Blockchain Fee Policies: Quantity- vs Price-Control

Design across Protocols

Abdoulaye Ndiaye*

NYU Stern School of Business

June 25, 2025

Abstract

This paper develops an economic theory to explain the optimal choice of blockchain transaction fee policies, contrasting quantity controls (as in Bitcoin) and price controls (as in Ethereum). I model the blockchain as a decentralized platform where validators, possessing temporary monopoly power and facing uncertain operational costs, choose transactions under varying user demand. The analysis finds that price controls outperform quantity controls when demand volatility is significant, the correlation between marginal costs and demand is low, and validators hold substantial bargaining power—conditions aligning closely with Ethereum's current structure following its proof-of-stake transition. Conversely, quantity controls are optimal when marginal costs positively correlate with demand and validators' bargaining power is limited, conditions characteristic of Bitcoin's proof-of-work model. Using recent transaction data from Ethereum's fee policy (EIP-1559), I identify that Ethereum's base fee adjusts more rapidly than optimal. Additionally, I derive bounds for block size targets to mitigate maximal extractable value (MEV) exploitation. These insights offer guidance for improving blockchain governance and optimizing fee policy design.

Keywords: Blockchain, Transaction Costs, Fee Policies, Bitcoin, Ethereum, Demand Fluctuations, Price Elasticity

^{*}Contact: andiaye@stern.nyu.edu. I thank Mohammad Akbarpour, Yannis Bakos, Joe Bonneau, Agostino Capponi, Davide Crapis, Jeremy Clark, Max Croce, Theo Diamandis, Joshua Gans, Pranav Garimidi, Itay Goldstein, Deeksha Gupta, Guillaume Haeringer, Hanna Halaburda, Urban Jermann, Kose John, Scott Kominers, Anthony Lee-Zhang, Jacob Leshno, Jiasun Li, Paolo Martellini, Ciamac Moallemi, Cyril Monnet, Barnabe Monnot, Dirk Niepelt, Mallesh Pai, Jonathan Payne, Max Resnick, Tim Roughgarden, Thomas Ruchti, Fahad Saleh, Elaine Shi, Philipp Strack, Ertem Nusret Tas, Aleh Tsyvinski, Quentin Vandeweyer, Venky Venkateswaran, Olivier Wang, Matt Weinberg, Ariel Zetlin-Jones, and seminar participants at a16z Crypto, Columbia Cryptoeconomics Conference, NYU, Carnegie Mellon Secure Blockchain Conference, the Office of Financial Research of the US Treasury Department, the CEPR Fintech and Digital Currencies RPN Workshop for helpful discussions and comments. A previous version of this paper circulated under the title "Why Bitcoin and Ethereum Differ in Transaction Costs: A Theory of Blockchain Fee Policies". First Version: July 27, 2023.

1 Introduction

Transaction costs play an essential role in the evolving areas of blockchain technology and decentralized finance. They manage the allocation of space on a blockchain—block space—which is a limited resource due to the economic and scaling limits of blockchain systems (Buterin, 2021; Budish, 2024). However, these transaction costs can induce illiquidity and create trade execution risk. As cryptocurrencies integrate into mainstream financial systems, exemplified by the approval of Bitcoin and Ethereum exchange-traded funds (ETFs) (SEC, 2024), inefficient fee policies can delay settlement of underlying crypto-assets and potentially affect financial stability. This paper provides a theory of how and why fee policies differ across blockchain platforms and what these fee policies should look like.

From a regulatory perspective, understanding the mechanisms behind how transaction services are priced and prioritized is crucial for developing effective oversight policies for blockchain platforms. Furthermore, from a blockchain design perspective, transaction costs directly affect economic efficiency and ultimately impact user adoption.

The Bitcoin and Ethereum blockchains are two leading examples that reflect opposite extremes in the choice of fee policies. For Bitcoin fees, a maximum quantity is set by the protocol (maximum block size), and users and transaction service providers independently choose fees. I call this the quantity-setting or quantity controls regime. On the other hand, Ethereum sets a minimum price while the block size adjusts in response to demand. I call this the price-setting or price controls regime. Most blockchains follow a variation of either the Bitcoin or the Ethereum fee policies, and some set constant fees by default or subsidize user fees until their network matures and faces congestion issues. In addition, there is an active debate on the future of the Bitcoin blockchain. As the payoff for transaction service providers programmatically decreases every four years, the role of fees becomes increasingly important for Bitcoin.

In this paper, I model a blockchain—similar to that underlying Bitcoin or Ethereum—as a distributed computing network where users submit transactions for inclusion by validators—

the transaction service providers. Transactions, representing data that modify 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 transaction processing. Using their limited resources, validators 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 impose a maximum block size, constraining block space supply.

On the supply side of transaction processing services, validators (also known as miners in proof-of-work systems) are responsible for processing transactions and adding them to the blockchain. While there are multiple validators in the network, each validator enjoys temporary monopoly power when they are selected to propose the next block due to the consensus mechanism (e.g., proof of work or proof of stake).¹ This grants them discretion over which transactions to include and how to prioritize them. Therefore, in our model, we focus on a representative validator with monopolistic characteristics during their turn to validate a block.²

The validator faces fluctuations in marginal costs due to varying operational costs of transaction processing on blockchains. In a proof-of-work protocol such as Bitcoin, miners use computational resources to solve a mathematical puzzle, and their marginal cost varies with the level of competition that they face before being selected to validate transactions. While competition among validators affects the probability of winning the right to produce a block (impacting expected benefits), we focus on the variability in validators' operational costs to capture the uncertainty in their marginal costs.³ In a proof-of-stake protocol such as the updated Ethereum, marginal costs are essentially constant due to the reduced reliance

¹See Saleh (2020) for an economic analysis of proof-of-stake.

 $^{^{2}}$ We interchangeably use the plural term 'validators' for the group of validators as a whole and the singular 'validator' when referring to the interim monopolist validator chosen to process transactions.

 $^{^{3}}$ In the case of monopolistic miners (such as a mining pool), Cong et al. (2021) show that the risk of being selected as a validator can be diversified.

on computational power.

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 reveal their maximum willingness-to-pay (their true valuations) for transaction inclusion. This can arise if waiting or transaction failure is costly and users rationally set bids equal to their utility cost of not transacting.⁴ This process forms the microfoundations of an aggregate user demand curve, with a demand shifter that captures low and high-demand situations.

In the face of these uncertainties, the protocol must commit to either a base fee—an exante price control—or a block size limit—an ex-ante quantity control.⁵ This choice highlights a disagreement between the blockchain designers' objective and the profit-maximization objective of the validator with significant interim monopoly power after they are chosen to validate a block. Indeed, a protocol designer who prefers full capacity utilization or has an exogenous technological block size target (that internalizes social benefits and costs) will choose the quantity-setting regime as it guarantees the desired block size irrespective of sources of uncertainty. On the other hand, a monopolist validator would choose either the price-setting or the quantity-setting regimes, depending on the relative degree of uncertainty in demand and marginal costs.

I model the resolution of the conflict between the blockchain designer's and monopolistic validator's preferences through Nash bargaining over the protocol profits. I show that, in this context, the choice of instruments is more nuanced than Weitzman's (1974) "prices vs. quantities" insight: namely, that price controls prove more effective when demand uncertainty is high and quantity controls more effective when uncertainty in marginal costs is high.

⁴Dominant-strategy incentive compatibility for all the fee policies considered in this paper is substantiated by the game-theoretic proofs provided by Roughgarden (2020, 2021).

⁵Since the fee policies are rule-based and encoded in the blockchain, the protocol must commit to a supply schedule before any uncertainty realizations. We analyze the two extremes of price-setting and quantitysetting supply schedules, motivated by the Bitcoin and Ethereum cases, and as a first step in the analysis of blockchain fee policies.

In general, the key determinants of the advantage of price-setting in blockchains are the validator's bargaining power, the elasticity of demand, the validator's uncertainty about demand, and the covariance of demand and marginal costs.

First, demand uncertainty favors price controls when the validator has high bargaining power, as block size adjustments provide the flexibility necessary to accommodate demand fluctuations. Second, a positive correlation between marginal costs and demand disfavors price controls. In this case, quantity adjustments would produce larger blocks when marginal costs are high, thereby decreasing efficiency. Third, a higher 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 her transaction in the next block—further amplifies the relative advantage of price controls over quantity controls and makes the choice over fee policies even more important. However, quantity controls become more effective if the monopolist validator has low bargaining power. Last, without uncertainty, the blockchain designer and validator remain indifferent between price and quantity controls.

This result helps us understand the differences in the policies determining fees and block space in Bitcoin and Ethereum. For both blockchains, the price elasticity of demand is high, and demand uncertainty is significant.

Ethereum can be viewed through the lens of our model as an instance where validators have high bargaining power (or are considered influential by the protocol designers).⁶ In addition, since its proof-of-stake upgrade, the marginal cost of a marginal block increase for Ethereum validators is virtually constant. In this case, my result points to price-setting as the most favorable policy for Ethereum. This helps explain the recent adoption of the Ethereum

⁶While the Ethereum Foundation (EF) maintains considerable authority over protocol upgrades, the validators' economic leverage in day-to-day block proposal can also be significant. In particular, validators may capture substantial amounts of fees and Maximum Extractable Value (MEV), and hence, they can indirectly influence or resist certain protocol-level fee adjustments if these undermine their profitability. Nonetheless, the Ethereum governance process—rooted in Ethereum Improvement Protocols (EIPs) and developer-led forks—often appears to grant the EF a strong position. Fracassi et al. (2024) find that 10 individuals are responsible for proposing 68% of all implemented Core EIPs, some of which are connected to Consensys, a private company, and other EF affiliates that have advised private companies. Recent comparisons of the Ethereum blockchain to its competitor Solana have also led to a new push to increase its block size limit https://pumpthegas.org/.

