Blockchain Price *vs.* Quantity Controls

Abdoulaye Ndiaye[0009−0000−7466−6444]

New York University, New York NY 10012, USA
andiaye@stern.nyu.edu
https://sites.google.com/view/abdoulaye-ndiaye

Abstract. Blockchains, the technology underlying cryptocurrencies, face large fluctuations in user demand. These fluctuations necessitate effective transaction fee mechanisms to manage service allocation. This paper models the blockchain designer's choice between price control and quantity control. We first derive an analytical expression for the advantage of a minimum fee over a rigid block size limit. When demand uncertainty is high, price controls are preferred. A large price elasticity of demand for block space amplifies this advantage. Using these insights, we provide novel results on the dynamics of optimal transaction fee mechanisms within a class of simple mechanisms, including that of the Ethereum blockchain. Further, we study optimal mechanisms that are resistant to complete value extraction by monopolistic blockchain service providers, known as validators. Using Ethereum data, we estimate the parameters of the optimal transaction fee mechanism and offer recommendations for blockchain designers.

Keywords: Blockchain, Price Controls, Quantity Controls, Transaction Fee Mechanisms, Cryptocurrencies, Ethereum, Monopolistic Competition

1 Introduction

Over the past decade, the blockchain industry has experienced remarkable growth, with the combined market capitalization of all cryptocurrencies reaching a peak of $3 trillion in 2021. The technology at the heart of these cryptocurrencies — blockchains — has emerged as a potent tool for facilitating efficient peer-to-peer transactions. However, the volatile nature of the market has led to significant fluctuations in demand for blockchain technology services, mirroring the instability of cryptocurrency prices. As a result, the market capitalization has since contracted to $1 trillion. In response to these demand uncertainties, blockchain designers are increasingly employing Transaction Fee Mechanisms (TFMs) to allocate the finite block space effectively. The implications of these mechanisms for the efficiency of blockchain systems underscore the need for rigorous economic analysis.

In this paper, we model a blockchain—like Bitcoin or Ethereum—as a distributed computing network where users submit transactions for inclusion by validators. Transactions, representing data that modifies the network's state (such as account balance transfers), are submitted by users with a bid to a publicly observable pool—known as the mempool—that indicates their willingness to pay for their transactions to be processed. Validators, using their limited resources, select a subset of transactions from the mempool to form a block. Each block, comprising an ordered sequence of transactions and a reference to the previous block, can be appended to the blockchain in every period. However, technological limitations enforce a maximum block size, thereby constraining the supply of block space.

On the supply side, validators — who provide block space inclusion services — face uncertainty in their marginal costs. On the demand side, atomistic users arrive at random and submit their transactions along with their willingness to pay for these transactions to be included in the next block. Given the paper's focus on the aggregate properties of the block space market, we operate under the assumption that users truthfully bid their valuations. Dominant-strategy incentive compatibility for all Transaction Fee Mechanisms (TFMs) considered in this paper is substantiated by the game-theoretic proofs provided by [15] and [16].

This process forms the micro-foundations of an aggregate user demand curve, which we study under both stable and unstable market conditions. In the face of these uncertainties, a blockchain designer opts for either a base fee—an *ex-ante* price control—or a block size limit—an *ex-ante* quantity control. This choice aims to maximize a welfare function that can encompass concerns for social welfare, monopolistic validator profits, or an externally set technological target.

The resolution of the blockchain designer's instrument choice casts [17]'s "prices *vs.* quantities" idea in a fresh light. Specifically, when user demand uncertainty is higher than uncertainty in the validators' marginal costs, price controls prove more effective than quantity controls. Within my model, we derive an analytical expression for the relative advantage of price controls over quantity controls under demand uncertainty and uncertain marginal costs.

Firstly, demand uncertainty favors price controls, as block size adjustments provide the flexibility necessary to accommodate demand fluctuations. Secondly, a negative correlation between marginal costs and demand also favors price controls. In such circumstances, quantity changes permit the production of larger blocks when marginal costs are low, thereby improving efficiency. Thirdly, the price elasticity of demand—the proportional change in block space demand in response to a proportional change in price for the marginal user seeking to include their transaction in the next block—further amplifies the relative advantage of price controls over quantity controls. Lastly, in the absence of uncertainty, the blockchain designer remains indifferent between price and quantity controls.

Building on these insights, we delve into the dynamics of optimal TFMs that maintain the simplicity of Ethereum's TFM, the most widely utilized public blockchain. Such mechanisms iteratively update base fees using information from the previous base fee and the deviation from the target block size. We refer to these as *Adaptive to a Deterministic Target* Transaction Fee Mechanisms (ADT-TFMs).

We establish that the optimal ADT-TFM is locally exponential for any user valuation distribution. Moreover, we prove that its adjustment parameter equates to the inverse price elasticity of demand at the target block size. Specifically, if the aggregate user demand curve is isoelastic, the optimal ADT-TFM is globally exponential with a fixed adjustment parameter. This provides theoretical justifications for a proposed upgrade to Ethereum's TFM from its current linear configuration [4].

Subsequently, we proceed to illustrate these results numerically. Estimating a user demand curve without random variation in the supply curve presents identification challenges, as it's difficult to distinguish between a movement along a fixed demand curve and a shift of the demand curve. To address this, we analytically show in my atomistic user model that the tail ratio of user valuations can serve as a proxy for the price elasticity of demand.

Using a random sample of Ethereum blockchain data, we show that the optimal adjustment parameter of the ADT-TFM is approximately 8%. At present, this ratio is fixed at an ad-hoc level of 12.5%, a rate capable of doubling prices within at least eight blocks.

On September 15, 2022, the Ethereum blockchain underwent an upgrade that, among other changes, resulted in a deterministic production rate of new blocks.[1] we observe that the adjustment rate was overshooting more prior to the proof-of-stake upgrade than following it, a finding that aligns with [12].

In addition, we investigate the optimal block size target for a monopolistic validator and offer tight bounds on its size relative to the block size limit. These bounds provide valuable insights for studies of TFMs involving monopolistic validators [13,10], as well as TFMs designed to prevent validators from monopolizing all surplus [2].

Numerically, we find that if the block size limit is set to accommodate all user transactions under average demand conditions and to resist total surplus extraction by validators, then the block size target should be less than $e^{-}1 \approx$ 37% of the block size limit. This bound is tight, as there exist user valuation distributions that lead to a ratio of the optimal monopolistic target to the block size limit that is arbitrarily close to this figure.

The block size target for the Ethereum blockchain is not set to align with the monopolist validator's target block size; instead, it is established at 50% of the block size limit. Nonetheless, there are widespread concerns regarding the potential for validators and other users to exploit their power to extract value through censoring, swapping, and front-running of mempool transactions, a practice colloquially referred to as *"maximal extractable value"* or MEV [6].

