

Optimal Dynamic Fees for Blockchain Resources

Davide Crapis
Robust Incentives Group
Ethereum Foundation
davide@ethereum.org

Ciamac C. Moallemi
Graduate School of Business
Columbia University
ciamac@gsb.columbia.edu

Shouqiao Wang
Graduate School of Business
Columbia University
shwang27@gsb.columbia.edu

Initial version: May 11, 2023

Current version: September 20, 2023

Abstract We develop a general and practical framework to address the problem of the optimal design of dynamic fee mechanisms for multiple blockchain resources. Our framework allows to compute policies that optimally trade-off between adjusting resource prices to handle persistent demand shifts versus being robust to local noise in the observed block demand. In the general case with more than one resource, our optimal policies correctly handle cross-effects (complementarity and substitutability) in resource demands. We also show how these cross-effects can be used to inform resource design, *i.e.* combining resources into bundles that have low demand-side cross-effects can yield simpler and more efficient price-update rules. Our framework is also practical, we demonstrate how it can be used to refine or inform the design of heuristic fee update rules such as EIP-1559 or EIP-4844 with two case studies. We then estimate a uni-dimensional version of our model using real market data from the Ethereum blockchain and empirically compare the performance of our optimal policies to EIP-1559.

1. Introduction

Users of public permissionless blockchains can modify the shared state of the network through *transactions* that are executed by a set of nodes with limited computational resources. To allocate resources among competing transactions most blockchains use *transaction fees*. Initial transaction fee mechanisms in the Bitcoin and Ethereum blockchains relied on users bidding for transaction inclusion as the main way of pricing congestion. Moreover, all computational resources were bundled into a unique virtual resource (“gas”) with fixed relative prices hardcoded in the protocol. Current R&D efforts are focused on improving transaction fee markets along two directions: (1) setting a minimum *dynamic base fee* (henceforth also called *price*) that is adjusted by the protocol as function of user demand and (2) *unbundling resources* so that different resources can be individually priced and their relative prices can also efficiently adjust with demand.

In this paper, we propose a new framework for choosing a resource pricing policy that makes significant progress across both directions. We consider the practical problem of a blockchain protocol that has to jointly update the prices of multiple resources at every block. We assume that the type of resources being metered and priced, as well as the block limits and sustainable targets for each resource, are pre-determined. These higher level decisions are the outcome of a design process that has interesting political, economic, and engineering considerations but are outside the current scope of our framework¹.

Our framework is both general and practical. Our main results characterize theoretically optimal policies in a realistic setting with multiple resources and time-varying demand. Our results can be used in two ways: (i) the policies can be *directly* implemented as we demonstrate, or (ii) insights from our main results can be used to construct and refine heuristics that approximate optimal policies. The latter point is particularly important in the blockchain environment, where, especially at Layer 1, the price computation itself is significantly resource constrained². We designed our framework with the following properties in mind:

¹ We note that some of our results do offer some insights on the resource design problem that we will briefly discuss.

² Layer 2s, depending on their architecture, can perhaps implement price policies that require more computation and are closer to the optimal ones.

- *Simplicity*: the price update algorithm should be simple enough that it can potentially be implemented without error and tested by multiple clients, the price computation should not consume significant network resources, and the algorithm should be future-proof (it should work with minimal to no upgrades required).
- *Robustness*: the price updates should be robust to sudden and substantial shifts in resource demand as well as fee manipulation attacks (*e.g.*, a full block attack).
- *Optimality*: the performance of the algorithm should satisfy some criterion of optimality.

Our main theoretical contributions are the design of a flexible modeling framework that satisfies most of these properties and the characterization of optimal dynamic pricing policies with multiple resources and time-varying demand under a diverse set of assumptions. In particular:

1. Our framework allows to specify an arbitrary number of resources to be priced, it models the aggregate resource demand dynamics as a system of linear equations and the realized block resource demand as a system of linear functions of resource prices. It uses as objective function the quadratic deviation of block usage from sustainable target at each block plus a regularization term for controlling price fluctuations. In this way, we cast the blockchain resource pricing problem into the linear-quadratic framework that has been studied in the optimal control literature and widely used in practical applications, from rocket control to congestion pricing in ride-sharing networks.
2. We derive the optimal pricing policy for the most generic formulation with n resources and price regularization. The price update optimally handles the trade-off between adjusting prices in response of persistent resource demand shifts versus being robust to local noise in the observed block demand. Moreover, it correctly takes into account the cross-effects in the joint price updates – *i.e.*, accounting that the demand for resource A will be suppressed (increased) as a result of a price increase of resource B when these resources are complementary (substitutable), and vice-versa.
3. We present a novel decomposition of the resource pricing problem into a set of independent problems, one for each resource. *We use the structure of the demand function to define what are effectively bundles of resources*, whose optimal prices can be computed independently. We call these bundles eigenresources and their prices eigenprices. Our result shows that when this decomposition is possible, optimal prices can be constructed from simple independent update rules. Moreover, the result offers actionable insights to protocol researchers that are thinking about market design for multiple resources.

We also take the optimal policies to real market data. We collect historical data on fourteen days and close to one hundred thousand blocks of the Ethereum blockchain. We estimate our model parameters using the EM algorithm and then compute the policies that we characterized in our theoretical results. We introduce an evaluation framework where we compare our policies to EIP-1559 across a set of performance metrics. The policies perform well across the set of metrics, with the regularized policy beating the benchmark on all metrics. We have reason to believe that our framework can perform even better in multi-dimensional settings with significant cross-effects between resources and we plan to extend this empirical analysis to two-dimensions as real market data on EIP-4844 becomes available.

On top of providing optimal pricing policies and algorithms to compute them, we also show how they can be used to inspire or refine simpler heuristics that are used in practice. We present two case studies for EIP-1559 and EIP-4844. In the case of EIP-1559, we derive a heuristic based on constant price sensitivity and show how the *price update* can be efficiently computed as a function of the resource *demand estimate*. The former is already computed by most blockchains that use an EIP-1559 type formula to update base fees at every block, where usually the latest block demand is used as a naïve single point estimate of demand. Our heuristic only requires to maintain one additional variable, the demand estimate, and both resource demands and prices can be updated after each block with a simple linear update that only uses local information.

A final thing to clarify is that the linear-quadratic optimization techniques we use to derive our optimal policy have been widely studied and used. Our contribution is to cast the blockchain resource pricing problem as a variant of the linear-quadratic framework, and define models for the resource demand dynamics and demand system that fit in the framework and yet can accommodate for many practical scenarios. We also note that the relation between EIP-1559 type rules and optimization of dynamical systems has been recently studied in the literature as we discuss in the next section. The novelty of our framework is to use explicit models for the time-varying demand dynamics, something that has not been considered before. This allows

us to derive dynamic pricing policies in a changing environment that depend on primitives of the demand system such as price sensitivity/cross-sensitivity and variance.

1.1. Literature Review

Blockchain transaction fee mechanisms have recently undergone major improvements and are the subject of current R&D in the Ethereum ecosystem. The first major set of improvements was originally researched by Buterin [2018] in a seminal paper that proposed a series of ground-breaking ideas, among which justifying the use of a dynamic *base fee* that is updated by the protocol to avoid congestion. This and related ideas such as burning base fees to align incentives were implemented in the Ethereum protocol³ via EIP-1559. Roughgarden [2020] provides a thorough game-theoretic analysis of this transaction fee mechanism, and also discusses in depth the base fee update rule focusing on desiderata, potential attack vectors, and alternative designs. Reijsbergen et al. [2022] provides a performance analysis using on-chain data and explores alternative designs.

The multi-dimensional transaction fee mechanism is a further upgrade that has recently been proposed by Buterin [2022] and has led to a new wave of interest on this topic. On the applied R&D side, EIP-4844 implements the first bi-dimensional fee mechanism for blockchains with the creation of the *Ethereum datagas* resource and market which is one of the keys to Ethereum scalability, see Crapis [2023] for an analysis and simulation.