Improvement Proposal 1559's (EIP-1559's) price-setting policy for the Ethereum blockchain. Furthermore, Ethereum's fee policies are found to perform better after Ethereum's proof-ofstake upgrade, consistent with the correlation between validators' marginal costs and demand being lower than miners' in a proof-of-work protocol.

For its part, Bitcoin is distinguished as the first blockchain with an anonymous founder, and its core developers strongly emphasize decentralization and censorship resistance. Most proposals to change Bitcoin's fees and block size policies have failed. Bitcoin can, therefore, be viewed through the lens of our model as an instance where validators (here miners) have low bargaining power (or are not attributed enough importance by protocol designers). In addition, under the proof-of-work protocol, the marginal cost of miners is largely positively correlated with demand. Our result then states that, in this case, price-setting is less effective than in the case of Ethereum.

Last, it is essential to consider that users (and validators) might value block space (and marginal costs) in dollars or real terms rather than in the native currency of the blockchain. To accommodate this, I expand the model to introduce uncertainty in the cryptocurrency price in US dollars or real terms. Notably, price controls are restricted to be expressed in units of the native currency. I show that volatility in the cryptocurrency price reduces the advantage of price controls over quantity controls.

Building on these insights, I briefly discuss the implications of my results for fee policies on Ethereum, the most widely utilized public blockchain, whose fee policy is a blueprint for many blockchains that follow the price-setting regime. I show that the rate at which Ethereum's fee policies readjust prices should be tightly linked to the price elasticity of demand to include a transaction in the next block. Using a random sample of Ethereum transaction data, I apply my framework to evaluate recent changes to Ethereum's fee policies. I find that the rate at which Ethereum's fees are changed is faster than optimal.

In addition, I characterize the optimal block size target (a quantity control) for a monopolistic validator and offer tight bounds on its size relative to the block size limit. There are widespread concerns regarding the potential for validators and other users to exploit their power to capture what is colloquially referred to as *"maximal extractable value"* (MEV) through the censoring, swapping, and front-running of mempool transactions (Daian et al., 2020). I find that Ethereum's current ratio of block size target to maximum block size leaves too much room for a monopolist validator to include her value-extracting transactions.

Blockchain designers need simple and robust economic insights to design their fee policies. With this need in mind, I conclude by suggesting open questions connected to this research.

1.1 Literature Review

The literature has generally explained blockchain fees arising because of competition for block space among heterogeneous, impatient users (Huberman et al., 2021). Easley et al. (2019) study the evolution of Bitcoin transaction fees, while Hinzen et al. (2022) highlight Bitcoin's limited adoption problem. Some studies highlight the strategic use of capacity on the Bitcoin blockchain (Lehar and Parlour, 2020; Malik et al., 2022). Catalini and Gans (2020) and Böhme et al. (2015) provide a simple primer on the economics of blockchains. This paper contributes to this literature by explaining the emergence of the two families of blockchain fee policies, those of Bitcoin and Ethereum.

Our modeling of demand uncertainty in blockchain usage and cryptocurrency investment is consistent with prior research. Aiello et al. (2023) documents the interaction of cryptocurrency price fluctuations with household investment behavior. Kogan et al. (2024) observe cryptocurrency usage and trades and identify momentum in cryptocurrency investments. In addition, fluctuations in uncertainty in Bitcoin block-production costs have been widely documented. Rehman and Kang (2021) show a negative correlation between the Bitcoin hash rate and energy commodity prices and Delgado-Mohatar et al. (2019) document increasing and varying electricity production cost of Bitcoin over time. However, since the proof-ofstake upgrade, the marginal production cost of Ethereum blocks is essentially fixed, and participation as a validator is determined by stake. Jermann (2023) studies the steady-state dynamics of Ethereum stake. We use these facts to motivate our modeling approach based on demand and supply uncertainties for blockchains.

This paper contributes to the literature on price versus quantity controls, a domain pioneered by Weitzman (1974). The issue of choosing a supply function under uncertainty has been explored in Klemperer and Meyer (1989). My approach aligns closely with that of Reis (2006) and Flynn et al. (2023), who study the choice of a supply function from a firm's perspective and its macroeconomic implications. While my work draws inspiration from Weitzman (1974), it diverges in that it contemplates the planner's (blockchain designer's) problem with a variety of goals, including purely technical objectives like those seen in practice, such as block size targets. The conclusions of this analysis are then applied to the design of transaction fee mechanisms (TFMs). Specifically, Ndiaye (2023) provides a summary of how the economic factors that affect the design of fee policies, as studied in this paper, correlate with the technical features of blockchains.

The literature taking a mechanism design perspective to examine fee policies is growing. Notably, Akbarpour and Li (2020) examine mechanisms immune to designer manipulations referred to as "credible mechanisms"—and demonstrate that the well-known second-price auction does not meet this credibility criterion. In the blockchain context, Roughgarden (2021) applies this credibility condition to fee policies and establishes that EIP-1559, Ethereum's TFM, which essentially acts as a first-price auction with a dynamic reserve price, and its variations are incentive-compatible for users and adhere to a form of myopic credibility for validators. These findings are further consolidated by Chung and Shi (2023). These papers provide game-theoretic foundations that guarantee that the ADT-TFMs (transaction fee mechanisms that are adaptive to a deterministic target) that I study are incentivecompatible. Ferreira et al. (2021) explore an alternative aspect of fee policies by investigating posted price mechanisms. The approach that I take in this paper is complementary to this strand of literature. Moreover, I study the dynamics of fee policies, offering insights into their updating rules. This paper also advances the literature examining the block space market from a macroeconomic perspective. The concepts formalized in Section 3 of this paper build upon and extend Buterin's (2018) reading of Weitzman (1974), with important differences characterizing the blockchain context because the protocol designer is limited in her capacity to enforce quantity or price controls. In other related work, Lavi et al. (2022) and Nisan (2023) investigate the monopolistic market for unlimited block space, in contrast to this paper, which takes the block size limit as an exogenous technological constraint.

Last, this paper contributes to studies of the dynamics of the Ethereum fee policies recently adopted in EIP-1559. Leonardos et al. (2021) studies the behavior of the dynamic system resulting from the fee policies, and Leonardos et al. (2022) uncovers numerous empirical properties for which this paper provides a theoretical explanation.

Outline: The paper is organized as follows: Section 2 gives an overview of the functioning of the Bitcoin and Ethereum protocols, makes the case that validators have had more bargaining power in the history of Ethereum than in that of Bitcoin, and explains how Ethereum fees are determined. Section 3 introduces the model and provides the main results of my analysis and an extension to cryptocurrency price fluctuations. Section 4 examines implications for Ethereum's fee policies. Last, Section 5 concludes the paper.

2 Transaction Fees and Block Space Utilization in Bitcoin and Ethereum Protocols

2.1 The Evolution Bitcoin and Ethereum Protocols and Fee Policies

A blockchain is a decentralized digital ledger that records transactions across a network of computers, called validators, in an immutable chained list of blocks. Bitcoin, introduced by Nakamoto (2008), was conceived as a peer-to-peer digital currency. In contrast, Ethereum, proposed by Buterin et al. (2013), extends the basic blockchain concept to include more

versatile functionalities such as self-executing agreements, known as "smart contracts".

In Bitcoin and Ethereum, each block has a fixed capacity because of technological constraints such as bandwidth and storage and the negative externalities associated with large blocks' propagation times. These limitations ensure that running a node remains accessible to a broad range of participants, thereby fostering decentralization. This limited capacity necessitates transaction fees, which serve as a market mechanism to allocate this scarce resource. Fees incentivize validators to prioritize and include transactions in a block.

Over the years, both Bitcoin and Ethereum have experienced changes in their block capacity and fee structures to adapt to changing network demands. On the price side, Bitcoin shifted from offering no-fee transactions to imposing transaction fees, as studied by Easley et al. (2019). On the block capacity side, since its inception, Bitcoin has given control of this capacity to the protocol developers. The original 1MB block size limit in the Bitcoin blockchain ignited intense debates within the community over scalability versus decentralization. On one side, proponents of a larger block size argued that increasing capacity would enable more transactions per block, alleviating congestion and lowering fees. On the other side, critics warned that larger blocks would raise the computational and storage requirements for running a validator, compromising the network's decentralization. This ideological divide peaked in 2017, leading to a "hard fork" that birthed Bitcoin Cash, a separate chain with an 8MB block size. Concurrently, Bitcoin adopted Segregated Witness (SegWit), effectively changing the block size limit to a more complex "block weight" limit of approximately 4 million units.

Ethereum, for its part, imposed transaction fees from the beginning. In Ethereum, "gas" serves as the unit of resources, and the "gas limit" dictates the maximum network capacity in a single block. Initially, Ethereum attempted a different paradigm and gave control over capacity to validators. Each validator was allowed to change capacity up or down by 0.1% for each new block. The underlying rationale was to transform the philosophical debate on block size into an economic decision, assuming that validators, heavily invested in the network,

would act individually in the long-term interest of the network. However, in practice, the history of Ethereum block size changes in Table A1 in Appendix A illustrates that validators have often abused their bargaining power and colluded to manipulate network capacity. This behavior is exemplified by an April 21, 2021, announcement from Sparkpool, a major validator based in China with 23.5% of the network's computational power at that time:

Osparkpool_eth: "We are raising gas limit to 15 million"

Obtcdentist: "did you get permission from the core devs?"

Osparkpool_eth: "What we need is advice from dev, not permission."

