

Hougaard, Jens Leth; Pourpouneh, Mohsen

Working Paper

Farsighted miners under transaction fee mechanism EIP1559

IFRO Working Paper, No. 2022/04

Provided in Cooperation with:

Department of Food and Resource Economics (IFRO), University of Copenhagen

Suggested Citation: Hougaard, Jens Leth; Pourpouneh, Mohsen (2022) : Farsighted miners under transaction fee mechanism EIP1559, IFRO Working Paper, No. 2022/04, University of Copenhagen, Department of Food and Resource Economics (IFRO), Copenhagen

This Version is available at:

<https://hdl.handle.net/10419/266990>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



IFRO Working Paper

Farsighted Miners under Transaction Fee Mechanism EIP1559

*Jens Leth Hougaard
Mohsen Pourpouneh*

IFRO Working Paper 2022 / 04

Farsighted Miners under Transaction Fee Mechanism EIP1559

Authors: Jens Leth Hougaard, Mohsen Pourpouneh

JEL-classification: D47, D53, L11, L17

Published: December 2022 (This paper is an update of an earlier version with the same number published August 2022)

See the full IFRO Working Paper series here:

https://ifro.ku.dk/english/publications/ifro_series/working_papers/

See other publications from the Center for Blockchains and Electronic Markets here:

<https://bcm.ku.dk/publications/>

Department of Food and Resource Economics (IFRO)

University of Copenhagen

Rolighedsvej 23

DK 1958 Frederiksberg DENMARK

<https://ifro.ku.dk/english/>

Farsighted Miners under Transaction Fee Mechanism EIP1559^{*}

Jens Leth Hougaard¹ and Mohsen Pourpouneh¹

University of Copenhagen, Copenhagen, Denmark {jlh, mohsen}@ifro.ku.dk

Abstract. We investigate the recent fee mechanism EIP1559 of the Ethereum network. Whereas previous studies have focused on myopic miners, we here focus on strategic miners in the sense of miners being able to reason k -blocks ahead. We derive expressions for optimal miner behavior (in terms of setting block sizes) in the case of 2-block foresight and varying degrees of hashing power. Results indicate that a sufficiently large mining pool will have enough hashing power to gain by strategic foresight. We further use a simulation study to examine the impact of both 2-block and 3-block foresight. In particular, the simulation study indicates that for realistic levels of hashing power, mining pools do not gain from the being able to reason more than 2 blocks ahead. Moreover, even though the presence of strategic miners increase the variation in block sizes and potentially empty blocks, overall system throughput tend to increase slightly compared to myopic mining.

Keywords: Blockchain · Ethereum · Transaction fee mechanism.

JEL Classification: D47, D53, L11, L17.

1 Introduction

On August 5, 2021, at block 12,965,000 the Ethereum network (Buterin et al. (2014)) implemented the so-called “London Hard Fork” (Buterin et al. (2019); Beiko (2021)), which changed its fee mechanism from what can loosely be described as a type of first price auction: blocks had a maximum (standard) size in terms of units of gas, and users attached a “bid” to their transaction indicating what they were willing to pay per unit of gas if their transaction was executed (i.e., included in a block). Miners typically fill up the blocks to standard size by selecting those transactions with the highest “bid”, hence maximizing their revenue for the current block. Consequently, this fee mechanism inherits the typical issues with first price auctions, Buterin (2021); Maskin et al. (2001), e.g., lack of user incentive compatibility, instability of the blockchain in the absence of the block reward, and over-bidding for faster inclusion in the block.

^{*} We wish to thank Jens Gudmundsson, Haiting Han and Luyao Zhang for helpful comments on earlier version of this paper. Support by the Center for Blockchains and Electronic Markets, funded by the Carlsberg Foundation under grant no. CF18-1112, is gratefully acknowledged.

