2 (L2) blockchains that have significantly enhanced the network’s capacity and efficiency. Major L2 solutions like Optimism, Arbitrum, and others have attracted assets worth billions of dollars and are processing millions of transactions daily. These platforms offload transaction processing from the main Ethereum blockchain (L1), enabling higher throughput and lower fees for users while maintaining the security guarantees of the underlying network.

A pivotal development facilitating this growth is Ethereum Improvement Proposal 4844 (EIP-4844), which has unlocked ample data availability for L2 rollups. By introducing a new transaction type that carries ephemeral data (blobs), EIP-4844 allows L2 blockchains to post large amounts of data to L1 at significantly reduced costs. This enhancement has resulted in very low fees for L2 users, further incentivizing adoption and driving transaction volumes.

However, the burgeoning L2 ecosystem faces two critical challenges:

1. **Fluctuating L1 Fees and Posting Costs:** The cost of posting transactions from L2 to L1 is subject to significant variability due to fluctuating L1

gas prices. These fluctuations impact the operational costs of rollups, making it challenging to predict and manage expenses effectively. This issue is poised to become even more pronounced as data availability demand begins to match supply, potentially increasing the cost and variability of data blobs used in rollup operations.

2. **Limited Capacity and Congestion Pricing:** Despite their enhanced capacity, L2 rollups have finite resources and may experience congestion during periods of high demand. To manage this, they might need to implement congestion fees, adjusting transaction costs to regulate network usage and maintain performance standards.

Addressing these challenges is essential for the sustainable growth of L2 solutions. Current studies have predominantly focused on the problem of optimal posting strategies to L1 in isolation, without concurrently considering the L2 fee mechanism. Practical systems have employed heuristic approaches to set L2 fees. Notably, Arbitrum Nitro has introduced a decoupled fee structure that separates fees into components aimed at recovering L1 posting costs and managing L2 congestion, drawing inspiration from Ethereum’s EIP-1559 mechanism. However, these methods lack rigorous theoretical underpinnings, leaving a gap in understanding the optimal interplay between posting strategies and pricing mechanisms.

In this paper, we develop a comprehensive transaction management framework for L2 blockchain systems that captures the complex interplay between transaction queue dynamics, posting strategies, and gas price fluctuations. For the first time, we jointly consider the problems of posting and pricing, providing a unified approach to optimize both aspects simultaneously. Our model operates in discrete time, aligned with L1 block intervals, and incorporates realistic behaviors of both transaction arrivals influenced by L2 fees and L1 gas price movements.

We explicitly set up the objectives of balancing operational costs and controlling congestion within the L2 network. By formulating an objective function that incorporates both posting costs and queuing delays, we aim to derive strategies that optimize the cumulative operational cost over time while maintaining network performance. Furthermore, we establish the first theoretical results for L2 pricing mechanisms under this framework, subject to technical assumptions that we validate through robust simulations.

Our primary contributions are as follows:

1. **Joint Modeling of Posting and Pricing Strategies:** We introduce a dynamic model that simultaneously considers the optimal posting of transactions from L2 to L1 and the L2 fee mechanism. This joint approach captures the dependencies between transaction queue dynamics, L1 gas price fluctuations, and user behavior in response to L2 fees.
2. **Optimal Posting Strategy with Threshold Policy:** We derive an optimal posting strategy for the L2 sequencer using dynamic programming techniques. We prove that a threshold policy is optimal, where the sequencer

decides to post all pending transactions or none based on a critical queue length that depends on the current L1 gas price. This finding simplifies the decision-making process and provides practical guidelines for L2 operators.

3. **Theoretical Foundations for L2 Pricing Mechanisms:** We establish the first theoretical results for L2 pricing, explicitly formulating the objectives of cost recovery and congestion control. By modeling the transaction arrival rate as a function of the L2 fee, we develop a dynamic fee adjustment strategy that ensures budget balance and manages network congestion. We prove the existence and uniqueness of optimal fees and show that with our fee adjustment mechanism, fees converge to the optimal fees, under some technical conditions. We also demonstrate the robustness of our approach in more general and practical settings through simulations.

The remainder of the paper is structured as follows. In Section 2, we detail our transaction management model, including the dynamics of the transaction queue, L1 gas price modeling, cost structure, compensation mechanism, and the formulation of the objective function. Section 3 is dedicated to deriving the optimal posting strategy. We employ dynamic programming techniques to establish the Bellman equation for our system and prove that a threshold policy is optimal for the sequencer’s decision-making process. In Section 4, we focus on the design of the L2 fee mechanism. We analyze how the L2 fee influences transaction arrival rates and develop a dynamic fee adjustment strategy that achieves budget balance and congestion control. We establish the first theoretical results for L2 pricing under this framework and present an adaptive algorithm for fee updates. Section 5 provides an analysis of the proposed mechanisms, including the assumptions made, convergence results, and discussions on practical implications. We validate our theoretical findings through simulations, demonstrating the robustness of our approach under various realistic network conditions.

1.1 Literature Review

Existing literature on decentralized L1 blockchains such as Bitcoin and Ethereum has extensively examined fee market designs, often using game-theoretic models to investigate how auction-based or monopolistic pricing schemes affect miner incentives, user behavior, and network throughput [2,9,17,18]. Roughgarden [15] provides an economic analysis of Ethereum’s EIP-1559, highlighting dynamic base-fee adjustments for congestion control. Further refinements in L1 fee mechanisms appear in works by Leonardos et al. [10], who analyzed fee market dynamics and demonstrated that optimal fee mechanisms can be achieved despite inherent market chaos, and Crapis et al. [6], who investigated optimal dynamic fees for blockchain resources, introducing models that adjust to network conditions to optimize fee structures. Additionally, Crapis [4] has provided an analysis of the fee market under EIP-4844, discussing its implications for L1 fee dynamics. While these studies primarily focus on decentralized settings where miners or validators collectively enforce protocol rules, our work centers on a more centralized L2 context, where a single operator unilaterally defines both posting and pricing

policies, thereby departing from the standard L1-centric paradigm. Although congestion and blockspace constraints remain central concerns, L2 rollups face the additional challenge of covering volatile posting costs on L1, which substantially shapes the L2 fee design. By jointly modeling L2 queue dynamic and L1 gas prices dynamics, our framework emphasizes maximizing throughput while achieving budget balance on the L2, thus bridging a key gap in the existing literature.

Studies on optimal posting strategies for L2 rollups have emerged to tackle the challenges posed by unpredictable L1 gas prices and the imperative for timely data finalization. Mamageishvili and Felten [12] propose a Q-learning approach to determine the optimal moment for a rollup to publish transactions on L1, modeling a trade-off between waiting for favorable gas prices and incurring delay costs under a quadratic delay cost model. Although their work offers valuable insights into the batch posting decision, it does not address the design of L2 fee mechanisms or account for the dynamics of user demand driven by network congestion. In contrast, our work considers user demand dynamics via an L2 fee model while adopting a linear delay cost framework. By rigorously proving a threshold policy property, we implement a policy iteration algorithm that is significantly more efficient than the Q-learning method. Similarly, Bar-On and Mansour [1] build on related ideas by offering threshold-based policies for specific classes of cost functions, thereby providing analytical insights into optimal posting schedules. However, their analysis is confined solely to the posting decision and does not examine how user fees might be dynamically set to influence transaction arrivals. In contrast, our framework jointly addresses both posting and pricing decisions, capturing the intricate interplay between volatile L1 gas costs and user-driven congestion. Moreover, we design and rigorously prove key properties of our dynamic L2 fee mechanism under relaxed assumptions. Furthermore, Crapis et al. [5] investigate the economics of EIP-4844, focusing on blob posting strategies and equilibrium cost-sharing among rollups. While their analysis highlights important trade-offs for rollups operating under blob constraints, our work focuses specifically on batch posting. Additionally, we incorporate a dynamic L2 fee mechanism to actively manage network congestion. By unifying the optimal posting strategy and L2 fee mechanism, our approach enables L2 systems to simultaneously maximize throughput and maintain budget balance, thereby offering a comprehensive and robust solution that extends beyond the scope of the aforementioned works.

2 Model

In what follows, we describe the main components of our transaction management framework for a Layer 2 blockchain system. This framework not only captures the dynamics of the transaction queue and gas prices but also outlines a decision model for optimizing transaction posting costs and waiting costs over time. We consider a realistic transaction processing scenario where both the de-

mand for transaction posting and fluctuating gas prices are modeled in discrete time, indexed by t .

Queue Dynamics. The L2 queue length at time t , denoted by Q_t , evolves according to the following equation:

$$Q_{t+1} = Q_t - S_t + A_t,$$

where S_t is the number of transactions posted in the t_{th} L1 block, and A_t represents the number of incoming transactions between the t_{th} and $(t+1)_{\text{th}}$ L1 block time. While A_t is modeled as an independent and identically distributed process conditioned on a specific Layer 2 gas price g , the distribution of A_t can vary as g changes, which captures the dynamic interplay between the transactions processed and new ones arriving.

L1 Gas Price Dynamics. The L1 gas price P_t , which is exogenous to the L2 structures, follows a mean-reverting process, i.e., it tends to move back toward a long-term average over time, with dynamics given by

$$P_{t+1} = \theta\mu + (1 - \theta)P_t + \sigma\omega_t,$$

where θ controls the rate of mean reversion towards the long-term average μ , σ is the volatility parameter, and ω_t follows an i.i.d. standard normal distribution. This is also known as an autoregressive AR(1) process.

The mean-reverting behavior of P_t is particularly realistic in the context of Ethereum’s EIP-1559 update, which introduces a mechanism for adjusting transaction fees that inherently aims to stabilize block sizes around a target size. Specifically, the gas fee for the next block increases if the current demand exceeds the target, and decreases if it falls below. If current demand exceeds the target then prices will increase, but as the demand subsequently responds to the increased prices, it will decrease, and in turn drive down future prices. This inherently leads to a mean-reverting dynamic. See [15] for an economic analysis of EIP-1559. Using a mean-reverting process to model P_t captures the realistic features of EIP-1559 while simplifying the model for better analysis. This approach is also supported in the literature; for example, [13] employ a fractional Ornstein-Uhlenbeck process to model gas fees, which is also a type of mean-reverting process.