A few recent papers consider the question of designing optimal fee update rules. Perhaps the closest to our work is that of Diamandis et al. [2022] which focuses on the multi-dimensional case and derives optimal prices as the solution of a welfare maximization problem and a general class of convex functions for the network cost. While the above focuses on a *static* problem with full information, in our present work we model both demand dynamics and the uncertainty of the protocol, that has to make optimal decisions based on noisy signals of time-varying demand. Leonardos et al. [2022] study the optimality of the EIP-1559 price update rule and alternative formulations in the uni-dimensional case through the lens of dynamical systems analysis. Ferreira et al. [2021] study the stability and welfare optimality of a dynamic posted-price mechanism that conditions prices on past block utilization and observable bids, they focus on steady state analysis of a system with one resource and a probabilistic demand distribution that does not change over time.

2. Model

In what follows we describe the main components of our framework. We consider a setting with a realistic demand process for the general n -dimensional resource pricing problem. We model aggregate resource demand at every block, in discrete time indexed by k , and describe the protocol decision problem.

Demand dynamics. The aggregate resource demand d_k of all user transactions that can potentially be included in block k is governed by mean-reverting process dynamics with Gaussian noise given by

$$d_{k+1} = (I - A^d)\mu^d + A^d d_k + \epsilon_k^d. \quad (1)$$

Here, $d_k \in \mathbb{R}^n$ is the resource demand vector in block k , $\epsilon_k^d \sim \mathcal{N}(0, W^d)$ is i.i.d. Gaussian noise with covariance matrix $W^d \in \mathbb{R}^{n \times n}$, and $\mu^d \in \mathbb{R}^n$ and $A^d \in \mathbb{R}^{n \times n}$ are parameters of the demand process. This process captures the essential properties of blockchain resource demand. The long-term mean base load of resources is given by μ^d , but demand can spike and deviate to high load and revert to the base load, the matrix A^d regulates the permanence of demand spikes for different resources.⁴

Observed demand. The protocol observes the aggregate resource demand $y_k \in \mathbb{R}^n$ of all transactions that are included into a valid block. We assume that y_k is given by the demand model

$$y_k = d_k + B_k p_k + \epsilon_k^y, \quad (2)$$

³ It has also been forked and adopted by many other protocols such as Filecoin, NEAR, and Ethereum L2 chains.

⁴ The more general case in which μ^d and A^d are time-varying also fit in our framework, we consider the constant case for a clearer exposition.

where $p_k \in \mathbb{R}^n$ are resource prices (determined by the protocol), $B_k \in \mathbb{R}^{n \times n}$ are price-sensitivity parameters, and $\epsilon_k^y \sim \mathcal{N}(0, W^y)$ is i.i.d. Gaussian noise with covariance matrix $W^y \in \mathbb{R}^{n \times n}$. That is, the protocol only observes realized demand for all resources, which we model as a linear decreasing function of resource prices. We also allow for cross-sensitivity, *i.e.*, if the matrix entry $B_{k,12}$ is non-zero, the price of resource 2 can have a positive (respectively, negative) impact on demand of resource 1, if the resources are complements (respectively, substitutes).

Price sensitivity. The price-sensitivity matrix B_k determining the observed demand y_k via (2) evolves according to

$$\mathbf{vec}(B_{k+1}) = (I - A^B)\mu^B + A^B \mathbf{vec}(B_k) + \epsilon_k^B, \quad (3)$$

where $\mathbf{vec}(\cdot)$ is a function that reshapes a matrix into a column vector by stacking its columns, $A^B \in \mathbb{R}^{n^2 \times n^2}$, $\mu^B \in \mathbb{R}^{n^2}$, and $\epsilon_k^B \sim \mathcal{N}(0, W^B)$ is i.i.d. Gaussian noise with covariance matrix $W^B \in \mathbb{R}^{n^2 \times n^2}$. It follows a mean-reverting process, where the long-term mean is given by the vector μ^B , and the parameters A^B quantifies how quickly $\mathbf{vec}(B_k)$ reverts to the base μ^B .

Price update. The resource prices for block $k + 1$ are decided after observing demand at block k . Characterizing optimal prices is the subject of the following section.

Objective. Our objective function is a weighted sum of two loss functions. First, we aim to find a pricing policy that minimizes resource demand deviation from the protocol block resource target $t \in \mathbb{R}^n$. We choose the long-term average of infinite-horizon aggregation of a quadratic loss function,

$$\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \|y_k - t\|_2^2.$$

Second, we aim to find a price update policy that minimizes price fluctuations over time. We choose the long-term average of infinite-horizon quadratic variation of the price process as our second loss function,

$$\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \|p_k - p_{k-1}\|_2^2.$$

We set $\lambda \geq 0$ as the weight for these two loss functions. We can get our objective function to be

$$\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K (\|y_k - t\|_2^2 + \lambda \|p_k - p_{k-1}\|_2^2).$$

The total loss assigns increasingly high penalty to higher deviation from target demand and more significant fluctuations in price process.⁵ Note that this is a problem that involves uncertainty, in the next section we show how we can apply stochastic control techniques to find an optimal policy that minimizes the expected objective function.

3. Optimal Policy

In this section, we delve into a comprehensive analysis of the optimal pricing policy. We first describe the information structure and belief update techniques that are required to solve the stochastic control problem. We then present two approaches to solve the problem, under different assumptions.

Information structure. We first define the information structure at every block k . The *hidden state* is given by the vector

$$x_k^\top \triangleq [d_k^\top \mathbf{vec}(B_k)^\top].$$

⁵ We can also specify weights and penalize deviation of resources differently, but we assume uniform weights or alternatively that resource measurements are already properly re-scaled to focus on the fundamental trade-offs between resources.

It evolves according to

$$x_{k+1} = (\mathbf{I} - A^x)\mu^x + A^x x_k + \epsilon_k^x, \quad \epsilon_k^x \sim \mathcal{N}(0, W^x),$$

where the parameters (A^x, μ^x, W^x) can be determined from the parameters $(A^d, \mu^d, W^d, A^B, \mu^B, W^B)$ based on equations (1) and (3). Here, we allow the dependencies between ϵ_k^d and ϵ_k^B , so that the noise term ϵ_k^x is i.i.d. Gaussian noise with covariance matrix W^x . The *observation equation* is

$$y_k = d_k + B_k p_k + \epsilon_k^y, \quad \epsilon_k^y \sim \mathcal{N}(0, W^y), \quad (4)$$

where ϵ_k^y is i.i.d. and independent with ϵ_k^x . The *information available* to the controller is

$$I_k \triangleq \sigma(y_0, y_1, \dots, y_k, p_0, p_1, \dots, p_k).$$

We require that policies be adapted to the filtration $\{I_k\}$.

Belief updates. We use the Kalman filter to derive the formulas for the Bayesian update of beliefs on the state x_k and on future observation y_k . The observation equation (4) can also be expressed as

$$y_k = C_k x_k + \epsilon_k^y,$$

for an appropriate choice of the matrix C_k (depending on the price vector p_k). Using a conjugate Gaussian prior, denote the prior at k by

$$x_{k-1} \sim \mathcal{N}(\hat{x}_{k-1}, \hat{\Sigma}_{k-1}).$$

The *predictive distribution* of x_k is Gaussian $\mathcal{N}(a_k, S_k)$ with parameters

$$\begin{aligned} a_k &= \mathbf{E}(x_k | I_{k-1}) = (\mathbf{I} - A^x)\mu^x + A^x \hat{x}_{k-1}, \\ S_k &= \mathbf{Var}(x_k | I_{k-1}) = A^x \hat{\Sigma}_{k-1} (A^x)^\top + W^x. \end{aligned}$$

The *predictive distribution* of y_k is Gaussian $\mathcal{N}(f_k, F_k)$ with parameters

$$\begin{aligned} f_k &= \mathbf{E}(y_k | I_{k-1}) = C_k a_k, \\ F_k &= \mathbf{Var}(y_k | I_{k-1}) = C_k S_k C_k^\top + W^y. \end{aligned}$$

The *posterior distribution* of x_k is Gaussian $\mathcal{N}(\hat{x}_k, \hat{\Sigma}_k)$ with parameters

$$\begin{aligned} \hat{x}_k &= \mathbf{E}(x_k | I_k) = (\mathbf{I} - K_k C_k) a_k + K_k y_k, \\ \hat{\Sigma}_k &= \mathbf{Var}(x_k | I_k) = (\mathbf{I} - K_k C_k) S_k. \end{aligned}$$

where $K_k \triangleq S_k C_k^\top (C_k S_k C_k^\top + W^y)^{-1}$ is the Kalman gain. This is the standard Kalman filter update which computes the belief on the current hidden state given the current information.