The new fee mechanism - dubbed EIP1559 - is an attempt to regulate demand and allow flexibility by changing the maximum block size (Buterin (2018, 2016, 2014)). To achieve this, the fee of the coming block is determined by the size of the current block and thereby known to the users akin to a dynamically adjusted posted price. This fee is known as the *base (or network) fee*, which is “burned” by the network (and thus not paid to miners).¹ The base fee is updated algorithmically and depends on the on-chain data. That is, if the size of the current block exceeds standard size (also called the *target* size) T then the base fee increases for the coming block, and similarly, if the size of the current block is below T , the base fee will decrease for the coming block. Blocks are now allowed to be of a size between 0 and $2T$ units of gas and users now attach a cap on the total fee they are willing to pay, f . The total fee (i.e., user payment) consists of the base fee plus a so-called “miner tip” , p , that the miner receives for including the transaction. A transaction is eligible to be included in a block if the network fee does not exceed the cap. If an eligible transaction is included in a block the user pays the minimum of the base fee plus miner tips and the total fee cap. The miner only receives the minimum of the miner tip and the cap minus the base fee. A myopic miner will therefore try to fill up the block as much as possible, including transactions with the highest miner tips. However, since the current block size influences the network fee of the coming block, a farsighted miner may act strategically when deciding on which transactions they want to include in the block.

First, we consider a miner who is able to reason ahead when deciding which transactions to include in a two-block sequence (i.e., a miner with 2-block foresight). Assuming that users’ willingness to pay are i.i.d. and uniformly distributed, we identify the optimal (i.e., revenue maximizing) block sizes for a single (monopolist) miner with 2-block foresight. Unlike a myopic miner, a miner with 2-block foresight may find it optimal to publish empty blocks depending on the size of the initial base fee. Since miners in practice compete to mine blocks, we also consider the percentage of hashing power needed in order to make it profitable for a miner to mimic the optimal strategy of a monopolist rather than being myopic: assuming that the rest of the network is myopic. Using a simulation study, we then demonstrate that a large mining pool (e.g. such as Ethermine²) can have hashing power enough to prefer publishing empty blocks rather than to mine myopically. In a similar fashion we repeat the exercise for miners with 3-block foresight and find again that publishing empty blocks may increase revenues compared to being myopic. However, our simulations show that nothing is gained in terms of aggregate revenue to having 2-block foresight. We conjecture that this remains true for higher levels of foresight as well for mining pools of realistic sizes (below 50% hashing power). Intuitively, with realistic levels of hashing power, the probability of mining several blocks in a row decreases fast and thus the ability to harvest the benefits of lowering the base

¹ By itself this will have deflationary effect benefiting all agents holding the cryptocurrency.

² <https://etherchain.org/miner>

fee. For small mining pools, in effect there is very little, if anything, to gain from having foresight of more than two blocks ahead and we conjecture that miners with higher levels of foresight optimally would mimic the strategy of a 2-block farsighted miner.

Our simulation study demonstrates that strategic mining behavior leads to more variation in base fees and block sizes, but somewhat surprisingly does not seem to lower system through-put. We further show that by changing the rule which up-dates the base fee to a rule which updates on the basis of the average of the two most recent blocks, lowers the variance of the base fee significantly even for myopic miners. Consequently, a strategic farsighted miner will now need more hashing power to gain by strategic manipulation of the base fee. Our conclusions appear to be robust to changes in the distributional assumption on users’ willingness to pay.

1.1 Related literature

Our paper follows up on a recent line of papers analysing various aspects of EIP1559 and fee mechanisms in general.

The economic properties of the EIP1559, were first studied in Roughgarden (2020, 2021). In this work, it is assumed that the miners are myopic, in the sense that they will always fill the block as much as possible without including any “fake” transactions. The intuition is that the base fee is burned by the network, while manipulating the base fee is costly for the miners. Furthermore, it is shown that in cases where the base fee is high enough to limit the set of eligible transactions to be below the maximum block size, the users’ “obvious optimal equilibrium bid” is to set the fee cap equal to the true valuation of the user and a tip which is equal to the miners’ marginal cost of executing the user’s transaction. Consequently, EIP1559 is incentive compatible for both users and myopic miners. In Chung and Shi (2021), the authors prove a conjecture in Roughgarden stating that no non-trivial incentive compatible mechanism can prevent (off-chain) miner-user collusion. As a result, they propose a new mechanism dubbed as “burning second-price auction”. Alternatives to the EIP1559 mechanism has also been discussed in Ferreira et al. (2021). They suggest a new mechanism (dubbed the dynamic posted-price-mechanism) which not only takes the size of the previous blocks into account but also the bids from the previous blocks in order to compute a fee for the subsequent blocks. Moreover, in Roughgarden (2020), several variations of EIP1559, such as the EIP2593 (dubbed as the escalator) are discussed.