Cost Structure. The cost incurred at the t_{th} L1 block is given by:

$$c(Q_t, S_t, P_t) = a(Q_t - S_t) + (b_0 + b_1 S_t)P_t \mathbf{1}_{\{S_t > 0\}},$$

where a is the proportional penalty on the number of waiting transactions after posting, b_0 is the fixed cost when transactions are posted, and b_1 scales linearly with the number of posted transactions, reflecting the sensitivity of the cost to the number of transactions. Observe that charging a linear penalty on the number of waiting transactions in each time step is equivalent to charging a linear penalty in the total number of time steps each transaction waits [11]. To see this more clearly, note that at each time step, every waiting transaction incurs a cost

of a . If a transaction waits over multiple time steps, it thus accumulates the sum of these step-by-step penalties. Interchanging the summation over time and over transactions (an argument analogous to summation by parts) demonstrates that this is precisely the same as penalizing the total waiting time for each transaction. Hence, the first term in the cost structure is fundamentally a penalty on queuing delay.

Compensation Mechanism. To mitigate user dissatisfaction due to transaction delays and to align the incentives of the sequencer, the L2 platform could implement a compensation mechanism. For every L1 block of delay experienced by a transaction, the platform refunds the user an amount proportional to the delay, calculated as a times the number of delayed L1 blocks.

This compensation ensures that from the user’s perspective, the utility remains consistent irrespective of the transaction posting time. Simultaneously, this approach aligns the incentives of the L2 sequencer with the cost function $c(Q_t, S_t, P_t)$ they are optimizing. By providing such refunds, the L2 sequencer demonstrates a commitment to minimizing transaction delays, rather than exploiting these for potential profit. This transparent compensation also reinforces trust in the platform’s operational integrity.

Objective Function. The overall objective is to minimize the expected cumulative discounted cost, given by:

$$J(Q_t, P_t) = \min_{\{S_s\}} \mathbb{E} \left[\sum_{s \geq t} \gamma^{s-t} (a(Q_s - S_s) + (b_0 + b_1 S_s) P_s \mathbf{1}_{\{S_s > 0\}}) \middle| Q_t, P_t \right], \quad (1)$$

where γ is the discount factor that captures the present value of future costs. The aim is to strike a balance between reducing the delay cost and managing posting costs effectively over time.

3 Optimal Posting Strategy

We continue our analysis of the optimal posting strategy under the assumption that the arrivals A_t are i.i.d., given a fixed Layer 2 fee. This section builds upon the established objective function (1) and focuses on optimizing the transaction posting mechanism under dynamic gas price conditions. The decision variable, S_t , determines the number of transactions to post per L1 block. It is optimized by evaluating immediate costs and forecasting future system states to enhance the transaction processing efficiency on Layer 2.

Theorem 1. Threshold Policy. *For any state (Q_t, P_t) , the optimal posting strategy S_t^* satisfies:*

$$S_t^* = \begin{cases} Q_t & \text{if } Q_t > Q^*(P_t), \\ 0 & \text{if } Q_t \leq Q^*(P_t), \end{cases}$$

where $Q^*(P_t)$ is a critical threshold dependent on the gas price P_t .

Proof. The proof is provided in Appendix A.

The threshold policy mandates a binary decision: either clear the queue by posting all transactions or withhold all transactions when conditions do not favor posting, based on the queue length relative to a dynamic threshold $Q^*(P_t)$. This straightforward, binary approach greatly simplifies operational decisions by removing intermediate options, which streamlines the posting mechanism and enhances system efficiency by ensuring that decisions are consistently aligned with current economic conditions.

By leveraging the threshold structure of the optimal policy, we significantly simplify the decision-making process in our Markov Decision Problem¹ [14]. This structure transforms the action space from $\{0, 1, \dots, Q\}$ to just two discrete options 0 or Q . Utilizing this binary action space, we implemented an efficient policy iteration algorithm. At each iteration, we evaluate the expected total cost for both possible actions by calculating the immediate cost and the expected future cost based on the current value function estimate. By directly comparing these two options, we can promptly update the policy without exhaustively searching over all possible actions. In contrast, [12] employs a Q-learning algorithm for the optimal posting strategy without utilizing a threshold policy, resulting in a computation time of approximately 72 hours for a single run, whereas our method completes in just about 6 seconds.

4 Layer 2 Pricing

In our model, we assume a uniform L2 fee per transaction, simplifying the varying costs typically seen due to different gas usages per transaction in reality. The fee mechanism in our analysis, similar to many L2 platforms, fundamentally addresses two critical objectives: maintaining budget balance and managing congestion. Therefore, our strategy seeks to optimize network throughput within the framework established by these baseline conditions, ensuring that enhancements in throughput do not compromise our commitment to budget balance and effective congestion management.

The objectives of the fee mechanism in our study, and similarly in many L2 platforms, are twofold: to maintain budget balance and to manage congestion effectively. A similar idea is also mentioned in [3]. The goal is to satisfy the minimum level of these two objectives; under this condition, we maximize the chain throughput. Our mechanism is well-suited for implementation on most optimistic rollup platforms, such as Arbitrum, Base, and Optimism.

Arrival Rate. We denote the arrival rate of transactions by $\lambda(g)$, which represents the expected number of incoming transactions between L1 block times, $\mathbb{E}[A(t; g)]$, where $A(t; g)$ represents the random variable of the number of incoming transactions between L1 block times, whose distribution varies as a function

¹ A Markov Decision Problem models sequential decision-making, where a system transitions between states based on chosen actions, with the goal of minimizing expected costs or maximizing rewards.

of the fee g . In our model, we assume that $A(t; g)$ follows a Poisson distribution [8], which is commonly used to describe event counts that have independent increments and occur with a constant rate. We further assume a linear relationship between the arrival rate and the fee:

$$\lambda(g) = \lambda_0 - kg,$$

where λ_0 is the maximum potential arrival rate when the fee is zero, and k is a constant that captures the sensitivity of the arrival rate to changes in the fee. If the fee exceeds the threshold $g > \lambda_0/k$, the arrival rate becomes zero, i.e., no arrivals occur. The linear demand curve assumption is supported by empirical analysis of [16]. While this assumption simplifies the analysis of the root existence condition, our fee mechanism is robust and applicable to a broad range of arrival rate models.

Budget Balance. Budget balance is achieved when the total fees collected from transactions match the total operational costs of processing those transactions within the L2 framework. To establish this equilibrium, we define the fee f to satisfy the following condition under the stationary distribution:

$$\mathbb{E}[A(t; f)f - c(S_t^*, Q_t, P_t; f)] = 0, \quad (2)$$

where $c(S_t^*, Q_t, P_t; f)$ is the cost associated with the optimal posting strategies given the L2 fee f . The root of this equation, $f = f^*$, defines the fee level at which the network achieves budget balance, ensuring financial sustainability by perfectly aligning revenues with costs.

Property 1. $\mathbb{E}[c(S_t^*, Q_t, P_t; f)]$ is strictly monotonically decreasing with respect to f .

Proof. The proof is provided in Appendix A.

Corollary 1. *The expected profit and loss per L1 block, given by*

$$\mathbb{E}[A(s; f)f - c(S_s^*, Q_s, P_s; f)],$$

is strictly monotonically increasing for $f \in [0, \lambda_0/(2k)]$.

Proof. The corollary is straightforward, since

$$\begin{aligned} \mathbb{E}[A(s; f)f - c(S_s^*, Q_s, P_s; f)] &= \lambda(f)f - \mathbb{E}[-c(S_s^*, Q_s, P_s; f)] \\ &= (\lambda_0 - kf)f - \mathbb{E}[c(S_s^*, Q_s, P_s; f)], \end{aligned}$$

where the product $(\lambda_0 - kf)f$ is strictly monotonically increasing for $f \in [0, \lambda_0/(2k)]$, and the expected cost $\mathbb{E}[c(S_s^*, Q_s, P_s; f)]$ is strictly monotonically decreasing due to Property 1.

Property 2. The arrival rate $\lambda(f) = 0$ if and only if $\mathbb{E}[c(S_t^*, Q_t, P_t; f)] = 0$.

Property 3. The cost function $\mathbb{E}[c(S_t^*, Q_t, P_t; f)]$ can be bounded by a decreasing linear function. Specifically,

$$\mathbb{E}[c(S_t^*, Q_t, P_t; f)] \leq (b_0 + b_1 \lambda_0 - b_1 k f) \mu.$$

Proof. The proof is provided in Appendix A.

Theorem 2. Existence of Unique Budget Balance Fee. *If*

$$\frac{\lambda_0^2}{4k} \geq (b_0 + b_1 \lambda_0) \mu,$$

then there exists a unique fee $f^ \in [0, \lambda_0/(2k)]$ that achieves budget balance, being the root of the equation defined in equation (2).*

Proof. The proof is provided in Appendix A.

The condition outlined in the theorem ensures that there is a fee level, f^* , which precisely balances the revenues from transaction fees with the costs of transaction processing, thereby achieving budget balance. Setting the fee below f^* results in revenues that fail to cover the operational costs, leading to financial losses for the platform. Conversely, setting the fee above f^* may generate surplus revenue, potentially turning a profit.

Target Arrival Rate. In managing the L2 network, it is critical to recognize the system's maximum capacity for executing transactions. The target arrival rate, denoted as $\bar{\lambda}$, is set based on this maximum capacity to ensure the network operates efficiently without being overwhelmed by an excessive volume of transactions.

Congestion Control. To align the actual arrival rate of transactions with the target $\bar{\lambda}$, a congestion control fee p is utilized. The fee that precisely balances the incoming transaction rate with the network's capacity is p^* , defined by the root of the equation:

$$\mathbb{E}[\bar{\lambda} - A(t; p)] = 0. \quad (3)$$

The fee p^* ensures that the number of transaction arrivals matches the target arrival rate $\bar{\lambda}$.

