

Optimality Despite Chaos in Fee Markets

Stefanos Leonardos^{1*} Daniël Reijsbergen^{2*} Barnabé Monnot³ Georgios Piliouras²

¹King’s College, London, UK

² Singapore University of Technology and Design, Singapore

³ Ethereum Foundation, Berlin, Germany

*The first and second authors contributed equally.

Abstract. Transaction fee markets are essential components of blockchain economies, as they resolve the inherent scarcity in the number of transactions that can be added to each block. In early blockchain protocols, this scarcity was resolved through a first-price auction in which users were forced to guess appropriate bids from recent blockchain data. Ethereum’s EIP-1559 fee market reform streamlines this process through the use of a base fee that is increased (or decreased) whenever a block exceeds (or fails to meet) a specified target block size. Previous work has found that the EIP-1559 mechanism may lead to a base fee process that is inherently chaotic, in which case the base fee does *not* converge to a fixed point even under ideal conditions. However, the impact of this chaotic behavior on the fee market’s main design goal – blocks whose long-term average size equals the target – has not previously been explored. As our main contribution, we derive near-optimal upper and lower bounds for the time-average block size in the EIP-1559 mechanism despite its possibly chaotic evolution. Our lower bound is equal to the target utilization level whereas our upper bound is $\approx 6\%$ higher than optimal. Empirical evidence is shown in great agreement with these theoretical predictions. Specifically, the historical average was $\approx 2.9\%$ larger than the target range under Proof-of-Work and decreased to $\approx 2.0\%$ after Ethereum’s transition to Proof-of-Stake. We also find that an approximate version of EIP-1559 achieves optimality even in the absence of convergence.

1 Introduction

In the seminal Bitcoin whitepaper [20], the concept of a *blockchain* was introduced as a secure data structure maintained by *nodes* in a peer-to-peer network. A blockchain consists of elementary database operations called *transactions* that modify a global state – e.g., cryptocurrency ownership or the state of smart contracts. Nodes provide an essential service to the blockchain’s users by broadcasting their transactions and responding to queries about the global state [12]. As such, a large and diverse network of nodes enhances *decentralization* in the sense that the availability and integrity of blockchain-enabled services do not depend on a handful of entities. As nodes execute every new transaction to maintain their view of the latest global state, the *computational cost* of new transactions must be limited to avoid excluding all but the most powerful nodes. In Ethereum [4],

this computational cost is measured through the notion of *gas*, and each block has a *gas limit* that is decided by the nodes. In Ethereum’s original design, each transaction has a *gas price* that indicates how much its creator is willing to pay for its inclusion on the blockchain. This mechanism behaves like a *first-price auction*, and shares all of its drawbacks [21]: users tend to bid untruthfully relative to the true valuation of their transaction, which leads to guesswork and overbidding that is detrimental to the user experience.

Ethereum Improvement Proposal (EIP) 1559 [3] simplifies Ethereum’s fee market through a protocol-set *base fee*. Instead of aiming to fill each block to the limit, it aims to achieve a long-term average *target*, which is half the maximum size of each block. The base fee is automatically updated to reflect market conditions: if a block is larger than the target, then demand for transaction inclusion is too high at the current price so the base fee is increased (and vice versa for smaller blocks). The base fee hence aims to reflect the constantly-changing *market-clearing price*, which is the theoretical price at which demand for transactions is precisely such that block sizes equal the target. To add a transaction to the blockchain, users pay the base fee per unit of spent gas – this payment is permanently destroyed or *burned* [13]. This mechanism is provably incentive-compatible in the sense that users bid close to their true valuation unless demand is extremely high [23,24]. However, whether the protocol is *optimal* in the sense that it achieves its main design goal – a long-term average block size that equals the target – has not previously been explored. Previous work has found that the base fee may *not* converge to the market-clearing price [16] even when market conditions remain unchanged, as the base fee process exhibits (Li-Yorke) *chaos* in a wide range of market conditions. As the base fees need not converge to the market-clearing price, it is natural to ask whether the long-term average block sizes in fact converge to the target.

In the current work, we investigate whether optimality is possible in fee markets that exhibit non-convergent behavior. Specifically, we show that the default EIP-1559 mechanism is *approximately optimal* even if the block sizes are chaotic. We find that, unless market conditions are such that the base fee converges to a fixed point, EIP-1559 1) overshoots the target but 2) by at most $\approx 6.27\%$. These results hold *regardless* of the specific market conditions beyond convergence, or the nature of the block creation protocol. This suggests that we can still analyze the system even if it does not reach an equilibrium, which is a very stringent condition to be met in practice. Furthermore, our results allow us to quantify the maximum degree to which excessively large blocks impact nodes with limited processing power. We have validated our theoretical results using historical data, as displayed in Figure 1 (see Appendix C): since its introduction, blocks in EIP-1559 were initially $\approx 2.9\%$ larger than the target, and after a major change to Ethereum’s block creation protocol¹ (the PoS “Merge” [14]) this overshoot dropped to $\approx 2.0\%$. Both of our main theoretical findings – blocks

¹ In particular, Ethereum’s consensus protocol switched from Proof-of-Work to Proof-of-Stake (PoS), which, as a by-effect, caused the time between the creation of new blocks to become constant.

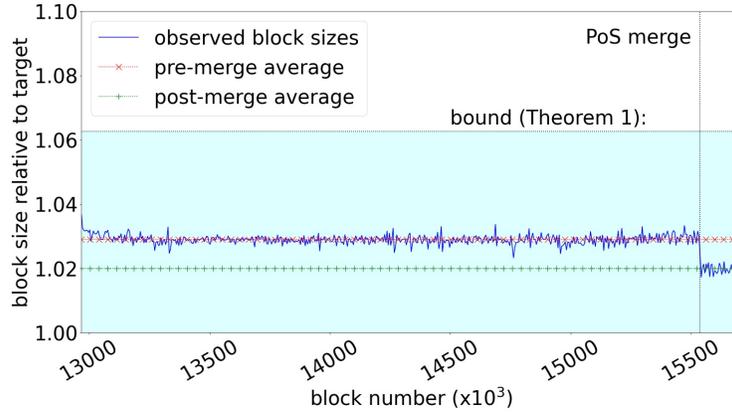


Fig. 1. Evolution of the relative block size since EIP-1559: the blue line without marks depicts the observed average block size over batches of 5000 consecutive blocks. The red and green lines with marks depict the averages over the periods before and after Ethereum’s switch to proof-of-stake, respectively. The colored region indicates the range of potential long-run averages covered by the bound of Theorem 1.

overshoot the target, but to a limited degree – have therefore been borne out in practice. Moreover, the persistence of excessive block sizes throughout the observation period suggests that this is not merely a honeymoon effect [22]. We also observe the tightness of our bound in a wide range of simulation experiments (see Appendix D).

In practice, the baseline variant of EIP-1559 is a linear approximation of an exponential update rule that is computationally inefficient to implement [9]. As a further contribution, we investigate the average-case performance of this ideal mechanism, deemed *exponential* EIP-1559. Our analysis suggests that exponential EIP-1559 *always* achieves the long-term average target, even if the base fee does not converge to the market-clearing price. We find that the manner in which EIP-1559 approximates an exponential function creates the observed overshoot – as such, this suggests an interesting direction for future protocol updates.²

Outline: The outline of our work is as follows. After a discussion of the context of our work (Section 2) and our model of a blockchain economy (Section 3), we present a unifying framework for the dynamics of different transaction fee market mechanisms, including the default EIP-1559 mechanism, its exponential variation, and alternative proposals (Section 4). In Section 5, we present our formal analysis. In Section 6, we discuss the generality of our results, and Section 7 concludes the paper.

² In Appendix B, we discuss further such designs including a recent proposal that sets base fees using the principles behind Automated Market Makers (AMMs) [5] and general dynamic posted price mechanisms [10].

2 Background and Related Work

In this section, we provide a high-level description of Ethereum’s fee market and EIP-1559, and present some concepts that we have not discussed previously. We conclude the section with an overview of related work on fee markets.

Ethereum: Ethereum is a cryptocurrency platform that supports *smart contracts*, i.e., software programs that are executed in a decentralized network. Ethereum’s global state consists of the state of all smart contracts and the amount of Ethereum’s native cryptocurrency token – *Ether* or ETH – in each user account. The purpose of Ethereum transactions is to transfer ETH from one account to another, or to create or call a smart contract. A selection of nodes have the ability to periodically group transactions into a new *block* and broadcast it to the network. Although the exact nature of these nodes depends on the consensus mechanism, we will refer to such nodes as “*miners*” for brevity. Each block points to a previous block, forming a *blockchain*. Each operation in a transaction consumes gas, and the amount of ETH that a user is willing to pay for each unit of gas depends on the user’s *valuation* of the transaction – i.e., how much utility she expects to derive from the transaction’s inclusion on the blockchain. Before EIP-1559, users would specify a gas price for each transaction that determines the amount of ETH spent per unit of gas. Demand for gas fluctuates over time, e.g., due to temporal patterns and events such as NFT drops, so non-expert users were forced to guess appropriate gas prices from recent data.

EIP-1559: EIP-1559 simplifies Ethereum’s original fee market design through the use of a dynamically adjusted base fee. The base fee at each time serves as a *posted price* that users need to pay to have their transaction processed at the next block. When users pay a transaction fee, an amount of ETH equal to the base fee is burned – however, a small amount of ETH can be awarded to the miner by the user in the form of a *miner’s tip*. Without the miner’s tip, miners would have no incentive to process transactions, which could cause them to create empty blocks instead. The base fee is continuously updated to reflect changing market conditions: if blocks are larger (smaller) than a fixed *target* size, then the base fee is increased (decreased) to reduce (increase) demand. The target has been set to roughly equal the maximum block size before EIP-1559 (i.e., 15M gas) – meanwhile, the maximum block size has been increased to 30M gas (i.e., twice the target). A higher maximum block size (relative to the target) increases the risk that nodes are overwhelmed during demand bursts.

Related work: In [23,24], three desirable properties for transaction fee markets are investigated: whether (1) users are incentivized to bid their true valuation, (2) miners are incentivized to follow the protocol’s inclusion rule, and (3) miners and users have no incentive to form cartels to subvert the protocol. It is shown that EIP-1559 always satisfies property (1) and (3), and (2) only when demand

is “stable” [24]. Ethereum’s original fee market (being a first-price auction) does not satisfy property (1), whereas property (3) would not hold if the base fee were awarded to miners instead of burned. The compatibility of these properties under general conditions is further explored by [7,25] and [11] in the context of transaction fee markets and by [18] in the context of NFT auctions.