Monnot et al. (2020); Leonardos et al. (2021), consider a dynamic system to evaluate the evolution of the base fee over time, and provide upper, and lower, bounds on the base fee given that miners and users are not speculating on the current base fee.

In Liu et al. (2022), the authors provide an empirical study to examine the effect of EIP1559 on the transaction fee dynamics, transactions waiting time, and security of the blockchain. The results show an improvement on the user experience by making fee estimation easier, mitigating intra-block difference of

gas price paid, and reducing users' waiting times. Finally, Reijsbergen et al. (2021) proposes a dynamic updating rule for the base fee. That is, the base fee rather than being updated based on only the previous block, updates by a sliding window which takes into account multiple previous blocks.

In comparison to previous studies, we do not assume from the outset that miners are myopic. On the contrary, we consider strategic miners with a foresight of up to k blocks ahead. We further consider situations where the base fee is so low that the set of eligible transactions is larger than the maximal block size.

2 Model and notation

A formal description of EIP1559 can be found several places, e.g. Roughgarden (2020, 2021). Here, we let T denote the standard block size in units of gas. A (block)chain with height t is defined by a profile $S = (s_1, \dots, s_t)$ where s_i is the block size of the i 'th block in the chain.

In EIP1559, $s_i \in [0, 2T]$ and the base fee at block t , denoted b_t , is determined by

$$b_t = b_{t-1} \left(1 + \frac{1}{8} \cdot \frac{s_{t-1} - T}{T} \right) = b_0 \prod_{j=1}^t \left(1 + \frac{1}{8} \cdot \frac{s_{j-1} - T}{T} \right). \quad (1)$$

Thus, when $s_i < T$ the base fee of the coming block $i + 1$ decreases and vice versa when $s_i > T$. This allows the users to predict the next block's "reserve price". In EIP1559 every transaction T must specify two parameters: a fee cap f , which determines the user's maximum willingness to pay per unit of gas for their transaction to be processed, and a tip p , which is the (maximum) amount that the user is willing to pay the miner to include their transaction.

Any transaction will therefore cost the user the minimum of her fee cap and the total fee payment, i.e.:

$$\text{user payment} = \min\{f, p + b_t\} \quad (2)$$

Consequently, the miner's revenue (payoff) from including a transaction in the block is given by:

$$\text{miner payoff} = \min\{f - b_t, p\} \quad (3)$$

To simplify our analysis we will, throughout the paper, consider legacy transactions only. After implementation of EIP1559, legacy transactions (type 0) are still allowed: here users only determine the fee cap, f and (for compatibility) we can consider the tip p as being set equal to f . Thus, users always pay f per unit of gas, and the miners receive $f - b_t$ per unit of gas since b_t is burnt. Standard economic logic seems to indicate that over time users' willingness to tip the miner will decrease (resp. increase) with increasing (resp. decreasing) base fees since users care only about their total payment. Moreover, even with constant tips, miner revenue is weakly increasing with decreasing base fees because the base fee affects the number of eligible transactions.

3 EIP1559: some preliminary observations

An immediate implication of Equation 1 is that the ordering of the blocks in a given chain S has no effect on the size of the base fee at a given height t . Formally,

Observation 1: Let $\pi : \{1, \dots, t-1\} \rightarrow \{1, \dots, t-1\}$ be a permutation of the indices $1, \dots, t-1$. Then $b_t(S) = b_t(\pi S)$.