The appropriate interpretation of these findings is that any ADT-TFM with a target block size exceeding 37% of the block size, which aligns with user demand under average demand conditions, would encourage validators to include

[1] This upgrade, known as the Merge and finalized at block number 15537393, transitioned the Ethereum network from proof-of-work to proof-of-stake. The specific details of the mechanism through which blocks are produced are not relevant to this analysis.

their own value-extracting transactions. This action would effectively reduce the supply available to users and allow validators to extract more surplus.

Lastly, it is important to consider that users (and validators) might value block space (and marginal costs) in USD or real-world terms as opposed to the native currency of the blockchain. To accommodate this, we expand the model to introduce uncertainty in the price of the cryptocurrency in USD or in real terms. Notably, price controls are restricted to be expressed in units of the native currency. we show that volatility in the cryptocurrency price reduces the advantage of price controls over quantity controls.

Blockchain designers require methods that are simple, robust, and principled for updating the parameters of their TFMs. Inspired by several robustness checks and quantitative explorations, we suggest potential ways to construct non-deterministic adaptive TFMs. These could better align with demand under uncertain conditions and resist commonly observed manipulations.

1.1 Literature Review

This paper contributes to the literature on price versus quantity controls, a domain pioneered by [17]. The issue of choosing a supply function under uncertainty has been explored in [9]. My approach aligns closely with [14] and [8], who scrutinize the choice between price and quantity from a firm's perspective and its macroeconomic implications. While my work draws inspiration from [17], it diverges in that it contemplates the planner's (blockchain designer's) problem with a variety of goals, including purely technical objectives seen in practice, such as block size targets. The conclusions of this analysis are then applied to the design of TFMs.

The literature on TFMs from a mechanism design perspective is growing. Notably, [1] examines mechanisms that are immune to designer manipulations—referred to as "credible mechanisms"—and demonstrates that the well-known second price auction doesn't meet this criterion. In the blockchain context, [16] applies this credibility condition to TFMs and establishes that EIP-1559, Ethereum's TFM, and its variations are incentive-compatible for users and adhere to a form of myopic credibility for validators. These findings are further consolidated by [5]. These papers provide game-theoretic foundations that guarantee that the ADT-TFMs we study are incentive-compatible. [7] explores an alternative aspect of TFMs by investigating posted price mechanisms. This paper takes a complementary approach to this strand of literature. Moreover, it delves into the dynamics of TFMs, offering insights into their updating rules.

This paper advances the literature that examines the block space market from a macroscopic viewpoint. The concepts formalized in Section 2 of this paper build upon and extend the work presented by [3]. Motivated by [10], [13] investigates the monopolistic market for unlimited block space, a contrast to this paper which views the block size limit as an exogenous technological constraint.

Lastly, this paper contributes to studies of the dynamics of EIP-1559, the Ethereum blockchain's TFM. [11] delves into the behavior of the dynamic system resulting from the TFM, and [12] uncovers numerous empirical properties that this paper provides a theoretical explanation for.

Outline: The paper is organized as follows: Section 2 introduces the model and provides an overview of blockchains and the economic principles of price and quantity controls. Section 2.1 is tailored for economists or readers unfamiliar with the workings of most blockchains. Section 3 examines the blockchain designer's instrument choices and identifies properties of optimal ADT-TFMs. Section 4 presents the data and numerical analysis. Section 5 extends the baseline model to accommodate uncertainty in cryptocurrency prices and includes some numerical robustness exercises. Lastly, Section 6 concludes the paper.

2 Model

2.1 Background

A *blockchain*, such as Bitcoin or Ethereum, is modeled as a distributed computer network where users submit *transactions* to be included in a chain of *blocks* by *validators*. The blockchain maintains a record of the network's state, such as account balances. A transaction t represents arbitrary data sent over the network to alter its state—for instance, to transfer a balance. Users submit transactions to a publicly observed pool of outstanding transactions (*mempool*), with a bid b_t, signifying their willingness to pay for transaction processing. Monopolistic validators, using quantities x_i of a finite number of resources $i \in [\![1, N]\!]$ (e.g., computation, bandwidth), select a subset of transactions from the mempool to form a block. A block of size q is an ordered sequence of transactions and a reference to the previous block. Validators add a block to the blockchain by a consensus mechanism (such as proof of work or proof of stake), a process irrelevant to this analysis. Technological constraints impose a maximum block size $q^{\max}$, thus limiting the supply of block space.[2]

Validators: Validators use a bundle of computational resources $x \in \mathbb{R}_+^N$ priced at per-unit resource rates $p_x \in \mathbb{R}_{++}^N$ to produce a block of size $q \leq q^{\max}$, as given by

$$q = \sum_{i=1}^{N} p_{x_i} x_i \tag{1}$$

Validators incur a technological cost of $C(q) = c(x_1, \ldots, x_N)$ when producing a block. These costs encompass validation operation costs and other costs associated with accessing and modifying the blockchain's state. From the blockchain designer's viewpoint, costs might also include delays in block propagation due to large blocks and other societal costs. These costs, denoted by $C(q; \eta)$, are subject to uncertainty represented by the distribution η.

[2] See [3] for the case of the Ethereum blockchain.

Users: User transactions populate the mempool according to a stochastic process. We posit that users, denoted by $j \in [0,1]$, are atomistic, with arrivals between two consecutive blocks, B_t, B_{t+1}, being independently and identically distributed according to a Poisson process X with a parameter of $\lambda q^{\max}$, where $\lambda \in \mathbb{R}_{++}$. For simplicity, we assume users leave the pool if their transaction is not included in the next block, only to return according to the arrival process.[3]

Each user j has a valuation v_j drawn from a common distribution f with a cumulative distribution function F, which is continuous and increasing.

Establishing a Demand Curve: Given that the transaction fee mechanisms under consideration are dominant-strategy incentive-compatible, as demonstrated by [16], it is reasonable to assume that users bid their true valuations, i.e., $b_t = v_j$. Consequently, for a given minimum bid for transaction inclusion, denoted by p, the number of users willing to pay the bid is $\lambda q^{\max}\bar{F}(p)$, where $\bar{F}(p) = 1 - F(p)$. The following lemma shows that this model serves as the microfoundation of an intuitive demand curve for block space, thereby linking demand parameters to model primitives.

Lemma 1. *The aggregate demand for block space can be represented as*

$$p = \left(\bar{F}\right)^{-1}\left(\frac{q}{\Psi}\right) \tag{2}$$

where the price elasticity of demand for the marginal user equals the tail ratio

$$\frac{pf(p)}{1-F(p)} \tag{3}$$

.