In [19,16], the behavior of the base fee under stable market conditions is investigated – it is found that the base fee exhibits Li-Yorke chaos [16], which in practice results in the prevalence of sequences of alternating full and empty blocks. This behavior was later confirmed to occur in practice [22]. In [10], the *social welfare* of fee market mechanism is investigated, and two alternative mechanisms to EIP-1559 are proposed that perform better on this metric. In [17], the impact of EIP-1559 on various user-centric measures such as average transaction fees, waiting times, and consensus security is investigated. In [8], an extension of EIP-1559 is considered in a setting in which the base fee depends on the availability of multiple fungible resources (i.e., beyond gas use). Finally, fee market design has been studied for other cryptocurrency platforms, e.g., Bitcoin [2,15].

The question of whether the fee market ensures that the long-term average block size indeed equals the target size has not been considered in these works, although [17] finds that the average size of blocks as measured in terms of the network load (which does not capture the gas use of, e.g., smart contract function calls) has increased from 64.05 to 78.01 kB.

3 Model and Notation

In this section, we introduce our notation to describe transaction fee markets mathematically. We have EIP-1559 in mind (cf. Section 2), but the description applies to variations of EIP-1559 and other related mechanisms as well.

Base fee and target block size: The main element of EIP-1559 like transaction fee markets is a dynamically adjusted *base fee*, b_n , that is updated after every block, $B_n, n \geq 0$. The goal of the mechanism is to update the base fee in such a way that blocks achieve a pre-defined target block size. Let T denote the target block size and let kT denote the maximum block size, for some integer $k \geq 1$. Currently, in the Ethereum blockchain, k is set at $k = 2$, i.e., the target is equal to half the maximum block size.

Users (transactions) and valuations: Users (transactions) arrive to the pool according to a stochastic process. Without any loss of generality, we assume that each transaction uses 1 unit of gas, and use the random variable N_n to describe the number of transactions that arrive between two consecutive blocks B_n, B_{n+1} for $n \geq 0$. Let $\lambda_n = \mathbb{E}(N_n)/T$ be the ratio of the expected value of N_n to the target T . We make no assumptions about the distribution of N_n beyond it having a finite mean, i.e., $\lambda_n < \infty$. To avoid trivial cases, we will assume that $\lambda_n > 1$, i.e., that the arrival rate is larger than the target block size. For the theoretical analysis, we will assume that users leave the pool if their transaction

is not included in the next block and return according to the specified arrival process.³ Whenever necessary, we will index users (transactions) with $i, j \in \mathbb{N}$.

Users' valuations: Each transaction, indexed by $i \in \mathbb{N}$, has a valuation, v_i . Valuations at time t are drawn from a distribution function F_n with support included in \mathbb{R}^+ . Typically, the support of F_n is bounded, i.e., there exists a maximum possible valuation $M \gg 0$.⁴ We will write $\bar{F}_n(x) := 1 - F_n(x)$ to denote the so-called *survival function* of the distribution F_n . For simplicity, we will assume that $\lambda_n \equiv \lambda$ and $F_n \equiv F$ for all $n \geq 0$, i.e., that λ and F are independent of the block height. We discuss how our results are straightforwardly generalized to a setting with time-dependent distributions in Section 6.

Bids and tips: User bids in EIP-1559 consist of two elements: (1) the *max fee*, f , which is the maximum amount per gas unit that the user is willing to pay for their transaction to be included and (2) the *max priority fee*, p , which is the maximum tip per gas unit that the user is willing to pay to the miner who includes their transaction. We assume that users bid truthfully and rationally, i.e., a user with valuation v_i will bid $(f, p) = (v_i, \epsilon)$ where $\epsilon > 0$ is the minimum amount that covers the miners' cost to process the transaction. Combining with the above, this generates the inclusion requirement: miner's tip = $\min \{f - b_n, p\} \geq \epsilon$.

Block Sizes: Let $g_n := g(b_n)$ denote the number of transactions that get included in block B_n when the base fee is equal to $b_n, n \geq 0$. Given a number of transactions $N_n = n$ and valuations v_1, \dots, v_n , we have that $g(b_n) = \min \{kT, \sum_{i=1}^n \mathbf{1}(v_i \geq b_n)\}$, i.e., the size of the block is equal to minimum between the block limit, kT , and the number of transactions whose max fee exceeds the base fee.⁵ We note that $\mathbf{1}(v_i \geq b_n)$ has a Bernoulli distribution with probability $\mathbb{P}(v_i \geq b_n) = \bar{F}(b_n)$ of observing 1. For our analysis, we will consider the mean-field approximation of the stationary demand, which results in

$$g(b_n) = \min \left\{ kT, \sum_{i=1}^{N_n} \mathbf{1}(v_i \geq b_n) \right\} = \min \{kT, \lambda T \bar{F}(b_n)\} \quad (1)$$

We denote the market-clearing price, i.e., the base fee for which $g(b^*) = T$, by b^* . From (1), we observe that $b^* = \bar{F}^{-1}(1/\lambda)$. Equation (1) also implies that $\lim_{b_n \downarrow 0} g(b_n) = kT$ since $\lambda > k$ by assumption, and $\lim_{b_n \rightarrow \infty} g(b_n) = 0$. To reflect practical settings, we will assume that b_n cannot become negative.

³ This assumption only reduces unnecessary complexities in the analysis and is relaxed in the simulations without significant effect in the results.

⁴ While unbounded valuations are unrealistic for practical purposes, we note that our results hold even for such cases.

⁵ To simplify notation, we henceforth assume that each user's priority fee, p , is equal to the miners' broken even cost ϵ . Equivalently, we only consider transactions that miners are willing to include and hence, we apply the indicator to b_n instead of $b_n + \epsilon$.

4 Fee Market Mechanisms: Base Fee Update Rules

Base fee update rules (BFURs), are functions $h : (0, \infty) \rightarrow (0, \infty)$. Their goal is to efficiently regulate block sizes via updates in the base fee, $b_{n+1} := h(b_n)$ with $b_0 > 0$. Intuitively, the base fee increases (decreases) whenever blocks are more (less) than the target and remains constant otherwise.

Design goal: Achieving the target block size in *each* block, however, turns out to be a very difficult [17,22] or even theoretically impossible goal [16]. To obtain a more tractable objective, it is still reasonable to ask whether the target block size is achieved *on average*. In symbols, let $G_N := \frac{1}{N} \sum_{n=1}^N g_n$ denote the average block size until block $N > 0$. Then, this requirement suggests that

$$\lim_{N \rightarrow \infty} G_N = T.$$

4.1 Proper Base Fee Update Rules

To avoid pathological cases, a base fee update rules h needs to satisfy some minimal regularity conditions. These are outlined in Definition 1.⁶

Definition 1 (Proper Base Fee Update Rules (PBFURs)). *Let T be the target block-size. An update rule $h : (0, \infty) \rightarrow (0, \infty)$ is called proper if it satisfies the following*

- (A.1) non-divergence: $h(b_n) \geq b_n$ if $g(b_n) \geq T$ and $h(b_n) \leq b_n$ if $g(b_n) \leq T$.
- (A.2) bounded relative differences: *there exists an $\alpha \geq 1$ such that $\alpha^{-1}b_n \leq h(b_n) \leq \alpha b_n$.*

Assumption (A.1) ensures that base fee updates are in the right direction, i.e., that the base fee does not decrease (increase) whenever the current block size is more (less) than the target. Assumption (A.2) excludes update rules with potentially unbounded updates. To include more general update rules, (A.2) can be relaxed to (the equivalent in flavor):

- (A.2') *there exists $\alpha > 1$ and $\beta > 0$ such that $\alpha^{-1}b_n - \beta \leq h(b_n) \leq \alpha b_n + \beta$.*

PBFURs have the desirable property that they generate a bounded sequence of base fees. This is established in Lemma 1 which can be proved by induction (cf. Appendix A.1).

Lemma 1. *If h is a PBFUR, and $0 < b_0 < \infty$, then*

$$\min\{b_0, \alpha^{-1}b^*\} \leq b_n \leq \max\{b_0, \alpha b^*\} \quad \text{for all } n \geq 0.$$

⁶ Apart from the current base fee, b_n , a base fee update rule may also depend on other parameters, θ_n , such as the target block size (time-independent) or the block size and efficient gas prices at time n (time-dependent). Whenever irrelevant, we will omit such parameters from the description of h . In Appendix B, we provide the generalized counterpart of Definition 1 that accounts for such dependencies.

4.2 Examples of PBFURs

EIP-1559: In the EIP-1559 transaction fee market [19,16,22], the base fee, b_n , is updated after every block (where blocks are indexed by their block height, $t > 0$) according to the following equation

$$b_{n+1} = b_n \left(1 + d \cdot \frac{g_n - T}{T} \right), \quad \text{for any } t \in \mathbb{N}, \quad (\text{EIP-1559})$$

where d denotes the *adjustment quotient* (or step-size or learning rate), currently set by default at $d = 0.125$. It will be convenient to use the notation $y_n := \frac{g_n - T}{T} \in [-1, 1]$, for the *normalized deviation at block t* , and

$$G_N := \frac{1}{N} \sum_{n=1}^N g_n \quad (2)$$

for the *average block size up to block N* . It is immediate to check that (EIP-1559) satisfies both A.1 and A.2 and is, thus, a PBFUR.

Exponential EIP-1559: Instead of the (EIP-1559) updates, we may consider the exponentially weighted updates

$$b_{n+1} = b_n (1 + d)^{\left(\frac{g_n - T}{T}\right)}, \quad \text{for any } t \in \mathbb{N}. \quad (\text{EXP-1559})$$

The standard (EIP-1559) update rule is the linear approximation (in the Taylor expansion of the function $d \mapsto (1 + d)^{y_n}$) of the update rule in equation (EXP-1559). Again, it is immediate to check that (EXP-1559) satisfies non-divergence (A.1) and bounded relative differences (A.2) and is, thus, a PBFUR. Using the generalized Bernoulli inequality, it is also straightforward to show that the updates of Equation (EXP-1559) are always less aggressive than (but in the same direction as) the updates of Equation (EIP-1559). In other words, if a block is congested ($y_n > 0$), then the next base fee will increase with both methods, but it will be higher with (EIP-1559). Similarly, if a block is not congested ($y_n < 0$), then the next base fee will decrease with both methods, and it will be lower with (EXP-1559).⁷

5 Analysis: Bounds on Average Block Sizes

5.1 EIP-1559

In this section, we are interested to obtain lower and upper bounds on the long-term average block sizes, G_N , $N \geq 0$, generated by the EIP-1559 update rule, cf. equation (EIP-1559). Our main result is summarized in Theorem 1.