To curb such practices, Ethereum introduced EIP-1559, a fee mechanism designed to limit validators' discretionary power over network capacity and enhance user fee predictability.

Figure 1 illustrates the time series of the capacity utilization rates of the Bitcoin and Ethereum blockchains. Considering blockchain upgrades over time and their variations, I define the block size target as the size below which the reserve price for pending transactions does not increase. The capacity utilization rate is, consequently, the daily average of the fraction of the realized block size to the block size target.

For Bitcoin, the block size target was 1MB pre-SegWit and is 4 million weights postupgrade. The top panel indicates an initial surge in Bitcoin's utilization rate during its nascent phase. However, as its use has become more widespread, the network has not continuously operated at full capacity. This fact is corroborated by Lehar and Parlour (2020), who attribute the prevalence of less-than-full blocks to strategic capacity management by monopolistic validators.

Before the London hard fork, Ethereum's block size target varied based on validator votes, ranging from 3.1 million to 15 million gas, as detailed in Appendix A. Post–London fork and EIP-1559 implementation, the target is a fixed 15 million gas, with an upper limit of 30 million and a dynamic reserve price adjustment for pending transactions to align closely with the target. The bottom panel reveals that, just as for Bitcoin, Ethereum's early-phase utilization rate rose steadily. However, following the introduction of EIP-1559 via the London



Figure 1: Time series of Bitcoin and Ethereum capacity utilization rates

hard fork, the blockchain has operated at full capacity. My analysis attributes this to EIP-1559's design, arguing that transaction fees now effectively shape and regulate the supply curve for monopolistic validators.

Such economic interactions between fee mechanisms and capacity have been overlooked in earlier studies. For instance, Lehar and Parlour (2020) focus on strategic capacity utilization in Bitcoin, while Huberman et al. (2021) theoretically ascribe full capacity utilization to the free entry of validators even though Bitcoin does not continuously operate at full capacity despite such entry. In my analysis in Appendix B.1, I examine factors affecting block space demand, including the number of active addresses, token prices, transaction fees, and residual demand through mempool size for Bitcoin. The data reveal strong correlations between traditional demand-side factors such as active addresses and transaction fees and the utilization rate for Bitcoin and Ethereum pre-EIP-1559. Post-EIP-1559, these factors became decoupled from block utilization.

In my study on block space supply in Appendix B.2, I evaluate factors such as computing power and mining pool concentration for both Bitcoin and Ethereum, adding the total Ethereum staked and validator pool concentration for Ethereum post-proof of stake. The results indicate that neither computing power nor Ethereum staked significantly impacted block utilization rates for Bitcoin or Ethereum. However, higher concentration rates among mining pools in Bitcoin, measured through the Herfindahl-Hirschman index (HHI), correlate with sub-100% block utilization, suggesting that miners may exercise market power to leave blocks less full. This pattern does not hold for Ethereum, where both before and after the transition to proof of stake, mining pool concentration shows no discernible effect on block utilization rates. The stabilization of Ethereum's block utilization post-London fork further corroborates this observation.

Advanced gas fee	×					
Max base fee (GWEI) 🚯						
33.973104442 ≈ 0.00071343 ETH	4 🗘	Next Block	Gas Estin	nator 🖲		EIP-1559 🧲
Current: 25.09 GWEI↑ 12hr: 24.09 - 38.35 GW	'EI	Pending block numb	er: 18016296			
Priority Fee (GWEI) 🚯		MORE LIKELY				LESS LIKI
0.1 ≈ 0.0000021 ETH		priority fee 0.19	priority fee 0.1	priority fee 0.1	priority fee 0.09	priority fee 0.08
2urrent: 0.1 - 15 GWEI↓ 12hr: 0.07 - 150.08 GW	/EI	max fee 31 99% probability	max fee 31 95% probability	max fee 31 90% probability	max fee 31 80% probability	max fee 31 70% probability
"Advanced"		Estimated number of	f transactions in pend	ding block: 83		
an limit 21000 Edit		Max price in pending	g block: 46			
		Base Fee: 24.53 GV	VEI Trending 🗡			
Course		Time since last upda	te: 4 seconds			
Save		powered by blocknative	>	View Charts	Explore ETH Gas	API → Share [

(a) User price choice from Metamask wallet

(b) Base fee and priority fee estimator from Blocknative

Figure 2: Ethereum fee settings and estimator

2.2 Transaction Fees on Ethereum

Ethereum fees initially operated on a bidding system constrained by a fixed block size. The fixed per-block gas limit and fluctuations in demand resulted in user delays, as the system lacked a "slack" mechanism to adjust block size to meet varying demand. In addition, since validators collected the fees, this system led to a first-price auction when demand was high. As first-price auctions are not incentive-compatible when user valuations are unobserved, this required complex fee estimation efforts from users.

With EIP-1559, Ethereum's fee structure underwent significant changes to address the issues with the first-price auction with a fixed block size limit. The system now operates with a target block size set at $q^{target} = 15M$ gas, a legacy from the pre-London hard fork settings, and a maximum block size of $q^{\max} = 2q^{target}$. Each transaction comes with a "base fee" on the user end, algorithmically adjusted based on network demand. The minimum gas

price, p_t , is adjusted based on the formula

$$p_t = p_{t-1} \cdot (1 + d \frac{q_{t-1} - q^{target}}{q^{target}})$$
(1)

where d is an adjustment parameter set to $\frac{1}{8}$, allowing the minimum price to double in 8 blocks when blocks are full. In addition to the base fee, users can include a "tip" to incentivize faster processing by validators. The dual-fee structure allows more predictability in transaction costs, as the base fee aligns closely with network congestion. Figure 2 shows fee settings and an estimator of the base fee for Ethereum users.

From a validator's standpoint, selection is stochastic, contingent on either the proofof-work or proof-of-stake mechanism. Though the selected validator enjoys a monopolistic position during her turn to validate a block, the protocol's fee structure regulates her behavior. Each transaction j carries a computational cost q_j , measured in gas units. Ethereum transaction senders pay an amount computed as $q_j \cdot \min\{p_t + \delta_j, c\}$, where δ_j is the tip and c the fee cap, with $c \ge p_t$. Incentive compatibility requires that, under normal demand conditions, the base fee adjusts upward to match the willingness to pay of the marginal user, and the tip is small in proportion. The aggregate base fee revenue over N transaction, $\sum_{j=1}^{N} q_j p_t$, is "burned", primarily to address the validator's off-chain incentives, as argued by Roughgarden (2021). Meanwhile, tips and a block reward go directly to the validator. By diverting a portion of the revenue away from the validator, the protocol has instruments to ensure that supply is not artificially restricted and that transactions of higher value are included.

EIP-1559 raises several questions regarding its rationale, potential for improvement, and the design of other fee policies that share its simplicity. Specifically, in the face of demand fluctuations, why might a protocol designer opt for a blockchain with an imposed quantity limit as in Bitcoin and Ethereum pre–London hard fork and dynamically adjust when demand fluctuates vs. one where price controls such as the EIP-1559 base fee are imposed and dynamically adjust when demand fluctuates? What determines the shape and parameters of such transaction fees? These questions will be the focus of subsequent sections in this paper.

3 Model

3.1 Environment

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 records 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. The monopolistic validator, using quantities x_i of a finite number of resources $i \in [1, N]$ (e.g., computation, bandwidth) selects 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.⁷

Validators: We model a setting with multiple validators; however, when a validator is selected to produce a block, she acts as a monopolist. This fact reflects the temporary monopoly power inherent in blockchain consensus mechanisms with a single block proposer. The validator uses a bundle of computational resources $x \in \mathbb{R}^N_+$ at exogenous per-unit

⁷See the related discussion in Buterin (2018) for the case of the Ethereum blockchain.

resource rates $g_x \in \mathbb{R}^N_{++}$ to produce a block of size $q \leq q^{\max}$, as given by

$$q = \sum_{i=1}^{N} g_{x_i} x_i \tag{2}$$

Each transaction consumes specific amounts of various resources (e.g., computation, storage, bandwidth). The total block size is the weighted sum of these resource usages across all included transactions.

The validator incurs a cost $C(q) = c(x_1, \ldots, x_N)$, based on the resources used to produce a block. These costs encompass validation operation costs and other costs associated with accessing and modifying the blockchain's state.⁸ The validator faces uncertainty in her marginal costs due to factors such as fluctuations in energy prices, hardware performance variability, and network conditions. Therefore, denote by $C(q; \eta)$ the validator costs of production and let η be a random variable with finite mean that represents the validator's marginal cost uncertainty.

Users: Users, denoted by $j \in [0, 1]$, are atomistic. We model user arrivals between two consecutive blocks, B_t, B_{t+1} as a Poisson process X with rate ζ . To facilitate comparative analysis under different block size limits, we introduce the notation $\zeta \equiv \lambda q^{\max}$, where λ reflects the arrival rate per unit of maximum block size. In our subsequent analysis, we consider λ a random variable capturing demand uncertainty. For simplicity, we assume that users leave the pool if their transaction is not included in the next block, only to return according to the arrival process.⁹

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.

⁸From the blockchain designer's viewpoint, costs might also include block propagation delays due to large blocks and other societal costs.

⁹Leonardos et al. (2021) confirm that this assumption does not significantly impact the dynamics of transaction fees, which are the focus of our analysis. In his research, Nisan (2023) accounts for residual demand in the mempool and finds transaction fee dynamics similar to those in Leonardos et al. (2022).