Specifically, when F is a Pareto distribution with scale p_m and shape α, the aggregate demand for block space is given by

$$\frac{p}{p_m} = \left(\frac{q}{\Psi}\right)^{-\frac{1}{\varepsilon}} \tag{4}$$

Here $p \in \mathbb{R}_+$ is the market price, $\Psi \equiv \lambda q^{\max}$ is a demand shifter, and $\varepsilon = \alpha$ is the price elasticity of demand for block space.

Proof. Refer to Appendix A.1 for the proof.

In Section 3, we will assume that the user valuation distribution is Pareto unless otherwise stated. The assumption of a Pareto distribution is not restrictive. For a general user valuation distribution, one can calculate the price elasticity of demand at all points from the tail ratio of the distribution. Using data from the Ethereum blockchain, we fit a Pareto distribution to the transaction in Section

[3] [11] confirms that this assumption does not significantly impact the dynamics of transaction fees, which are the focus of our analysis. In their research, [13] accounts for residual demand in the mempool and finds similar dynamics of transaction fees as [12].

4. This demand curve will prove useful when considering uncertainty in the user arrival rate λ, leading to uncertainty in the demand shifter Ψ. When $\lambda > 1$, we encounter a high-demand scenario where not all transactions can be included in the next block, whereas $\lambda < 1$ reflects a low-demand scenario where the block is not filled to capacity. Given this context and considering the uncertainties in both the cost $C(q;\eta)$ and demand Ψ, we explore the conditions under which a blockchain protocol would find it beneficial to introduce price or quantity controls.

3 The Blockchain Designer Problem

3.1 Price Controls and The Elasticity of Demand

The analysis conducted by Weitzman (1974) hinges on a Taylor approximation of the private benefit function and marginal cost. Consequently, it provides a local result around an equilibrium price-quantity pair, where shocks to demand and validator costs can occur. For analytical tractability, we approach the blockchain designer's problem within the framework of a demand-specified model. This model is governed by equations (1) and (4), featuring uncertainty in the demand shifter, Ψ, and in the cost to validators, indicated by η.

Blockchain Designer Objective: Given a price p per unit of block space and a fixed block reward R, the blockchain protocol profits are

$$\Pi = R + \mathbb{E}\left[pq - C(q;\eta)\right]. \tag{5}$$

If the blockchain designer were to maximize the profits of monopolist validators, they would then maximize equation (5). However, protocol revenue can be, in part, diverted by the protocol treasury, burned, or rebated to users for incentivization purposes beyond the scope of this paper.[4] Therefore, the blockchain designer's objective balances social welfare and technological considerations while ensuring validators have enough profit to be willing to provide their validator services.

We consider the following blockchain designer objective function that captures these considerations:

$$\mathcal{V} = \mathbb{E}\left[S(q)\right]^{1-\beta}\,\mathbb{E}\left[pq - C(q;\eta)\right]^{\beta}. \tag{6}$$

In this objective, the parameter $\beta \in [0;1]$ captures the bargaining power of validators. $\beta = 1$ means that the blockchain designer maximizes validator profits, while $\beta = 0$ means that the designer only optimizes for other common considerations captured by the function $S(q)$. Here are a few examples:

Example 1. (Social Welfare) Denote u, the utility of the representative user of the blockchain. Then:

$$S(q) = u(q) \tag{7}$$

[4] See [16].

captures a concern for the utility of users, which may differ and conflict with validator profits.

Example 2. (Technological Block Size Targets) Suppose the designer has a target block size, denoted by q^{target}, which they aim to achieve. There might be technical constraints that necessitate reaching this specific block size target. Consequently, the designer might only be satisfied when $q = q$, which can be represented as

$$S(q) = \delta\{q - q^{target}\} \tag{8}$$

where δ symbolizes the Dirac function, which equals zero everywhere except at the origin.

However, it might also be plausible for the designer to allow some degree of deviation from the target block size. In such a case, we can represent this as $S(q) = \ell\{q - q^{target}\}$ where ℓ stands for some loss function.

As a specific instance, consider a loss function that behaves as the narrowing limit of a normal distribution centered at zero, which provides an approximation to the Dirac function:

$$S(q) = \phi\left(\frac{q - q^{target}}{a}\right) \tag{9}$$

Here, ϕ denotes the density of the standard normal distribution, and a is a parameter that brings the function closer to the Dirac function as it approaches zero.

Example 3. (Validator Profits) Ignoring the block reward R as it remains constant, we can consider the scenario where the designer aims to maximize the monopolistic validator's profit. This can be represented as:

$$S(q) = 1 \tag{10}$$

In this case, given a price p per unit of block space, the block size q adjusts to optimize profit. Furthermore, when the block size limit $q^{\max}$ becomes binding, the equilibrium price adapts to meet user demand.

Assuming users and smart contract writers optimize resource usage, the cost to the validator becomes

$$C(q; p_x, \eta) = \min_{x \in \mathbb{R}_+^N} c(x_1, \ldots, x_N; \eta) \tag{11}$$

subject to (1). The bundle x can be interpreted as the various resources that constitute a user transaction, such as bandwidth, computational operations, etc. Let's consider that c is homogeneous of degree 1 in x.[5] We then have:

$$C(q; p_x, \eta) = \Gamma(\eta, p_x) q \tag{13}$$

[5] A typical example is the "constant-scale" cost function

$$c(x_1, \ldots, x_N; \eta) = \eta \prod_{i=1}^{N} x_i^{\varepsilon_i} \quad \text{such that} \quad \sum_{i=1}^{N} \varepsilon_i = 1 \tag{12}$$

The expression for the marginal cost Γ is derived in Appendix A.2.

Price Controls: Assuming the block space demand follows the isoelastic demand curve derived in (4), with a price elasticity of demand $\varepsilon > 1$and that the block size limit is not binding, the equilibrium block size lies on the demand curve, i.e., $q = \Psi\left(\frac{p}{p_m}\right)^{-\varepsilon}$. The blockchain designer's problem of settings optimal prices then becomes:

$$\mathcal{V}^p = \max_{p\in\mathbb{R}_+} \mathbb{E}\left[S\left(\Psi\left(\frac{p}{p_m}\right)^{-\varepsilon}\right)\right]^{1-\beta}\left[(p-\Gamma)\times\Psi\left(\frac{p}{p_m}\right)^{-\varepsilon}\right]^{\beta} \tag{14}$$

Quantity Controls: If the monopolistic validator sets a block size $q \leq q^{\max}$, it sells transaction inclusion at the price that clears markets ex-post: $p = p_m\left(\frac{q}{\Psi}\right)^{-\frac{1}{\varepsilon}}$. Then the blockchain designer's value of setting the optimal block space is

$$\mathcal{V}^q = \max_{q\in\mathbb{R}_+} \mathbb{E}\left[S\left(q\right)\right]^{1-\beta}\left[\left(p_m\left(\frac{q}{\Psi}\right)^{-\frac{1}{\varepsilon}}-\Gamma\right)\times q\right]^{\beta} \tag{15}$$

The log-difference between the values of price controls and block space controls can be defined as follows:

$$\Delta^{\log} = \log\mathcal{V}^p - \log\mathcal{V}^q. \tag{16}$$

To provide insight into the choice between price and quantity, consider a blockchain designer objective that accounts for social welfare with utility $S(q) = u(q) = q^{\nu}$ for $\nu > 0$ and an arbitrary bargaining power β for validators. The following proposition establishes the relationship between the relative value of price controls, the price elasticity of demand, and other moments of the shock to demand and marginal costs.