Thus, assume for convenience that block sizes are increasingly ordered $s_1 \leq \dots \leq s_t$. Consider two increasingly ordered chains of the same height t , $S = (s_1, \dots, s_t)$ and $S' = (s'_1, \dots, s'_t)$, with the same through-put, i.e., $\sum_{j=1}^t s_j = M = \sum_{j=1}^t s'_j$. Denote by $\mathcal{S}(M)$ the set of such chains with height t and through-put M . Now, S is said to Lorenz-dominate S' (written $S \succ S'$) iff $\sum_{j=1}^k s_j \geq \sum_{j=1}^k s'_j$ for all $k = 1, \dots, t-1$. In other words, the transactions are more equally distributed between blocks in S than in S' . Clearly, if $M = tT$ then $s_i = T$, for all i , is the unique Lorenz maximal chain whereas (for t even) the chain where $s_i = 0$ for $i = 1, \dots, t/2$ and $s_i = 2T$ for $i = t/2 + 1, \dots, t$ is the unique Lorenz minimizer.

A real valued differentiable function $b : [0, 2T]^t \rightarrow \mathbb{R}$ is said to be Shur-concave if it preserves the Lorenz ordering, i.e., if $S \succ S' \Rightarrow b(S) \geq b(S')$.

It is well-known (see e.g., Theorem 4, page 89 in Marshall et al. (1979)) that b is Shur-concave iff the partial derivatives are decreasingly ordered, i.e., $b'_1 \geq \dots \geq b'_n$. It is clear from Equation 1 that b_t has decreasingly ordered partial derivatives and thus the base fee b_t is minimized for the most unequally distributed block sizes. Formally,

Observation 2: If S is a Lorenz minimizer on $\mathcal{S}(M)$ then $b_t(S) = \min\{b_t(S') \mid S' \in \mathcal{S}(M)\}$.

Now, a legacy transaction is characterized by its fee cap f . Thus, since the base fee is burned, Observation 2 indicates a revenue maximizing miner prefer Lorenz minimizing chains. More precisely, at any given height t a miner maximizes the revenue of block t if the total amount of throughput is distributed as unevenly as possible over the blocks in the chain. Of course, the miner of block t is not necessarily in control of the previous block sizes, but it is clear that the choice of block size at a given height influences the revenue in later blocks. In this way a farsighted miner will not automatically fill up the blocks as much as possible, but may choose the block size strategically. For instance, consider a rational miner in control of a block chain of height $t = 2$. When she is deciding on how many transactions to process in block 1, she aims at maximizing the total revenue obtained from both block 1 and 2. If the initial base fee is 0 and the number of incoming transactions is large enough compared to the maximal block size, it will be revenue maximizing to fill up both blocks to full size (i.e., $s_1 = s_2 = 2T$). However, if the initial base fee is given by the steady state base fee (i.e., the clearing price which makes exactly T transactions eligible) it may be revenue maximizing to let $s_1 = 0$ and $s_2 = 2T$. Obviously, the picture becomes much more complicated taking into account that there are multiple competing miners (mining pools) as we will show in the coming sections.

In the following sections, we will analyze the potential for miners' strategic behaviour in further detail both analytically (for a single miner with 2-block foresight) and by simulations (for miners with up to 3-block foresight and varying degrees of hashing power).

4 Myopic versus strategic miners

As mentioned in the previous section, whether a given transaction τ , is *eligible* to be executed depends on the user's cap f and the base fee b_t at the current time t . In particular, a transaction is eligible if $f - b_t \geq 0$. We assume that users willingness to pay (i.e., their cap) are i.i.d and follow a uniform distribution $f \sim U[0, F]$. Moreover, suppose that a fixed number n of new transactions arrive at every time interval t (i.e., for every block).

Proposition 1. *Assuming that there are n transactions uniformly distributed on $[0, F]$, the expected number of transactions that are eligible for inclusion in block t is $n \frac{F-b_t}{F}$.*

Proof. See Appendix 8.

If $n \frac{F-b_t}{F} < 2T$ the miner can at most include $n \frac{F-b_t}{F}$ (eligible) transactions in the block at height t . This will give the miner an expected payoff of $n \frac{F-b_t}{F} (\frac{F-b_t}{2})$.

If $n \frac{F-b_t}{F} \geq 2T$ the maximal block size restricts the number of eligible transactions that can be executed. In this case the miner can at most expect a payoff of $2T(F - \frac{TF}{n} - b_t)$.

A *myopic* miner, is a miner that maximizes expected payoff per block without taking into account that the current block size influences the base fee of the next block (as given by Equation 1) which in turn influences the set of eligible transactions for the next block etc. In other words, a myopic miner always fills up the block as much as possible (and do not add fake transactions).