Theorem 3. Existence of Unique Congestion Control Fee. *If*

$$\frac{\lambda_0}{2} \leq \bar{\lambda} \leq \lambda_0,$$

then there exists a unique fee $p^ \in [0, \lambda_0/(2k)]$ that achieves congestion control, being the root of the equation defined in equation (3).*

Proof. This proof is straightforward, because we only need to solve

$$0 = \mathbb{E}[\bar{\lambda} - A(t; p^*)] = \bar{\lambda} - \lambda(p^*) = \bar{\lambda} - \lambda_0 + k p^*.$$

The equation has a unique root $p^* = (\lambda_0 - \bar{\lambda})/k$. The value $p^* \in [0, \lambda_0/(2k)]$ if and only if $\lambda_0/2 \leq \bar{\lambda} \leq \lambda_0$.

The condition outlined in the theorem ensures that there is a fee level, p^* , which precisely aligns the actual transaction arrival rate with the target rate λ , effectively managing congestion. Setting the fee below p^* may lead to network overload, while setting it above p^* can result in underutilization of network capacity.

Optimal Fee Strategy. Our fee mechanism is strategically designed to simultaneously achieve budget balance and manage congestion, fundamental conditions for the stable operation of an L2 platform. To this end, the fee must be set at least as high as f^* to cover operational costs and ensure financial sustainability, and at least as high as p^* to regulate the flow of transactions and prevent system overload.

Given our goal to maximize network throughput, which ideally involves keeping fees as low as possible, the optimal fee charged is $\max(f^*, p^*)$. This strategy ensures that fees are not set higher than necessary to meet the foundational requirements, allowing the platform to process the maximum number of transactions without compromising financial viability or operational stability. By charging $\max(f^*, p^*)$, we maintain a balance that supports the highest possible throughput within the constraints of budget balance and congestion management.

4.1 L2 Fee Mechanism

In this part, we introduce a dynamic fee mechanism designed to achieve both budget balance and congestion control while maximizing throughput. When considering only budget balance, we have a fee mechanism based on updating the fee f . Similarly, when considering only congestion control, we have a fee mechanism based on updating the fee p . Since we may not know in advance whether the budget balance fee f^* is greater than the congestion control fee p^* , we propose an adaptive approach that updates both fees based on observed network conditions to determine the optimal fee.

Fee Update Mechanism for Budget Balance. For the budget-balancing fee f , we define its update rule as:

$$f_{t+1} = \Pi_{[0, \lambda_0/(2k)]}(f_t - a \cdot X(t; f_t)),$$

where $\Pi_{[a,b]}(x)$ projects x onto the interval $[a, b]$ to ensure the fee remains within feasible bounds, and $a > 0$ is a step size parameter controls the magnitude of each fee update, i.e., larger values lead to more aggressive adjustments, while smaller values yield more conservative changes. Here, $X(t; f_t)$ represents the cumulative profit or loss over the t -th posting period when using fee f_t , defined as:

$$X(t; f_t) = \sum_s (A(s; f_t)f_t - c(S_s^*, Q_s, P_s; f_t)).$$

In this expression, $A(s; f_t)$ is the observed number of arrivals during the s -th L1 block time after applying fee f_t , and $c(S_s^*, Q_s, P_s; f_t)$ is the cost associated with

the optimal posting strategy at the s -th L1 block given fee f_t . The summation over s aggregates over multiple L1 block intervals within the t -th posting period. Specifically, suppose that after the $(t-1)$ -th posting, the queue is empty, thanks to Theorem 1, the Threshold Policy. During the t -th posting period, transactions arrive over several L1 block intervals. We may decide not to post immediately, accumulating transactions and observing the profit and loss in each interval (the revenue from fees minus the cost). We sum these values until we decide to post, obtaining $X(t; f_t)$.

The update rule aims to achieve financial sustainability by aligning fees with operational costs. If $X(t; f_t) < 0$, it suggests the current fee is not covering costs, prompting an increase in f_{t+1} . Conversely, if $X(t; f_t) > 0$, there is room to reduce the fee in f_{t+1} without undermining financial health, potentially boosting transaction volume.

Fee Update Mechanism for Congestion Control. For the congestion control fee p , we define its update rule as:

$$p_{t+1} = \Pi_{[0, \lambda_0/(2k)]} (p_t - b \cdot Y(t; p_t)),$$

where $b > 0$ is a step size parameter. Here, $Y(t; p_t)$ represents the cumulative difference between the target arrival rate and the actual arrivals over the t -th posting period when using fee p_t , defined as:

$$Y(t; p_t) = \sum_s (\bar{\lambda} - A(s; p_t)).$$

In this expression, $A(s; p_t)$ is the observed number of arrivals during the s -th L1 block time after applying fee p_t , and $\bar{\lambda}$ is the target arrival rate that the network aims to maintain to avoid congestion. The summation over s captures the total effect over multiple L1 block intervals within the t -th posting period.

The congestion control fee is adjusted to align the actual transaction flow with the target arrival rate. A negative $Y(t; p_t)$ indicates congestion, suggesting a need to increase p_{t+1} to reduce the incoming transaction rates. Conversely, a positive $Y(t; p_t)$ suggests the capacity to handle more transactions, allowing a fee reduction.

Adaptive Fee Selection Mechanism. Since we do not know whether $f^* \leq p^*$ or $f^* > p^*$, we adopt an adaptive mechanism that updates both fees based on observed network performance to determine the optimal fee.

Let g_t represent the fee applied during the t -th posting period, and let δ_t be an indicator variable, representing whether g_t is updated based on the fee update mechanism for budget balance or congestion control. If g_t is updated based on the budget balance fee update mechanism, the indicator variable $\delta_t = 1$; otherwise, the indicator variable $\delta_t = 0$.

Define $\zeta(t) = \max\{s \leq t : \delta_s = 1\}$ and $\eta(t) = \max\{s \leq t : \delta_s = 0\}$, which represent the last time we have updated the budget balance fee and congestion control fee up to time t respectively. The decision variable δ_{t+1} and the fee g_{t+1} for the next posting period are updated according to the following rules:

– If $\delta_t = 1$:

- If $Y(t; g_t) < 0$, indicating network congestion, we set $\delta_{t+1} = 0$ to switch to the congestion control fee sequence. The corresponding fee update rule

$$g_{t+1} = \Pi_{[0, \lambda_0/(2k)]} (g_{\eta(t)} - b \cdot Y(\eta(t); g_{\eta(t)})) .$$

- If $Y(t; g_t) \geq 0$, we keep $\delta_{t+1} = 1$, continuing with the budget balance fee sequence. The corresponding fee update rule

$$g_{t+1} = \Pi_{[0, \lambda_0/(2k)]} (g_{\zeta(t)} - a \cdot X(\zeta(t); g_{\zeta(t)})) .$$

– If $\delta_t = 0$:

- If $X(t; g_t) < 0$, indicating that revenues fail to cover operational costs, we set $\delta_{t+1} = 1$. This decision switches back to the budget balance fee update mechanism for the next period

$$g_{t+1} = \Pi_{[0, \lambda_0/(2k)]} (g_{\zeta(t)} - a \cdot X(\zeta(t); g_{\zeta(t)})) .$$

- If $X(t; g_t) \geq 0$, indicating financial stability, we keep $\delta_{t+1} = 0$, continuing with the congestion control fee sequence. The corresponding fee update for the next period

$$g_{t+1} = \Pi_{[0, \lambda_0/(2k)]} (g_{\eta(t)} - b \cdot Y(\eta(t); g_{\eta(t)})) .$$

This method shares similarity with the multi-armed bandit problem, where we iteratively select the fee update mechanism between different objectives based on observed performance. It enables the system to adaptively switch between fee mechanisms as necessary, optimizing for either budget balance or congestion control in response to real-time conditions. This mechanism is particularly robust, adept at handling non-stationary network conditions and ensuring that fees stay close to the ideal level despite fluctuations. Through this adaptive approach, the system maintains crucial controls over budget and congestion, thereby enhancing throughput and ensuring consistent operational efficiency.

4.2 Analysis

We analyze the theoretical results for the L2 fee mechanism, considering some relaxations such as decreasing step size and i.i.d. Layer 1 gas fee prices. This analysis aims to establish a robust understanding of the dynamics underpinning the fee update mechanism within a theoretical framework.

Proposition 1. *For any feasible fee g , where $0 \leq g \leq \lambda_0/(2k)$, the expected value of $X(t; g)$ under the stationary distribution can be expressed as:*

$$\mathbb{E}[X(t; g)] = \mathbb{E}[\tau(g)] \cdot \mathbb{E}[A(s; g)g - c(S_s^*, Q_s, P_s; g)],$$

where $\tau(g)$ is the number of L1 blocks between two consecutive postings, which empty the queue according to Theorem 1. Therefore, if f^* is the fee level such

that $\mathbb{E}[A(f^*)f^* - c(S_s^*, Q_s, P_s; f^*)] = 0$, then $\mathbb{E}[X(t; f^*)] = 0$. Similarly, for the congestion control metric:

$$\mathbb{E}[Y(t; g)] = \mathbb{E}[\tau(g)] \cdot \mathbb{E}[\bar{\lambda} - A(s; g)].$$

Consequently, if p^* is the fee level where $\mathbb{E}[\bar{\lambda} - A(s; p^*)] = 0$, then $\mathbb{E}[Y(t; p^*)] = 0$.

Proof. The proof is provided in Appendix A.

To underpin a well-founded theoretical result, we incorporate the following considerations:

1. **Decreasing Step Size.** The step size parameters a_t and b_t decrease over time, defined as $a_t = a/t$ and $b_t = b/t$.
2. **Independent and Identically Distributed L1 Gas Fee.** We assume that the L1 gas fee sequence $\{p_s\}$ follows an independent and identically distributed distribution.
3. **Multiple Observations.** For each fee update, we consider κ postings and use the observations of all these postings to select the fee update mechanism.