Proposition 1. *Assume that (Ψ, η) follows a joint log-normal distribution. Then, the relative value of price controls over quantity controls is given by:*

$$\Delta^{\log} = \frac{1}{2}\left(\left(\hat{\nu} - \frac{\bar{\nu}}{\varepsilon}\right)\sigma_{\Psi}^2 - 2(\varepsilon\nu - \beta)\sigma_{\Psi,\Gamma}\right) \tag{17}$$

where $\bar{\nu} = (1-\beta)\nu + \beta$ and $\hat{\nu} = (1-\beta)\nu^2 + \beta$

Proof. Refer to Appendix A.2 for the proof.

This equation can be interpreted by setting $\beta = 1$, which gives

$$\Delta^{\log} = \frac{1}{2}\left(\varepsilon - 1\right)\left(\frac{1}{\varepsilon}\sigma_{\Psi}^2 - 2\sigma_{\Psi,\Gamma}\right) \tag{18}$$

In this case, price-setting is preferable to quantity-setting when (i) demand volatility is high, and (ii) the covariance between demand and real marginal costs is low. Uncertainty in demand favors price controls as block size adjustments can flexibly respond to demand fluctuations. Additionally, a negative correlation between marginal costs and demand favors price controls, as this allows for the production of larger blocks when marginal costs are low, enhancing efficiency. The price elasticity of demand, which dictates how quickly prices react to changes, mediates the degree to which the firm values (i) and (ii). A larger price elasticity of demand favors price controls. In general, these comparative statistics remain valid as long as demand is relatively elastic, i.e., $\varepsilon > \max\{\beta/\nu, \bar{\nu}/\hat{\nu}\}$.

Lastly, note that the advantage of price-setting in equation (17) could be empirically estimated if the elasticity of demand, validators' uncertainty about demand, and marginal costs can be measured. This insight will guide us in estimating the appropriate magnitude of the adjustment parameter in a transaction fee mechanism.

3.2 Optimal Transaction Fee Mechanisms

In this section, we define a family of simple TFMs that include EIP-1559, Ethereum's TFM. We study their dynamics, determine the shape and adjustment rate of the optimal mechanism within this family, and provide bounds on the target block size that align with the incentives of a monopolistic validator.

Definition 1. *(ADT-TFM) A transaction fee mechanism is called Adaptive to a Deterministic Target (ADT) if there exists a deterministic block size, q^{target} (the target), and a deterministic function, f (the adjustment function), such that the base fee satisfies:*

$$\frac{p_{t+1}}{p_t} = g\left(\frac{q_t - q^{target}}{q^{target}}\right) \tag{19}$$

Example 4. (EIP-1559) The base fee in EIP-1559 is ADT with linear adjustment function $g(x) = 1 + d \times x$ where the adjustment parameter is $d = \frac{1}{8}$.

Let q^* denote the optimal quantity control or the block size that a blockchain designer aims to achieve for a specific technological target, $S(q) = \delta\{q - q^{target}\}$. In this case, $q^* = q^{target}$. We consider a general demand curve, with price elasticity of demand $\varepsilon(q^{target})$ and uncertainty in demand represented by λ. In this context, a TFM is considered to be robustly optimal if, following a sudden change or shock in demand, it manages to bring the realized quantity as close as possible to the targeted level in the worst-case scenario. The following proposition determines the shape and slope of the optimal ADT-TFM.

Proposition 2. *Suppose the demand curve is log-convex, the robustly optimal ADT-TFM is an exponential function with an adjustment parameter equal to the inverse price elasticity of demand, $d = \frac{1}{\varepsilon(q^{target})}$. In other words, $g(x) = \exp(d \cdot x)$ and*

$$p_{t+1} = p_t \exp\left(\frac{1}{\varepsilon(q^{target})} \frac{q_t - q^{target}}{q^{target}}\right) \tag{20}$$

The intuition of the proof in Appendix A.3 goes as follows. Let f denote the adjustment function of the optimal TFM and $p(q_t)$ represent the price that matches demand at the block size q_t. The base fees then satisfy:

$$\ln p_{t+1} - \ln p_t = \ln g\left(\underbrace{\frac{q_t - q^{target}}{q^{target}} \frac{p(q^{target})}{p(q_t) - p(q^{target})}}_{g(q_t)} \frac{p(q_t) - p(q^{target})}{p(q^{target})}\right) \tag{21}$$

Near q^{target}, we have:

$$g(q_t) \xrightarrow[q^{target}]{} \frac{q'(p(q^{target}))p(q^{target})}{q^{target}} = \varepsilon(q^{target}) \tag{22}$$

And if p_{t+1} maintains the block size near q^{target}, then

$$\ln p_{t+1} - \ln p_t \sim \frac{p(q_t) - p(q^{target})}{p(q^{target})} \tag{23}$$

Let x denote this price growth. From (21), we get

$$x = \ln\left(g\left(\varepsilon(q^{target}) \cdot x\right)\right) + o(x) \tag{24}$$

for all x in the neighborhood of zero.

Solving this functional equality yields $g(x) = \exp\left(\frac{x}{\varepsilon(q^{target})}\right)$ in the neighborhood of zero. The proof extends this argument with uncertainty in demand and shows that this function is optimal for the worst-case demand scenario in demand fluctuations.

In particular, when the elasticity of demand is constant, expression (22) becomes an equality everywhere. The adjustment parameter is a constant and equals the inverse of the price elasticity of demand. Moreover, the adjustment function takes the form of an exponential function everywhere. Similarly, if we restrict the adjustment function to be linear or of the form $(1+d)^{\frac{q_t - q^{target}}{q^{target}}}$ as studied by [12], the adjustment rate remains the inverse of the price elasticity of demand. Now, let's determine the block size target that aligns with the optimal target of a monopolistic validator.