⁷ Ethereum researchers [9], also study an alternative exponential EIP-1559 form which relies on the exponential approximation $1 + dy_n \approx e^{dy_n}$ for dy_n small enough. Again, it is immediate to see that this rule satisfies (A.1) and (A.2) and is, thus, a PBFUR.

Theorem 1. Let $(g_n)_{t>0}$ denote the sequence of block sizes generated by the base fee $(b_n)_{t>0}$ of the EIP-1559 update rule (EIP-1559) with learning rate $d \in (0, 1)$ for an arbitrary valuation distribution on \mathbb{R}_+ . Then, the long-term average block size $\lim_{N \rightarrow +\infty} G_N$ satisfies

$$T \leq \lim_{N \rightarrow +\infty} G_N \leq \left[1 - \frac{\ln(1+d)}{\ln(1-d)} \right]^{-1} 2T. \quad (3)$$

In words, Theorem 1 implies that the EIP-1559 update rule either meets or *slightly* overshoots the target of T . The extent of possible overshooting (see upper bound in (3) of Theorem 1) depends on the choice of the learning rate d . For instance, at the current default level of $d = 0.125$, this yields a bound of approximately 1.0627 or 106% of the block size T . The upper bound for values of $d \in (0, 0.5]$ (which are of practical interest) is visualized in Figure 2. As we see, this bound grows approximately linearly in d . This pattern continues for larger values of d till eventually growing exponentially fast and approaching 2 (or 200% of T) in the limit $d \uparrow 1$ (but such values are, at least currently, only of theoretical interest). The proof of Theorem 1 utilizes the following upper and lower linear bounds on the natural logarithm.

Lemma 2. (i) Let $x > -1$. Then it holds that

$$\ln(1+x) \leq x,$$

with equality if and only if $x = 0$. (ii) Let $d \in (0, 1)$ and let $|x| \leq d$. Then, it holds that

$$\ln(1+x) \geq \alpha x + \beta,$$

with $\alpha = \frac{1}{2d} \cdot [\ln(1+d) - \ln(1-d)]$ and $\beta = \frac{1}{2} \cdot [\ln(1+d) + \ln(1-d)]$.

Using Lemma 2, we can now prove Theorem 1.

Proof (Proof of Theorem 1). Let $d \in (0, 1)$. By taking the natural logarithm on both sides of equation (EIP-1559), we obtain that

$$\ln(b_{n+1}/b_n) = \ln\left(1 + d \frac{g_n - T}{T}\right)$$

Recall that $\frac{g_n - T}{T} \in [-1, 1]$ by definition. Thus, by applying Lemma 2 on $\ln\left(1 + d \frac{g_n - T}{T}\right)$ with $d \frac{g_n - T}{T} \in [-d, d]$, we obtain that

$$\alpha d \cdot \frac{g_n - T}{T} + \beta \leq \ln(b_{n+1}/b_n) \leq d \cdot \frac{g_n - T}{T},$$

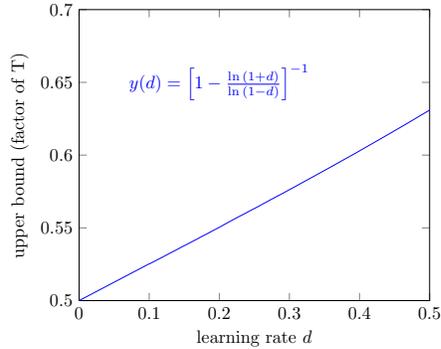


Fig. 2. Upper bound (scaling factor of maximum block size if $k = 2$) in equation (3) of Theorem 1. The upper bound grows almost linearly for the relevant values of d .

with α, β as above. Combining the above and solving for g_n , we obtain that

$$\frac{T \cdot \ln(b_{n+1}/b_n)}{d} + T \leq g_n \leq \frac{T \cdot \ln(b_{n+1}/b_n)}{\alpha d} + \left(1 - \frac{\beta}{\alpha d}\right) T.$$

Observe that if we sum up all terms from 1 to N , then the term involving $\ln(b_{n+1}/b_n)$ on both sides telescopes to

$$\sum_{n=1}^N \ln(b_{n+1}/b_n) = \sum_{n=1}^N (\ln b_{n+1} - \ln b_n) = \ln b_N - \ln b_1.$$

Thus, summing up all terms from 1 to N in the previous inequality, and using the notation $G_N = \frac{1}{N} \sum_{n=1}^N g_n$, we obtain that

$$G_N \leq \frac{T(\ln b_N - \ln b_1)}{N\alpha d} + \left(1 - \frac{\beta}{\alpha d}\right) T, \quad (4)$$

and

$$\frac{T(\ln b_N - \ln b_1)}{Nd} + T \leq G_N, \quad (5)$$

for the upper and lower bounds respectively. Concerning the term on the right hand side of the upper bound, observe that after some standard algebraic manipulation, we can write $\left(1 - \frac{\beta}{\alpha d}\right) \cdot T = \left[1 - \frac{\ln(1+d)}{\ln(1-d)}\right]^{-1} \cdot 2T$. To conclude observe that $\lim_{N \rightarrow +\infty} b_N < M$ for some $M > 0$, since $(b_n)_{n \geq 0}$ is bounded by Lemma 1.⁸ This implies that

$$\lim_{N \rightarrow +\infty} \frac{T \ln b_N - \ln b_1}{N} = 0.$$

Thus, taking the limit $N \rightarrow +\infty$ on both sides of the inequalities in (4), (5), we obtain that

$$T \leq \lim_{N \rightarrow +\infty} G_N \leq \left[1 - \frac{\ln(1+d)}{\ln(1-d)}\right]^{-1} \cdot 2T,$$

as claimed. \square

5.2 Visualizations: Bifurcation Diagrams and Long-Term Averages

To gain more intuition on the previous results, we proceed to visualize the individual trajectories of the base fee dynamics and the resulting block sizes for a simulated demand realization. This is done in the *bifurcation diagrams* of Figure 3.

⁸ The limit is also bounded from above if transactions are no longer included by miners when the base fee becomes so high that the computational cost of processing transactions outweighs any potential miner's tip.

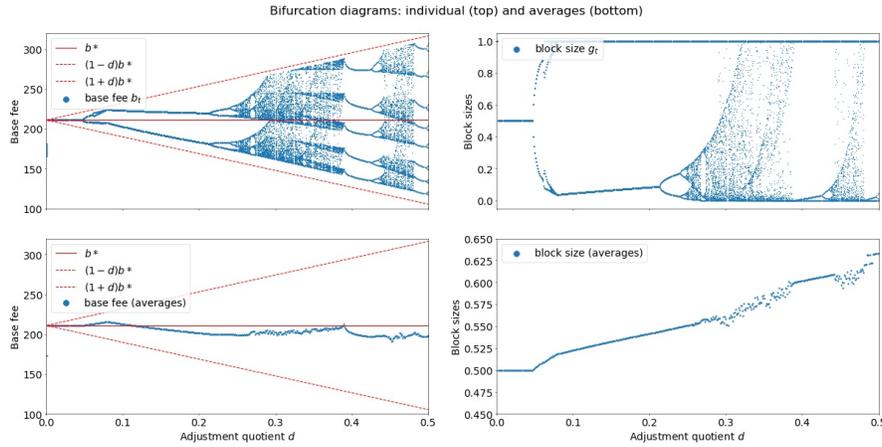


Fig. 3. EIP-1559: Individual trajectories (100 iterations after skipping 200 iterations) of the base fee dynamics (top left) and block sizes (top right) for various values of the adjustment quotient, d , in $(0, 0.5]$ (horizontal axis). The bottom panels show averages of the trajectories in the top panels. Demand (users-valuations) has been simulated from an exponential distribution on $[205, +\infty)$ with mean $\mu = 210$ and standard deviation $\sigma = 5$. The results are robust to different distributions and initializations (currently $b_0 = 100$), cf. Appendix D. Despite the chaotic behavior of the individual trajectories (top panels), the long-term averages (bottom panels) exhibit predictable and mathematically tractable patterns (cf. upper bound of block-sizes in Figure 4). Note: Unlike Figure 1, the scale of the y -axis in the block-size panels (top-right and bottom-right) is between 0 (empty block) and 1 (full block) and the target is equal to 0.5.

Individual trajectories (top panels) The bifurcation diagram in the top left panel shows the individual trajectories of the base fee dynamics (blue dots) for different value of the adjustment quotient, d , (horizontal axis). Recall that the default value of d is 0.125. To generate the plots, we have drawn the user-valuations from a gamma distribution with mean $\mu = 220$ and standard deviation $\sigma = 10$. The depicted blue dots show 1000 updates after a burn-in of 500 updates (higher numbers of iterations suggest that the dynamics have converged to the depicted attractors). For low values of d (below 0.08), we see that the base fee dynamics converge to a single value, close to or exactly at the theoretical optimum, b^* . For larger values of d (approximately between 0.1 and 0.3), the dynamics oscillate between two values and for most remaining values of d (larger than 0.3 and for a small regime roughly between 0.08 and 0.1), the dynamics exhibit chaotic behavior (multiple dots dispersed over the whole interval between the two diagonal red lines). In all cases, the dynamics remain within the bounded region $[(1-d)b^*, (1+d)b^*]$.

The bifurcation diagram in the top right panel shows the resulting block sizes. Following similar patterns to the base fee dynamics, block sizes converge to the target value (0.5 or half-full) for low values of d , oscillate between full

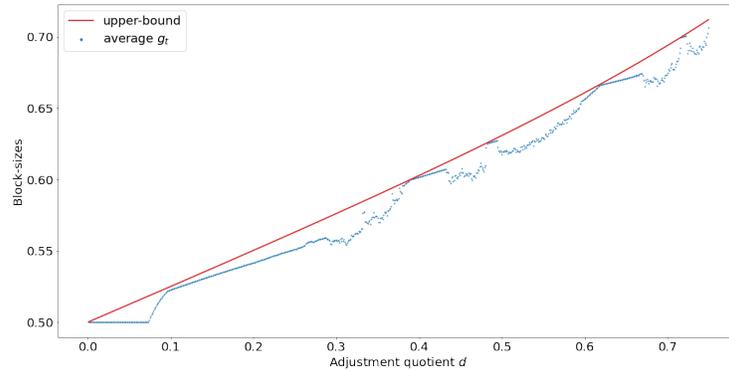


Fig. 4. Estimated upper bound (red line) of the average block-sizes as given by equation (3) in Theorem 1. The blue dots represent the average block-sizes for normally distributed simulated demand (user-valuations) with mean $\mu = 210$ and standard deviation $\sigma = 2.5$ at different values of the adjustment quotient, d , (horizontal axis). The long-term average block sizes grow linearly with d consistent with the estimated upper bound. Moreover, the upper bound is *tight*. Qualitatively equivalent results obtain for arbitrary demand distributions, e.g., uniform and gamma with arbitrary parameters (not presented here).

and (almost) empty for intermediate values of d and become chaotic for larger values of d .