Optimization problem. In the following results, we will characterize control policies that seek to minimize the expected long-term average cost, solving the problem

$$\begin{aligned} &\underset{p}{\text{minimize}} \quad J \triangleq \lim_{K \rightarrow \infty} \mathbf{E} \left\{ \frac{1}{K} \sum_{k=1}^K (\|y_k - t\|_2^2 + \lambda \|p_k - p_{k-1}\|_2^2) \right\} \\ &\text{subject to } p_{k+1} \text{ is } I_k\text{-adapted.} \end{aligned} \quad (5)$$

3.1. Target Loss Objective

Our first result focuses on the case where we only aim to find a price update policy that minimizes resource demand deviation from the target t . In particular, we do not add any penalty for fluctuations in the price process and we make the following assumption.

Assumption 1. $\lambda = 0$.

The following theorem characterizes the optimal policy under Assumption 1.

Theorem 1. *Consider Problem (5) and suppose Assumption 1 holds. Then, the optimal policy updates prices according to the rule*

$$p_{k+1}^* = [\mathbf{E}(B_{k+1}^\top B_{k+1} | I_k)]^{-1} \mathbf{E}(B_{k+1}^\top (t - d_{k+1}) | I_k).$$

Proof. The proof is provided in Appendix A.

Under Assumption 1, our focus is directed solely towards achieving the primary objective: making the observed, or realized, demand align as closely as possible with our predefined target t . This approach can be viewed as a pursuit of optimal demand estimation. We use Kalman filter to facilitate joint updates across our potential demand d_k and the price sensitivity matrix B_k . This dynamic filtering technique enables us to incorporate new information into our estimation process. The optimal vector of resource prices for block k is the one that matches the expected demand of each resource to the respective target. Note that in general, the optimal policy necessarily takes into account the expected demand, price sensitivity, and target of resource j when updating the price of resource i . Whenever demand-side cross-effects are present the optimal prices need to be set jointly.

3.2. General Setting: Model Predictive Control

We now turn to the general problem, in which we have a penalty for fluctuations in the price process, with weight $\lambda > 0$. This is important because large pricing fluctuations between adjacent blocks directly degrade user experience. Thus, controlling the magnitude of price update is a requirement for every policy that can be adopted in practice and it was one of the main design goals of EIP-1559.

Our framework allows us to flexibly control price fluctuations with the additional regularization term. However, this increases the computational complexity of the optimal policy. One particular challenge is that the price-sensitivity matrix B_k is stochastic. Since this multiplies the decision variable p_k , our setting does not fall into the standard framework of linear-quadratic-Gaussian (LQG) control.

This leads us to the following heuristic procedure: first, we observe that if the evolution of price-sensitivity matrix B_k is deterministic, then the problem can be optimally solved using standard LQG control methods. Then, in the stochastic setting, we repeatedly apply the optimal deterministic policy, always using the then current estimate for B_k , but pretending (for the purpose of control) that the future evolution of B_k is deterministic — this is the idea of model predictive control.

Optimal policy under deterministic price-sensitivity. For the moment, we will make the following assumption:

Assumption 2. *The price-sensitivity matrix B_k is deterministic or, in other words, the noise term $\epsilon_k^B \equiv 0$, for all $k \geq 0$.*

Under this assumption, the dynamics for B_k are given by

$$\text{vec}(B_{k+1}) = (I - A^B)\mu^B + A^B \text{vec}(B_k). \quad (6)$$

Using LQG control methods, we are able to derive an optimal policy explicitly, and the following theorem characterizes the optimal policy:

Theorem 2. *Consider problem (5) and suppose Assumption 2 holds. Then, the optimal policy updates prices according to the rule*

$$p_{k+1}^* = \left(I - \frac{Q_k}{\lambda}\right) p_k + \frac{Q_k}{\lambda} \text{aim}_k. \quad (7)$$

Here, $Q_k \in \mathbb{R}^{n \times n}$ is a positive definite matrix capturing the quadratic coefficient of the price vector p_k the cost-to-go function (we provide an explicit derivation for this as the solution of a Riccati equation in the Appendix A). aim_k is the “aim” price defined by

$$\text{aim}_k = \sum_{s=k+1}^{\infty} (I - Z_s) Z_{s-1} Z_{s-2} \cdots Z_{k+1} \bar{p}_s, \quad (8)$$

where $Z_s \triangleq (Q_s + B_s^\top B_s)^{-1} Q_s$ is a positive definite matrix, and \bar{p}_s is the market clearing price that satisfies

$$\mathbf{E}(y_s | I_k, p_s = \bar{p}_s) = t.$$

The sum of the coefficient matrices of the aim price satisfies

$$\sum_{s=k+1}^{\infty} (I - Z_s) Z_{s-1} Z_{s-2} \cdots Z_{k+1} = I.$$

Proof. The proof is provided in Appendix A.

We can understand Theorem 2 through an interpretation analogous to that developed by Gârleanu and Pedersen [2013], who consider different LQG control problem in a quantitative trading setting. First, observe that according to (7), the optimal price in block $k + 1$ can be expressed as the (matrix-weighted) convex combination of the current price p_k and a future “aim” price \mathbf{aim}_k . This captures a trade-off between the two objectives of not changing prices much (by staying close to the current price) and ensuring that future observed demand aligns with the target t (by targeting the aim price).

The aim price \mathbf{aim}_k , in particular, is a (matrix-weighted) average of future market clearing prices $\{\bar{p}_s\}$. Each \bar{p}_s is the price for block s that so that the expected demand in block s matches the target t , conditional on information known up to block k . The matrix weights in (8) specify exactly how much to weight to put on different future periods. These weights guide the controller in how much to emphasize targeting demand in different future blocks, when making the pricing decision for the next block. This captures the idea that one should set immediate prices in a way that not only works well for the next block, but also leaves the controller in a good position for subsequent blocks downstream.

Model predictive control. Of course, the original problem (5) allows for stochastic and time-varying price-sensitivity, and this is an important feature in the empirical analysis of Section 6. Therefore, we will adapt the policy of Theorem 2 to this more general setting using the idea of model predictive control (MPC):

1. After each block k , we estimate the current price-sensitivity according to $\hat{B}_k = \mathbf{E}(B_k | I_k)$, using all available information available at that time, under the original dynamics (3). This can be accomplished using a standard Kalman filter, as described in the belief updates of Section 2.
2. Given the estimate \hat{B}_k , we solve for the optimal policy under Theorem 2, that is, assuming that $B_k = \hat{B}_k$ and that the future evolution of B_k is deterministic as per (6).

While the MPC policy is no longer optimal for the problem (5) under stochastic price-sensitivity, we shall see in the empirical results of Section 6 that nevertheless this procedure is a powerful heuristic.

4. Resource Design and Separable Eigenresources

In this section, we present a pricing approach which holds significant economic implications. Instead of pricing individual resources as explicitly specified, we focus on pricing the linear combinations of these resources, which we refer to as eigenresources. Our main result also has implications for protocol designers thinking about how to structure supply-side resources and create markets for pricing them. The idea here is that, in the policies of Section 3, in general resources must be *jointly* priced. This is because there could be cross-price elasticities, since some resources can be complements or substitutes. In this section, we discuss a setting where, through the definition of appropriate resources, the problem nevertheless separates into different control problems for each resource.

We will derive our result under the assumption that the price-sensitivity parameter B_k can be decomposed into principal components consistently over time:

Assumption 3. *There exist fixed, orthogonal matrices $U, V \in \mathbb{R}^{n \times n}$ such that, for all k ,*

$$B_k = U \text{diag}(\delta_k) V^\top.$$

Here, $\delta_k \in \mathbb{R}^n$ is the sensitivity vector for eigenresources governed by a mean-reverting process

$$\delta_{k+1} = (I - A^\delta)\mu^\delta + A^\delta\delta_k + \epsilon_k^\delta. \quad (9)$$

where $A^\delta \in \mathbb{R}^{n \times n}$, $\mu^\delta \in \mathbb{R}^n$, and $\epsilon_k^\delta \sim \mathcal{N}(0, W^\delta)$ represents the i.i.d. Gaussian noise with covariance matrix $W^\delta \in \mathbb{R}^{n \times n}$.