Establishing a Demand Curve: Given that the fee policies under consideration are incentive-compatible in dominant strategies, as demonstrated by Roughgarden (2021), it is reasonable to assume that users bid their accurate 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 is 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 $p = (\bar{F})^{-1} \left(\frac{q}{\Psi}\right)$, where the price elasticity of demand for the marginal user equals the tail ratio $\frac{pf(p)}{1-F(p)}$.

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

$$\frac{p}{p_m} = \left(\frac{q}{\Psi}\right)^{-\frac{1}{\varepsilon}} \tag{3}$$

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 C.1 for the proof.

In Section 3.3, assume that users' valuations for transaction inclusion follow a Pareto distribution. This choice is motivated by empirical observations in blockchain networks, where transaction fees often exhibit heavy-tailed distributions due to a small number of users willing to pay significantly higher fees for urgent processing. For instance, data from the Ethereum network shows that transaction fees follow a heavy-tailed distribution (Guo et al., 2019, See), justifying the Pareto assumption.

While the Pareto distribution is a simplification, it allows for analytical tractability. 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. The demand curve from the Pareto distribution will be helpful when I consider uncertainty in the user arrival rate λ , leading to uncertainty in the demand shifter Ψ . By modeling λ as a random variable, we incorporate demand uncertainty into our analysis and can study the effect of stochastic demand on the optimal fee policies. For realization $\lambda > 1$, 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. Given this context and considering the uncertainties in both the cost $C(q; \eta)$ and demand Ψ , we study the conditions under which a blockchain protocol designer would find introducing price or quantity controls beneficial.

3.2 Tension Between Protocol Objectives and The Validator's Incentives

Given that each validator has monopoly power during their assigned turn, we model the interaction between a single validator and the protocol designer. This allows us to analyze the strategic choices and conflicts in setting fee policies.

3.2.1 Protocol Designer's Preference for Maximum Capacity Utilization

The designer selects the control variable (price or quantity) before uncertainties are realized, aiming to maximize expected social welfare while considering the strategic responses of validators and users.

If the protocol designer has a preference over the block size, then she will prefer quantity controls over price controls as the former insulates this variable from demand fluctuations. To illuminate the potential conflict between the protocol designer's objectives and the miners' incentives, consider the following examples of objective S(q) for the protocol designer.

Example 2. (Social Welfare) Let u be an increasing function representing the utility of a

representative user on the blockchain. The social welfare function S(q) can be defined as

$$S(q) = \begin{cases} u(q) & \text{if } q \le q^{\max} \\ -\infty & \text{if } q > q^{\max} \end{cases}$$

$$\tag{4}$$

This example reflects a protocol designer's concern for maximizing user utility, an objective that can diverge from the profit-maximizing motive of the validator.

Claim 3. Suppose the protocol designer aims to choose either a base fee (price control) or a block size limit (quantity control) to maximize social welfare as defined in Equation (4). If 0 is in the support of Ψ , then the protocol designer would prefer to have blocks at full capacity. The only price that aligns with the designer's objective is zero.

Proof. If the protocol designer can choose the block size, then she can guarantee a maximal utility at q^{limit} . With a price choice, however, we have $q = \max\{\Psi\left(\frac{p_m}{p}\right)^{\varepsilon}, q^{limit}\}$. This means that for any price p, there is a deadweight loss when demand intensity is low, $\Psi \leq \left(\frac{p}{p_m}\right)^{\varepsilon} q^{limit}$, which makes the expectation lower when 0 is in the support of Ψ . Therefore, the only transaction fee that ensures the designer achieves the maximum block utilization under all demand scenarios is a zero price. All users are willing to submit transactions at this price, and the block is filled regardless of demand fluctuations.

Setting a very low price eliminates deadweight loss from less than full blocks in lowdemand conditions. This claim highlights a point of conflict with the validator, who may have different objectives, such as profit maximization.

Example 4. (Technological Block Size Targets) The protocol designer may have a target block size q^{target} to optimize network performance, security, or decentralization. For example, setting a target block size can help manage the trade-off between transaction throughput and the time it takes for blocks to propagate through the network. She might also allow some degree of deviation from the target block size, which we can represent as $S(q) = l\{q - q^{target}\}$, where

l stands for some loss function. For instance, a square loss function could be used to penalize deviations from the target:

$$S(q) = -\left(q - q^{target}\right)^2 \tag{5}$$

Claim 5. Under a technological block size target, the protocol designer prefers quantity controls over price controls. In particular, for the square loss function, the expected block size under optimal price controls $\mathbb{E}[q]$ is less than the target q^{target} , and the loss is given by

$$\mathbb{E}[\left(q(\Psi) - q^{target}\right)^2] = \frac{Var(\Psi)}{\mathbb{E}[\Psi^2]} (q^{target})^2 \tag{6}$$

Proof. Refer to Appendix C.2.

This loss quantifies the variance in the block size, relative to the target, that arises because of the fluctuations in demand Ψ when the price is optimally chosen ex-ante. This highlights that the loss for the protocol designer is more significant when the demand shifter has a high variance.

3.2.2 The Validator's Monopoly Power

While validators compete ex-ante to be selected for block production, the selected validator effectively becomes a monopolist for that block. This model captures the essential features of the validator market, where competition determines which validator gets to produce a block, but monopoly power in, so a monopolist validator's profit Π consists of the block reward R, the transaction fee revenue pq, and the cost $C(q; \eta)$ associated with producing a

block of size q.¹⁰

$$\Pi = \mathbb{E}\left[R + pq - C(q;\eta)\right].$$
(7)

Application developers and users design transactions to be as efficient as possible to minimize costs, i.e., the cost to the validator is¹¹

$$C(q; g_x, \eta) = \min_{x \in \mathbb{R}^N_+} c(x_1, \dots, x_N; \eta)$$
(8)

subject to (2). The bundle x can be interpreted as the various resources that constitute a user transaction, such as bandwidth and computational operations. Let us consider that c is homogeneous of degree 1 in x.¹² We then have

$$C(q; g_x, \eta) = \Gamma(\eta, g_x)q \tag{10}$$

The expression for the marginal cost Γ is derived in Appendix C.3.

A monopolist validator would then maximize profits in equation (7). In contrast to the protocol designer, who systematically prefers to set quantities ex-ante, the monopolist validator might choose either the price-setting or the quantity-setting regime, depending on

$$c(x_1, \dots, x_N; \eta) = \eta \prod_{i=1}^N x_i^{\varepsilon_i} \quad \text{such that } \sum_{i=1}^N \varepsilon_i = 1$$
(9)

¹⁰For simplicity, I abstract from the fact that not all fees collected go to the validator, and the monopolist validator can capture extra revenue from MEV. In Ethereum's EIP-1559 mechanism, the base fee is burned rather than paid to validators. Hence, validators only receive the priority tip component of users' payments. This burning mechanism is critical: without it, validators could insert empty transactions to inflate future block fees at no net cost, undermining the price-adjustment process. In addition to transaction fees, validators can capture MEV by reordering or selectively including transactions to exploit on-chain arbitrage opportunities (e.g., DEX front-running). On Ethereum, a sophisticated ecosystem of searchers and builders helps validators systematically capture such profits.

¹¹In blockchain networks, transaction fees are often proportional to the computational and storage resources required. Therefore, developers and users are incentivized to optimize their transactions to reduce costs.

¹²A typical example is the Cobb–Douglas cost function

the degree of demand uncertainty.

3.3 Bargaining Problem

I model the resolution of the potential conflict between the protocol designer and the validator by assuming that the price- or quantity-setting regime is determined by Nash bargaining over the protocol profits.¹³ Before uncertainty is realized, the protocol designer and the validator commit ex-ante to a fixed base fee (price-setting) or fixed block space (quantity-setting). In practice, such a process is implemented by a set of rules in the fee policies of the blockchain to which validators and blockchain designers agree. Even though the validator has ultimate control over which transactions, of what s,ize and at what pr, ice can be included in a block, sophisticated mechanisms can be encoded to enforce prices and quantities. Protocol revenue can be diverted in part to the protocol treasury, burned, or rebated to users to implement the fee policies; this design feature is beyond the scope of this paper.¹⁴ The planner's objective balances social welfare and technological considerations while ensuring that the validator has enough excess profits to be willing to provide her services.

The following objective captures the planner's problem:

$$\mathcal{V} = \mathbb{E}\left[S(q)\right]^{1-\beta} \mathbb{E}\left[pq - C(q;\eta)\right]^{\beta}$$
(11)

In this objective, the parameter $\beta \in [0; 1]$ captures the bargaining power of the validator. $\beta = 1$ means that the outcome maximizes the excess profits of the validator,¹⁵ while $\beta = 0$ means that the outcome optimizes for the protocol designer's objective captured by the

¹³Even without a formal bargaining game, one could analyze the fee mechanism preferences of the interim monopolist validator. The main takeaways on price vs. quantity-setting can stand alone as a simpler exercise. ¹⁴See Roughgarden (2021).

¹⁵For analytical simplicity, we focus on transaction fees as the primary variable component of the validator's profit. While the block reward is important, it is often a fixed amount per block and does not directly influence the choice between price and quantity controls in our model. Another interpretation of our analysis is to look at the long run, where blockchain rewards are often programmed to decay exponentially to zero in order to avoid overinflation. An analysis involving the full profit of the validator $\mathbb{E}[R + p\phi q - C(q; \eta)]$ where ϕ is a stochastic share of priority fee as in Jermann (2023) would yield similar insights.

function S(q). The Nash bargaining problem is defined over expected utilities, taking into account uncertainty in deman λ , an marginal costs η . We assume that both parties have common knowledge of the distributions of these random variables.