In summary, the bifurcation plots in the two top panels show the effect of the adjustment parameter, d , on the individual trajectories of the base fees and the block sizes. These plots illustrate how small changes in the adjustment quotient, i.e., in the studied *bifurcation parameter*, can cause dramatic changes in the observable trajectories of both base fees and block sizes.

Averages (bottom panels) The bottom panels show the averages of the trajectories that are depicted in the top panels; base fees (bottom left) and block sizes (bottom right). We can see that the base fee slightly undershoots the ideal value of b^* and that the block sizes (slightly) overshoot the target of 0.5. Moreover, the deviation from this target grows linearly in the adjustment quotient d . At the current level, i.e., $d = 0.125$, the averages are approximately at 0.53.

Tight upper-bound on block sizes: The figures in the averages in Figure 3 are not specific to the simulated demand and generalize to arbitrary demand distributions. Figure 4 shows the estimated upper-bound (red line) of Equation (3) and the realized average block-sizes for user-valuations drawn from a normal distribution with mean $\mu = 210$ and standard deviation $\sigma = 10$. The upper-bound is tight and approximates very well the actual evolution of the block-size averages for various values of the adjustment quotient d . Qualitatively equivalent results obtain for arbitrarily parameterized uniform, normal and gamma user-valuation distributions (not presented here).

5.3 Exponential EIP-1559

As we show in Theorem 2, (EXP-1559) achieves time average convergence exactly to the target block sizes of T . However, (EIP-1559) is more relevant from practical purposes since it requires integer rather than floating point calculations.

Theorem 2. *For the dynamical system in equation (EXP-1559), it holds that*

$$\lim_{N \rightarrow +\infty} G_N = T,$$

i.e., the time average of the block sizes (or block occupancies), $(G_N)_{N \geq 1}$, converges to the target value T as the number, N , of updates grows to infinity.

The proof of Theorem 2 mirrors the steps in the proof of Theorem 1 and is therefore deferred to Appendix A. However, it is worth noting that there is nothing special about T ; the time-average of the dynamic in equation (EXP-1559) would converge to any given target block-size (as it appears in the numerator of the exponent) within, of course, the admissible limits.

Convergence Rates From the proof of Theorem 2, we can also reason about the convergence rate. After N time-steps, we have that the distance $d(G_N, T)$ between $G_N := \frac{1}{N} \sum_{t=0}^N g_n$ and the target T is equal to

$$d(G_N, T) = \frac{T (\ln b_N - \ln b_1)}{2 \ln(1 + d)} \cdot \frac{1}{N}$$

which drops linearly in N , i.e., $\mathcal{O}(1/N)$. The constant factor depends on $T, \ln 1 + d$ and the error due to initialization, $\ln b_N - \ln b_1$. Since b_N is bounded (cf. Lemma 1, we also know that $\ln b_N$ cannot grow (in absolute value) beyond certain bounds. Thus, all these terms vanish at a rate of $1/N$. Note, that in a similar fashion, we can get similar rates for the linear-EIP1559 update, cf. Theorem 1. In this case, we have that

$$\frac{T (\ln b_N - \ln b_1)}{2d} \cdot \frac{1}{N} \leq d(G_N, T) \leq \frac{T (\ln b_N - \ln b_1)}{2\alpha d} \cdot \frac{1}{N} - \frac{\beta T}{\alpha d 2}.$$

Again, the convergence rate is linear (by the same reasoning as above for $\ln b_N$ and the only thing that changes is the constant error term in the upper bound.

Visualizations for Exponential EIP-1559 Figure 5 shows the same data as Figure 3 but for the exponential EIP-1559 update rule (cf. (EXP-1559)). The updates of both base fees (top left panel) and block sizes (top right panel) are more smooth (albeit not entirely non-chaotic). However, as can be seen from the left panels, the base fee dynamics tend to overshoot the actual ideal value. This behavior is slightly dependent on the initialization of the dynamics and is not consistent among all possible simulations in many of which the base fee dynamics meet b^* (not presented here). However, in all cases, the important observation

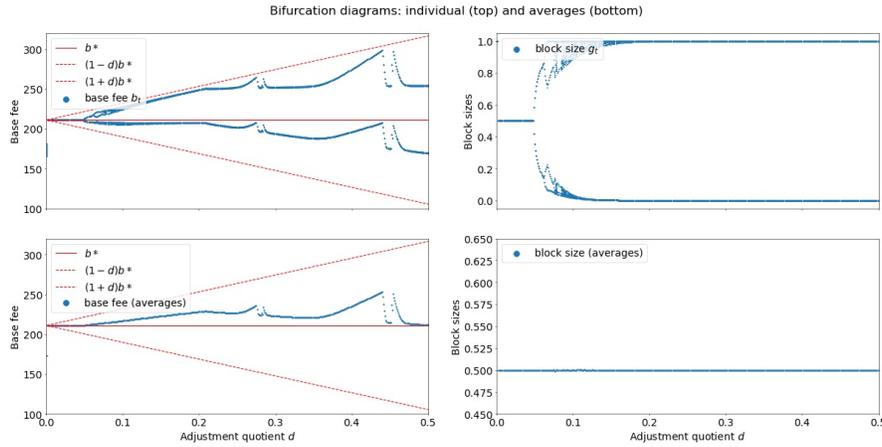


Fig. 5. Exponential EIP-1559: Individual trajectories of the base fee dynamics (top left) and block sizes (top right) for various values of the adjustment quotient, d (horizontal axis). The panels are the same as in Figure 3. Exponential updates result in more regular (individual) trajectories and exactly achieve the block-size target on average. While the exact base fee trajectory is sensitive to b_0 , it is not clear why the average base fee does not approach the market-clearing price, b^* (horizontal red line).

concerns the block sizes and the fact that the target is met exactly (horizontal line at exactly 0.5 in the bottom right panel) regardless of the initialization and the value of the adjustment quotient d . This outcome is consistent with the theoretical prediction of Theorem 2.

6 Discussion

Overshoot The inequality $\ln\left(1 + d \cdot \frac{g_n - T}{T}\right) \leq d \cdot \frac{g_n - T}{T}$ that is used to derive the lower bound of T in Theorem 1 (cf. inequality (i) in Lemma 2) holds with equality *if and only if* $d \cdot \frac{g_n - T}{T} = 0$, i.e., if and only if $g_n = T$. This implies that the long-term average block sizes will be equal to the target, T , if and only if the system equilibrates at a fixed point. However, the base fee dynamics are provably chaotic for almost all market conditions (demand and user valuations) that can be met in practice [16], and, thus, individual blocks will generally deviate from the target. Consequently, the lower bound will hold with *strict* inequality, implying a positive overshoot in EIP-1559 (almost) regardless of market conditions.

Generality & Robustness We note that the proof of Theorem 1 does not in any way rely on F_n or the distribution of N_n , but only on the learning rate d . As such, our results hold regardless of the exact distribution of valuations or the block creation protocol. In fact, the only technical requirement is that for Lemma 1, we need the existence of a market-clearing price b^* to express our

bounds. However, if we make the market-clearing price b^* time-dependent i.e., b_n^* instead of b^* , we only need to require that there exist $b_{\min}^* > 0$ and $b_{\max}^* < \infty$ such that $b_{\min}^* \leq b_n^* \leq b_{\max}^*$ for all $n \geq 0$ to obtain similar bounds. These are loose restrictions in practice. Our results also do not rely on k , e.g., if the maximum block size were set to 4 the target, then this would have no impact on the theoretical bounds. Interestingly, the upper bound in (3) is less than twice the target *regardless* of how small the target is relative to the maximum.

Although our theoretical results establish general bounds on the long-term average block size, the precise average values may depend on many factors, including the distribution of user valuations and the block creation protocol. We do observe from Figure 1 that the observed block size averages over 5000-block batches exhibit remarkably consistent behavior both before and after the PoS merge. Since EIP-1559, market conditions have changed considerably: the base fee itself has changed from around 100 GWei (1 GWei = 10^{-9} ETH) in Jan. 2022 to around 10 Gwei in Aug. 2022. However, the average block size per batch has remained around 102.9% of the target throughout this period. Interestingly, the average block size dropped to around 102.0% immediately after the switch to PoS. The reason behind this drop is a stimulating direction for future research. One hypothesis is that before PoS, the inter-block times had an (approximately) exponential distribution, whereas they are constant in Ethereum’s PoS protocol [6]. As such, the variance of N_n is smaller and block sizes are more regular. Another hypothesis is that blocks are less congested after the switch due to a decrease in inter-block times from roughly 13 seconds on average⁹ to 12 seconds.

7 Conclusions

In this paper, we have formally analyzed the long-term performance of the standard EIP-1559 transaction fee market mechanism and its closely related variants; exponential EIP-1559 among others. Our findings provide a theoretical justification for the anecdotal evidence that blocks are slightly more full than normal. As our main contribution, we have found that both designs, the baseline EIP-1559 and its exponential variant, are approximately and exactly optimal, respectively, even under the prevailing chaotic conditions in inter-block sizes. Importantly, this implies that these mechanisms can still achieve their goals even if the underlying system does not equilibrate, a condition that is rarely met in practice. The empirical data since the launch of EIP-1559 suggest that our results accurately capture reality: observable average block sizes are within the sharp approximation bounds predicted here and this is, in fact, robust to the underlying protocol functionality including both pre- and post-PoS merge periods.

Concerning future work, the current paper provides a framework to evaluate the performance and analyze the stability of transaction fee or other related cryptoeconomic mechanisms. Practical use cases suggest that blockchain economies are systems with complex dynamics: when these economies are close

⁹ https://ycharts.com/indicators/ethereum_average_block_time

to their equilibrium state, they can re-adjust their parameters and self-stabilize. However, once they are pushed further away and/or lose their peg to their fundamentals, they start to spiral away and eventually collapse (e.g., the Terra/Luna crypto network). Determining the limits in which these instabilities emerge already before such mechanisms are launched in practice, is critical to improve their efficiency and avoid future financial catastrophes.