Given the matrices (U, V) , we can change coordinates as follows:

Eigendemand. The eigendemand of the eigenresources $\tilde{y}_k \in \mathbb{R}^n$ that are included into a valid block is given by

$$\tilde{y}_k = U^{-1}y_k.$$

Here, the columns of U^{-1} define linear combinations of the original resources. These new eigenresources are what will be prices.

Eigenprice. The eigenprice of eigenresources $\tilde{p}_k \in \mathbb{R}^n$ is given by

$$\tilde{p}_k = V^\top p_k.$$

Eigentarget. The eigentarget of eigenresources $\tilde{t} \in \mathbb{R}^n$ is given by

$$\tilde{t} = U^{-1}t.$$

The orthogonal property of matrix U ensures the cost for the demand deviation from target t equals to the cost for the eigendemand deviation from eigentarget \tilde{t} , i.e. $\|y_k - t\|_2^2 = \|\tilde{y}_k - \tilde{t}\|_2^2$, which means it is equivalent for us to minimize the eigendemand deviation from eigentarget \tilde{t} .

Eigenprice sensitivity. We can get the eigendemand \tilde{y}_k satisfies

$$\tilde{y}_k = \tilde{d}_k + \text{diag}(\delta_k)\tilde{p}_k + \tilde{\epsilon}_k^y,$$

where $\tilde{d}_k \triangleq U^{-1}d_k$ and $\tilde{\epsilon}_k^y \triangleq U^{-1}\epsilon_k^y$ by $y_k = d_k + B_k p_k + \epsilon_k^y$. We can see that δ_k is the eigenprice-sensitivity vector indicating the price-sensitivity of each eigenresource. Equation (9) shows that the long-term mean base eigenprice-sensitivity is given by the vector μ^δ , but it can also fluctuate around the base eigenprice-sensitivity vector, and the parameters A^δ quantifies how quickly the eigenprice-sensitivity vector reverts to the base eigenprice-sensitivity vector.

Optimal policy. The following theorem characterizes the optimal policy under Assumption 1 and Assumption 3.

Theorem 3. Consider the objective of minimizing the expected long-term average cost,

$$\begin{aligned} & \underset{p}{\text{minimize}} \quad J \triangleq \lim_{K \rightarrow \infty} \mathbf{E} \left\{ \frac{1}{K} \sum_{k=1}^K \|y_k - t\|_2^2 \right\} \\ & \text{subject to } p_k \text{ is } I_{k-1}\text{-adapted.} \end{aligned}$$

Then, the optimal policy updates prices according to the rule

$$p_{k+1}^* \triangleq V \tilde{p}_{k+1}^*,$$

where the i th element of eigenprice \tilde{p}_{k+1} satisfies

$$\tilde{p}_{k+1,i}^* = \frac{\mathbf{E}(\delta_{k+1,i}(\tilde{t}_i - \tilde{d}_{k+1,i})|I_k)}{\mathbf{E}(\delta_{k+1,i}^2|I_k)},$$

$\delta_{k+1,i}, \tilde{t}_i, \tilde{d}_{k+1,i}$ are the i th element of $\delta_{k+1}, \tilde{t}, \tilde{d}_{k+1}$ respectively.

Proof. The proof is provided in Appendix A.

The optimal price in block $k + 1$ is the product of orthogonal matrix V and the eigenprice \tilde{p}_{k+1} . The explicit expression of each eigenprice $\tilde{p}_{k+1,i}^*$ shows that optimal prices for eigenresources depend only on the its own expected demand, price-sensitivity, and target. This implies that when one can leverage this decomposition, one can devise pricing strategies that are simple yet optimal. In particular, they do not require the joint determination of all prices that can be more computationally intensive.

Note that Theorem 3 is analogous to Theorem 1, in that they both handle the $\lambda = 0$ case (Assumption 1). There is a similar analogous result for the deterministic price-sensitivity case of Theorem 2: in that case, the optimal policy will similarly separate across eigenresources. We omit an explicit statement and proof of this result since it is relatively obvious given the developments thus far.

Beyond pricing, these results also holds implications for resource market design. In particular, if the protocol designer can define virtual resources that have minimal demand-side interaction between them, then simple heuristics that implement independent price updates between resources will have a smaller efficiency loss versus the optimal pricing policy. We further expand on this insight in Section 5 when we discuss EIP-4844.

5. Implications

In this section we study particular instances of our framework (*e.g.*, uni-dimensional or bi-dimensional resources) and develop practical implications by comparing optimal policies to heuristics that are used in practice, such as the EIP-1559 and the EIP-4844 price update rules. For a more direct comparison with these rules, in this section we consider the optimal policy characterized in Theorem 1 when $\lambda = 0$, and we focus on the special case in which the price sensitivity B_k is deterministic and constant.

5.1. Uni-dimensional Resource and Ethereum Gas

Suppose there is only one resource, for example a virtual resource such as *gas* in Ethereum. Define $\beta \triangleq \partial y_k / \partial p_k < 0$ to be the (deterministic, constant) price sensitivity of demand in this one dimensional model, substituting $p_{k+1} = p_k + u_k$ in the optimal pricing policy of Theorem 1 yields the following equation for the optimal price update

$$u_k = -\beta^{-1}(a_{k+1} + \beta p_k - t), \quad (10)$$

where the predictive estimate of demand can be expressed as

$$a_{k+1} = (1 - \alpha)\mu^d + \alpha \left(\frac{\sigma_{\epsilon^y}^2}{\alpha^2 \hat{\sigma}_k^2 + \sigma_{\epsilon^d}^2 + \sigma_{\epsilon^y}^2} a_k + \frac{\alpha^2 \hat{\sigma}_k^2 + \sigma_{\epsilon^d}^2}{\alpha^2 \hat{\sigma}_k^2 + \sigma_{\epsilon^d}^2 + \sigma_{\epsilon^y}^2} (y_k - \beta p_k) \right).$$

Here, α is the uni-dimensional entry of the matrix A^d and the variances $\hat{\sigma}_k^2, \sigma_{\epsilon^d}^2, \sigma_{\epsilon^y}^2$ are the uni-dimensional entries of the matrices $\hat{\Sigma}_k, W^d, W^y$, respectively. The *optimal price update* has two main properties:

- (i) *Maintains the optimal balance between responsiveness to change and robustness to noise.* It does this by using the signal-to-noise ratio $\sigma_{\epsilon^y}^2 / (\alpha^2 \hat{\sigma}_k^2 + \sigma_{\epsilon^d}^2)$ to optimally weigh in new information – the higher the idiosyncratic variance $\sigma_{\epsilon^y}^2$ of block size observation, *vis-à-vis* a more persistent state change, the less responsive the optimal controller will be to local deviation from the current estimate.
- (ii) *Optimally computes the next period price to match predicted and target demand.* First computing the predicted demand at the current price $a_{k+1} + \beta p_k$ and then adjusting with the price sensitivity. The *optimal price update is increasing in predicted demand and decreasing in price sensitivity* — *i.e.*, the more sensitive users are to prices the lower the update required to control demand.

Optimal Policy vs. EIP-1559. We compare the optimal price update policy to the EIP-1559 update rule and distill insights that can guide future design and improvements. Consider the generic EIP-1559 style linear price update rule

$$p_{k+1} = p_k \left(1 + \gamma_k \frac{y_k - t}{t} \right), \quad (11)$$

where EIP-1559 has $\gamma_k^{\text{eip1559}} \triangleq 1/8$. We can apply a few simple transformations to the optimal policy to obtain a comparable form. Note that the price elasticity of block k demand at the target point t is $\eta_k = \beta p_k / t$.