Price Controls: Assuming that the block space demand follows the isoelastic demand curve derived in (3), 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 problem of setting optimal prices then becomes

$$\mathcal{V}^{p} = \max_{p \in \mathbb{R}_{+}} \mathbb{E} \left[S \left(\Psi \left(\frac{p}{p_{m}} \right)^{-\varepsilon} \right) \right]^{1-\beta} \mathbb{E} \left[(p-\Gamma) \times \Psi \left(\frac{p}{p_{m}} \right)^{-\varepsilon} \right]^{\beta}$$
(12)

Quantity Controls: If the block size is set at q, the transaction is included in the block at the price that clears markets ex-post: $p = p_m \left(\frac{q}{\Psi}\right)^{-\frac{1}{\varepsilon}}$. Then, the value of setting the optimal block space is

$$\mathcal{V}^{q} = \max_{q \in \mathbb{R}_{+}} \mathbb{E} \left[S(q) \right]^{1-\beta} \mathbb{E} \left[\left(p_{m} \left(\frac{q}{\Psi} \right)^{-\frac{1}{\varepsilon}} - \Gamma \right) \times q \right]^{\beta}$$
(13)

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

$$\Delta^{\log} = \log \mathcal{V}^p - \log \mathcal{V}^q \tag{14}$$

To derive some insight into the choice between price and quantity controls, consider a blockchain designer's objective that accounts for social welfare with utility $S(q) = u(q) = q^{\nu}$ for $\nu > 0$. 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 given an arbitrary bargaining power β for the validator.

Proposition 6. Assume that (Ψ, η) follows a joint log-normal distribution with log-variances

 $\sigma_{\Psi}^2, \sigma_{\eta}^2$ and log-covariance $\sigma_{\Psi,\Gamma}$. Then, for any $\beta > 0$, the relative value of price controls over quantity controls is given by:

$$\Delta^{\log} = \frac{1}{2} \left(\left(\hat{\nu} - \frac{\bar{\nu}}{\varepsilon} \right) \sigma_{\Psi}^2 - 2(\varepsilon \nu - \beta) \sigma_{\Psi,\Gamma} \right)$$
(15)

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

Proof. Refer to Appendix C.3.

We can interpret this equation first by looking at the limit for $\beta \to 1, \nu \to 1$, which gives

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

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 positive correlation between marginal costs and demand disfavors price controls, as this would lead to the production of larger blocks when marginal costs are high. 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 indicate that price-setting has an advantage as long as demand is relatively elastic, i.e., $\varepsilon > \bar{\nu}/\hat{\nu}$, and the validator has enough bargaining power, $\beta > \varepsilon \nu$.

Taking Stock: The above result helps us understand the differences in the policies that determine fees and block space in Bitcoin and Ethereum. We can see that both blockchains are instances where price elasticity of demand is high $\varepsilon > 1$ and demand uncertainty is significant $\sigma_{\Psi}^2 \gg 0$.

The background provided in Section 2.1 suggests that Ethereum can be viewed through the lens of this model as an instance where the bargaining power of the validator is high $\beta \to 1$

(or is considered high by the protocol designers). Thus, our use of 'high bargaining power' for validators should be interpreted in the sense of their economic incentive alignment in the transaction-fee market, rather than an overriding governance control of the protocol overall. In addition, since the proof-of-stake upgrade, the marginal cost for Ethereum validators has been virtually constant, so $\sigma_{\Psi,\Gamma} \approx 0$. Proposition 6 states that, in this case, price-setting is the most favorable policy. This explains the adoption of the EIP-1559 price-setting policy for the Ethereum blockchain. Furthermore, Figure 1 shows that EIP-1559 has performed better since Ethereum's proof-of-stake upgrade, consistent with the correlation between the validator's marginal costs and demand being lower relative to miners' in a proof-of-work protocol.

Bitcoin is distinguished as the first blockchain with an anonymous founder, and its core developers strongly emphasize decentralization and censorship resistance. As discussed in Section 2.1, most proposals to change Bitcoin's fees and block size policies have failed. Bitcoin can, therefore, be viewed through the lens of our model as an instance where the bargaining power of validators (here miners) is low $\beta \rightarrow 0$ (or is not attributed enough importance by protocol designers, given the low block space utilization rate still observed). In addition, it is safe to assume that under the proof-of-work protocol, miners' marginal cost is largely positively correlated with demand $\sigma_{\Psi,\Gamma} \gg 0$. Proposition 6 states that, in this case, price-setting is less effective.

In general, the key determinants of the advantage of price-setting in blockchains are the validator's bargaining power, the elasticity of demand, the validator's uncertainty about demand, and marginal costs. This insight will guide us in Section 4 in studying the implications of my results for Ethereum transaction fees.

3.4 Effect of Cryptocurrency Price Fluctuations

Cryptocurrency prices can be volatile, while users value transaction processing services in dollars or real terms. This section examines the impact of cryptocurrency price volatility on the choice between price and quantity controls. Let P denote the exchange rate between 1 USD and the cryptocurrency (equivalently, the inverse of the cryptocurrency price 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 cryptocurrency price and fiat currency's value. Users pay transaction fees in the cryptocurrency at a nominal price p, implying that the dollar (equiv. real) value 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 6 in this context:

Proposition 7. Suppose (Ψ, η, P) is jointly log-normal distributed. Then, the relative value of price adjustments over quantity adjustments is

$$\Delta^{\log} = \frac{1}{2} (\varepsilon - 1) \left(\frac{1}{\varepsilon} \sigma_{\Psi}^2 - 2\sigma_{\Psi,\Gamma} - \varepsilon \sigma_P^2 - 2\varepsilon \sigma_{P,\Gamma} \right)$$
(16)

In addition to the findings of Proposition 6, 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.

4 Implications for Ethereum Transaction Fees

In this section, I explore how my results can inform the design of Ethereum fee policies. Proposition 6 implies that, under the conditions of (15), any fixed block size target q^{target} can be improved upon by setting prices ex-ante and letting quantities adjust. As discussed in Section 2, Ethereum's fee policies follow such an approach to guarantee blocks of average size q^{target} . How can prices be set iteratively and heuristically? Below, I define a family of simple TFMs, including EIP-1559, Ethereum's TFM. I 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 8.** (ADT-TFM) A TFM 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

$$\frac{p_{t+1}}{p_t} = g\left(\frac{q_t - q^{target}}{q^{target}}\right) \tag{17}$$

Example 9. (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 $\varepsilon(q^{target})$ that is assumed to be fixed and uncertainty in demand represented by λ . In this context, a TFM is considered 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 worstcase scenario. The following proposition determines the shape and slope of the optimal ADT-TFM.

Proposition 10. Suppose that 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

$$p_{t+1} = p_t \exp\left(\frac{1}{\varepsilon(q^{target})} \frac{q_t - q^{target}}{q^{target}}\right)$$
(18)

The intuition of the proof in Appendix C.4 goes as follows. Let f denote the adjustment function of the optimal TFM and $p(q_t)$ the price that matches demand at block size q_t . The base fees then satisfy

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

Near q^{target} , we have:

$$g(q_t) \xrightarrow[q^{target}]{} \frac{q'(p(q^{target}))p(q^{target})}{q^{target}} = \varepsilon(q^{target})$$
(20)

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

$$\ln p_{t+1} - \ln p_t \sim \frac{p(q_t) - p(q^{target})}{p(q^{target})}$$
(21)

Let x denote this price growth. From (19), we obtain $x = \ln (g (\varepsilon(q^{target}) \cdot x)) + o(x)$ 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 (20) 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 Leonardos et al. (2022), the adjustment rate remains the inverse of the price elasticity of demand. This elasticity, however, captures the price elasticity of inclusion of a transaction in the next block. Given the user interface shown in Figure 2, users are less likely to change their price in a short period, so this elasticity will be larger than the price elasticity of inclusion of a transaction within 5–10 minutes, for which users can substitute waiting in the mempool.

Adjustment Parameter for Ethereum: I now approximate Ethereum's adjustment rate through the lens of my analysis. I use Ethereum data because of its widespread availability, the simplicity of its ADT-TFM, and the global usage of its blockchain.

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) are also analyzed.¹⁶

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. I do not have random variation in supply to trace the demand curve here. Instead, informed by Lemma 1, I calculate the Pareto tail of transaction gas prices—as the nearest proxy for user valuations—for blocks of size around q^{target} to obtain the price $\varepsilon(q^{target})$. This approach has limitations that I discuss in Appendix D.

My favorite number is a Pareto coefficient of 12.62 for blocks of size within $\pm 5\%$ of the block size target (over 7252 blocks and 12965717 transactions), which yields an optimal adjustment rate of 7.92%. This is significantly below Ethereum's current adjustment rate of 12.5%. This result aligns with the finding of Leonardos et al. (2022), who simulate the dynamic system of EIP-1559 and find stability around the target block size only for adjustment parameters below 8%. Several alternative choices in Appendix D yield a range between 6.14% and 11% for the Ethereum optimal adjustment rate. My contribution clarifies that the adjustment rate encapsulates the economic concept of inverse price elasticity of demand, which must be measured or approximated on-chain.

¹⁶The results are consistent when I separately sample 100,000 blocks from before and after the merge.

Target Block Size and MEV: Now, let us determine the block size target that aligns with the optimal target of a monopolistic validator. The goal of such a block size target is to make the blockchain immune to a simple form of MEV—that is, to prevent the validator from including her value-extracting transactions while reducing the effective supply available to users. The following definition makes this notion explicit.

Definition 11. (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 12. Given any isoelastic demand curve, the target block size aligning with the monopolist validator's optimal target block size is expressed as

$$\frac{q^{target}}{q^{\max}} = \left(\frac{\varepsilon}{\varepsilon - 1}\right)^{-\varepsilon} \mathbb{E}\left[\lambda^{\frac{1}{\varepsilon}}\right]^{\varepsilon}$$
(22)

Proof. Refer to Appendix C.5.