Given these considerations, our fee update rules can be described as follows:

- If $\delta_t = 1$ and $\sum_{s=\kappa t+1}^{s=\kappa(t+1)} Y(s; g_t) \geq 0$, or if $\delta_t = 0$ and $\sum_{s=\kappa t+1}^{s=\kappa(t+1)} X(s; g_t) < 0$, we set $\delta_{t+1} = 1$, and the fee update rule

$$g_{t+1} = \Pi_{[0, \lambda_0/(2k)]} \left(g_{\zeta(t)} - a \cdot \frac{\sum_{s=\kappa\zeta(t)+1}^{s=\kappa(\zeta(t)+1)} X(s; g_{\zeta(t)})}{\kappa} \right).$$

- If $\delta_t = 1$ and $\sum_{s=\kappa t+1}^{s=\kappa(t+1)} Y(s; g_t) < 0$, or if $\delta_t = 0$ and $\sum_{s=\kappa t+1}^{s=\kappa(t+1)} X(s; g_t) \geq 0$, we set $\delta_{t+1} = 0$, and the fee update rule

$$g_{t+1} = \Pi_{[0, \lambda_0/(2k)]} \left(g_{\eta(t)} - b \cdot \frac{\sum_{s=\kappa\eta(t)+1}^{s=\kappa(\eta(t)+1)} Y(s; g_{\eta(t)})}{\kappa} \right).$$

Theorem 4. *Assuming the conditions and relaxations defined above hold, consider the dynamic fee update mechanisms with δ_t indicating whether the fee update is for budget balance or congestion control. Define:*

- $i(t) = \sum_{s=1}^t \mathbf{1}\{\delta_s = 1\}$, the total number of budget balance fee updates up to t ,
- $j(t) = t - i(t)$, the total number of congestion control fee updates up to t ,
- $f_{i(t)} = g_{\zeta(t)}$, linking the budget balance fee to the last update up to t ,
- $p_{j(t)} = g_{\eta(t)}$, linking the congestion control fee to the last update up to t .

Suppose the condition $\lambda_0^2/(4k) \geq (b_0 + b_1\lambda_0)\mu$ for the existence of unique budget balance fee in Theorem 2, and the condition $\lambda_0/2 \leq \bar{\lambda} \leq \lambda_0$ for the existence of unique congestion control fee in Theorem 3 hold. Then, the following properties hold almost surely as $t \rightarrow \infty$:

1. Both $i(t) \rightarrow \infty$ and $j(t) \rightarrow \infty$.
2. The sequences $f_t \rightarrow f^*$ and $p_t \rightarrow p^*$.
3. The long-run average proportions of updates,

$$\frac{i(t)}{t} \rightarrow \pi_f \quad \text{and} \quad \frac{j(t)}{t} \rightarrow \pi_p,$$

where (π_f, π_p) is the stationary distributions under the transition matrix P . Here, P governs the transition of δ_t when the fee is fixed at the optimal levels, i.e., $g_t \equiv f^*$ for $\delta_t^* = 1$ and $g_t \equiv p^*$ for $\delta_t^* = 0$. Moreover, for any initial distribution, the distribution of δ_t converges to the stationary distribution (π_f, π_p) .

4. If $f^* \neq p^*$, as the number of observations $\kappa \rightarrow \infty$, the stationary distribution

$$(\pi_f, \pi_p) \rightarrow (\mathbf{1}\{f^* > p^*\}, \mathbf{1}\{f^* < p^*\}).$$

As $\kappa \rightarrow \infty$, the convergence rate

$$|\pi_f - \mathbf{1}\{f^* > p^*\}| = |\pi_p - \mathbf{1}\{f^* < p^*\}| = O(1/\sqrt{\kappa}).$$

Proof. The proof is provided in Appendix A.

According to Theorem 4, the sensitivity of our adaptive fee selection mechanism to network fluctuations can be moderated by adjusting the parameter κ , which determines the number of postings observed before a fee update. By increasing κ , the mechanism aggregates more data across multiple L1 block intervals before reconsidering fee changes. This extended observation window tends to smooth out short-term volatility and reduces the frequency of switching between fee update rules, thereby decreasing the likelihood of erratic fee adjustments driven by transient network effects.

In practical terms, the fee mechanism described in Section 4.1 sets $\kappa = 1$ because real-world systems are inherently non-stationary. Updating the fee for each posting, similar to how EIP-1559 updates fees for each block, allows the mechanism to quickly adapt to changing network conditions and improve robustness. Updating fees after observing a larger set of postings (i.e., increasing κ) presents a trade-off. On one hand, it improves stability by reducing sensitivity to short-term fluctuations, leading to more consistent fee decisions. On the other hand, it may diminish the mechanism's responsiveness to rapid changes in network conditions. This reduced agility may compromise robustness, as fees might not adapt swiftly to evolving demand patterns during extended observation periods.

5 Simulation

In this section, we empirically evaluate our L2 fee mechanism through a series of simulations designed to test its performance under varying conditions. Initially corroborating the theoretical model, we progressively introduce more realistic scenarios shifting from i.i.d., which assumes the L1 gas fees are independent

and identically distributed across time, providing a simplified theoretical setting, to non-i.i.d., which models the fees via a mean-reverting process to capture temporal dependence and better reflect real-world dynamics, price distributions and from decreasing to constant step sizes. These simulations aim to assess whether the fees f_n and p_n stay close to their theoretical optimal values f^* and p^* under less idealized conditions, providing insight into the mechanism’s robustness and practical applicability.

5.1 Simulation Setup

We conduct simulations to evaluate the performance of our L2 fee mechanism across four distinct scenarios, which vary by the nature of L1 gas fee conditions, either i.i.d. or non-i.i.d., and the approach of step sizes, either decreasing or constant. For each scenario, we assess how well the fee mechanism approaches the theoretical optimal fees f^* and p^* , as well as analyze the long-term frequency of selecting either the budget balance or congestion control fee update mechanisms. These simulations are crucial for understanding how well the mechanism can maintain its efficiency and effectiveness in different market conditions.

Price Generation. For i.i.d. scenarios, prices are generated from a predefined normal distribution. For non-i.i.d. scenarios, prices are generated using the mean-reverting AR(1) process, as defined in Section 2.

Step Size Configuration: For decreasing step sizes, we use $a_t = a/t$ and $b_t = b/t$, where a and b are predefined. For constant step sizes, a_t and b_t are fixed throughout the simulation.

Parameters Setup. Our parameters are calibrated as follows:

- **L1 gas fee parameter.** We use the L1 base fee data during the second half of February 2024 to get the mean and variance. Then for the i.i.d. case, we set its mean $\mu = 3.86 \times 10^{-8}$ and standard deviation $v = 1.93 \times 10^{-8}$. For the non-i.i.d. case, we set $\mu = 3.86 \times 10^{-8}$, $\theta = 0.1$ and $\sigma = 8.41 \times 10^{-9}$.
- **Arrival rate parameter.** The typical L2 transaction fee is around 0.09 USD, equivalent to 3.6×10^{-5} ETH at an ETH price of 2500 USD. Some historical data shows an average of 120 transactions per L1 block time before EIP-4844. We model the transaction decline by setting a fee cap at 0.27 USD, or 1.08×10^{-4} ETH, above which no transactions occur. This setup defines a linear arrival rate model with parameters $\lambda_0 = 180$ and $k = 1.67 \times 10^6$, based on these two fee thresholds.
- **Single observation.** We set $\kappa = 1$, which means that for each fee update, we only consider the current posting to determine the fee update mechanism.

5.2 Simulation Results

Decreasing step size and i.i.d. L1 fee case. The simulation validates the theoretical results for decreasing step sizes with i.i.d. L1 fee, as shown in Appendix B.1. Each graph demonstrates the corresponding parameter converges

in the long run. These outcomes support the use of the decreasing step size approach in achieving target efficiency in a stationary environment.

Decreasing step size and non-i.i.d. L1 fee case. In the absence of a formal theorem for almost surely convergence of the sequence f_t , p_t , $i(t)/t$ and $j(t)/t$ under non-i.i.d. conditions, our simulations serve as a critical empirical test. The results, as shown in Appendix B.2, indicate that despite the complexity introduced by the non-i.i.d. nature of L1 fees, all of these parameters still demonstrate a tendency towards convergence over an extended period. This performance suggests that our fee mechanism is robust even in more complicated and realistic L1 gas fee conditions.

Constant step size and non-i.i.d. L1 fee case. This scenario arguably presents the most practical and realistic conditions for the implementation of our L2 fee mechanism. While we lack a formal convergence theorem for this case, the simulation results, as shown in Appendix B.3, are promising, showing that the fees f_n and p_n tend to stabilize close to f^* and p^* over time. The histograms indicate that this stabilization occurs within a reasonable neighborhood given the simulation environment. This observation aligns with the theorems presented in Chapter 8 of [7], which discusses the convergence behavior for constant step-size cases, indicating asymptotic convergence within a neighborhood of the root. Moreover, the simulation indicates that the relative frequencies $i(t)/t$ and $j(t)/t$ of selecting different update mechanisms appear to converge. Although almost surely convergence may not hold in highly stable environments, the robustness and adaptability of our fee mechanism are essential for effectively managing the dynamic and often unstable conditions encountered in real-world applications.

6 Conclusion

This paper presents a framework to optimize transaction posting and dynamic fee mechanisms in L2 blockchain systems. By integrating models for transaction arrivals, queue dynamics, and cost structures, we showed that an optimal threshold policy can dictate when to post transactions to L1, balancing operational costs and system performance.

On the pricing side, we developed an L2 fee mechanism that achieves both budget balance and congestion control. Our analysis established the existence and uniqueness of fee levels that satisfy these objectives and examined the convergence properties of the adaptive update mechanism. Simulations confirm that the method efficiently adjusts fees to maintain stability and manage congestion.

Overall, our work bridges transaction posting strategies and fee design, offering a principled method that enhances operational efficiency, financial viability, and congestion management. This integrated perspective paves the way for more resilient and scalable L2 platforms and opens avenues for future research. For example, extending the framework to include blob transactions and other complex settings could further refine our approach to handle diverse transaction types and evolving network demands.