Solving for β , substituting into the scaling factor of the optimal update (10), and substituting the resulting u_k^* into the full price update equation $p_{k+1} = p_k + u_k^*$ yields

$$p_{k+1} = p_k \left(1 + \frac{1}{|\eta_k|} \frac{a_{k+1} + \beta p_k - t}{t} \right).$$

Thus, the optimal policy is similar to the EIP-1559 rule (11), with two key differences:

- EIP-1559 uses the current block size y_k as a naïve estimate for the next period demand, while the optimal policy uses the predictive estimate;
- EIP-1559 divides the adjustment by a fixed parameter, set to 8, instead of the optimal parameter which is equal to the absolute value of the current demand elasticity $|\eta_k|$.

Another way of comparing is by noting that the optimal policy is actually equivalent to the EIP-1559 style rule (11) with step-size

$$\gamma_k^{\text{opt}} = \frac{1}{|\eta_k|} \frac{a_{k+1} + \beta p_k - t}{y_k - t},$$

which shows how, while EIP-1559 updates prices with a constant learning rate of $\gamma_k^{\text{eip1559}} = 1/8$, this policy optimally adapts the learning rate to the predicted demand.

5.2. Bi-dimensional Resource and Ethereum Datagas

Consider now the case of two resources. This is akin to how the Ethereum fee market once EIP-4844 introduces an additional market for *datagas*. We expand on this shortly, for now consider the generic bi-dimensional $n = 2$ resource case, the optimal update at step k has the form

$$\begin{bmatrix} u_{k,1} \\ u_{k,2} \end{bmatrix} = \frac{1}{\beta_{12}\beta_{21} - \beta_{11}\beta_{22}} \begin{bmatrix} -\beta_{22} & \beta_{12} \\ \beta_{21} & -\beta_{11} \end{bmatrix} \begin{bmatrix} \Delta_{k,1} \\ \Delta_{k,2} \end{bmatrix},$$

where $\Delta_{k,i} = (a_{k+1} + Bp_k - t)_i$ is the predicted deviation of resource i demand from its target and β_{ij} the element in row i and column j of B . The update to price 1 can thus be expressed as

$$u_{k,1} = \frac{\Delta_{k,1} - \beta_{12}\beta_{22}^{-1}\Delta_{k,2}}{\beta_{11} - \beta_{21}\beta_{22}^{-1}\beta_{12}},$$

which computes the optimal joint resource price update taking into account both resource demand balances and cross-sensitivities. We illustrate how this works by considering three scenarios of increasing sophistication:

- (i) if resource 1 is independent of price 2 ($\beta_{12} = 0$) then the optimal update is similar to the uni-dimensional case, $u_{k,1} = \Delta_{k,1}/\beta_{11}$;
- (ii) if resource 1 is dependent of price 2 ($\beta_{12} \neq 0$) but resource 2 is independent of price 1 ($\beta_{21} = 0$) then the optimal update is $u_{k,1} = (\Delta_{k,1} - \beta_{12}\beta_{22}^{-1}\Delta_{k,2})/\beta_{11}$, where the second term in the numerator accounts for the joint update of price 2. Suppose $\Delta_{k,2} > 0$, then the optimal update of price 1 will be adjusted *downwards if the resources are complement* ($\beta_{12} < 0$) and *upwards if the resources are substitute* ($\beta_{12} > 0$);
- (iii) if both cross sensitivities are positive then the optimal update also accounts for the *indirect effect* of changing price 1 on resource 1 via its impact on the demand and optimal update of resource 2 (i.e., the pathway $\beta_{21} \rightarrow 1/\beta_{22} \rightarrow \beta_{12}$ in the demand system represented in Figure 1). The denominator in the full update formula above is the “net sensitivity” of resource 1 to price 1.

Optimal Policy vs. EIP-4844. We now compare the optimal policy to the rule that will be used in Ethereum’s fee market after EIP-4844 is implemented. As mentioned before, this EIP introduces a new resource called *datagas* on top of the standard *gas*. This new resource is consumed by a particular type of transaction called *BlobTransaction* which is likely going to be used by most Layer 2 systems to commit data to Layer 1. The




Figure 1: Example of prices-resources relationship in the demand system with cross-sensitivity matrix B (bi-dimensional case).

EIP also introduces a new fee market for the new resource with a price update rule that mimics EIP-1559.⁶ Thus at every block the prices of both resources are updated as follows

$$p_{k+1,\text{dat}} = p_{k,\text{dat}} \left(1 + \frac{1}{8} \frac{y_{k,\text{dat}} - t_{\text{dat}}}{t_{\text{dat}}} \right), \quad p_{k+1,\text{gas}} = p_{k,\text{gas}} \left(1 + \frac{1}{8} \frac{y_{k,\text{gas}} - t_{\text{gas}}}{t_{\text{gas}}} \right).$$

One important thing to note is that the updates calculated via these rules are independent from one another, while they are coupled in the optimal update policy. For example, consider the first update rule for datagas above and compare it to the optimal update policy,

$$p_{k+1,\text{dat}} = p_{k,\text{dat}} + \frac{a_{k+1,\text{dat}} + \beta_{\text{dat,dat}} p_{k,\text{dat}} + \beta_{\text{dat,gas}} p_{k,\text{gas}} - t_{\text{dat}}}{\beta_{\text{dat,dat}} - \beta_{\text{gas,dat}} \beta_{\text{gas,gas}}^{-1} \beta_{\text{dat,gas}}} - \beta_{\text{dat,gas}} \beta_{\text{gas,gas}}^{-1} \frac{a_{k+1,\text{gas}} + \beta_{\text{gas,dat}} p_{k,\text{dat}} + \beta_{\text{gas,gas}} p_{k,\text{gas}} - t_{\text{gas}}}{\beta_{\text{dat,dat}} - \beta_{\text{gas,dat}} \beta_{\text{gas,gas}}^{-1} \beta_{\text{dat,gas}}},$$

which properly accounts for the effect of an update in gas price on the future demand of datagas and, moreover, it adjusts the datagas price update taking into account the joint update of the gas price. This difference is important because *BlobTransactions* consume both resources and thus the demands are not independent. Failing to take this into account will likely result in cycles of over/under correction of prices. Inefficiencies of this type will be even more prevalent when moving beyond two resources, so we believe that future designs for multi-dimensional resource pricing should have some of the good properties of the optimal pricing policy highlighted here.

6. Empirical Analysis

We now demonstrate how, besides generating theoretical insights and benchmarks, our framework can easily be used in applied settings and calibrated or estimated using onchain data. We focus on a uni-dimensional application, for this we can use historical onchain data for gas used and the associated base fees that are computed via the EIP-1559 rule that is live on Ethereum mainnet. In the future, once the data on other resources will be available, for example the data market of EIP-4844 we can extend the analysis to the multi-dimensional case.

6.1. Empirical Setup

Data. We use historical market data for 99,547 blocks that were added to Ethereum in the span of two weeks (starting on February 24, 2023). The first block is 16,694,514 and the last block is 16,794,061, we rescale the

⁶ In reality an exponential form is used for *datagas* instead of the linear form used for *gas*. Since the two forms are approximately equal we use the linear form for both resources to make the presentation clearer.

first block number in our dataset to 0 for clarity. We also note that this is a generic 14-day period that does not represent any particular market conditions. For every block we observe the prevailing EIP-1559 base fee and we compute the total gas consumed aggregating over all the transactions included in the block.

Parameter Estimation. We estimate the uni-dimensional version of our model using the Expectation-maximization (EM) method on the observed sample. In particular, given the sample observations and the full unobserved state, respectively

$$\mathcal{I}_T \triangleq \{(y_0, p_0), (y_1, p_1), \dots, (y_T, p_T)\} \quad \text{and} \quad \mathcal{S}_T \triangleq \{(d_0, \beta_0), (d_1, \beta_1), \dots, (d_T, \beta_T)\},$$

we want an estimator for our model $\theta \triangleq (\mu^d, \mu^\beta, \alpha^d, \alpha^\beta, \sigma_{\epsilon^d}, \sigma_{\epsilon^\beta}, \rho_{\epsilon^d, \epsilon^\beta}, \sigma_{\epsilon^y}, d_0, \beta_0, \sigma_{d_0}, \sigma_{\beta_0}, \rho_{d_0, \beta_0})$, where α^d and α^β are the one dimensional entries of the respective matrices A^d and A^β , σ denotes the standard deviation and ρ the correlation coefficient of the respective noise terms, and the last five parameters are the prior means, standard deviations, and correlation of the state variables at time 0. The EM algorithm finds the MLE of the marginal likelihood by iteratively applying two steps:

E step. Compute conditional likelihood $p(\mathcal{S}_T | \mathcal{I}_T; \theta_t)$.