This proposition states that the maximum block size should have an adequate buffer above the block size; otherwise, if the block size target is too high relative to the maximum block size, the monopolist validator will have incentives to fill the block up to the target size with her value-extracting transactions. In particular, if q^{max} is adjusted to coincide with user demand under average demand conditions, then $\frac{q^{\text{max}}}{q^{target}} > e$. For the price elasticity found above, I find that the ratio of the maximum block size to the block size target should be greater than 2.83 for Ethereum, which is currently set at 2.

5 Conclusion

In this paper, I have studied different blockchain fee policies to allocate block space efficiently when various sources of uncertainty exist. Namely, I have compared the quantity-setting regime used by Bitcoin and the price-setting regime used by Ethereum under different sources of uncertainty. Price controls are optimal in an environment with low validator bargaining power, high price elasticity of demand, high demand uncertainty, and high marginal costs during the validation process. In addition, a high variance of the cryptocurrency price mitigates the advantages of price controls. These insights help explain the difference between Bitcoin's and Ethereum's fee policies. Next, I have applied these insights to a family of fee policies to which the method used by Ethereum belongs. Under mild assumptions on the shape of the demand for block space, I find that a crucial parameter of the mechanism, the rate of adjustment of prices, equals the inverse price elasticity of demand for inclusion in the next block. Some calculations using Ethereum transaction data suggest that the rate at which Ethereum's fees change is faster than optimal.

Henceforth, understanding the economics of multidimensional fees is crucial, as transactions use various types of resources. The question of designing fee policies for durable resources is growing in importance as blockchain states expand significantly over time, particularly due to resources such as storage. Additionally, as some blockchain states may experience less congestion than others, it is imperative to explore fee policies that differentiate pricing based on varying demand levels rather than solely pricing block space.¹⁷ While a comprehensive examination of these complex issues lies beyond the scope of this paper, they offer promising opportunities for future research and further refinement of my analysis.

References

Aiello, Darren, Scott R Baker, Tetyana Balyuk, Marco Di Maggio, Mark J Johnson, and Jason D Kotter, "The effects of cryptocurrency wealth on household consumption and

¹⁷Ndiaye (2024) studies state-contingent fee policies.

investment," Technical Report, National Bureau of Economic Research 2023. 7

- Akbarpour, Mohammad and Shengwu Li, "Credible auctions: A trilemma," *Econometrica*, 2020, 88 (2), 425–467. 8
- Böhme, Rainer, Nicolas Christin, Benjamin Edelman, and Tyler Moore, "Bitcoin: Economics, technology, and governance," *Journal of economic Perspectives*, 2015, 29 (2), 213–238.
- Budish, Eric, "Trust at Scale: The Economic Limits of Cryptocurrencies and Blockchains^{*†}," *The Quarterly Journal of Economics*, 10 2024, p. qjae033. 2
- Buterin, Vitalik, "Blockchain resource pricing," URL: https://ethresear. ch/uploads/default/original X, 2018, 2. 9, 16
- $_$, "The Limits to Blockchain Scalability," 5 2021. 2
- _ et al., "Ethereum white paper," *GitHub repository*, 2013, 1, 22–23. 9
- Catalini, Christian and Joshua S. Gans, "Some Simple Economics of the Blockchain," Commun. ACM, jun 2020, 63 (7), 80–90. 7
- Chung, Hao and Elaine Shi, "Foundations of transaction fee mechanism design," in "Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)" SIAM 2023, pp. 3856–3899. 8
- Cong, Lin William, Zhiguo He, and Jiasun Li, "Decentralized mining in centralized pools," The Review of Financial Studies, 2021, 34 (3), 1191–1235. 3
- Daian, Philip, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels, "Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability," in "2020 IEEE Symposium on Security and Privacy (SP)" IEEE 2020, pp. 910–927. 7
- Delgado-Mohatar, Oscar, Marta Felis-Rota, and Carlos Fernández-Herraiz, "The Bitcoin mining breakdown: Is mining still profitable?," *Economics Letters*, 2019, 184, 108492.
- Easley, David, Maureen O'Hara, and Soumya Basu, "From mining to markets: The evolution of bitcoin transaction fees," *Journal of Financial Economics*, 2019, 134 (1), 91–109. 7, 10
- Ferreira, Matheus VX, Daniel J Moroz, David C Parkes, and Mitchell Stern, "Dynamic posted-price mechanisms for the blockchain transaction-fee market," in "Proceedings of the 3rd ACM Conference on Advances in Financial Technologies" 2021, pp. 86–99. 8
- Flynn, Joel P, George Nikolakoudis, and Karthik A Sastry, "A Theory of Supply Function Choice and Aggregate Supply," 2023. 8
- Fracassi, Cesare, Moazzam Khoja, and Fabian Schär, "Decentralized Crypto Governance? Transparency and Concentration in Ethereum Decision-Making," January 2024. 5
- Guo, Dongchao, Jiaqing Dong, and Kai Wang, "Graph structure and statistical properties of Ethereum transaction relationships," *Information Sciences*, 2019, 492, 58–71. 18
- Hinzen, Franz J, Kose John, and Fahad Saleh, "Bitcoin's limited adoption problem," Journal of Financial Economics, 2022, 144 (2), 347–369. 7
- Huberman, Gur, Jacob D Leshno, and Ciamac Moallemi, "Monopoly without a monopolist: An economic analysis of the bitcoin payment system," *The Review of Economic Studies*, 2021, 88

(6), 3011 - 3040. 7, 13

- Jermann, Urban J., "A Macro Finance Model for Proof-of-Stake Ethereum," January 2023. 7, 23
- Klemperer, Paul D and Margaret A Meyer, "Supply function equilibria in oligopoly under uncertainty," *Econometrica: Journal of the Econometric Society*, 1989, pp. 1243–1277. 8
- Kogan, Shimon, Igor Makarov, Marina Niessner, and Antoinette Schoar, "Are cryptos different? evidence from retail trading," Journal of Financial Economics, 2024, 159, 103897. 7
- Lavi, Ron, Or Sattath, and Aviv Zohar, "Redesigning Bitcoin's Fee Market," ACM Trans. Econ. Comput., may 2022, 10 (1). 9, 46
- Lehar, Alfred and Christine A Parlour, "Miner collusion and the bitcoin protocol," Available at SSRN 3559894, 2020. 7, 11, 13, 37
- Leonardos, Stefanos, Barnabé Monnot, Daniël Reijsbergen, Efstratios Skoulakis, and Georgios Piliouras, "Dynamical analysis of the eip-1559 ethereum fee market," in "Proceedings of the 3rd ACM Conference on Advances in Financial Technologies" 2021, pp. 114–126. 9, 17
- _ , Daniël Reijsbergen, Daniël Reijsbergen, Barnabé Monnot, and Georgios Piliouras, "Optimality Despite Chaos in Fee Markets," arXiv preprint arXiv:2212.07175, 2022. 9, 17, 29, 30, 48
- Malik, Nikhil, Manmohan Aseri, Param Vir Singh, and Kannan Srinivasan, "Why bitcoin will fail to scale?," *Management Science*, 2022, 68 (10), 7323–7349. 7
- Nakamoto, Satoshi, "Bitcoin: A peer-to-peer electronic cash system," Decentralized business review, 2008. 9
- Ndiaye, Abdoulaye, "Blockchain Price vs. Quantity Controls," Working Paper, 2023. 8
- __, "Relative Pricing and Efficient Allocation in Blockchains," Research Paper, NYU Stern School of Business 11 2024. Available at SSRN. 32
- Nisan, Noam, "Serial Monopoly on Blockchains," 2023. 9, 17, 46
- Rehman, Mobeen Ur and Sang Hoon Kang, "A time-frequency comovement and causality relationship between Bitcoin hashrate and energy commodity markets," *Global Finance Journal*, 2021, 49, 100576. 7
- Reis, Ricardo, "Inattentive producers," The Review of Economic Studies, 2006, 73 (3), 793–821. 8
- Roughgarden, Tim, "Transaction fee mechanism design for the Ethereum blockchain: An economic analysis of EIP-1559," arXiv preprint arXiv:2012.00854, 2020. 4
- ____, "Transaction fee mechanism design," ACM SIGecom Exchanges, 2021, 19 (1), 52–55. 4, 8, 15, 18, 23
- Saleh, Fahad, "Blockchain without Waste: Proof-of-Stake," The Review of Financial Studies, 07 2020, 34 (3), 1156–1190. 3
- SEC, "Statement on the Approval of Spot Bitcoin Exchange-Traded Products," https://www.sec. gov/newsroom/speeches-statements/gensler-statement-spot-bitcoin-011023 2024. Accessed: 2024-10-22. 2
- Weitzman, Martin L, "Prices vs. Quantities," The Review of Economic Studies, 1974, 41 (4), 477–491. 4, 8, 9

FOR ONLINE PUBLICATION

Appendix

A History of Block Size Changes in Ethereum

Dates	Changes in Gas	Reasons and Context		
	Limit			
March 4th, 2016	3,141,592 to 4,712,388	Max gas increased by 1.5x, and minimum gas price		
		reduced from 50 to 20 gwei (a denomination of the		
		Ethereum cryptocurrency) to improve affordabil-		
		ity and throughput.		
September 22nd,	4,712,388 to 1,000,000	Network experienced a DDoS attack, leading to a		
2016		drastic reduction in block size for security mea-		
		sures.		
September 22nd,	1,000,000 to 1,500,000	DDoS attack was partially mitigated, allowing a		
2016		limited increase in block size.		
October 15th,	1,500,000 to 500,000	Block size was decreased as a precautionary step		
2016		before implementation of a hard fork to address		
		security vulnerabilities.		
October 19th,	500,000 to 3,000,000	Hard fork successfully implemented, security issues		
2016		resolved, leading to a substantial increase in block		
		size.		
October 20th,	3,000,000 to 1,500,000	Another DDoS attack occurred, prompting a re-		
2016		duction in block size to safeguard the network.		