Acknowledgements

This research/project is supported in part by the National Research Foundation, Singapore and DSO National Laboratories under its AI Singapore Program (AISG Award No: AISG2-RP-2020-016), NRF 2018 Fellowship NRF-NRFF2018-07, NRF2019-NRF-ANR095 ALIAS grant, grant PIESGP-AI-2020-01, AME Programmatic Fund (Grant No.A20H6b0151) from the Agency for Science, Technology and Research (A*STAR) and Provost’s Chair Professorship grant RGEPPV2101. It is also supported by the National Research Foundation (NRF), Prime Minister’s Office, Singapore, under its National Cybersecurity R&D Programme and administered by the National Satellite of Excellence in Design Science and Technology for Secure Critical Infrastructure, Award No. NSoE DeST-SCI2019-0009.

References

1. Google BigQuery. <https://cloud.google.com/bigquery>.
2. Soumya Basu, David Easley, Maureen O’Hara, and Emin Gün Sirer. StableFees: A predictable fee market for cryptocurrencies. *Available at SSRN: <https://ssrn.com/abstract=3318327> or <http://dx.doi.org/10.2139/ssrn.3318327>*, 2019.
3. V. Buterin, E. Conner, R. Dudley, M. Slipper, I. Norden, and A. Bakhta. EIP1559: Fee market change for ETH 1.0 chain. [Online], 2019.
4. Vitalik Buterin. Ethereum: A next-generation smart contract and decentralized application platform. [Online], 2014.
5. Vitalik Buterin. Make EIP 1559 more like an AMM curve. Ethereum Research [Online], 2021.
6. Vitalik Buterin, Diego Hernandez, Thor Kampefner, Khiem Pham, Zhi Qiao, Danny Ryan, Juhyeok Sin, Ying Wang, and Yan X Zhang. Combining GHOST and Casper. *arXiv preprint arXiv:2003.03052*, 2020.
7. Hao Chung and Elaine Shi. Foundations of transaction fee mechanism design. *arXiv preprint arXiv:2111.03151*, 2021.
8. Theo Diamandis, Alex Evans, Tarun Chitra, and Guillermo Angeris. Dynamic pricing for non-fungible resources. *arXiv preprint arXiv:2208.07919*, 2022.
9. Dankrad Feist. Exponential eip-1559. [Online], 2022.
10. Matheus VX Ferreira, 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*, pages 86–99, 2021.
11. Yotam Gafni and Aviv Yaish. Greedy Transaction Fee Mechanisms for (Non-)myopic Miners, 2022.

12. Adem Efe Gencer, Soumya Basu, Ittay Eyal, Robbert van Renesse, and Emin Gün Sirer. Decentralization in Bitcoin and Ethereum networks. In *International Conference on Financial Cryptography and Data Security*, pages 439–457. Springer, 2018.
13. Kostis Karantias, Aggelos Kiayias, and Dionysis Zindros. Proof-of-burn. In *International Conference on Financial Cryptography and Data Security*, pages 523–540. Springer, 2020.
14. Sam Kessler. The Ethereum Merge is done, opening a new era for the second-biggest blockchain. [Online], 2022.
15. Ron Lavi, Or Sattath, and Aviv Zohar. Redesigning Bitcoin’s fee market. *ACM Transactions on Economics and Computation*, 10(1):1–31, 2022.
16. Stefanos Leonardos, Barnabé Monnot, Daniël Reijnders, 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*, AFT ’21, page 114–126, New York, NY, USA, 2021. Association for Computing Machinery.
17. Yulin Liu, Yuxuan Lu, Kartik Nayak, Fan Zhang, Luyao Zhang, and Yinhong Zhao. Empirical analysis of EIP-1559: Transaction fees, waiting time, and consensus security. *arXiv preprint arXiv:2201.05574*, 2022.
18. Jason Milionis, Dean Hirsch, Andy Ardit, and Pranav Garimidi. A Framework for Single-Item NFT Auction Mechanism Design. *arXiv e-prints*, page arXiv:2209.11293, September 2022.
19. Barnabé Monnot, Qingze Hum, Chuan Sen Marcus Koh, Stefanos Leonardos, and Georgios Piliouras. Ethereum’s transaction fee market reform of EIP 1559. *Proceedings of the WINE 2020 Workshop on Blockchain*, 2020.
20. Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. [Online], 2008.
21. Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, USA, 2007.
22. Daniël Reijnders, Shyam Sridhar, Barnabé Monnot, Stefanos Leonardos, Stratis Skoulakis, and Georgios Piliouras. Transaction Fees on a Honeymoon: Ethereum’s EIP-1559 One Month Later. In *2021 IEEE International Conference on Blockchain (Blockchain)*, pages 196–204, 2021.
23. Tim Roughgarden. Transaction fee mechanism design for the ethereum blockchain: An economic analysis of eip-1559, 2020.
24. Tim Roughgarden. Transaction fee mechanism design. *ACM SIGecom Exchanges*, 19(1):52–55, 2021.
25. Elaine Shi, Hao Chung, and Ke Wu. What can cryptography do for decentralized mechanism design? *arXiv preprint arXiv:2209.14462*, 2022.

A Appendix

A.1 Omitted Proofs

Proof (Proof of Lemma 1). We prove the assertion using induction.

Base case: clearly, $b_0 \geq \min\{b_0, \alpha^{-1}b^*\}$ and $b_0 \leq \max\{b_0, \alpha b^*\}$ by the definition of min and max, so the lemma holds for the base case.

Induction step: we consider two cases:

1. $b_n \leq b^*$, then $b_{n+1} = h(b_n) \geq b_n \geq \min\{b_0, \alpha^{-1}b^*\}$ where the first inequality follows from h being a proper update rule (non-divergence) and the second from the induction hypothesis. Furthermore,

$$b_{n+1} = h(b_n) \leq ab_n \leq \alpha b^* \leq \max\{b_0, \alpha b^*\},$$

where the first inequality follows from h being a proper update rule (bounded relative differences) and the second from the case assumption.

2. $b_n \geq b^*$: similar as above. \square

Proof (Proof of Lemma 2). Part (i) is well known and follows directly from the Taylor expansion of $\ln(1+x)$, i.e.,

$$\ln(1+x) = \sum_{n=1}^{\infty} (-1)^{n+1} \frac{x^n}{n} \leq (-1)^{1+1} \cdot \frac{x^1}{1} = x.$$

To obtain part (ii), i.e., to lower bound $\ln(1+x)$ by a linear function $y(x) = ax + b$ for $|x| \leq d$, we only need to require that $y(-d) = \ln(1-d)$ and $y(d) = \ln(1+d)$. This follows directly from the fact that $\ln(1+x)$ is concave, and thus, its graph lies above the line that connects any two of its points for all intermediate values between these points. Solving the system of these two equations, yields

$$y(x) = \frac{\ln(1+d) - \ln(1-d)}{2d}x + \frac{\ln(1+d) + \ln(1-d)}{2},$$

as claimed. \square

Proof (Proof of Theorem 2). We start from equation (EXP-1559) and take the natural logarithm of both sides

$$\ln(b_{n+1}/b_n) = \left(\frac{g_n - T}{T} \right) \cdot \ln(1+d),$$

which after some rearranging of the terms yields

$$g_n = \frac{T}{2} \cdot \frac{\ln(b_{n+1}/b_n)}{\ln(1+d)} + \frac{T}{2}.$$

Summing up all terms from 1 to N , we obtain that

$$\frac{1}{N} \sum_{n=1}^N g_n = \frac{T}{2N \ln(1+d)} \cdot \sum_{n=1}^N \ln(b_{n+1}/b_n) + \frac{1}{N} \cdot \frac{N \cdot T}{2}. \quad (6)$$

The sum on the left hand side telescopes to

$$\sum_{n=1}^N \ln(b_{n+1}/b_n) = \sum_{n=1}^N (\ln b_{n+1} - \ln b_n) = \ln b_N - \ln b_1.$$

Thus, (6) becomes

$$\frac{1}{N} \sum_{n=1}^N g_n = \frac{T(\ln b_N - \ln b_1)}{2N \ln(1+d)} + \frac{T}{2}.$$

To conclude observe that $\lim_{N \rightarrow +\infty} b_N < M$ for some $M > 0$, since $(b_n)_{n \geq 0}$ is bounded. This implies that

$$\lim_{N \rightarrow +\infty} \frac{T(\ln b_N - \ln b_1)}{2N \ln(1+d)} = 0.$$

Thus, taking the limit $N \rightarrow +\infty$ on both sides of the previous equation, we obtain that

$$\lim_{N \rightarrow +\infty} \frac{1}{N} \sum_{n=1}^N g_n = T,$$

as claimed. \square

From the proof of Theorem 2, we can also reason about the convergence rate. After N time-steps, we have that the distance $d(G_N, T)$ between $G_N := \frac{1}{N} \sum_{t=0}^N g_t$ and the target T is equal to

$$d(G_N, T) = \frac{T(\ln b_N - \ln b_1)}{2 \ln(1+d)} \cdot \frac{1}{N}$$

which drops linearly in N , i.e., $\mathcal{O}(1/N)$. The constant factor depends on $T, \ln 1+d$ and the error due to initialization, $\ln b_N - \ln b_1$. Also these terms vanish at a rate of $1/N$.

B General BFURs

In this section, we discuss some general BFURs which have been proposed by the research community or which have appeared in the relevant literature. To allow dependencies on more parameters (other than the current base fee), we first provide some additional definitions and notation.

Let $\Theta \subseteq \mathbb{R}^d$ for some integer $d > 0$ denote a *parameter* space. Generalized base fee update rules (GBFURs), are functions $h : (0, \infty) \times \Theta \rightarrow (0, \infty)$. In other words, GBFURs condition their updates on both the current base fee, b_n , and on additional time-dependent or time-independent information, e.g., the current block-size, g_n , or the initial base fee, b_0 etc. The notion of *proper* GBFURs can be adapted accordingly (cf. Definition 1).

Definition 2 (Generalized Proper Base Fee Update Rules (GPBFURs)).

Let T be the target block-size and let Θ be a parameter space. An generalized update rule $h : (0, \infty) \times \Theta \rightarrow (0, \infty)$ is called *proper* if it satisfies the following

- (GA.1) non-divergence: $h(b_n, \theta_n) \geq b_n$ if $g(b_n) \geq T$ and $h(b_n, \theta_n) \leq b_n$ if $g(b_n) \leq T$.
- (GA.2) bounded relative differences: there exists an $\alpha \geq 1$ such that $\alpha^{-1}b_n \leq h(b_n, \theta_n) \leq \alpha b_n$.