Definition 2. *(Myopic Miner Incentive Compatibility) A quantity target is Myopic Miner Incentive Compatible (MMIC), if a myopic miner, by creating no fake transactions and adhering to the suggested block size target q^{target}, maximizes her profit.*

The MMIC definition implies that a miner who aims to maximize her revenue should be motivated to comply with the proposed quantity target when choosing her block size ex-ante.

Proposition 3. *Given any isoelastic demand curve, the target block size aligning with the monopolist validator's optimal target block size is expressed as:*

$$\frac{q^{target}}{q^{\max}} = \left(\frac{\varepsilon}{\varepsilon - 1}\right)^{-\varepsilon} \mathbb{E}\left[\lambda^{\frac{1}{\varepsilon}}\right]^{\varepsilon} \tag{25}$$

Notably, if $q^{\max}$ *is adjusted to coincide with user demand under average demand conditions, then:*

$$\frac{q^{target}}{q^{\max}} < \frac{1}{e} \approx 37\% \tag{26}$$

Proof. Refer to Appendix A.4 for proof.

4 Numerical Analysis

This section presents a numerical analysis aimed at quantifying the parameters of optimal ADT-TFMs. The analysis uses Ethereum data due to its widespread availability, the simplicity of its ADT-TFM, and the global usage of its blockchain.

Data: A random sample of 100,000 blocks, encompassing 16,881,386 transactions, was extracted from the complete set of Ethereum blocks. This sample spans from the introduction of EIP-1559 at the London hard fork (block number 12965000, August 5, 2021) to block number 17731768 (July 20, 2023). Additional random block subsamples from before and after the Ethereum merge (block number 15537393, September 15, 2022) were also analyzed.[6]

The median block in the sample contains 143 transactions. Each block is associated with a number and a timestamp, the total gas used by all transactions in the block (equivalent to q in the model), and an array of transactions. Each transaction includes information on its "gas" unit, gas price, and other metadata.

Methodology: The inference of a demand curve requires random variation in supply to distinguish between shifts along the demand curve and shifts in the demand curve itself. Nevertheless, such random variation in supply is rare due to the programmatically defined rules of blockchains.

Instead, a different strategy, informed by the user demand model presented in Section 2, is adopted. Lemma 1 shows that for any density f of user valuations, the price elasticity of demand is the tail ratio $\frac{pf(p)}{1 - F(p)}$. Specifically, if the distribution is Pareto, the price elasticity is its Pareto tail coefficient. Knowing the Pareto tail of the distribution of user valuations allows the determination of the optimal adjustment parameter from Proposition 2 as the inverse of the Pareto tail coefficient.

[6] The results were consistent when sampling 100,000 blocks before and after the merge separately.

In this analysis, transaction gas prices serve as the nearest proxy for user valuations, given the available data. A Pareto distribution is fitted to the empirical distribution of effective transaction gas prices in each randomly selected block to calculate the optimal adjustment rate. However, this approach is not without limitations.[7] Firstly, gas prices censor user valuations at the lower end of the distribution due to the base fee. Secondly, pending transactions in the mempool, which carry lower base fees, are not included. Consequently, the estimate will reflect a heavier tail than the actual user valuation distribution. Therefore, the estimate of the Pareto tail coefficient is an underestimate, meaning its inverse—the optimal adjustment rate—will be overestimated. The adjustment parameters identified here should thus be considered as an upper bound.[8] To address these limitations, several robustness exercises are performed in Section 5, which include restricting the estimation to blocks with a minimum gas price below a certain threshold (to limit the censoring of low valuations) and to blocks that are less than full (to limit the censoring of mempool transactions). These robustness exercises do not alter the primary conclusions of this numerical analysis.

Results: Given that the exponential shape with an adjustment parameter equal to the inverse price elasticity of demand in Proposition 2 is local for a general user valuation distribution, the preferred estimate focuses on blocks with sizes close to the target. Estimating the Pareto tail for blocks of size within ±5% of the block size target (over 7252 blocks and 12965717 transactions) yields a Pareto coefficient of 12.62 and an optimal adjustment rate of 7.92%. This aligns with [12], who simulate the dynamic system of EIP-1559 and find stability around the target block size only for adjustment parameters below 8%. Our contribution clarifies that the adjustment rate encapsulates the economic concept of inverse price elasticity of demand, which must be measured or approximated on-chain.

5 Extensions and Robustness

This section introduces an extension to the model, accounting for fluctuations in cryptocurrency prices, and presents some robustness checks for the numerical estimates. The results suggest that the volatility of cryptocurrency prices may favor a block size limit over base fees.

5.1 Implications of Cryptocurrency Price Fluctuations

This extension examines the impact of cryptocurrency price volatility on the choice between price and quantity controls. Let P denote the exchange rate

[7] As users demand different amounts of block space, each transaction is weighed by the gas units it uses to fit the Pareto distribution.

[8] Given that the optimal adjustment rate found here is lower than its current value of 12.5%, this upper bound estimate offers valuable insights for the design of Ethereum's TFM.

between 1 USD and the cryptocurrency (equivalently, the inverse of the cryptocurrency price when expressed in dollars). Another way to interpret P is the exchange rate between 1 unit of consumption goods and the cryptocurrency. This real model accommodates variations in both the cryptocurrency price and the value of fiat currency. Users pay transaction fees in the cryptocurrency at a nominal price p, implying that the dollar value (equiv. real) of these payments is $\frac{p}{P}$. Meanwhile, Γ represents the dollar (equiv. real) marginal cost.

The following proposition, formulated for simplicity with $\beta = 1$ (though similar insights apply for other parameters), provides an equivalent to Proposition 1 in this context:

Proposition 4. *Suppose (Ψ, η, P) is jointly log-normal distributed. Then the relative value of price adjustments over quantity adjustments is:*

$$\Delta^{\log} = \frac{1}{2}(\varepsilon - 1)\left(\frac{1}{\varepsilon}\sigma_{\Psi}^2 - 2\sigma_{\Psi,\Gamma} - \varepsilon\sigma_P^2 - 2\varepsilon\sigma_{P,\Gamma}\right) \tag{27}$$

In addition to the findings of Proposition 1, the variance of the cryptocurrency price and the covariance between the cryptocurrency price and the dollar (equiv. real) marginal cost both decrease the relative advantage of price controls over quantity controls.

5.2 Numerical Robustness

This section presents the results of supplementary robustness checks to validate the primary numerical analysis findings. These checks were performed under various conditions to address potential concerns highlighted in Section 4, such as the base fee causing censoring of low valuations and the exclusion of pending mempool transactions. To investigate the consistency of the Pareto tail coefficient and the optimal adjustment rate, adjustments were made to the selection of blocks within a size range of $q^{target}(1 \pm \delta\%)$ and different limits to base fees.