October 23rd,	1,500,000 to 2,000,000	Successful mitigation of attacks led to increased
2016		block size, signaling a return to stability.
November 24th,	2,000,000 to 3,300,000	Continued stability and growing user base justified
2016		another increase in block size.
December 5th,	3,300,000 to 4,000,000	Network achieved consistent stability, allowing a
2016		further increase in block size.
June 3rd, 2017	4,000,000 to 4,712,388	Increase in network activity and user engagement
		necessitated a rise in block size.
June 29th, 2017	4,712,388 to 6,283,184	New target gas limit set at 4,700,000 to better
		match growing ecosystem demands.
December 10th,	6,283,184 to 8,000,000	Rise in transaction volume driven by Cryptokitties
2017		NFT required an increase in block size.
September 19th,	8,000,000 to	Tether's migration to the Ethereum blockchain led
2019	10,000,000	to increased transaction demands, prompting a
		block size increase.
June 19th, 2020	10,000,000 to	Miner consensus to increase block size was reached,
	12,000,000	supporting the growing Ethereum ecosystem.
July 25th, 2020	12,000,000 to	Minor increase following another round of miner
	12,500,000	agreements aimed at fine-tuning network perfor-
		mance.
April 21st, 2021	12,500,000 to	Berlin hard fork led to efficiency improvements,
	15,000,000	enabling a substantial increase in block size.
August 5th,	15,000,000 to	London hard fork brought about major improve-
2021	30,000,000	ments in transaction fee predictability and network
		efficiency, justifying a block size doubling.

B Capacity Utilization

B.1 Demand Factors

For my analysis, I identified the number of active addresses, token prices, and transaction fees as potential determinants of block space demand. For Bitcoin (BTC) only, I also added the size of the mempool, i.e., the pool of transactions that has yet to be confirmed and included in a block.

For BTC, the top graph of Figure A1 shows that the number of active addresses has a strong correlation with the block utilization rate, as it acts as a good proxy for demand. Similarly, transaction fees show a tight link with the utilization rate, capturing the fact that users must offer a higher fee to be picked by miners in periods of block congestion. As explained in Lehar and Parlour (2020), mempool size is not a perfect predictor of block utilization, as miners can leave blocks empty, even when there are transactions waiting, to extract more profits. Finally, the BTC price has a loose relationship with the utilization rate, possibly because of its high volatility due to speculation.

For Ethereum (ETH), the bottom graph of Figure A1 shows that, before the London fork, the number of active addresses once again had the highest correlation with block utilization rates, representing best users' demand for transactions. Both transaction fees and the ETH price have stronger correlations with block fullness than their BTC counterparts. After the implementation of EIP-1559, gas levels stabilized consistently at the target level, suggesting that the implementation of the new fee system successfully achieved its objective. As a consequence, the block utilization rate became uncorrelated with its previously valid demandside determinants.

B.2 Supply Factors

For my analysis, I identified the computing power of miners and the concentration rate of mining pools as potential determinants of block space supply. For ETH only, I considered the



(a) Bitcoin network demand factors per day against utilization rate



(b) Ethereum network demand factors per day against utilization rate

Figure A1: Time series of Bitcoin and Ethereum demand factors



(b) Ethereum network supply factors against utilization rateFigure A1: Time series of Bitcoin and Ethereum supply factors

total amount of ETH staked and the concentration of validator pools after the proof-of-stake update.

For BTC, the top graph of Figure A1 shows that the amount of computing power invested into mining does not affect the block utilization rate. I measure it using the hash rate, which measures how many guesses are made per second to solve the code to mine the next block. Regarding mining pool concentration, computed as the Herfindahl–Hirschman index (the sum of the squares of individual market shares) from the shares of blocks mined in a onemonth period, the graph seems to support the hypothesis that miners exercise their market power by leaving blocks less than full, as periods of higher HHI are usually accompanied by utilization rates below 100%.

For ETH, the bottom graph of Figure A1 shows that, under both proof of work and proof of stake, the computing power (or the amount of ETH staked, of which 32 ETH are required to activate a validator software), does not affect the utilization rate, similarly to what we observe for BTC. In contrast to the BTC case, however, the pool concentration does not affect the block utilization rate either. The HHI measures (from the share of blocks mined in a one-month period under proof of work and from the share of ETH staked over a one-month period under proof of stake) are not correlated with the utilization rate, which is particularly noticeable after the stabilization of the utilization rate post–London fork.

C Analytic Proofs

C.1 Proof of Lemma 1

At a price p, demand for block space is the measure of users' willingness to pay p for transaction inclusion, i.e., $\lambda q^{\max} \bar{F}(p)$. This yields the demand curve $p = (\bar{F})^{-1} \left(\frac{q}{\lambda q^{\max}}\right)$. 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 $q = \lambda q^{\max} \bar{F}(p)$ and $dq/dp = -\lambda q^{\max} f(p)$, the demand elasticity is $\frac{pf(p)}{1-F(p)}$.

When F is a Pareto distribution with scale p_m and shape $\alpha,$

$$\bar{F}(p) = Pr(v > p) = \begin{cases} \left(\frac{p_m}{p}\right)^{\alpha} & \text{for } p \ge p_m \\ 1 & \text{for } p < p_m \end{cases}$$
(23)

so that, above the minimum price p_m , demand is

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

Therefore, we obtain

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

C.2 Proof of Claim 5

The optimal ex-ante quantity choice is just $q = q^{target}$ for which the loss is zero. The optimal ex-ante price choice solves

$$\min_{p \in \mathbb{R}^+} E\left[\left(\Psi \cdot \left(\frac{p}{p_m}\right)^{-\varepsilon} - q^{target} \right)^2 \right]$$
(24)

Denote $x \equiv \left(\frac{p}{p_m}\right)^{-\varepsilon}$. The first-order condition for the choice of p (resp. x) in (24) is

$$x = \frac{\mathbb{E}[\Psi]q^{target}}{\mathbb{E}[\Psi^2]} \tag{25}$$

The expected block size is then

$$\mathbb{E}[\Psi x] = \frac{\mathbb{E}[\Psi]^2 q^{target}}{\mathbb{E}[\Psi^2]} \le q^{target}$$
(26)

Replacing (25) in the value of the loss function (24) yields

$$\frac{\mathbb{E}[\Psi^2] - \mathbb{E}[\Psi]^2}{\mathbb{E}[\Psi^2]} (q^{target})^2 = \frac{\operatorname{Var}(\Psi)}{\mathbb{E}[\Psi^2]} (q^{target})^2$$

C.3 Proof of Proposition 6

We first find the expression of the marginal cost Γ in (10). The first-order condition is

$$c_i(x_1, \dots, x_N; \eta) = \gamma g_{x_i} \tag{27}$$

where γ is the Lagrangian of constraint (2). Since c is homogeneous of degree 1, we have

$$c(x_1, \dots, x_N; \eta) = \sum_{i=1}^{N} c_i(x_1, \dots, x_N; \eta) x_i = \gamma \sum_{i=1}^{N} g_{x_i} x_i = \gamma q$$
(28)

Therefore, $\Gamma = \gamma$. Evaluating at q = 1 yields

$$\Gamma(\eta, g_x) = c(x_1(1, g_x, \eta), \dots, x_N(1, g_x, \eta); \eta)$$
(29)

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

$$\frac{\varepsilon(1-\beta)\nu}{p^*} + \frac{\varepsilon\beta}{p^*} = \frac{\beta\mathbb{E}[\Psi]}{p^*\mathbb{E}[\Psi] - \mathbb{E}[\Gamma\Psi]}$$
(30)

$$\frac{(1-\beta)\nu}{q^*} + \frac{\beta}{q^*} = \frac{\beta/\varepsilon p_m(q^*)^{-1-1/\varepsilon} \mathbb{E}[\Psi^{1/\varepsilon}]}{p_m(q^*)^{-1/\varepsilon} \mathbb{E}[\Psi^{1/\varepsilon}] - \mathbb{E}[\Gamma]}$$
(31)

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

$$p^* = \frac{\varepsilon \bar{\nu}}{\varepsilon \bar{\nu} - \beta} \frac{\mathbb{E}[\Gamma \Psi]}{\mathbb{E}[\Psi]}$$
(32)

$$q^* = \left(\frac{\varepsilon\bar{\nu}}{\varepsilon\bar{\nu} - \beta} \frac{1}{p_m} \frac{\mathbb{E}[\Gamma]}{\mathbb{E}[\Psi^{1/\varepsilon}]}\right)^{-\varepsilon}$$
(33)

The log values of price controls and quantity controls are then

$$\log \mathcal{V}^p = -\varepsilon \bar{\nu} \log(p^*/p_m) + (1-\beta) \log(\mathbb{E}[\Psi^{\nu}]) + \beta \log(p^* \mathbb{E}[\Psi] - \mathbb{E}[\Psi\Gamma])$$
(34)

$$\log \mathcal{V}^q = \bar{\nu} \log(q^*) + \beta \log(p_m(q^*)^{-1/\varepsilon} \mathbb{E}[\Psi^{1/\varepsilon}] - \mathbb{E}[\Gamma])$$
(35)