Using the above, we can now proceed to the study of general BFURs.

AMM-based mechanisms: Another alternative to (EIP-1559) that has been proposed is the AMM-based mechanism which is defined as follows [5]. Let x_n denote the stochastic process of the *excess gas* issued. The evolution of x_n is given by

$$x_{n+1} = \max\{0, x_n + g(b_n) - T\} \quad (7)$$

with $x_0 = 0$. Then, the AMM base fee update rule is given by

$$b_{n+1} = qe^{qx_{n+1}} \quad (\text{AMM})$$

where q is the *adjustment quotient* which determines how the base fee is chosen based on the current excess gas (q is similar to parameter d in the standard EIP-1559 update). From equations (7) and (AMM), we have that

$$b_{n+1} = b_n \cdot e^{q \max\{-x_n, g(b_n) - T\}}. \quad (\text{AMM-closed})$$

The (AMM-closed) base fee update rule depends on both the current base fee, b_n , and the current excess gas, x_n . Thus, it is a generalized update rule. Since $x_n \geq 0$, we have that $\max\{-x_n, g(b_n) - T\} = g(b_n) - T$ whenever $g(b_n) \geq T$. This implies that $h(b_n, x_n) \geq b_n$ if this is the case. Similarly, if $g(b_n) < T$, then $\max\{-x_n, g(b_n) - T\} < 0$ which implies that $b_{n+1} = h(b_n, x_n) < b_n$. Thus, (AMM-closed) satisfies (GA.1). Since e^y is continuous for any $y \in \mathbb{R}$, it is also straightforward to show that (AMM-closed) satisfies (GA.2). Thus, (AMM) is a GPBFURs.

Upper bound on average block sizes: We next derive an upper bound for the average block sizes with the (AMM) mechanism. Equation (7) implies that $x_{n+1} \geq x_n + g(b_n) - T$ or equivalently that $g(b_n) \leq x_{n+1} - x_n + T$. Thus,

$$\frac{1}{N} \sum_{n=1}^N g(b_n) \leq \frac{1}{N} \sum_{n=1}^N (x_{n+1} - x_n) + T = \frac{1}{N} (x_{N+1} - x_1) + T$$

If $\lim_{N \rightarrow \infty} x_N < X$ for some X , i.e., if x_N stays bounded, then $\lim_{N \rightarrow \infty} \frac{1}{N} (x_{N+1} - x_1) = 0$ and therefore

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N g(b_n) \leq T.$$

Alternatively, from equation (AMM-closed), we have that

$$\frac{1}{N} \sum_{n=1}^N \max\{T - x_n, g(b_n)\} = T + \frac{1}{qN} (\ln b_N - \ln b_1).$$

Again, if b_N is bounded (which is always the case from below by $L = 7$ and also from above, under the assumption that users' valuations are bounded), then

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N \max\{T - x_n, g_n\} = T.$$

Sufficient condition for convergence of AMM to T . If $x_n \geq T$ for some $n \geq 0$, then,

$$\max \left\{ \frac{T}{2} - x_n, g_n \right\} = g_n.$$

To deal with the case $0 \leq x_n \leq T$, one sufficient condition is to assume that for each such x_n , it holds $g_n(b_n) > T - x_n$, where $b_n = \frac{d}{T} e^{\frac{d}{T} x_n}$. To obtain a simple sufficient condition, we may require that

$$\min_{x_n \in [0, T]} g_n(b_n) > \max_{x_n \in [0, T]} \{T - x_n\} = T.$$

Using that g_n is decreasing in b_n , and that b_n is strictly increasing in x_n , this reduces to requiring that $g_n(b_n(T)) > T$, where $b_n(T) = \frac{d}{T} e^d$.

In other words, if there are sufficient users (i.e., more than T) with valuations above $\frac{d}{T} e^d$, then the max in the previous summation is always resolved by g_n and we obtain convergence to T . In symbols, we can formulate the sufficient condition as

$$\lambda \bar{F} \left(\frac{d}{T} e^d \right) > 1$$

or equivalently (if we solve for d),

$$d e^d < T \bar{F}^{-1}(\lambda/2)$$

where the inequality was reversed, since the inverse of a decreasing function (here \bar{F}) is also decreasing.

Finally, if $b^* = \bar{F}^{-1} \left(\frac{T}{2\lambda} \right)$, then by (1) we have that $g(b^*) = T$. Moreover, if for some $n > 0$, it holds that $x_n > 0$ and that $b_n = b^*$, then

$$x_{n+1} = \max\{0, x_n\} = x_n,$$

which implies that there is no change in the excess gas. In this case, it holds that

$$x_n = q^{-1} \ln(q^{-1} b^*).$$

This can be also seen in the example of Figure 6, where the valuations are drawn uniformly at random from $[200, 230]$ with $\lambda = 1$, and we use an AMM-based mechanism with $q = \frac{1}{10}$. The market clearing price $b^* = 215$, and the excess gas converges to $10 \cdot \ln(2150) \approx 76.73$.

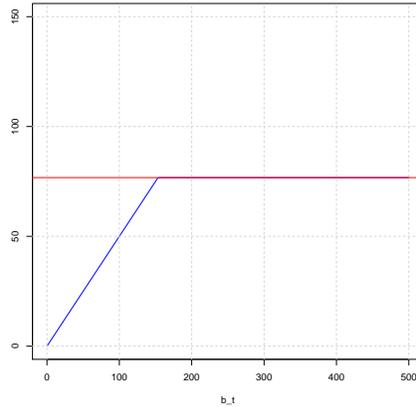


Fig. 6. Excess gas over time in the AMM mechanism for uniform demand distribution.

Dynamic Posted-Price Mechanisms: Another approach to transaction fee mechanisms that has been proposed in the literature [10] is the following valuation-based or *welfare-based* approach

$$b_{n+1} = \frac{\alpha}{kT} \sum_{i=1}^{g_n} v_{i,n} + (1 - \alpha)b_n \quad (\text{WEL})$$

and the corresponding *truncated welfare-based* alternative:

$$b_{n+1} = \begin{cases} \alpha(1 + \delta)b_n + (1 - \alpha)b_n, & \text{if } g_n = kT \\ \frac{\alpha}{kT} \sum_{i=1}^{g_n} \min\{v_{i,n}, (1 + \delta)b_n\} + (1 - \alpha)b_n, & \text{else.} \end{cases} \quad (\text{T-WEL})$$

where kT denotes the block capacity. These Dynamic Posted Price mechanisms condition their updates also on the $v_{i,n}$, i.e., the valuations of the included transactions at time n . However, their scope is to maximize welfare rather than meeting a specific target T . Accordingly, condition (GA.1) is irrelevant in this case. However, assuming that valuations are bounded, it is again straightforward to show that these mechanisms satisfy (GA.2).

Effective Gas Price Correction Update Rule: In some cases, when demand shifts abruptly upwards, the EIP-1559 transaction fee market reverts back to a first price auction (FPA). During such times, e.g., during demand bursts due to NFT drops or other reasons, the prevailing base fee is too low to control the changing demand. This results in a series of full blocks in which all included transactions pay unreasonably high tips (FPA conditions) before the base fee reaches the appropriate level to achieve half-full blocks.

To counter the above phenomenon several alternative mechanisms have been proposed [10,22]. However, these mechanisms treat deviations from the target block size equally (treat all kinds of deviations equally) or rely their correction terms again on the previous block-sizes.

However, FPA conditions are unlike any other deviations from target block-sizes due to the distinctive characteristic of very high tips to miners, i.e., of very high effective gas prices in comparison to the prevailing base fee. Next, we use this intuition to design a modified EIP-1559 mechanism that precisely aims to mitigate such FPA conditions.

Let m_n denote the minimum effective gas price paid at block height t , and let $\delta_n := m_n - b_n$. By definition, $\delta_n > 0$. If δ_n is large, e.g., $\delta_n > b_n$, then this means that all transactions in block t were selected under first price auction (FPA) conditions. By contrast, if δ_n is small, e.g., $\delta_n < 0.1b_n$, then this means that the current base fee is correct for at least some users. In the first case, (large δ_n), the mechanism needs to aggressively update b_n , whereas in the second case, the mechanism is ok with a normal update. This motivates the following modification of the standard EIP-1559 update rule

$$b_{n+1} = b_n \left(1 + d \cdot \frac{g_n - T}{T} \right) + \lambda \cdot \text{ReLU}(\delta_n - \gamma b_n),$$

where $\text{ReLu}(x) := \max(0, x)$. Since $\delta_n = m_n - b_n$, we can rewrite the above as

$$b_{n+1} = \underbrace{b_n \left(1 + d \cdot \frac{g_n - T}{T} \right)}_{\text{EIP-1559}} + \underbrace{\lambda \cdot \text{ReLu}(m_n - (1 + \gamma)b_n)}_{\text{EGP Correction}},$$

where m_n is the minimum effective gas price of block t . We will refer to this update rule as the *Effective Gas Price Correction Update Rule* (EGP-CURE). In comparison to EIP-1559, EGP-CURE has two additional (learnable) parameters: (i) γ which is the threshold that triggers the effective gas price correction mechanism and λ which calibrates the intensity of the updates in case that the correction mechanism has been triggered. Note that, similar to g_n , m_n can be directly estimated from the observable data that is included in block t .

EGP-CURE is not PFUR If we assume that user valuations are bounded by some constant M , then the EGP-CURE update rule, $b_{n+1} := h(b_n)$ defined above, can be easily shown to satisfy the generalized (A.2') condition.

Lemma 3. *Assume that user valuations are bounded by some constant M . Then, the EGP-CURE defined above by $b_{n+1} := h(b_n)$ satisfies generalized bounded relative differences condition, (A.2').*

Proof. If $m_n \leq (1 + \gamma)b_n$, then EGP-CURE degenerates to EIP-1559 which has been shown to satisfy the more restrictive (A.2). Thus, it remains to treat the case $m_n > (1 + \gamma)b_n$. In this case, EGP-CURE becomes

$$b_{n+1} = b_n \left(1 + d \cdot \frac{g_n - T}{T} \right) + \lambda \cdot (m_n - (1 + \gamma)b_n).$$

and we have that

$$b_{n+1} = b_n \left(1 + d \frac{g_n - T}{T} - \lambda(1 + \gamma) \right) + \lambda m_n \leq b_n(1 + d - \lambda(1 + \gamma)) + \lambda M$$

where the last inequality holds because of the bounded valuations assumption, i.e., $m_n < M$ for any $t \geq 0$. This means that EGP-CURE satisfies condition (A.2') with constants $\alpha \equiv (1 + d - \lambda(1 + \gamma))$ and $\beta \equiv \lambda \cdot M$. Showing that $h(b_n)$ is bounded from below is trivial (since the correction term is always nonnegative) which concludes the proof.