References

1. Yogev Bar-On and Yishay Mansour. Optimal publishing strategies on a base layer. *arXiv preprint arXiv:2312.06448*, 2023.
2. Soumya Basu, David Easley, Maureen O’Hara, and Emin Gün Sirer. Towards a functional fee market for cryptocurrencies. *arXiv preprint arXiv:1901.06830*, 2019.
3. Lee Bousfield, Rachel Bousfield, Chris Buckland, Ben Burgess, Joshua Colvin, Edward W Felten, Steven Goldfeder, Daniel Goldman, Braden Huddleston, H Kalonder, et al. Arbitrum nitro: A second-generation optimistic rollup, 2022.
4. Davide Crapis. Eip-4844 fee market analysis. <https://ethresear.ch/t/eip-4844-fee-market-analysis/15078>, March 2023.
5. Davide Crapis, Edward W Felten, and Akaki Mamageishvili. Eip-4844 economics and rollup strategies. *arXiv preprint arXiv:2310.01155*, 2023.
6. Davide Crapis, Ciamac C. Moallemi, and Shouqiao Wang. Optimal dynamic fees for blockchain resources. *CoRR*, abs/2309.12735, 2023.
7. J Harold, G Kushner, and George Yin. Stochastic approximation and recursive algorithm and applications. *Application of Mathematics*, 35(10), 1997.
8. SK Katti and A Vijaya Rao. Handbook of the poisson distribution, 1968.
9. Ron Lavi, Or Sattath, and Aviv Zohar. Redesigning bitcoin’s fee market. *ACM Transactions on Economics and Computation*, 10(1):1–31, 2022.
10. Stefanos Leonardos, Daniël Reijsbergen, Barnabé Monnot, and Georgios Piliouras. Optimality despite chaos in fee markets. *CoRR*, abs/2212.07175, 2022.
11. John DC Little and Stephen C Graves. Little’s law. *Building intuition: insights from basic operations management models and principles*, pages 81–100, 2008.
12. Akaki Mamageishvili and Edward W. Felten. Efficient Rollup batch posting strategy on Base Layer. *CoRR*, abs/2212.10337, 2022. [arXiv:2212.10337](https://arxiv.org/abs/2212.10337), doi: 10.48550/arXiv.2212.10337.
13. Bernhard K Meister and Henry CW Price. Gas fees on the ethereum blockchain: from foundations to derivative valuations. *Frontiers in Blockchain*, 7:1462666, 2024.
14. Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
15. Tim Roughgarden. Transaction fee mechanism design for the ethereum blockchain: An economic analysis of eip-1559. *arXiv preprint arXiv:2012.00854*, 2020.
16. Matt Stephenson and Ally Zach. Getting the pricing right in crypto, 2024. URL: <https://panteracapital.com/research-getting-the-pricing-right-in-crypto/>.
17. Aviv Yaish and Aviv Zohar. Correct cryptocurrency asic pricing: Are miners overpaying? In *5th Conference on Advances in Financial Technologies (AFT 2023)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2023.
18. Andrew Chi-Chih Yao. An incentive analysis of some bitcoin fee designs. *arXiv preprint arXiv:1811.02351*, 2018.

A Proofs

Proof of Theorem 1. Consider the objective function

$$J(Q_t, P_t) = \min_{\{S_s\}} \mathbb{E} \left[\sum_{s \geq t} \gamma^{s-t} [a(Q_s - S_s) + (b_0 + b_1 S_s) P_s \mathbf{1}\{S_s > 0\}] \right].$$

We first show that When the state is (Q_t, P_t) , we have either $S_t^* = 0$ or $S_t^* = Q_t$. We then show the following Lemmas.

Lemma 1. *For any feasible Q, P , we have*

$$2J(Q+1, P) \geq J(Q, P) + J(Q+2, P).$$

Proof. Suppose for any given sequence $\{(A_s, P_s)\}_{s \geq t}$, the optimal strategy for $J(Q+1)$ is $\{S_s^*\}_{s \geq t}$, and $u = \min\{s \geq t \mid S_s^* > 0\}$. We construct the following strategy for $J(Q)$ and $J(Q+2)$. For $J(Q)$, we let

$$S_s^Q = \begin{cases} S_s^*, & \text{when } s \neq u, \\ S_u^* - 1, & \text{when } s = u. \end{cases}$$

For $J(Q+2)$, we let

$$S_s^{Q+2} = \begin{cases} S_s^*, & \text{when } s \neq u, \\ S_u^* + 1, & \text{when } s = u. \end{cases}$$

Define the random variable

$$J_\omega(Q, P) = \sum_{s \geq t} \gamma^{s-t} [a(Q_s - S_s) + (b_0 + b_1 S_s) P_s \mathbf{1}\{S_s > 0\}] \Big| (Q_t, P_t) = (Q, P).$$

Then we can show that

$$\begin{aligned} & 2J_\omega(Q+1, P \mid \{(A_s, P_s)\}_{s \geq t}) \\ &= 2J_\omega(Q+1, P \mid \{(A_s, P_s)\}_{s \geq t}, \{S_s^*\}_{s \geq t}) \\ &\geq J_\omega(Q, P \mid \{(A_s, P_s)\}_{s \geq t}, \{S_s^Q\}_{s \geq t}) \\ &\quad + J_\omega(Q+2, P \mid \{(A_s, P_s)\}_{s \geq t}, \{S_s^{Q+2}\}_{s \geq t}) \\ &\geq J_\omega(Q, P \mid \{(A_s, P_s)\}_{s \geq t}) + J_\omega(Q+2, P \mid \{(A_s, P_s)\}_{s \geq t}), \end{aligned}$$

where the first inequality holds because

$$\begin{aligned} & 2c(Q_s, S_s^*, P \mid Q_t = Q+1) \\ &= c(Q_s, S_s^Q, P \mid Q_t = Q) + c(Q_s, S_s^{Q+2}, P \mid Q_t = Q+2) \end{aligned}$$

for all $s \neq u$ and

$$\begin{aligned} & 2c(Q_u, S_u^*, P | Q_t = Q + 1) \\ & \geq c(Q_u, S_u^Q, P | Q_t = Q) + c(Q_u, S_u^{Q+2}, P | Q_t = Q + 2). \end{aligned}$$

Hence, we have

$$\begin{aligned} 2J(Q + 1, P) &= \mathbb{E}[2J_\omega(Q + 1, P) | \{(A_s, P_s)\}_{s \geq t}] \\ &\leq \mathbb{E}[J_\omega(Q, P) | \{(A_s, P_s)\}_{s \geq t}] + \mathbb{E}[J_\omega(Q + 2, P) | \{(A_s, P_s)\}_{s \geq t}] \\ &= J(Q, P) + J(Q + 2, P). \end{aligned}$$

By Bellman equation, we have

$$\begin{aligned} & J(Q_t, P_t) \\ &= \min_{0 \leq S_t \leq Q_t} a(Q_t - S_t) + (b_0 + b_1 S_t) P_t \mathbf{1}\{S_t > 0\} \\ & \quad + \gamma \mathbb{E}[J(Q_t - S_t + A_t, P_{t+1}) | Q_t - S_t, P_t] \\ & \triangleq \min_{0 \leq S_t \leq Q_t} J^Q(Q_t, P_t, S_t). \end{aligned}$$

Lemma 2. *If $J^Q(Q_t, P_t, S_t + 1) < J^Q(Q_t, P_t, S_t)$ holds, then we always have $J^Q(Q_t, P_t, S_t + 2) < J^Q(Q_t, P_t, S_t + 1)$.*

Proof. After plugging in the definition of J^Q , this Lemma is equivalent to if

$$\begin{aligned} & -a + (b_0 \mathbf{1}\{S_t=0\} + b_1) P_t - \gamma (\mathbb{E}[J(Q_t - S_t + A_t, P_{t+1}) \\ & \quad - J(Q_t - S_t - 1 + A_t, P_{t+1}) | Q_t - S_t, P_t]) \leq 0, \end{aligned}$$

then we have

$$\begin{aligned} & -a + b_1 P_t - \gamma (\mathbb{E}[J(Q_t - S_t + A_t - 1, P_{t+1}) \\ & \quad - J(Q_t - S_t - 2 + A_t, P_{t+1}) | Q_t - S_t, P_t]) \leq 0. \end{aligned}$$

Comparing these two inequalities, we only need to prove that

$$\begin{aligned} & -a + b_1 P_t - \gamma (\mathbb{E}[J(Q_t - S_t + A_t - 1, P_{t+1}) \\ & \quad - J(Q_t - S_t - 2 + A_t, P_{t+1}) | Q_t - S_t, P_t]) \\ & \leq -a + (b_0 \mathbf{1}\{S_t=0\} + b_1) P_t - \gamma (\mathbb{E}[J(Q_t - S_t + A_t, P_{t+1}) \\ & \quad - J(Q_t - S_t - 1 + A_t, P_{t+1}) | Q_t - S_t, P_t]), \end{aligned}$$

which means we only need to show that

$$\begin{aligned} & \mathbb{E}[2J(Q_t - S_t + A_t - 1, P_{t+1}) - J(Q_t - S_t + A_t, P_{t+1}) \\ & \quad - J(Q_t - S_t - 2 + A_t, P_{t+1}) | Q_t - S_t, P_t] \geq 0. \end{aligned}$$

This can be directly derived from the Lemma 1.

By the Lemma 2, we know that if for any feasible S_t , we have

$$J^Q(Q_t, P_t, S_t + 1) \geq J^Q(Q_t, P_t, S_t),$$

then the optimal policy is $S_t^* = 0$. Otherwise, suppose the minimum feasible S_t such that

$$J^Q(Q_t, P_t, S_t + 1) < J^Q(Q_t, P_t, S_t)$$

is S , which also means that

$$J^Q(Q_t, P_t, 0) \leq J^Q(Q_t, P_t, 1) \leq \dots \leq J^Q(Q_t, P_t, S).$$

Then we iteratively use the Lemma to get

$$J^Q(Q_t, P_t, S) > J^Q(Q_t, P_t, S + 1) > \dots > J^Q(Q_t, P_t, Q_t).$$

Combing these two inequalities, we can get

$$\min_{0 \leq S_t \leq Q_t} J^Q(Q_t, P_t, S_t) = \min\{J^Q(Q_t, P_t, 0), J^Q(Q_t, P_t, Q_t)\},$$

which means that the optimal value S_t^* is either 0 or Q_t .