In contrast, a *strategic* miner, is a miner which maximizes the total expected payoff thinking k blocks ahead when she determines which transactions to include in the current, and up-coming $k - 1$ blocks. We say that a strategic miner has a foresight of k -blocks in that case.

4.1 A monopolist miner with 2-block foresight

The problem of a single (monopolist) miner with 2-block foresight is already surprisingly complex. Since the miner can only reason one block ahead of the current block, it is optimal for the miner to fill up the second block as much as possible given the set of eligible transactions at block 2 (which in turn depends on the size of the first block). Thus, the problem boils down finding the optimal size of the first block. The following intermediate result turns out to be convenient.

Proposition 2. *Let (B_t, B_{t+1}) be two consecutive blocks, such that $s_t \neq 0$ and $s_{t+1} \neq 2T$. The total expected miner payoff of (B_t, B_{t+1}) is smaller than that of $(\bar{B}_t, \bar{B}_{t+1})$ when $\bar{s}_t = 0$ and $\bar{s}_{t+1} \neq 0$ or $\bar{s}_t \neq 0$ and $\bar{s}_{t+1} = 2T$.*

Proof. See Appendix 9.

We can now determine the payoff maximizing size of the first block for a 2-block farsighted miner as a function of the initial base fee: the second block should always be full as possible given the set of eligible transactions.

Proposition 3. *The optimal size of the first block for a 2-block farsighted miner is:*

$$s_t(b_t) = \max\{0, \min\{2T, (\frac{5}{4}b_t + \frac{TF}{n} - F)(\frac{-2n}{F})\}\} \quad (4)$$

Proof. See Appendix 10.

Consequently, whereas a myopic miner will always fill up every block as much as possible, a 2-block farsighted miner may optimally leave the first block empty if the first block's base fee is sufficiently high. Clearly, this hinges on the fact that the same miner gets to mine two blocks in a row: leaving the current block empty relies on the ability to harvest the benefits of a decreasing base fee for the second block.

Figure 1 below illustrates the optimal size of the first block for a 2-block farsighted monopolist miner in the particular case where $n = 100$, $T = 15$ and $F = 10$.

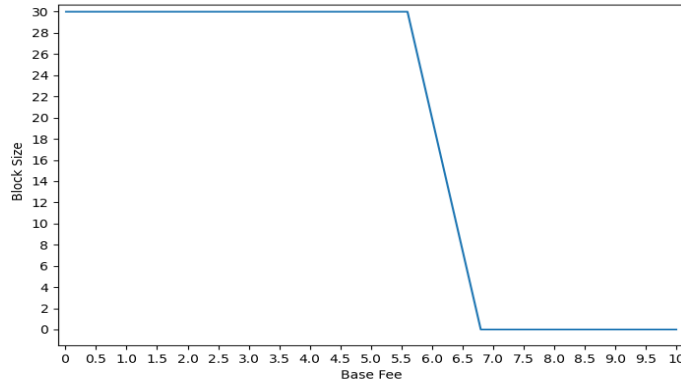


Fig. 1: The optimal size of the first block, with $T = 15$, $F = 10$, and $n = 100$.

When the base fee is low (i.e., $b_t \in [0, 5.6]$) the miner will fill up the first block to max size $2T$, whereas when the base fee is high (i.e., $b_t \in [6.8, 10]$) the first block is optimally be left empty: in between the optimal block size will monotonically decrease.

4.2 Competition: A 2-block farsighted miner with hashing power α

In practice, miners are competing to verify blocks so we therefore turn to the case of a miner with different levels of hashing power. We start out to examine

how much hashing power a 2-block farsighted miner needs in order to do better than a myopic miner (taking into account that there is a probability proportional to the hashing power of getting to mine the second block). Note, that while it generally matters for the optimal strategy of a k -block farsighted miner whether the rest of the network is myopic or farsighted as well, there is no difference in the particular case of $k = 2$.