Estimations were performed on three samples: full, pre-, and post-merge. Tables 1, 2, and 3 in Appendix A.5 provide detailed results. The estimated optimal adjustment rates are consistently lower than the 12.5% adjustment rate and remain stable under various data partitionings. The estimate is smaller before than after the merge, as previously found. The adjustment rates decrease as the window around the target block size widens and the maximum base fee increases. Ethereum transaction fees are paid in units of gwei, a denomination of the Ethereum cryptocurrency. The increase in rates as the max base fee limit becomes more restrictive (blocks with a base fee less than 30 gwei) can only produce thicker tails due to a restricted range and censoring. Variations based on δ, the window of the target block size, are quite robust, with an adjustment rate in the full sample estimation ranging from 7.33% to 8.03% for a more accommodating max base fee of 200 gwei. Estimates from pre-merge and post-merge suggest the optimal adjustment rate lies within the 6% to 10% window, serving as an upper limit.

6 Conclusion

This paper analyzes the market for user transactions on blockchains, focusing on the balance between price and quantity controls under uncertainty. We present a user model that establishes a demand curve for block space, linking demand parameters to model primitives. The model reveals that the price elasticity of demand for the marginal user is a simple statistic of the distribution of user valuations. If the user valuation distribution follows a Pareto distribution, the aggregate demand for block space takes a simple form.

We explore the welfare implications of price controls in the face of demand and social cost uncertainty, drawing on ideas proposed by [17]. We also study the optimal quantity control for a monopolistic miner, showing that volatility in cryptocurrency prices diminishes the benefits of price controls over quantity controls.

A key conclusion drawn from the analysis is that optimal Transaction Fee Mechanisms *Adaptive to a Deterministic-Target* (ADT-TFM) are exponential and adjust log prices with a slope equal to the inverse price elasticity of demand. Numerical analysis suggests that Ethereum's adjustment parameter might be above the range of our estimates. In addition, we show that a target block size exceeding 37% of the block size that matches average user demand would incentivize validators to include their own value-extracting transactions.

Blockchain designers must adopt simple, robust, and principled methods for updating TFM parameters. Yet, updating parameters and the "rules of the game" as we go might not be best for scalability. The consideration of non-deterministic adjustment rates for TFMs could provide a solution. Preliminary quantitative explorations using adjustment rates derived from prices and quantities of the preceding two blocks have shown promising results in reflecting market conditions during those periods. A more stable and manipulation-resistant approach could involve calculating an average elasticity over a range of prior blocks. Alternatively, introducing noise to the target block size, effectively the "supply curve," could help infer demand fluctuations under normal conditions. However, these economists' dream of ideal experiments must be balanced against the practical realities of managing a system where hundreds of billions of dollars are at stake. This conclusive remark underscores the complexity and nuanced nature of designing effective TFMs for blockchains.

References

1. Akbarpour, M., Li, S.: Credible auctions: A trilemma. Econometrica **88**(2), 425–467 (2020)
2. Bahrani, M., Garimidi, P., Roughgarden, T.: Transaction fee mechanism design with active block producers. arXiv preprint arXiv:2307.01686 (2023)
3. Buterin, V.: Blockchain resource pricing. URL: https://ethresear.ch/uploads/default/original X **2** (2018)
4. Buterin, V.: Make eip 1559 more like an amm curve. Ethereum Research [Online] (2021)

5. Chung, H., Shi, E.: Foundations of transaction fee mechanism design. In: Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA). pp. 3856–3899. SIAM (2023)
6. Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., Breidenbach, L., Juels, A.: Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In: 2020 IEEE Symposium on Security and Privacy (SP). pp. 910–927. IEEE (2020)
7. Ferreira, M.V., Moroz, D.J., Parkes, D.C., Stern, M.: Dynamic posted-price mechanisms for the blockchain transaction-fee market. In: Proceedings of the 3rd ACM Conference on Advances in Financial Technologies. pp. 86–99 (2021)
8. Flynn, J.P., Nikolakoudis, G., Sastry, K.: Prices vs. quantities: A macroeconomic analysis. Quantities: A Macroeconomic Analysis (June 24, 2023) (2023)
9. Klemperer, P.D., Meyer, M.A.: Supply function equilibria in oligopoly under uncertainty. Econometrica: Journal of the Econometric Society pp. 1243–1277 (1989)
10. Lavi, R., Sattath, O., Zohar, A.: Redesigning bitcoin's fee market. ACM Trans. Econ. Comput. **10**(1) (May 2022). https://doi.org/10.1145/3530799
11. Leonardos, S., Monnot, B., Reijsbergen, D., Skoulakis, E., Piliouras, G.: Dynamical analysis of the eip-1559 ethereum fee market. In: Proceedings of the 3rd ACM Conference on Advances in Financial Technologies. pp. 114–126 (2021)
12. Leonardos, S., Reijsbergen, D., Reijsbergen, D., Monnot, B., Piliouras, G.: Optimality despite chaos in fee markets. arXiv preprint arXiv:2212.07175 (2022)
13. Nisan, N.: Serial monopoly on blockchains (2023)
14. Reis, R.: Inattentive producers. The Review of Economic Studies **73**(3), 793–821 (2006)
15. Roughgarden, T.: Transaction fee mechanism design for the ethereum blockchain: An economic analysis of eip-1559. arXiv preprint arXiv:2012.00854 (2020)
16. Roughgarden, T.: Transaction fee mechanism design. ACM SIGecom Exchanges **19**(1), 52–55 (2021)
17. Weitzman, M.L.: Prices vs. quantities. The Review of Economic Studies **41**(4), 477–491 (1974)

A Appendix

A.1 Proof of Lemma 1

At a price p, demand for block space is the measure of users willing to pay p for transaction inclusion, i.e., $\lambda q^{\max}\bar{F}(p)$. This yields the demand cure $p = \left(\bar{F}\right)^{-1}\left(\frac{q}{\lambda q^{\max}}\right)$. The price elasticity of demand is defined by (negative) the percentage change in quantity demanded over the percentage change in price, i.e., $-\frac{dq/q}{dp/p}$. Since from the demand curve

$$q = \lambda q^{\max}\bar{F}(p) \tag{28}$$

and

$$dq/dp = -\lambda q^{\max} f(p) \tag{29}$$

the demand elasticity is $\frac{pf(p)}{1-F(p)}$.