We have already proved that the optimal posting strategy S_t^* is either 0 or Q_t . We still need to prove the existence of $Q^*(P_t)$. Then we only need to prove that if $S_t^*(Q_t, P_t) = Q_t$, we always have $S_t^*(Q_t + 1, P_t) = Q_t + 1$, because in this way, the threshold

$$Q^*(P_t) = \min\{Q : S_t^*(Q, P_t) = Q\}.$$

The condition of $S_t^*(Q_t, P_t) = Q_t$ means that

$$J^Q(Q_t, P_t, 0) \geq J^Q(Q_t, P_t, Q_t).$$

This is equivalent to

$$aQ_t + \gamma \mathbb{E}[J(Q_t + A_t, P_{t+1}) | Q_t, P_t] \geq (b_0 + b_1 Q_t)P_t + \gamma \mathbb{E}[J(A_t, P_{t+1}) | P_t].$$

By Lemma 1, given Q_t , we can show that $\forall A_t = A, P_{t+1}$, there is

$$J(Q_t + A, P_{t+1}) - J(A, P_{t+1}) \leq \frac{Q_t}{Q_t + 1} (J(Q_t + 1 + A, P_{t+1}) - J(A, P_{t+1})),$$

which means that

$$\begin{aligned} & \mathbb{E}[J(Q_t + A, P_{t+1}) - J(A, P_{t+1}) | Q_t] \\ & \leq \frac{Q_t}{Q_t + 1} \mathbb{E}[J(Q_t + 1 + A, P_{t+1}) - J(A, P_{t+1}) | Q_t]. \end{aligned}$$

Thus, we can derive

$$\begin{aligned}
 & \gamma \mathbb{E} [J(Q_t + 1 + A, P_{t+1}) - J(A, P_{t+1}) | Q_t] \\
 & \geq \frac{Q_t + 1}{Q_t} \gamma \mathbb{E} [J(Q_t + A, P_{t+1}) - J(A, P_{t+1}) | Q_t] \\
 & \geq \frac{Q_t + 1}{Q_t} (-aQ_t + (b_0 + b_1 Q_t) P_t) \\
 & = (Q_t + 1) \left(-a + \left(\frac{b_0}{Q_t} + b_1 \right) P_t \right) \\
 & \geq (Q_t + 1) \left(-a + \left(\frac{b_0}{Q_t + 1} + b_1 \right) P_t \right) \\
 & = -a(Q_t + 1) + (b_0 + b_1(Q_t + 1)) P_t.
 \end{aligned}$$

By rearranging the inequality and Bellman equation, we can get

$$J^Q(Q_t + 1, P_t, 0) \geq J^Q(Q_t + 1, P_t, Q_t + 1),$$

which means that $S_t^*(Q_t + 1, P_t) = Q_t + 1$.

Proof of Property 1. For all $0 \leq f_1 < f_2 \leq \lambda_0/k$, we can get $\lambda_1 > \lambda_2$, where λ_1 and λ_2 are the transactions arrival rate per L1 block time with the L2 fee f_1 and f_2 respectively.

Leveraging the properties of Poisson processes, the arrival process with rate λ_1 can be decomposed into two independent Poisson processes: one with rate λ_2 and another with rate $\lambda_3 = \lambda_1 - \lambda_2$. Suppose f_3 satisfies $\lambda(f_3) = \lambda_1 - \lambda_2$. Let $A(t; f_1)$, $A(t; f_2)$, and $A(t; f_3)$ denote the number of arrivals up to time t under fees f_1 , f_2 , and f_3 , respectively, such that

$$A(t; f_1) = A(t; f_2) + A(t; f_3).$$

Assume $S_t^*(f_2)$ and $S_t^*(f_3)$ are the optimal policies for processing the arrival processes $A(t; f_2)$ and $A(t; f_3)$, respectively. Define a combined policy $\tilde{S}_t = S_t^*(f_2) + S_t^*(f_3)$ for the total arrival process $A(t; f_1)$. Since $S_t^*(f_1)$ is the optimal policy for $A(t; f_1)$, it satisfies

$$\mathbb{E} [c(S_t^*(f_1), Q_t, P_t; f_1)] \leq \mathbb{E} [c(\tilde{S}_t, Q_t, P_t; f_1)].$$

The cost function c is additive over independent processes due to the independence of $A(t; f_2)$ and $A(t; f_3)$. Therefore, we have

$$\mathbb{E} [c(\tilde{S}_t, Q_t, P_t; f_1)] = \mathbb{E} [c(S_t^*(f_2), Q_t^{(2)}, P_t; f_2)] + \mathbb{E} [c(S_t^*(f_3), Q_t^{(3)}, P_t; f_3)],$$

where $Q_t^{(2)}$ and $Q_t^{(3)}$ denote the queue lengths corresponding to the arrival processes $A(t; f_2)$ and $A(t; f_3)$, respectively.

Combining these results, we obtain

$$\mathbb{E} [c(S_t^*(f_1), Q_t, P_t; f_1)] \leq \mathbb{E} [c(S_t^*(f_2), Q_t^{(2)}, P_t; f_2)] + \mathbb{E} [c(S_t^*(f_3), Q_t^{(3)}, P_t; f_3)].$$

This concludes the proof of the monotonicity.

Proof of Property 3. Consider a posting strategy that we post all the transactions at each block, we have

$$\begin{aligned}
\mathbb{E}[c(S_t^*, Q_t, P_t; f)] &\leq \mathbb{E}[c(Q_t, Q_t, P_t; f)] \\
&= \mathbb{E}[(b_0 + b_1 Q_t) P_t \mathbf{1}\{Q_t > 0\}] \\
&\leq \mathbb{E}[(b_0 + b_1 Q_t) P_t] \\
&= (b_0 + b_1 \lambda(f)) \mu \\
&= (b_0 + b_1 \lambda_0 - b_1 k f) \mu.
\end{aligned}$$

Proof of Theorem 2. Suppose $C(f) = (b_0 + b_1 \lambda_0 - b_1 k f) \mu$, which is the upper bound of $\mathbb{E}[c(S_t^*, Q_t, P_t; f)]$ by *Property 3*. We first consider the sufficient condition that there is at least one root for $\lambda(f)f = C(f)$. Since

$$C(\lambda_0/k + b_0/(b_1 k)) = 0$$

and $C(f)$ is monotonically decreasing, the sufficient condition that $\lambda(f)f = C(f)$ has at least one root can be

$$\max\{\lambda(f)f\} \geq C(0).$$

This is equivalent to

$$\frac{\lambda_0^2}{4k} \geq (b_0 + b_1 \lambda_0) \mu.$$

If $\lambda_0^2/(4k) \geq (b_0 + b_1 \lambda_0) \mu$, we know that there exists $\hat{f} \in [0, \lambda_0/(2k)]$ such that $\lambda(\hat{f})\hat{f} = C(\hat{f})$. When $f = 0$, we have

$$\lambda(f)f = 0 < \mathbb{E}[c(S_t^*, Q_t, P_t; 0)].$$

When $f = \hat{f}$, we have

$$\lambda(\hat{f})\hat{f} = C(\hat{f}) \geq \mathbb{E}[c(S_t^*, Q_t, P_t; \hat{f})].$$

Thus, there exists a root $f^* \in [0, \hat{f}]$, such that

$$\lambda(f^*)f^* = \mathbb{E}[c(S_t^*, Q_t, P_t; f^*)].$$

The uniqueness of the root is derived by the strictly monotonically increasing property of $\mathbb{E}[A(s; f)f - c(S_s^*, Q_s, P_s; f)]$ when $f \in [0, \lambda_0/(2k)]$.

Proof of Proposition 1.

Proof. By ergodic theorem and the strong law of large number, we have

$$\begin{aligned}
 \mathbb{E}[X(t; g)] &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \left[\sum_{s=1}^{\tau_i(g)} A_s^{(i)}(g)g - c^{(i)}(S_s^*, Q_s, P_s; g) \right] \\
 &= \lim_{N \rightarrow \infty} \frac{\sum_{i=1}^N \tau_i(g)}{N} \frac{1}{\sum_{i=1}^N \tau_i(g)} \sum_{i=1}^N \sum_{s=1}^{\tau_i(g)} A_s^{(i)}(g)g - c^{(i)}(S_s^*, Q_s, P_s; g) \\
 &= \lim_{N \rightarrow \infty} \frac{\sum_{i=1}^N \tau_i(g)}{N} \\
 &\quad \cdot \lim_{N \rightarrow \infty} \frac{1}{\sum_{i=1}^N \tau_i(g)} \sum_{i=1}^N \sum_{s=1}^{\tau_i(g)} A_s^{(i)}(g)g - c^{(i)}(S_s^*, Q_s, P_s; g) \\
 &= \mathbb{E}[\tau(g)] \mathbb{E}[A(g)g - c(S_s^*, Q_s, P_s; g)].
 \end{aligned}$$

Here $A_s^{(i)}(g) = A(l; g)$ and $c^{(i)}(S_s^*, Q_s, P_s; g) = c(S_l^*, Q_l, P_l; g)$, where

$$l = \sum_{j=1}^{i-1} \tau_j(g) + s.$$

After plugging in f^* , we will have $\mathbb{E}[X(s; f^*)] = 0$.

Similarly, we can also get the result for $\mathbb{E}[Y(t; g)]$.