Concerning condition (A.1), it is immediate to show that EGP-CURE satisfies $h(b_n) \geq b_n$ if $g_n \geq T$. However, it does not necessarily hold that $h(b_n) \leq b_n$ if $g_n \leq T$. For instance, if there are only few transactions in a block (less than half-full block), but they all pay very high tips (above the threshold), then the correction mechanism will be triggered and, under certain mathematical conditions, it will not allow the base fee to decrease (despite a less than full block).

Such a case may only be of theoretical interest, since it implies a clearly suboptimal behavior in which users pay unreasonably very high fees to get their

transactions included in an otherwise non-full block. Nevertheless, if one wants to exclude such a case, one may additionally impose that the EGP correction term is triggered only if $g_n > T$, i.e., only if the block is more than half-full.

C Empirical Evaluation

In this section we describe the blockchain data from after EIP-1559 used to construct Figure 1.

Description of Data. To collect blockchain data, we have used Google’s BigQuery data service [1]. Our sample includes all blocks b_n whose index t is in $\{t^-, \dots, t^+ - 1\}$ where $t^- = 12965000$ and $t^+ = 15200000$. In particular, t^+ is the index of the first block after EIP-1559, and block t^+ was mined on 23 July, 2022. We group the blocks of our sample into $M = 447$ consecutive batches, such that each batch consists of $N = (t^+ - t^-)/M = 5000$ blocks. For each batch $i = 1, \dots, M$, let $t_i^- = t^- + Ni$ and $t_i^+ = t^- + N(i + 1) - 1$ be the indices of the first and last block in batch i , respectively. We then define $y_i \triangleq \frac{1}{N} \sum_{t=t_i^-}^{t_i^+} g_n/T$ as the average relative gas use in the blocks of batch i . We note that $g_n/T \in [0, 1]$ for all t , so it must hold that $y_i \in [0, 1]$ as well. To calculate y_i for each batch, we use the query from Figure 7. Our approach is fully reproducible – anyone with a BigQuery account can use the query from Figure 7 to reproduce the data in this section.

Analysis. We have visualized the data created by the query of Figure 7 in Figure 1. The solid blue line in Figure 1 depicts the evolution of y_i across the observation period, whereas the red dashed line depicts the sample average $\mu_y = \frac{1}{M} \sum_{i=1}^M y_i$, which equals 0.5145. We observe that the batch averages occur in a narrow band around μ_y . From Figure 4, we observe that the theoretical upper bound of the block size for $d = 0.125$ equals $\left(1 - \frac{\ln(1.125)}{\ln(0.875)}\right)^{-1} \approx 0.5313$. This shows that, as predicted by our analysis, the average block sizes both exceed the target relative block size of 0.5, but not to a degree at which they would exceed our theoretical upper bound.

D Extensive Simulations

In this part, we present systematic simulations of the (EIP-1559) and (EXP-1559) BFURs. We consider two bifurcation parameters: (1) the adjustment quotient, d , and (2) the range of valuations, which is parameterized by w (see Table 1).

Adjustment quotient: For the adjustment quotient, we allow values in the $(0, 0.5]$ range, since lower values are not admissible and larger values are not relevant to practice (they result in extremely aggressive updates). We consider three types of distributions for user valuations which are presented in Table 1. To define the

```

SELECT
  AVG(gas_used/gas_limit) AS block_size,
  MIN(number) AS min_number,
FROM (
  SELECT
    gas_used,
    gas_limit,
    number,
    FLOOR(number/5000) as bucket
  FROM `bigquery-public-data.crypto-ethereum.blocks`
  WHERE number >= 12965000 AND number < 15200000
)
GROUP BY bucket ORDER BY bucket

```

Fig. 7. The SQL query used in Google BigQuery to produce the data for Figure 1.

Distribution	Parameters (μ, σ^2)	Support
Uniform	m $w^2/12$	$[m - w/2, m + w/2]$
Normal	m $w^2/16$	$(m - w, m + w)$
Gamma	$m - aw$ w^2a	$(m - aw, +\infty)$

Table 1. Distributions of users valuations that have been used in the simulations. For the normal distribution, the support refers to the 4σ interval ($> 99.99\%$ of all values). For the gamma distribution, special cases occur when a is an integer (Erlang distribution) and when $a = 1$ (exponential distribution).

distributions, we use parameters, m (equal or related to the mean), w (related to standard deviation) and a (only necessary for gamma distributions). For the simulations below, we fix $m = 210, w = 20$ and for the gamma distribution, $a = 0.5$, and allow d to vary between $(0, 0.5]$ as mentioned above. In all cases, we assume mean arrival rate 4 times the target, T , i.e., two times the block size. The results are shown in Figures 8 and 9. The blue dots (trajectories of base fees and block sizes) show data from $N_{iter} = 100$ iterations after skipping $N_{skip} = 200$ iterations. Allowing more iterations (either N_{skip} or N_{iter}) does not produce different attractors which suggests that the plots show the steady state, i.e., what happens after the mechanism has converged. The initial base fee is set to $b_0 = 170$.¹⁰

Qualitatively similar results are obtained for different values of parameters m, w and a (not presented here).

Range of valuations: For this part, we fix d to its default value, $d = 0.125$, and vary w which parameterizes the standard deviation (variance) of the distribution of users valuations. The range of users valuations has a significant effect on the stability of both BFURs since more (less) concentrated valuations tend to

¹⁰ The Jupyter notebooks that have been used to generate these plots are available as supplementary material (or upon request).

produce more (less) aggressive changes in block sizes. The rest of the parameters, i.e., $Niter$, $Nskip$, b_0 , are the same as above. For all cases, we allow w to vary between 0 (very concentrated) and 20 (less concentrated). The results are presented in Figures 10 and 11.

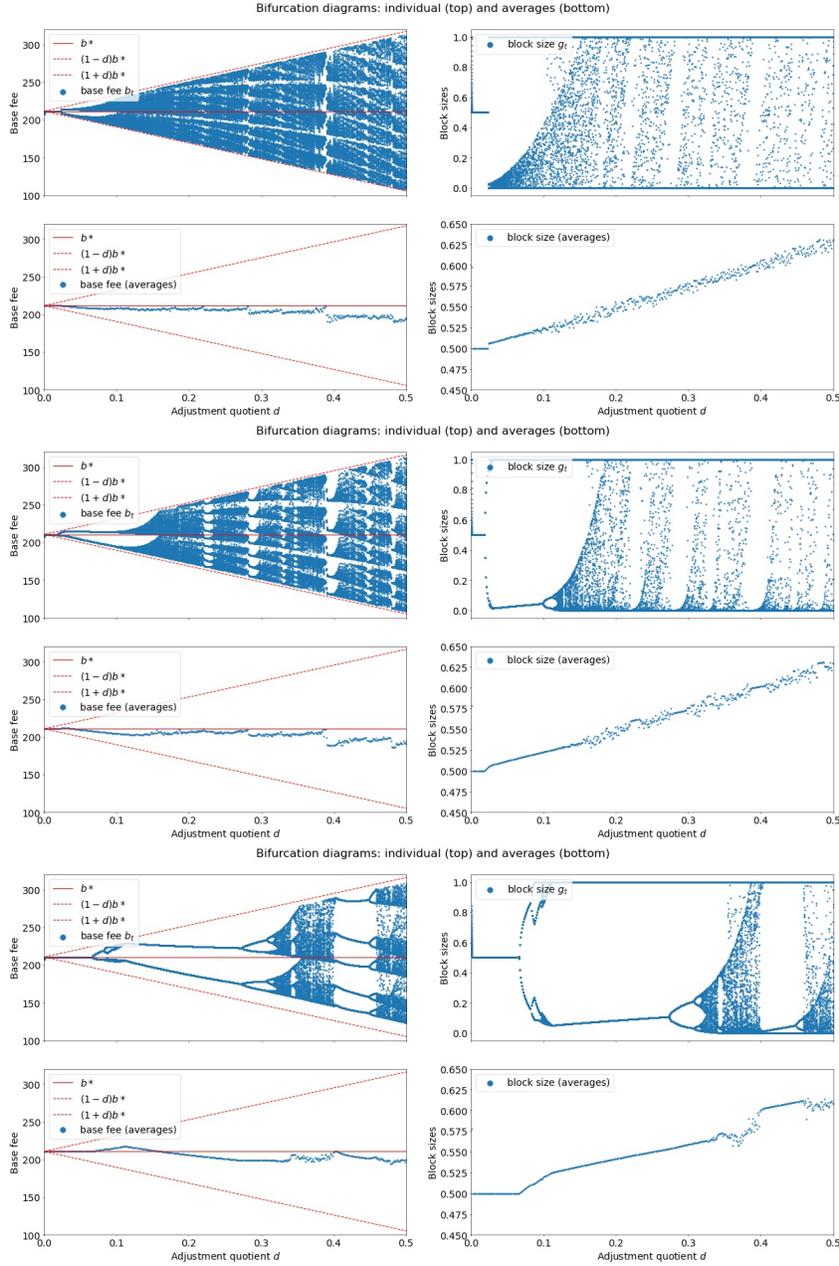


Fig. 8. Simulations of (EIP-1559) for different values of the adjustment quotient (bifurcation parameter, d) for uniform (top), normal (middle) and gamma (bottom) distributions. Note: similar to Figures 3 and 5, the scale of the y -axis in block-size panels (top- and bottom-right) are between 0 and 1 (target block size is 0.5).