M step. Compute $\theta_{t+1} = \operatorname{argmax}_\theta \mathbb{E}[\log p(\mathcal{S}_T, \mathcal{I}_T; \theta) | p(\mathcal{S}_T | \mathcal{I}_T; \theta_t)]$.

In particular we use the Kalman smooth to compute the likelihood in the E step and then compute optimal solution of the M step where we have a convex objective function (Appendix B reports a detailed description of the algorithm). The result is the following MLE of our model parameters:

Parameter	Value
μ^d	7.30×10^7
μ^β	-2.38×10^{-3}
α^d	9.97×10^{-1}
α^β	9.98×10^{-1}
σ_{ϵ^d}	1.04×10^6
σ_{ϵ^β}	5.35×10^{-5}
$\rho_{\epsilon^d, \epsilon^\beta}$	-5.48×10^{-1}
σ_w	7.21×10^6
d_0	4.47×10^7
β_0	-9.62×10^{-4}
σ_{d_0}	1.47×10^6
σ_{β_0}	3.05×10^{-4}
ρ_{d_0, β_0}	-2.57×10^{-1}

Table 1: MLE Estimates of Model Parameters via EM Algorithm with Kalman smoother.

Simulation. We now use our model to simulate the optimal policy and then compare to the EIP-1559 policy for our main results. The main outcome of our simulation is the state estimates that, together with the model parameters, fully characterize our optimal policy. We use the Kalman filter to update state estimates online and then compute the optimal policies as characterized in Section 3. The following figure reports the state estimates. We can see how the state estimates vary during the day, adapting to the patterns of observed demand. The mean for the demand sensitivity is $\bar{\beta} = -0.0025$, we can use the prevailing price to compute the implied demand elasticity $\eta_k = \beta_k p_k / t$ at the target for every block. The mean implied elasticity over the entire sample is $\bar{\eta} \approx -4$.

6.2. Empirical Results

We compute performance metrics for three policies and summarize them in the following table. The first policy is the optimal policy for the problem without price regularization ($\lambda = 0$) that we characterized in




Figure 2: State estimates using the Kalman filter, potential demand \hat{d}_k (left) and demand sensitivity $\hat{\beta}_k$ (right), for the first 7,080 blocks in our sample. The orange line reports the mean for the respective parameter estimate computed over *all blocks in our sample*.

Theorem 1, we call this policy LINDY(0). The second policy is the policy with price regularization ($\lambda > 0$), computed according to the MPC method that we derived from Theorem 2, we call this policy LINDY(λ) with $\lambda = 10^{-7}$ in this case. The last row reports the benchmark EIP-1559 policy. The main results are reported in Table 2, for each policy we compute the following metrics:

- **Gas Used Bias: $\text{Bias}(y, t)$.** Average deviation of gas used from target t over all blocks in our sample. Unit: gas.
- **Gas Used Standard Deviation: $\text{SD}(y)$.** Sample standard deviation of gas used. Unit: gas.
- **Root Mean Squared Deviation (of gas used from target): $\text{RMSD}(y, t)$.** Square root of sum of squared difference between gas used and 15M gas target over all blocks. It can also be expressed as function of bias and standard deviation as follows

$$\text{RMSD}(y, t) = \sqrt{\frac{1}{T} \sum_{k=1}^T (y_k - t)^2} = \sqrt{\text{Bias}(y, t)^2 + \text{SD}(y)^2}.$$

Unit: gas.

- **Fraction of near-full blocks: $\phi_{0.95}(y)$.** Fraction of blocks that are more than 95% full, i.e. blocks that use more than 28.5M gas. Unit: percent.
- **Root Mean Squared Update: $\text{RMSU}(p)$.** Square root of sum of squared updates, i.e. the difference between p_{k+1} and p_k , for every block k in our sample, i.e.,

$$\text{RMSU}(p) = \sqrt{\frac{1}{T} \sum_{k=1}^T (p_k - p_{k-1})^2}.$$

Unit: gwei.

The first column shows that the EIP-1559 policy overshoots the fifteen million gas target by 1% on average in our sample while the LINDY policies are both extremely accurate, overshooting by fraction of a percent. Using column two, we can compute the *RMSD* which shows that while the average deviation for the EIP-1559 policy is about 5.75M the LINDY policies deviate on average by about 5.4M, a *6% improvement*.

Moreover, EIP-1559 results in *2.65 times more near-full blocks compared to the LINDY policies*, suggesting that the performance improvement at times of high demand might be higher than the average. Finally, note that adding a bit of regularization yields a policy, LINDY(λ), that is much smoother on average than EIP-1559 with an average update of only 1.1 gwei, a *28% improvement in the RMSU*. In what follows we analyze these properties with additional analyses on the performance of LINDY(λ) versus the benchmark.

Figure 3 reports two histograms to visualize the observed and simulated values for gas used (left) and fee updates (right) for all 99,547 blocks in our sample, comparing the EIP-1559 with the LINDY(λ) policy.

	Bias(y, t)	SD(y)	RMSD(y, t)	$\phi_{0.95}(y)$	RMSU(p)
LINDY(0)	102 (99.9%)	5,404,967 (5.9%)	5,404,967 (6.0%)	2.0% (62%)	1.9 (-35.7%)
LINDY(λ)	3,932 (97.2%)	5,404,474 (5.9%)	5,404,476 (6.0%)	2.0% (62%)	1.1 (21.4%)
EIP-1559	140,495	5,745,604	5,747,321	5.3%	1.4

Table 2: Performance metrics for the LINDY policies and EIP-1559, the regularized policy has $\lambda = 10^{-7}$. In parenthesis below each metric the percentage improvement of the respective LINDY policy versus the EIP-1559 benchmark. The unit of the first three columns is Ethereum *gas* and the unit of the last column is *gwei*s.

Note that we plot the fee update $u_k = p_{k+1} - p_k$ instead of the absolute fee. The plot on the left shows how our policy is able to achieve a distribution of gas used that is much less concentrated around the limit of 30M gas and closer to the gas target in general. At the same time, our policy shows lower variability in fees in aggregate, the distribution of updates is concentrated around the value of 0 in the right plot.




Figure 3: Histogram of gas used (left) and update size (right); comparing LINDY(λ) policy to EIP-1559.

Finally, we look at how our policy performs in different market regimes. In particular, we identify intervals of blocks that are in a *demand spike* and blocks that are in a *stable demand* regime. We use a heuristic that is solely based on the observed gas used. We compute a symmetric moving average of 25 blocks (roughly 5 minutes) for every block in our sample, then we classify as demand spikes the blocks that have a moving average of more than 20 million gas used (or $2/3$ of the block limit) and stable demand the blocks that have a moving average between 13.5 and 16.5 million gas used (or within $1/10$ of the block target).

Figure 4 shows an example of a demand spike and an example of stable and slightly increasing demand. We can appreciate how the LINDY(λ) policy is more responsive to demand spikes, with fees surging more quickly than EIP-1559 and also going down more quickly after the demand spike (left). Also, computing the RMSD over demand spike blocks only shows that the optimal policy performs significantly better than EIP-1559, with an improvement of almost 12%. At the same time, our policy is less variable and less reactive to noise during the many periods of stable demand (right).