Proof of Theorem 4. First, we show that $i(t) \rightarrow \infty$ and $j(t) \rightarrow \infty$ almost surely. If $i(t) < \infty$, there must $\exists M \geq 0$, such that $\forall t > M$, we have $\delta_t = 0$. Consider

$$\begin{aligned}
 P(i(t) < \infty) &\leq \sum_{l=0}^{\infty} P(M = l) P(\forall t > M, \delta_t = 0 | M = l) \\
 &\leq \sum_{l=0}^{\infty} P(\forall t > M, \delta_t = 0 | M = l) \\
 &= \sum_{l=0}^{\infty} P(\forall t > l, \delta_t = 0 | M = l).
 \end{aligned}$$

For all $l \geq 0$, we have

$$\begin{aligned}
& P(\forall t > l, \delta_t = 0 | M = l) \\
&= \prod_{t=l+1}^{\infty} P(\delta_t = 0 | M = l, \delta_{t-1} = \dots = \delta_{l+1} = 0) \\
&\leq \prod_{t=l+2}^{\infty} P(\delta_t = 0 | M = l, \delta_{t-1} = \dots = \delta_{l+1} = 0) \\
&\leq \prod_{t=l+2}^{\infty} \max_{p \in [0, \lambda_0/(2k)]} P(\delta_t = 0 | M = l, \delta_{t-1} = \dots = \delta_{l+1} = 0, p_{j(t-1)} = p) \\
&= \prod_{t=l+2}^{\infty} \max_{p \in [0, \lambda_0/(2k)]} P(\delta_t = 0 | \delta_{t-1} = 0, p_{j(t-1)} = p) \\
&= \prod_{t=l+2}^{\infty} \max_{p \in [0, \lambda_0/(2k)]} P\left(\sum_{s=\kappa(t-1)+1}^{s=\kappa t} X(s; p) \geq 0\right) \\
&\leq \prod_{t=l+2}^{\infty} (1 - \rho) \\
&= 0,
\end{aligned}$$

where

$$\rho = 1 - \max_{p \in [0, \lambda_0/(2k)]} P\left(\sum_{s=\kappa(t-1)+1}^{s=\kappa t} X(s; p) \geq 0\right),$$

and it is easy to show that $\rho > 0$ by using the compact property of interval $[0, \lambda_0/(2k)]$. Hence, we have

$$\begin{aligned}
P(i(t) < \infty) &\leq \sum_{l=0}^{\infty} P(\forall t > l, \delta_t = 0 | M = l) \\
&= \sum_{l=0}^{\infty} 0 \\
&= 0,
\end{aligned}$$

which means that $P(i(t) = \infty) = 1$. Similarly, we have $P(j(t) = \infty) = 1$. Thus, $i(t) \rightarrow \infty$ and $j(t) \rightarrow \infty$ almost surely.

Second, we show that $f_t \rightarrow f^*$ and $p_t \rightarrow p^*$ almost surely. Here, we can think of the sequence $\{g_t\}$ is generated as the following equivalent way. We first generate two sequences $\{f_t\}$ and $\{p_t\}$ independently. Then $\{g_t\}$ is a merged sequence of the sequence $\{f_t\}$ and $\{p_t\}$. The way to choose from either $\{f_t\}$ or $\{p_t\}$ is based on our adaptive fee selection mechanism. Since we have already proved that $i(t) \rightarrow \infty$ and $j(t) \rightarrow \infty$ almost surely, we only need to show that the sequences $\{f_t\}$ and $\{p_t\}$ that we generated at the beginning converge to f^*

and p^* almost surely respectively. This can be derived straightforwardly using the Theorem 2.1 at Page 127 of [7] with Theorem 2 and Theorem 3.

Third, we show that the long-run average proportions $i(t)/t \rightarrow \pi_f$ and $j(t)/t \rightarrow \pi_p$ almost surely. Since we already show that $i(t) \rightarrow \infty$, $j(t) \rightarrow \infty$, $f_t \rightarrow f^*$, $p_t \rightarrow p^*$ and when $f_t \equiv f^*$, $p_t \equiv p^*$, the sequence δ_t becomes a discrete time Markov chain with transition matrix P , we can also get the transition matrix in our scenario, which only based on $f_{i(t)}$ and $p_{j(t)}$ will converge to P almost surely. For all $\epsilon > 0$, there $\exists \alpha > 0$, such that for any matrix P' which satisfies $|P_{uv} - P'_{uv}| < \alpha$ for any $u, v \in 0, 1$, we always have the stationary distribution of P' , denoted by (π'_f, π'_p) , satisfies $|\pi_f - \pi'_f| < \epsilon$. Note that $|\pi_f - \pi'_f| < \epsilon$ also means $|\pi_p - \pi'_p| < \epsilon$. Since the real transition matrix at time t , denoted by $P_t \rightarrow P$ almost surely. There exists a finite time N , such that $\forall t > N$, we have $|P_{uv} - (P_t)_{uv}| < \alpha$.

In order to show that $i(t)/t \rightarrow \pi_f$, we only need to show that

$$\frac{\sum_{l=1}^n t_l^{(f)}}{n} \rightarrow \frac{1}{\pi_f},$$

where $t_l^{(f)}$ denotes the number of time period until the sequence δ_t change to the $(l+1)_{\text{th}}$ 1 from the l_{th} 1. We consider a matrix \hat{P} , such that $\hat{P}_{11} = P_{11} - a$ and $\hat{P}_{01} = P_{01} - a$. For all $t > N$, we know that the transition matrix $(P_t)_{01} \geq \hat{P}_{01}$ and $(P_t)_{11} \geq \hat{P}_{11}$. We construct the following two dimensional i.i.d. random variable ξ_{uv} follows uniform distribution from the interval $[0, 1]$ to illustrate the transition of state. At any time t , suppose the state $\delta_t = 1$, we check the $i(t)_{\text{th}}$ row of i.i.d. sequence $\{\xi_{i(t)s}\}_{s \geq 1}$. If $\xi_{i(t)1} \leq (P_t)_{11}$, we have $\delta_{t+1} = 1$, otherwise, $\delta_{t+1} = 0$. We will always stay on the $i(t)_{\text{th}}$ row until δ visits the state 1 again. Thus, $t_{i(t)}^{(f)}$ represents the total amount of time one spends on the row of $i(t)$. Then, we construct the other random variable based on the $i(t)_{\text{th}}$ row of $\{\xi_{i(t)s}\}_{s \geq 1}$. At the beginning, we let another sequence β_t starts from 1, if $\xi_{i(t)1} \leq \hat{P}_{11}$, we have $\beta_{t+1} = 1$, otherwise, $\beta_{t+1} = 0$. The sequence β performs like a Markov chain with transition matrix \hat{P} . The random variable $\hat{t}_{i(t)}$ is the total number of β sequence return back to 1. Hence, we know that $\hat{t}_{i(t)} \leq t_{i(t)}^{(f)}$. Thus, we have

$$\liminf \frac{\sum_{l=1}^n t_l^{(f)}}{n} \geq \liminf \frac{\sum_{l=1}^n \hat{t}_l}{n} = \lim \frac{\sum_{l=1}^n \hat{t}_l}{n} \geq \frac{1}{\pi_f + \epsilon}.$$

Similarly, we can also construct another matrix and get

$$\limsup \frac{\sum_{l=1}^n t_l^{(f)}}{n} \leq \frac{1}{\pi_f - \epsilon}.$$

By the arbitrary property of $\epsilon > 0$, we let $\epsilon \rightarrow 0$ and get

$$\frac{\sum_{l=1}^n t_l^{(f)}}{n} \rightarrow \frac{1}{\pi_f},$$

Thus, we get $i(t)/t \rightarrow \pi_f$, and we can also get $j(t)/t \rightarrow \pi_p$ in the same way.

Then we show that for any initial distribution, the distribution of δ_t converges to the stationary distribution (π_f, π_p) . It is not hard (just two dimension) to show that for any two dimensional vector $x \neq 0$ such that $x^\top \mathbf{1} = 0$, we have $\|Px\| \leq \gamma \|x\|$, where $\gamma = |P_{00} + P_{11} - 1| < 1$. Suppose P_t is the transition matrix from time period δ_t to δ_{t+1} . Since we know that $P_t \rightarrow P$, there must exist a $N > 0$, such that $\forall t > N$, we have $\|P_t - P\|$ is sufficiently small such that we have $\forall x^\top \mathbf{1} = 0, x \neq 0$, we have

$$\|P_t x\| \leq \frac{\gamma + 1}{2} \|x\|.$$

Then, for any initial distribution u_1 , suppose v is the stationary distribution under transition matrix P , we have

$$\begin{aligned} \|u_{t+1} - v\| &= \|P_t u_t - v\| \\ &\leq \|P_t u_t - P_t v\| + \|P_t v - P v\| \\ &\leq \frac{\gamma + 1}{2} \|u_t - v\| + \|P_t - P\| \|v\| \\ &= \frac{\gamma + 1}{2} \|u_t - v\| + \|P_t - P\|. \end{aligned}$$

Suppose $\|P_t - P\| = \alpha_t \rightarrow 0$, and $\hat{\gamma} = (\gamma + 1)/2 < 1$. we have

$$\frac{\|u_{t+1} - v\|}{\hat{\gamma}^{t+1}} \leq \frac{\|u_t - v\|}{\hat{\gamma}^t} + \frac{\alpha_t}{\hat{\gamma}^{t+1}}.$$

After summing up from the both side, we can get

$$\|u_t - v\| \leq \sum_{l=1}^{t-1} \hat{\gamma}^{t-1+l} \alpha_l + \hat{\gamma}^{t-1} \|u_1 - v\|.$$

For all $\epsilon > 0$, since $\alpha_n \rightarrow 0$, there exists $N(\epsilon)$, such that $\forall n > N(\epsilon)$, we have $\alpha_n < \epsilon$. Also, there exists $M(\alpha)$ to be the upper bound of sequence $\{\alpha_n\}$. Then, we have

$$\begin{aligned} \|u_t - v\| &\leq \sum_{l=1}^{t-1} \hat{\gamma}^{t-1+l} \alpha_l + \hat{\gamma}^{t-1} \|u_1 - v\| \\ &\leq \sum_{l=1}^{N(\epsilon)} \hat{\gamma}^{t-1+l} M(\alpha) + \sum_{l=1}^{\infty} \hat{\gamma}^{l-1} \epsilon + \hat{\gamma}^{t-1} \|u_1 - v\| \\ &\leq \frac{\hat{\gamma}^{N(\epsilon)}}{1 - \hat{\gamma}} M(\alpha) + \frac{\epsilon}{\hat{\gamma}} + \hat{\gamma}^{t-1} \|u_1 - v\|. \end{aligned}$$

Then, we let $\epsilon \rightarrow 0$, which also means $N(\epsilon) \rightarrow \infty$, so that we can get $\|u_t - v\| \rightarrow 0$. Hence, $u_t \rightarrow v$, where $v = (\pi_f, \pi_p)$.