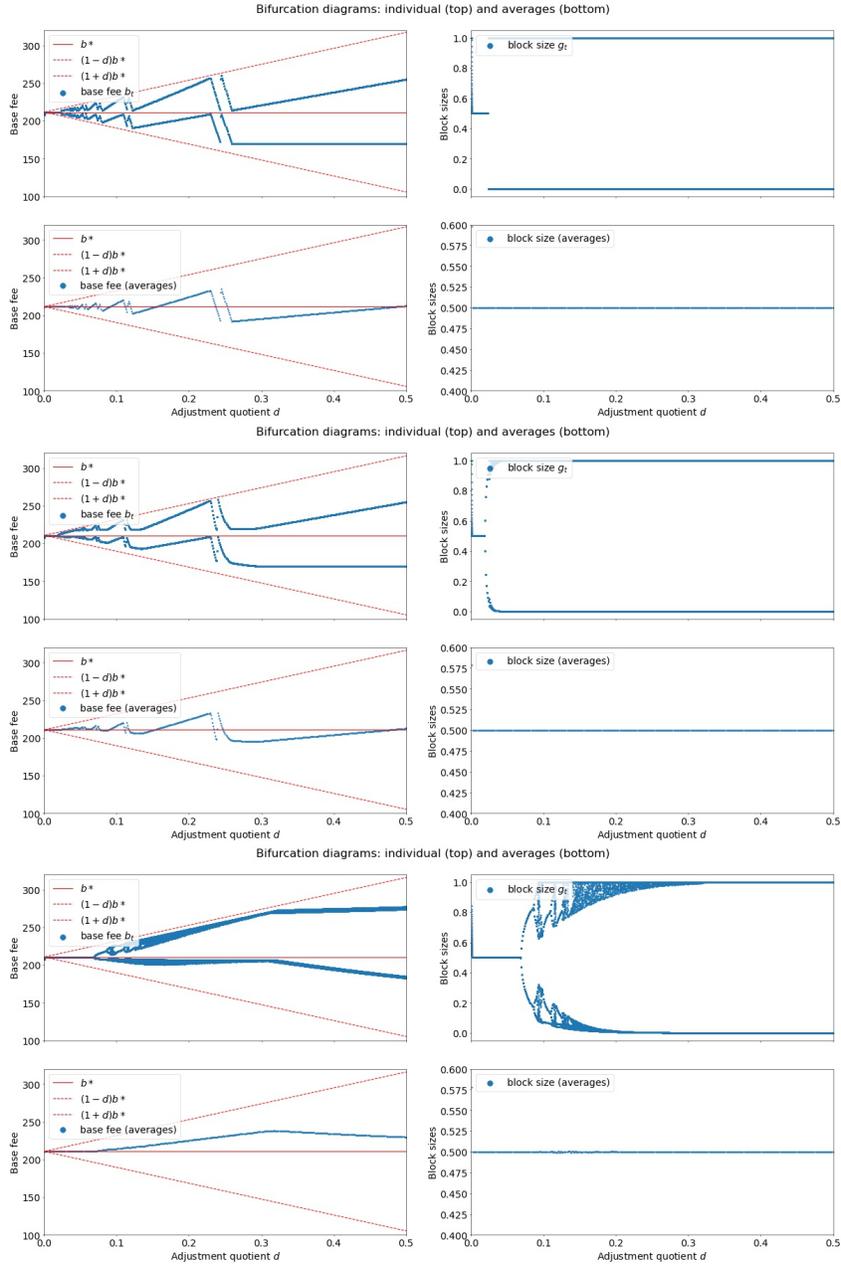


Fig. 9. Simulations of (EXP-1559) for different values of the adjustment quotient (bifurcation parameter, d) for uniform (top), normal (middle) and gamma (bottom) distributions. The individual base fee trajectory depends on the initial conditions, but averages and block sizes are robust regardless of the exact base fee realization.

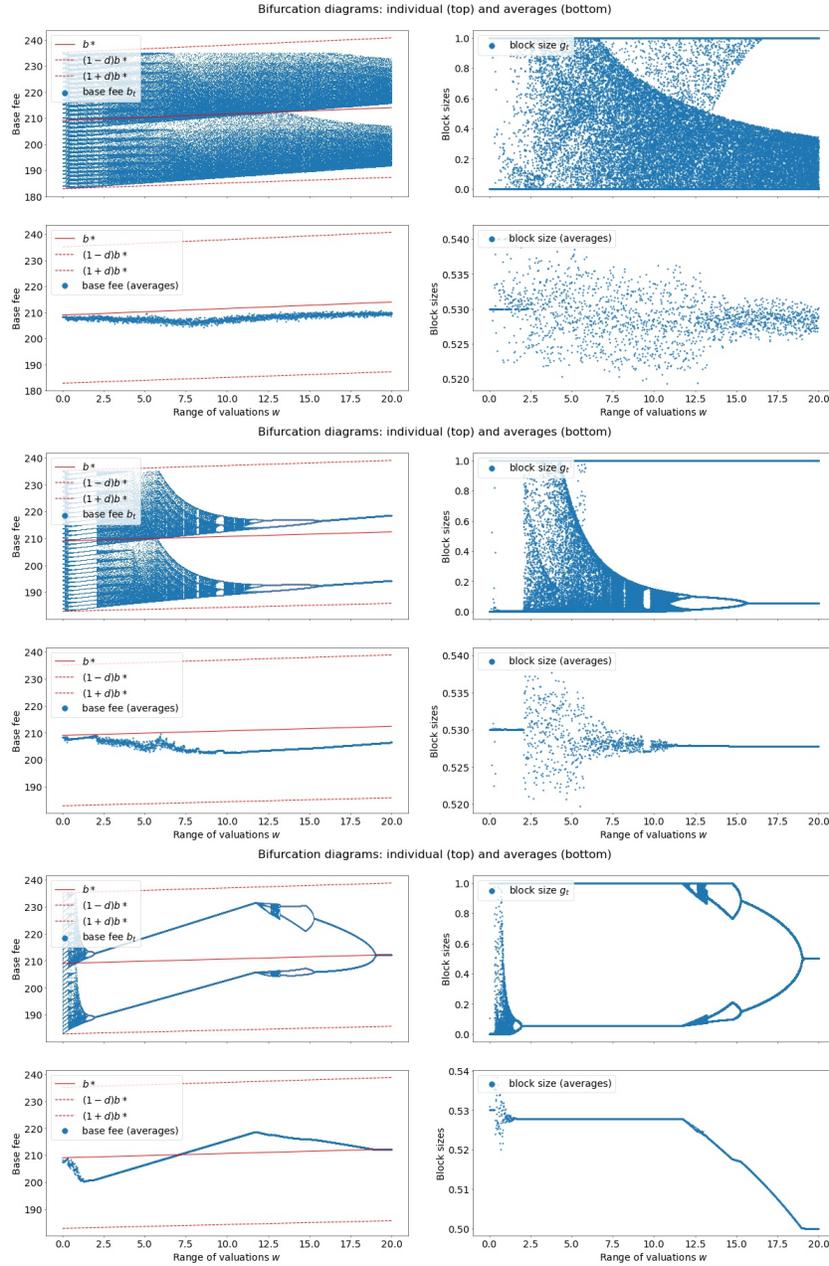


Fig. 10. Simulations of (EIP-1559) for default $d = 0.125$ and uniform (top), normal (middle) and gamma (bottom) distributions with varying range of valuations (bifurcation parameter, w). The individual base fee trajectory depends on the initial conditions, but averages and block sizes are robust regardless of the exact base fee realization.

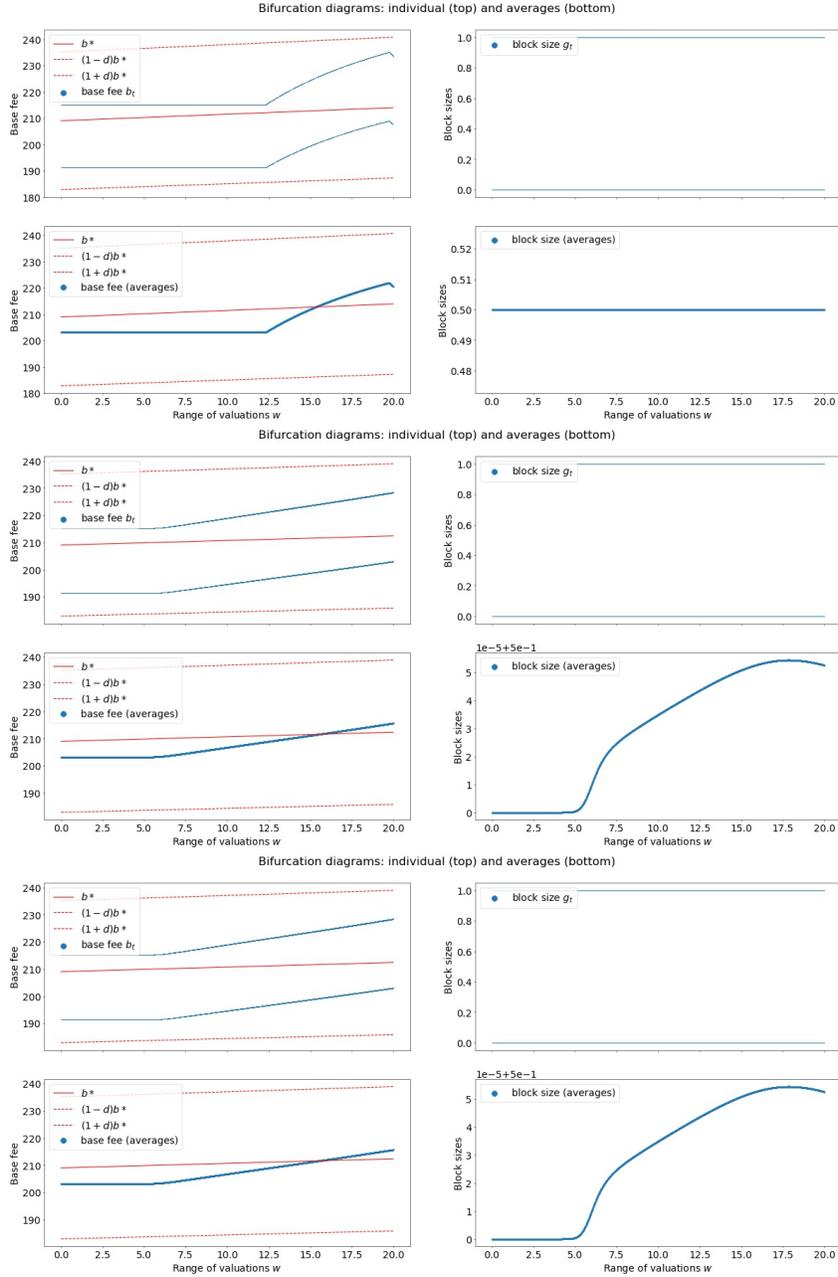


Fig. 11. Simulations of (EXP-1559) for default $d = 0.125$ and uniform (top), normal (middle) and gamma (bottom) distributions with varying range of valuations (bifurcation parameter, w). For the average block size plots (bottom left panels), note that $1e - 5 + 5e - 1 = 0.5001$.