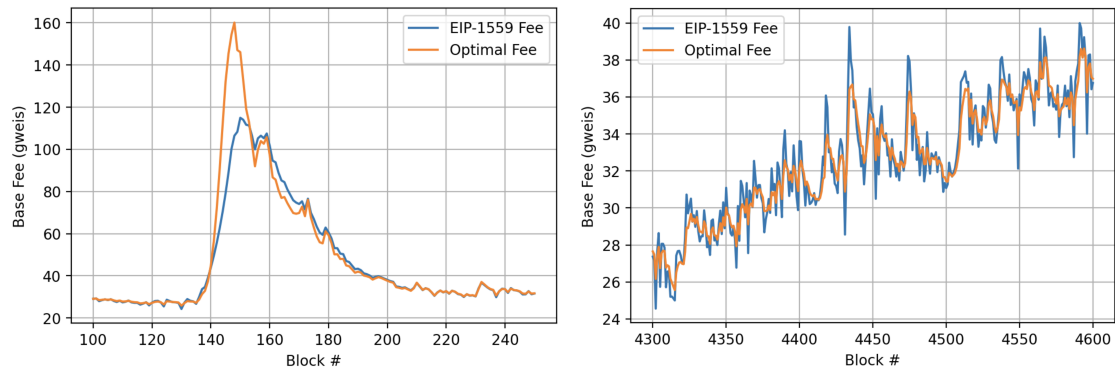


Figure 4: Base fee for EIP-1559 vs LINDY(λ) policy during a demand spike (left, blocks 120 to 200) and a stable period with increasing demand (right, blocks 4300 to 4600).

Bibliography

- Vitalik Buterin. Blockchain resource pricing. <https://ethresear.ch/uploads/default/original/2X/1/197884012ada193318b67c4b777441e4a1830f49.pdf>, August 2018.
- Vitalik Buterin. Multidimensional eip-1559. <https://ethresear.ch/t/multidimensional-eip-1559/11651>, 2022.
- Davide Crapis. Eip-4844 fee market analysis. <https://ethresear.ch/t/eip-4844-fee-market-analysis/15078>, March 2023.
- Theo Diamandis, Alex Evans, Tarun Chitra, and Guillermo Angeris. Dynamic pricing for non-fungible resources: Designing multidimensional blockchain fee markets, 2022. URL <https://arxiv.org/abs/2208.07919>.
- Matheus V. X. Ferreira, Daniel J. Moroz, David C. Parkes, and Mitchell Stern. Dynamic posted-price mechanisms for the blockchain transaction-fee market. Proceedings of the 3rd ACM Conference on Advances in Financial Technologies, 2021. URL <https://doi.org/10.1145/2F3479722.3480991>.
- Nicolae Gârleanu and Lasse Heje Pedersen. Dynamic trading with predictable returns and transaction costs. <https://onlinelibrary.wiley.com/doi/epdf/10.1111/jofi.12080>, 2013.
- Stefanos Leonardos, Daniël Reijbergen, Barnabé Monnot, and Georgios Piliouras. Optimality despite chaos in fee markets, 2022. URL <https://arxiv.org/abs/2212.07175>.
- Daniël Reijbergen, Shyam Sridhar, Barnabé Monnot, Stefanos Leonardos, Stratis Skoulakis, and Georgios Piliouras. Transaction fees on a honeymoon: Ethereum’s eip-1559 one month later, 2022. URL <https://arxiv.org/abs/2110.04753>.
- Tim Roughgarden. Transaction fee mechanism design for the ethereum blockchain: An economic analysis of eip-1559. <https://timroughgarden.org/papers/eip1559.pdf>, 2020.

A. Proofs

Proof of Theorem 1. We consider the objective function as a finite sum

$$\begin{aligned} J_K &\triangleq \mathbf{E} \left\{ \sum_{k=1}^K \|y_k - t\|_2^2 \right\} \\ &= \sum_{k=1}^K \mathbf{E} \left\{ \mathbf{E} \left[\|d_k + B_k p_k + \epsilon_k^y - t\|_2^2 | I_k \right] \right\} \end{aligned}$$

to calculate the optimal policy $p_{k+1,K}^*$. Note that the objective J_K separates into a sum of quadratic terms. Equating the first-order conditions for p_{k+1} to 0 for every k yields the control policy

$$p_{k+1,K}^* = \left[\mathbf{E} \left(B_{k+1}^\top B_{k+1} | I_k \right) \right]^{-1} \mathbf{E} \left(B_{k+1}^\top (t - d_{k+1}) | I_k \right).$$

Then we let K goes to infinity, and we can get

$$\begin{aligned} p_{k+1}^* &= \lim_{K \rightarrow \infty} p_{k+1,K}^* \\ &= \left[\mathbf{E} \left(B_{k+1}^\top B_{k+1} | I_k \right) \right]^{-1} \mathbf{E} \left(B_{k+1}^\top (t - d_{k+1}) | I_k \right). \end{aligned}$$

□

Proof of Theorem 2. We consider the objective function as a finite sum

$$J_K = \min_p \mathbf{E} \left\{ \sum_{k=1}^K \left(\|y_k - t\|_2^2 + \lambda \|p_k - p_{k-1}\|_2^2 \right) \right\}$$

to calculate the optimal policy $p_{k+1,K}^*$. Using the techniques in dynamic programming, we further suppose

$$J_{k,K} \triangleq \min_p \mathbb{E} \left\{ \sum_{s=k+1}^K \left(\|y_s - t\|_2^2 + \lambda \|p_s - p_{s-1}\|_2^2 \right) \middle| I_k \right\}.$$

Lemma 1. We can use induction to show that there exists matrices $Q_{k,K}, R_{k,K} \in \mathbb{R}^{n \times n}$ and vector $\tau_{k,K} \in \mathbb{R}^n$ such that

$$J_{k,K} = p_k^\top Q_{k,K} p_k - 2\tau_{k,K}^\top p_k + 2(a_{k+1}^d)^\top R_{k,K} p_k + H_{k,K},$$

where $a_{k+1}^d \triangleq \mathbb{E}(d_{k+1} | I_k)$ and $H_{k,K}$ is independent with the choice of decision variables p . Moreover, $Q_{k,K}, R_{k,K}, \tau_{k,K}$ satisfy

$$\begin{aligned} Q_{k-1,K} &= \lambda I - \lambda^2 (\lambda I + Q_{k,K} + B_k^\top B_k)^{-1}, \\ R_{k-1,K}^\top &= \lambda (\lambda I + Q_{k,K} + B_k^\top B_k)^{-1} (B_k^\top + R_{k,K}^\top A^d), \\ \tau_{k-1,K} &= \lambda (\lambda I + Q_{k,K} + B_k^\top B_k)^{-1} \left(B_k^\top t + \tau_{k,K} - R_{k,K}^\top (I - A^d) \mu^d \right). \end{aligned}$$

Proof. For $k = K$, $J_{K,K} = 0$, we can let $Q_{K,K} = 0$, $R_{K,K} = 0$, and $\tau_{K,K} = 0$. Suppose for $k = s$, we have

$$J_{s,K} = p_s^\top Q_{s,K} p_s - 2\tau_{s,K}^\top p_s + 2(a_{s+1}^d)^\top R_{s,K} p_s + H_{s,K}.$$

For $k = s - 1$, we can first get the optimal policy

$$\begin{aligned} p_{s,K}^* &= \operatorname{argmin}_{p_s} \mathbf{E} \left\{ \|y_s - t\|_2^2 + \lambda \|p_s - p_{s-1}\|_2^2 + J_{s,K} | I_{s-1} \right\} \\ &= \operatorname{argmin}_{p_s} \mathbf{E} \left\{ \|d_s + B_s p_s + \epsilon_s^y - t\|_2^2 + \lambda \|p_s - p_{s-1}\|_2^2 + p_s^\top Q_{s,K} p_s - 2\tau_{s,K}^\top p_s + 2(a_{s+1}^d)^\top R_{s,K} p_s | I_{s-1} \right\}, \end{aligned}$$

by the expression of $J_{s,K}$. It is a quadratic function with respect to the decision variable p_s . We can get the optimal policy

$$p_{s,K}^* = -(\lambda I + Q_{s,K} + B_s^\top B_s)^{-1} \left(B_s^\top a_s^d - B_s^\top t - \lambda p_{s-1} - \tau_{s,K} + R_{s,K}^\top [(I - A^d)\mu^d + A^d a_s^d] \right),$$

by equating its first-order condition to 0. After plugging in and comparing the coefficient, the optimal value

$$J_{s-1,K} = p_{s-1}^\top Q_{s-1,K} p_{s-1} - 2\tau_{s-1,K}^\top p_{s-1} + 2(a_s^d)^\top R_{s-1,K} p_{s-1} + H_{s-1,K},$$

where

$$\begin{aligned} Q_{s-1,K} &= \lambda I - \lambda^2 (\lambda I + Q_{s,K} + B_s^\top B_s)^{-1}, \\ R_{s-1,K}^\top &= \lambda (\lambda I + Q_{s,K} + B_s^\top B_s)^{-1} (B_s^\top + R_{s,K}^\top A^d), \\ \tau_{s-1,K} &= \lambda (\lambda I + Q_{s,K} + B_s^\top B_s)^{-1} \left(B_s^\top t + \tau_{s,K} - R_{s,K}^\top (I - A^d) \mu^d \right). \end{aligned}$$