Fourth, we only need to show the convergence of the transition matrix P . Since we are studying the property of P , we should assume that the fee sequence $f_n \equiv f^*$ and $p_n \equiv p^*$. Suppose $f^* > p^*$, by Central Limit Theorem, we know that

$$P(\delta_{n+1} = 0 | \delta_n = 1) = P\left(\sum_{s=\kappa t+1}^{s=\kappa(t+1)} Y(s; f^*) < 0\right) \rightarrow 0$$

as $\kappa \rightarrow \infty$. Similarly, we also have $P(\delta_{n+1} = 0 | \delta_n = 0) \rightarrow 0$. Hence, the stationary distribution $\pi_f \rightarrow 1$ and $\pi_p \rightarrow 0$ as $\kappa \rightarrow \infty$. Conversely, when $f^* < p^*$, we have the stationary distribution $\pi_f \rightarrow 0$ and $\pi_p \rightarrow 1$ as $\kappa \rightarrow \infty$.

The convergence rate can be get by the Berry–Esseen Theorem. If $f^* > p^*$, by Berry–Esseen Theorem, we can show that $P(\delta_{n+1} = 0 | \delta_n = 1) = O(1/\sqrt{\kappa})$ and $P(\delta_{n+1} = 0 | \delta_n = 0) = O(1/\sqrt{\kappa})$. Then, after expressing the stationary distribution as a function of each element of P , we can directly get the convergence rate

$$|\pi_f - \mathbf{1}\{f^* > p^*\}| = |\pi_p - \mathbf{1}\{f^* < p^*\}| = O(1/\sqrt{\kappa}).$$

The same method also works for $f^* < p^*$.

B Figures

B.1 Decreasing step size and i.i.d. L1 fee case

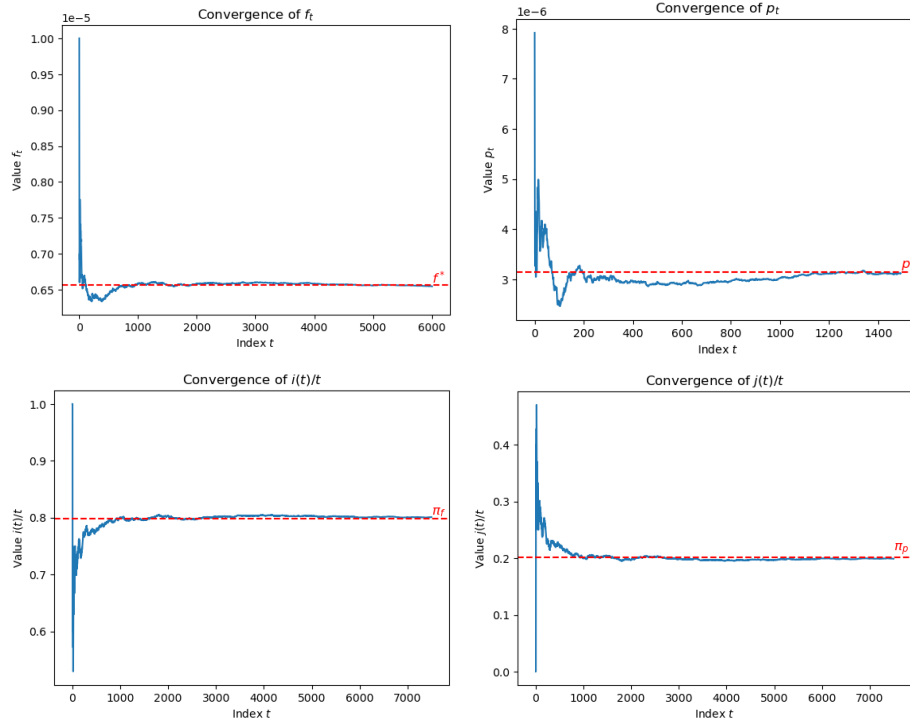


Fig. 1. Illustration of the fee update mechanism under decreasing step size and i.i.d. L1 fees. The top row displays the evolution of the posting fees over time, with red dashed lines marking the theoretical optimal levels f^* and p^* . In the bottom row, the plots show the proportion of times each type of fee update mechanism (budget balance vs. congestion control) is selected. All four trajectories exhibit convergence behavior in line with the theoretical results.

B.2 Decreasing step size and non-i.i.d. L1 fee case

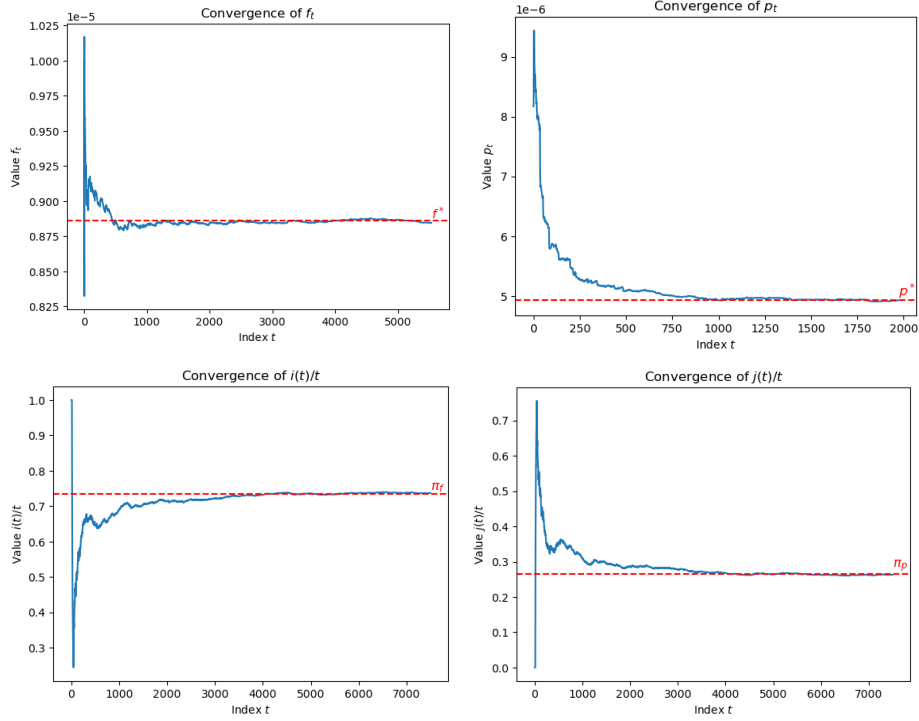


Fig. 2. Illustration of the fee update mechanism under decreasing step size and non-i.i.d. L1 fees. The top row displays the evolution of the posting fees over time, with red dashed lines marking the theoretical optimal levels f^* and p^* . In the bottom row, the plots show the proportion of times each type of fee update mechanism (budget balance vs. congestion control) is selected. All four trajectories exhibit convergence behavior in line with the theoretical results.

B.3 Constant step size and non-i.i.d. L1 fee case.

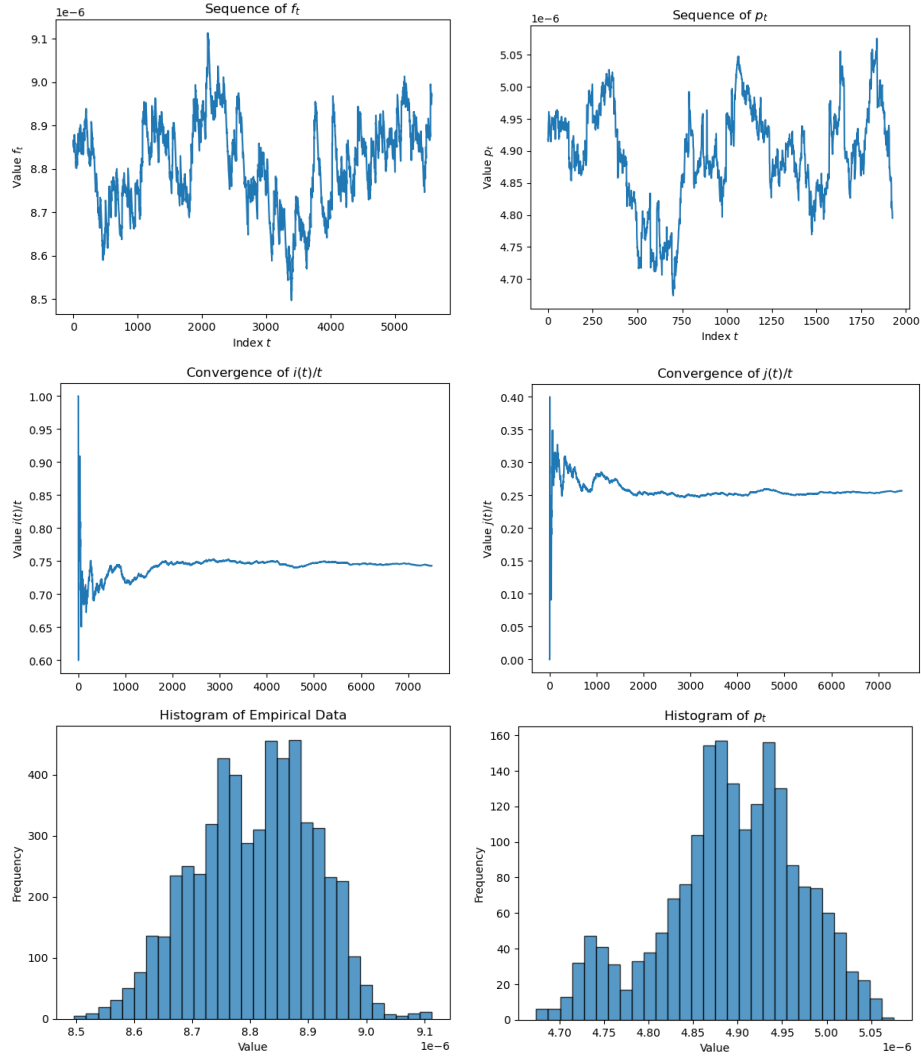


Fig. 3. Illustration of the fee update mechanism under constant step size and non-i.i.d. L1 fees. The top row shows time-series of the fee updates f_t and p_t under a mean-reverting L1 fee process. In the middle row, the plots track the long-run fraction of times each update rule is chosen, indicating how often the budget-balance versus congestion-control mechanism is active. The bottom row presents empirical histograms of f_t and p_t . Although no formal convergence theorem applies in this setting, the simulation demonstrates that both fees remain near their targets, and that the selection frequencies also stabilize, highlighting the mechanism's robustness under realistic market dynamics.