Replacing the optimal choices with their values in (32) and (33), we obtain

$$\log \mathcal{V}^p = -\bar{\nu}\log(p^*/p_m) + (1-\beta)\log(\mathbb{E}[\Psi^\nu]) + \beta\log(\frac{\beta}{\varepsilon\bar{\nu}-\beta}\mathbb{E}[\Psi\Gamma])$$
(36)

$$\log \mathcal{V}^{q} = \bar{\nu} \log(q^{*}) + \beta \log(\frac{\beta}{\varepsilon \bar{\nu} - \beta} \mathbb{E}[\Gamma])$$
(37)

Simplifying yields

$$\log \mathcal{V}^{p} - \log \mathcal{V}^{q} = \bar{\nu} \varepsilon \log \mathbb{E}[\Psi]$$

$$+ (\bar{\nu} \varepsilon - \beta) (\log \mathbb{E}[\Gamma] - \log \mathbb{E}[\Psi\Gamma])$$

$$- \varepsilon \bar{\nu} \log \mathbb{E}[\Psi^{1/\varepsilon}] + (1 - \beta) \log \mathbb{E}[\Psi^{\nu}]$$
(38)

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

$$\log \mathbb{E}[\Psi] = \mu_{\Psi} + \frac{1}{2}\sigma_{\Psi}^2$$
(39)

$$\log \mathbb{E}[\Gamma] - \log \mathbb{E}[\Psi\Gamma] = -\mu_{\Psi} - \frac{1}{2}\sigma_{\Psi}^2 - \sigma_{\Psi,\Gamma}$$
(40)

$$\log \mathbb{E}[\Psi^{1/\varepsilon}] = \frac{1}{\varepsilon} \mu_{\Psi} + (\frac{1}{\varepsilon^2}) \frac{1}{2} \sigma_{\Psi}^2$$
(41)

$$\log \mathbb{E}[\Psi^{\nu}] = \nu(\mu_{\Psi} + \frac{1}{2}\nu\sigma_{\Psi}^2)$$
(42)

Putting them together, we obtain the result

$$\log \mathcal{V}^p - \log \mathcal{V}^q = \frac{1}{2} \left(\left(\hat{\nu} - \frac{\bar{\nu}}{\varepsilon} \right) \sigma_{\Psi}^2 - 2(\varepsilon \nu - \beta) \sigma_{\Psi,\Gamma} \right)$$
(43)

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

C.4 Proof of Proposition 10

Let us start by defining the notion of optimality in this dynamic context. Let q^t be the equilibrium quantity at time t. Denote as $x_t = \frac{q_t - q^{target}}{q^{target}}$ the percentage deviation from the target at time t. Denote as 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 to a protocol price p_t that is not properly set so that q_t deviates from 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

$$p_{t+1} = \bar{F}^{-1}\left(\frac{q^{target}(1+x_t)}{q^{\max}\lambda_t}\right)g(d \times x_t)$$
(44)

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

$$x_{t+1} = \frac{\bar{F}\left(\bar{F}^{-1}\left(\frac{q^{target}(1+x_t)}{q^{\max\lambda_t}}\right)g(d \times x_t)\right)}{\bar{F}(p^{target})} - 1$$
(45)

We can see that by setting

$$g(d \times x_t) = \frac{\bar{F}^{-1}(\frac{q^{target}}{q^{\max}\lambda})}{\bar{F}^{-1}(\frac{q^{target}(1+x_t)}{q^{\max}\lambda_t})}$$
(46)

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 . Thus, we need to evaluate the deviation that is the closest to the worst-case value of z_t . We have

$$\ln \bar{F}^{-1}\left(\frac{q_{t+1}}{q^{target}}\right) = \ln \bar{F}^{-1}\left(\frac{q^{target}(1+x_t)}{q^{\max}\lambda_t}\right) - \ln \bar{F}^{-1}\left(\frac{q^{target}}{q^{\max}\lambda}\right) + \ln g(d \times x) \tag{47}$$

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

$$\ln \bar{F}^{-1}\left(\frac{q^{target}}{q^{\max}\lambda}\right) - \ln \bar{F}^{-1}\left(\frac{q^{target}(1+x_t)}{q^{\max}\lambda_t}\right) \ge \frac{1}{\varepsilon(q^{target})}(x_t + z_t)$$
(48)

From (47), 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 $f = \exp(q^{target})$.

C.5 Proof of Proposition 12

From equation (33), the optimal quantity for a monopolistic miner is, for $\beta = 1$,

$$q^* = \left(\frac{\varepsilon}{\varepsilon - 1} \frac{1}{p_m} \frac{\mathbb{E}[\Gamma]}{\mathbb{E}[\Psi^{1/\varepsilon}]}\right)^{-\varepsilon}$$
(49)

With $\Psi = \lambda q^{\text{max}}$. 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 obtains a positive value in expectation. Therefore, MMIC requires that $p_m \leq \mathbb{E}[\Gamma]$; thus,

$$q^* \leq q^{\max} \left(\frac{\varepsilon - 1}{\varepsilon}\right)^{\varepsilon} \mathbb{E}[\lambda^{1/\varepsilon}]^{\varepsilon}$$
 (50)

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

$$\frac{q^*}{q^{\max}} \leq \left(\frac{\varepsilon - 1}{\varepsilon}\right)^{\varepsilon} \tag{51}$$

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

These bounds provide valuable insights for studies of fee policies involving a monopolistic validator (Nisan, 2023; Lavi et al., 2022) and fee policies designed to prevent the validator from monopolizing all the surplus (?).

Figure A1 illustrates how the upper bound of the ratio q^*/q^{max} changes with ε . This upper bound approaches an asymptote of $e^{-1} \approx 37\%$, and it reaches 95% of this limit when $\varepsilon = 10.42$. Ethereum currently sets the target block size to half the block size limit. The correct interpretation of this result is that any ADT-TFM with a target block size exceeding 37% of the block size that meets user demand in an average demand scenario would not be invulnerable to a simple form of MEV. This is because MMIC necessitates that the validator has no incentive to include her own value-extracting transactions while simultaneously reducing the effective supply available to users.



Figure A1: Optimal target block size for a monopolist validator as a function of ε

D Numerical Examples

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 because of the programmatically defined rules of blockchains.

Instead, a different strategy, informed by the user demand model presented in Section 3, 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 10 as the inverse of the Pareto tail coefficient.

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.¹⁸ First, gas prices censor user valuations at the lower end of the distribution because of the base fee. Second, 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 an upper bound.¹⁹ To address these limitations, several robustness exercises are performed, which include restricting the estimation to blocks with a minimum gas price below a certain threshold (to limit the censoring of low valuations). These robustness exercises do not alter the primary conclusions of this numerical analysis.

Discussion: The estimate from a sample prior to the merge suggests an adjustment rate of 6.57% (tail coefficient 15.22), while the optimal adjustment rate for the blocks following the merge is 8.68% (tail coefficient 11.53). The observation that the current adjustment rate of 12.5% overshoots more before the merge than after aligns with the findings of Leonardos et al. (2022). The authors postulate that this discrepancy is because inter-block times were (approximately) exponentially distributed prior to the proof-of-stake upgrade whereas they are constant in Ethereum's proof-of-stake protocol. A constant block arrival rate more accurately reflects our model, suggesting that an estimate around 8% is more suitable for the current blockchain. One might wonder whose price elasticity our estimate represents and why it is so high (over 10). The adjustment rate represents the inverse elasticity of the marginal user submitting her transaction for inclusion in the next block. Given the

¹⁸As users demand different amounts of block space, each transaction is weighted by the gas units that it uses to fit the Pareto distribution.

¹⁹Given 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.

low block size limit, which is set because of technological constraints, it is not implausible that a marginal decentralized finance (DeFi) user with sensitive transactions would exhibit significant demand elasticity, similar to how high-frequency traders are sensitive to spreads.

Another consideration is that this analysis does not account for users whose transactions remain in the mempool for several blocks before confirmation. Notably, under EIP-1599, base fees do not decrease fast enough after a surge in demand. Fee policies that respond to intra-day or intra-hourly demand variations would not need to adjust prices based on the size of the previous block. Instead, it could maintain fixed prices during high-demand periods, fill blocks, and monitor the mempool for price adjustments. However, mempool data are typically not recorded on-chain (i.e., as part of the immutable blockchain) and can be easily manipulated.

Alternative Calculations: This section presents the results of supplementary robustness checks to validate the primary numerical analysis findings. These checks are performed under various conditions to address potential concerns highlighted above, 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, I adjust the selection of blocks within a size range of $q^{target}(1 \pm \delta\%)$ and different limits on base fees.

The estimations are performed on three samples: the full sample, the pre-merge sample, and the post-merge sample. Tables A2, A3, and A4 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 for the pre- than for the post-merge sample, 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. The increase in rates as the max base fee limit becomes more restrictive (for blocks with a base fee of less than 30 gwei) can only produce thicker tails

Parameters		Outputs		Observations
δ	max base fee (gwei)	shape α	adjustment rate \boldsymbol{d}	number of blocks
	200	12.45	8.03	7189
E 07	100	11.83	8.45	6645
3%	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 A2: 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)

because of 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 the pre-merge and post-merge samples suggest that the optimal adjustment rate lies within the 6% to 10% window, where the latter serves as an upper limit.

Blockchain designers must adopt simple, robust, and principled methods for updating TFM parameters. However, updating parameters and the "rules of the game" as we go might not be the best track for fostering scalability. Considering nondeterministic 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 terms of how well these rates reflect market conditions during the relevant 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 us infer demand fluctuations under normal

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
99(7	200	11.48	8.71	30029
	100	11.29	8.86	29454
33 70	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 A3: 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 A4: 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)

conditions.