□

Suppose $Q_k = \lim_{K \rightarrow \infty} Q_k^K$, $R_k = \lim_{K \rightarrow \infty} R_k^K$, $\tau_k = \lim_{K \rightarrow \infty} \tau_k^K$, $H_k = \lim_{K \rightarrow \infty} (H_{k,K} - H_{0,K})$, and $\tilde{J}_k = p_k^\top Q_k p_k - 2\tau_k^\top p_k + 2(a_{k+1}^d)^\top R_k p_k + H_k$. Taking limit in Lemma 1, we have

$$\begin{aligned} p_{k+1}^* &= \operatorname{argmin}_{p_{k+1}} \mathbf{E} \left\{ \|y_{k+1} - t\|_2^2 + \lambda \|p_{k+1} - p_k\|_2^2 + \tilde{J}_{k+1} | I_k \right\}, \\ \tilde{J}_k &= \min_{p_{k+1}} \mathbf{E} \left\{ \|y_{k+1} - t\|_2^2 + \lambda \|p_{k+1} - p_k\|_2^2 + \tilde{J}_{k+1} | I_k \right\}. \end{aligned} \quad (12)$$

We take the partial derivative with respect to p_k on both sides of the equation 12. Then we have

$$2Q_k p_k - 2\tau_k + 2R_k^\top a_{k+1}^d = -2\lambda(p_{k+1}^* - p_k),$$

which means that

$$p_{k+1}^* = \left(I - \frac{Q_k}{\lambda} \right) p_k + \frac{Q_k}{\lambda} \operatorname{aim}_k,$$

where $\operatorname{aim}_k = Q_k^{-1}(\tau_k - R_k^\top a_{k+1}^d)$.

We can also get

$$\begin{aligned} Q_k &= \lambda I - \lambda^2 (\lambda I + Q_{k+1} + B_{k+1}^\top B_{k+1})^{-1}, \\ R_k^\top &= \lambda (\lambda I + Q_{k+1} + B_{k+1}^\top B_{k+1})^{-1} (B_{k+1}^\top + R_{k+1}^\top A^d), \\ \tau_k &= \lambda (\lambda I + Q_{k+1} + B_{k+1}^\top B_{k+1})^{-1} (B_{k+1}^\top t + \tau_{k+1} - R_{k+1}^\top (I - A^d) \mu^d), \end{aligned}$$

after taking limit. By some algebraic manipulations, we can get

$$Q_k^{-1} = \lambda^{-1} (Q_{k+1} + B_{k+1}^\top B_{k+1})^{-1} (\lambda I + Q_{k+1} + B_{k+1}^\top B_{k+1}),$$

$$\tau_k - R_k^\top a_{k+1}^d = \lambda (\lambda I + Q_{k+1} + B_{k+1}^\top B_{k+1})^{-1} (B_{k+1}^\top (t - a_{k+1}^d) + \tau_{k+1} - R_{k+1}^\top ((I - A^d)\mu^d + A^d a_{k+1}^d)).$$

Thus, we can get

$$\begin{aligned} \operatorname{aim}_k &= Q_k^{-1} (\tau_k - R_k^\top a_{k+1}^d) \\ &= (Q_{k+1} + B_{k+1}^\top B_{k+1})^{-1} (B_{k+1}^\top (t - a_{k+1}^d) + \tau_{k+1} - R_{k+1}^\top ((I - A^d)\mu^d + A^d a_{k+1}^d)) \\ &= (B_{k+1}^\top B_{k+1} + Q_{k+1})^{-1} (B_{k+1}^\top B_{k+1} \bar{p}_{k+1} + Q_{k+1} \mathbf{E}(\operatorname{aim}_{k+1} | I_k)), \end{aligned}$$

where $\bar{p}_{k+1} \triangleq B_{k+1}^{-1} (t - a_{k+1}^d)$ satisfying $\mathbf{E}(y_{k+1} | I_k, p_{k+1} = \bar{p}_{k+1}) = t$. is the market clearing price at block $k+1$. We can obtain the result stated in Theorem 2 by employing the same technique to represent α_{k+1} and iterating this process repeatedly. □

Proof of Theorem 3. By Theorem 1, the optimal policy

$$\begin{aligned} p_{k+1}^* &= [\mathbf{E}(B_{k+1}^\top B_{k+1} | I_k)]^{-1} \mathbf{E}(B_{k+1}^\top (t - d_{k+1}) | I_k) \\ &= [\mathbf{E}(V \text{diag}(\delta_{k+1}) U^\top U \text{diag}(\delta_{k+1}) V^\top | I_k)]^{-1} \mathbf{E}(V \text{diag}(\delta_{k+1}) U^\top (U\tilde{t} - U\tilde{d}_{k+1}) | I_k) \\ &= V [\mathbf{E}(\text{diag}(\delta_{k+1})^2 | I_k)]^{-1} \mathbf{E}(\text{diag}(\delta_{k+1})(\tilde{t} - \tilde{d}_{k+1}) | I_k). \end{aligned}$$

By the eigenprice $\tilde{p}_{k+1}^* = V^\top p_{k+1}^*$, we can obtain

$$\tilde{p}_{k+1}^* = [\mathbf{E}(\text{diag}(\delta_{k+1})^2 | I_k)]^{-1} \mathbf{E}(\text{diag}(\delta_{k+1})(\tilde{t} - \tilde{d}_{k+1}) | I_k),$$

which is equivalent to Theorem 3. \square

B. Empirical Models

We introduce the Expectation-Maximization (EM) algorithm to estimate the parameters in our model using real data. The EM algorithm is a two-step iterative method designed for finding maximum likelihood estimates of parameters in probabilistic models with incomplete data. It is especially useful in situations where direct optimization of the likelihood is computationally challenging or not feasible.

In our study, the EM algorithm is applicable to both our baseline model and eigenresources model. For clarity and brevity, we will take the baseline model as an illustrative example to introduce how the EM algorithm works.

Model review. Demand dynamics

$$d_{k+1} = (I - A^d)\mu^d + A^d d_k + \epsilon_k^d.$$

Price-sensitivity matrix

$$\text{vec}(B_{k+1}) = (I - A^B)\mu^B + A^B \text{vec}(B_k) + \epsilon_k^B.$$

Observed demand

$$y_k = d_k + B_k p_k + \epsilon_k^y.$$

Hidden state and observed state. The hidden state in our model is $x_k^\top = [d_k^\top \text{vec}(B_k)^\top]$. It evolves according to

$$x_{k+1} = (I - A^x)\mu^x + A^x x_k + \epsilon_k^x, \quad \epsilon_k^x \sim \mathcal{N}(0, W^x).$$

The observed state in our model is y_k . It satisfies

$$y_k = C_k x_k + \epsilon_k^y, \quad \epsilon_k^y \sim \mathcal{N}(0, W^y).$$

Parameters. The parameter that need to be estimated in our model is

$$\theta \triangleq (\mu^x, A^x, W^x, W^y, a_0, S_0),$$

where x_0 and S_0 represent the normal-mean and normal-variance for the initial hidden state respectively.

EM algorithm.

1. **Initialization:** Choose initial values θ_0 for the parameters.
2. **E-Step (Expectation):** Given the current parameters θ_t , compute the expectation of the log likelihood

$$l(\theta | \theta_t) \triangleq \mathbf{E} \{ \log(x_{0:T}, y_{0:T}; \theta) | x_{0:T}, y_{0:T} \sim p(\theta_t), I_T \}$$

with respect to the hidden state and observed state.

3. **M-Step (Maximization):** Maximize the expected log likelihood found in the E-step with respect to the parameter θ , and update the parameter to this new value, i.e., $\theta_{t+1} = \arg\max_{\theta} l(\theta | \theta_t)$.

Repeat the E-Step and M-Step until the algorithm converges.