When F is a Pareto distribution with scale p_m and shape α,

$$\bar{F}(p) = Pr(v > p) = \begin{cases} \left(\frac{p_m}{p}\right)^{\alpha} & \text{for } p \geq p_m \\ 1 & \text{for } p < p_m \end{cases} \tag{30}$$

So that above the minimum price p_m demand is

$$q = \lambda q^{\max}\left(\frac{p_m}{p}\right)^{\alpha} \tag{31}$$

therefore we obtain

$$\frac{p}{p_m} = \left(\frac{q}{\lambda q^{\max}}\right)^{-\frac{1}{\alpha}} \tag{32}$$

A.2 Proof of Proposition 1

We first find the expression of the marginal cost $\mathit{\Gamma}$ in (13). The first order condition is:

$$c_i(x_1, \ldots, x_N; \eta) = \gamma p_{x_i} \tag{33}$$

where γ is the Lagrangian of the constraint (1). Since c is homogeneous of degree 1, we have:

$$c(x_1, \ldots, x_N; \eta) = \sum_{i=1}^{N} c_i(x_1, \ldots, x_N; \eta)x_i = \gamma \sum_{i=1}^{N} p_{x_i} x_i = \gamma q. \tag{34}$$

Therefore, $\mathit{\Gamma} = \gamma$. Evaluating at $q = 1$ yields:

$$\mathit{\Gamma}(\eta, p_x) = c(x_1(1, p_x, \eta), \ldots, x_N(1, p_x, \eta); \eta) \tag{35}$$

To prove the proposition, we now take logarithms of the blockchain designer's objective and maximize over p and q. The first-order conditions for the price choice and quantity choice are:

$$\frac{\varepsilon(1-\beta)\nu}{p^*} + \frac{\varepsilon\beta}{p^*} = \frac{\beta\mathbb{E}[\Psi]}{p^*\mathbb{E}[\Psi] - \mathbb{E}[\Gamma\Psi]} \tag{36}$$

$$\frac{(1-\beta)\nu}{q^*} + \frac{\beta}{q^*} = \frac{\beta/\varepsilon p_m (q^*)^{-1-1/\varepsilon}\mathbb{E}[\Psi^{1/\varepsilon}]}{p_m (q^*)^{-1/\varepsilon}\mathbb{E}[\Psi^{1/\varepsilon}] - \mathbb{E}[\Gamma]} \tag{37}$$

Denote $\bar{\nu} = (1-\beta)\nu + \beta$. Then we obtain,

$$p^* = \frac{\varepsilon\bar{\nu}}{\varepsilon\bar{\nu} - \beta}\frac{\mathbb{E}[\Gamma\Psi]}{\mathbb{E}[\Psi]} \tag{38}$$

$$q^* = \left(\frac{\varepsilon\bar{\nu}}{\varepsilon\bar{\nu} - \beta}\frac{1}{p_m}\frac{\mathbb{E}[\Gamma]}{\mathbb{E}[\Psi^{1/\varepsilon}]}\right)^{-\varepsilon} \tag{39}$$

The log value of price controls and quantity controls are then

$$\log \mathcal{V}^p = -\varepsilon\bar{\nu}\log(p^*/p_m) + (1-\beta)\log(\mathbb{E}[\Psi^\nu]) + \beta\log(p^*\mathbb{E}[\Psi] - \mathbb{E}[\Psi\Gamma]) \tag{40}$$

$$\log \mathcal{V}^q = \bar{\nu}\log(q^*) + \beta\log(p_m(q^*)^{-1/\varepsilon}\mathbb{E}[\Psi^{1/\varepsilon}] - \mathbb{E}[\Gamma]) \tag{41}$$

Replacing optimal choices with their values in (38) and (39) we get

$$\log \mathcal{V}^p = -\bar{\nu}\log(p^*/p_m) + (1-\beta)\log(\mathbb{E}[\Psi^\nu]) + \beta\log(\frac{\beta}{\varepsilon\bar{\nu} - \beta}\mathbb{E}[\Psi\Gamma]) \tag{42}$$

$$\log \mathcal{V}^q = \bar{\nu}\log(q^*) + \beta\log(\frac{\beta}{\varepsilon\bar{\nu} - \beta}\mathbb{E}[\Gamma]) \tag{43}$$

Simplifying yields:

$$\begin{aligned}\log \mathcal{V}^p - \log \mathcal{V}^q = {} & \bar{\nu}\varepsilon\log\mathbb{E}[\Psi] \\ & +(\bar{\nu}\varepsilon - \beta)(\log\mathbb{E}[\Gamma] - \log\mathbb{E}[\Psi\Gamma]) \\ & -\varepsilon\bar{\nu}\log\mathbb{E}[\Psi^{1/\varepsilon}] + (1-\beta)\log\mathbb{E}[\Psi^\nu]\end{aligned} \tag{44}$$

For the joint log- (Ψ, Γ), with mean

$$\mu = \begin{pmatrix}\mu_\Psi \\ \mu_\Gamma\end{pmatrix}$$

And variance-covariance matrix

$$\Sigma = \begin{pmatrix}\sigma_\Psi^2 & \sigma_{\Psi,\Gamma} \\ \sigma_{\Psi,\Gamma} & \sigma_\Gamma^2\end{pmatrix}$$

We have

$$\log \mathbb{E}[\Psi] = \mu_\Psi + \frac{1}{2}\sigma_\Psi^2 \tag{45}$$

$$\log \mathbb{E}[\Gamma] - \log \mathbb{E}[\Psi\Gamma] = -\mu_\Psi - \frac{1}{2}\sigma_\Psi^2 - \sigma_{\Psi,\Gamma} \tag{46}$$

$$\log \mathbb{E}[\Psi^{1/\varepsilon}] = \frac{1}{\varepsilon}\mu_\Psi + (\frac{1}{\varepsilon^2})\frac{1}{2}\sigma_\Psi^2 \tag{47}$$

$$\log \mathbb{E}[\Psi^\nu] = \nu(\mu_\Psi + \frac{1}{2}\nu\sigma_\Psi^2) \tag{48}$$

Putting them together, we get the result

$$\log \mathcal{V}^p - \log \mathcal{V}^q = \frac{1}{2}\left(\left(\hat{\nu} - \frac{\bar{\nu}}{\varepsilon}\right)\sigma_\Psi^2 - 2(\varepsilon\nu - \beta)\sigma_{\Psi,\Gamma}\right) \tag{49}$$

where $\hat{\nu} = (1-\beta)\nu^2 + \beta$

A.3 Proof of Proposition 2

Let's start by defining the notion of optimality in this dynamic context. Let q^t the equilibrium quantity at time t. Denote $x_t = \frac{q_t - q^{target}}{q^{target}}$ the percentage deviation from the target at time t. Denote λ the arrival rate in normal times (i.e. the expected arrival rate). Consider a shock to the demand curve $\lambda_t^{-1} \equiv \lambda^{-1} + z_t$. At the protocol set price p_t, the quantity lies on the demand curve $q_t = \lambda_t q^{\max}\bar{F}(p_t)$. The deviation from target x_t can be due to the shock to demand z_t or a protocol price p_t that is not properly set so that q_t deviates from not the target.

We have the expression $p_t = \bar{F}^{-1}(\frac{q_t}{q^{\max}\lambda_t})$. From the expression of the ADT-TFMs, we have:

$$p_{t+1} = \bar{F}^{-1}(\frac{q^{target}(1+x_t)}{q^{\max}\lambda_t})g(d \times x_t) \tag{50}$$

Without loss of generality at time $t+1$, demand returns to normal so that $\lambda_{t+1} = \lambda$.

Known Intensity of Demand: Suppose for now that the realization of z_t i.e. λ_t is known, then the deviation from the target quantity at time $t+1$ is

$$x_{t+1} = \frac{\bar{F}\left(\bar{F}^{-1}(\frac{q^{target}(1+x_t)}{q^{\max}\lambda_t})g(d \times x_t)\right)}{\bar{F}(p^{target})} - 1. \tag{51}$$

We can see that by setting

$$g(d \times x_t) = \frac{\bar{F}^{-1}(\frac{q^{target}}{q^{\max}\lambda})}{\bar{F}^{-1}(\frac{q^{target}(1+x_t)}{q^{\max}\lambda_t})} \tag{52}$$

We guarantee that The quantity at time $t+1$ is at the target. The issue is that demand λ_t is uncertain so we look at the function f that performs in the worst-case scenario.

Unknown Intensity of Demand: Because demand is uncertain, the adjustment function can depend on the gap from target x_t but not on λ_t. So we need to evaluate the deviation that is closest for the worst case value of z_t. We have:

$$\ln \bar{F}^{-1}\left(\frac{q_{t+1}}{q^{target}}\right) = \ln \bar{F}^{-1}(\frac{q^{target}(1+x_t)}{q^{\max}\lambda_t}) - \ln \bar{F}^{-1}(\frac{q^{target}}{q^{\max}\lambda}) + \ln g(d \times x) \quad (53)$$

Suppose $-\ln \bar{F}^{-1}$ is concave, then for any x_t and $z_t = \lambda_t^{-1} - \lambda^{-1}$

$$\ln \bar{F}^{-1}(\frac{q^{target}}{q^{\max}\lambda}) - \ln \bar{F}^{-1}(\frac{q^{target}(1+x_t)}{q^{\max}\lambda_t}) \geq \frac{1}{\varepsilon(q^{target})}(x_t + z_t) \quad (54)$$

From (53), we have that, $\ln g(d \times x)$ is the closest approximation of $\ln \bar{F}^{-1}(\frac{q^{target}}{q^{\max}\lambda}) - \ln \bar{F}^{-1}(\frac{q^{target}(1+x_t)}{q^{\max}\lambda_t})$ that is independent of z_t. Therefore, $f(d \times x) = x/\varepsilon(q^{target})$,i.e. $d = \varepsilon(q^{target})$ and $f = \exp$.

A.4 Proof of Proposition 3

From equation (39), the optimal quantity for a monopolistic miner is for $\beta = 1$,

$$q^* = \left(\frac{\varepsilon}{\varepsilon - 1}\frac{1}{p_m}\frac{\mathbb{E}[\Gamma]}{\mathbb{E}[\Psi^{1/\varepsilon}]}\right)^{-\varepsilon} \quad (55)$$

With $\Psi = \lambda q^{\max}$. Now suppose that the minimum user valuation is greater than the expected marginal cost $p_m > \mathbb{E}[\Gamma]$. Then by including her own transactions up to the block size limit $q^{\max}$ and paying the base fee to herself, the validator gets a positive value in expectation. Therefore, MMIC requires that $p_m \leq \mathbb{E}[\Gamma]$, thus

$$q^* \leq q^{\max}\left(\frac{\varepsilon - 1}{\varepsilon}\right)^{\varepsilon} \mathbb{E}[\lambda^{1/\varepsilon}]^{\varepsilon}. \quad (56)$$

By Jensen's inequality, $\mathbb{E}[\lambda^{1/\varepsilon}]^{\varepsilon} \leq \mathbb{E}[\lambda]$. So the block size limit is set to match user demand in expectation, then $\mathbb{E}[\lambda] = 1$, thus

$$\frac{q^*}{q^{\max}} \leq \left(\frac{\varepsilon - 1}{\varepsilon}\right)^{\varepsilon}. \quad (57)$$

The right-hand side is an increasing function for $\varepsilon > 1$ with limit e^{-1}

A.5 Tables

Parameters		Outputs		Observations
δ	max base fee (gwei)	shape α	adjustment rate d	number of blocks
5%	200	12.45	8.03	7189
	100	11.83	8.45	6645
	60	11.01	9.08	5867
	30	9.50	10.53	4130
33%	200	12.77	7.83	43718
	100	12.09	8.27	40471
	60	11.26	8.88	35649
	30	9.48	10.55	25193
87.5%	200	13.64	7.33	78881
	100	12.62	7.92	70769
	60	11.53	8.67	60141
	30	9.50	10.53	41140

Table 1: shape of Pareto fit α and optimal adjustment rate s for the different maximum gas prices (base fee) in units of gwei and selection of blocks within size $q^{target} \pm \delta\%$. Full sample estimation.

Parameters		Outputs		Observations
δ	max base fee (gwei)	shape α	adjustment rate d	number of blocks
5%	200	11.46	8.73	5088
	100	11.24	8.89	4962
	60	10.84	9.23	4688
	30	9.55	10.47	3547
33%	200	11.48	8.71	30029
	100	11.29	8.86	29454
	60	10.89	9.18	27884
	30	9.55	10.47	21288
87.5%	200	11.32	8.83	42042
	100	11.15	8.97	41348
	60	10.78	9.28	39426
	30	9.47	10.56	30890

Table 2: shape of Pareto fit α and optimal adjustment rate s for the different maximum gas prices (base fee) in units of gwei and selection of blocks within size $q^{target} \pm \delta\%$. Post-merge sample estimation.

Parameters		Outputs		Observations
δ	max base fee (gwei)	shape α	adjustment rate d	number of blocks
5%	200	14.85	6.73	2101
	100	13.57	7.37	1683
	60	11.70	8.55	1179
	30	9.15	10.93	583
33%	200	15.59	6.41	13689
	100	14.22	7.03	11017
	60	12.57	7.96	7765
	30	9.09	11.01	3905
87.5%	200	16.29	6.14	36839
	100	14.69	6.81	29421
	60	12.96	7.71	20715
	30	9.58	10.44	10250

Table 3: shape of Pareto fit α and optimal adjustment rate s for the different maximum gas prices (base fee) in units of gwei and selection of blocks within size $q^{target} \pm \delta\%$. Pre-merge sample estimation.