Let the networks total hashing power be normalized to 1, and let a given miner have hashing power $\alpha \in (0, 1)$. In Figure 2 below, we illustrate the miner's decision tree: with probability $(1 - \alpha)^2$ she will not get to mine any of the two blocks; with probability α^2 she will get to mine both; and with probability $\alpha(1 - \alpha)$ she will get to mine one block, either the first or the second.

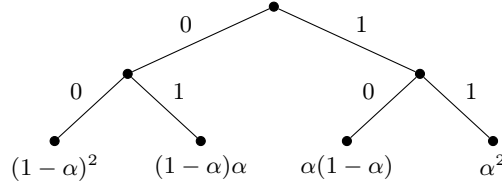


Fig. 2: The decision tree of a 2-block farsighted miner with hashing power α . An edge labelled 1 indicates that the farsighted miner creates a block, and edge labelled 0 indicates otherwise.

Note that for the left branch of the tree there is no difference between the payoff for a myopic and a 2-block farsighted miner. Therefore, we focus on the right branch. Furthermore, by Proposition 3, even as a monopolist it is optimal for a farsighted miner to act as if myopic when the base fee is sufficiently low. Therefore, we focus on base fee values $b_t \geq \frac{4}{5}F(1 - \frac{T}{n})$ where a (monopolist) 2-block farsighted miner optimally sets $s_t = 0$. Hence, we consider the minimum required computational power for a farsighted miner such that producing an empty first block is profitable in expectation.

Let R_1^M denote the payoff of a myopic miner from the first block on the right branch of the tree in Figure 2, and R_2^M be the payoff of a myopic miner from the second block of the right branch. Similarly, let R_1^F and R_2^F be the payoff of a farsighted miner from the first and second blocks on the right branch, respectively. So, the expected payoff, of a myopic miner from the right branch is $\alpha(1 - \alpha)R_1^M + \alpha^2(R_1^M + R_2^M)$. If the farsighted miner decides to leave the first block empty, i.e., $R_1^F = 0$, so the expected payoff of a farsighted miner is $\alpha^2 R_2^F$. Therefore, a farsighted miner is better off to produce an empty first block whenever, $\alpha^2 R_2^F \geq \alpha(1 - \alpha)R_1^M + \alpha^2(R_1^M + R_2^M)$, which simplifies to:

$$\alpha > \frac{R_1^M}{R_2^F - R_2^M} \quad (5)$$

As a special case, consider a myopic miner and assume that the initial base fee b_t is at steady state level: leaving exactly T eligible transactions. At block

t , the miner therefore fill the block with all eligible transactions. In this case, the eligible transactions are distributed uniformly on $[b_t, F]$. Given that from every transaction the amount of b_t is burnt, the payoff of the myopic miner from each transaction is uniformly distributed on $[0, F - b_t]$. Therefore the average revenue of the miner, i.e., R_1^M , of including eligible transactions in the block is $R_1^M = T \left(\frac{F-b_t}{2} \right)$. Since $b_{t+1} = b_t$, the average revenue of the miner is the same for block $t + 1$, that is, $R_2^M = R_1^M = T \left(\frac{F-b_t}{2} \right)$.

Now, if the miner has 2-block foresight, she will leave the first block empty if she gets to mine it (with probability α): since the farsighted miner sets $s_t = 0$, we get $b_{t+1} = \frac{7}{8}b_t$, and the miner can choose the same transactions and earn $\frac{1}{8}b_t$ more for each transaction with a block size of $2T$. So the revenue of the second block becomes $R_2^F = T(F - b_t) + \frac{2T}{8}b_t$. Plugging, R_1^M , R_2^M and R_1^F into Equation 5, a level-2 farsighted miner will do better than a myopic miner if:

$$\alpha > \frac{\frac{T}{2}(F - b_t)}{T(F - b_t) + \frac{T}{4}b_t - \frac{T}{2}(F - b_t)} \Rightarrow \alpha > \frac{F - b_t}{F - \frac{b_t}{2}}$$

At steady state $b_t = F(1 - \frac{T}{n})$, so the above equation simplifies to $\alpha > \frac{2T}{n+T}$. Say, $n = 100$ new transactions arrive at time t and $t + 1$ and that target size is $T = 15$, then when the base fee is at the steady state, the 2-block farsighted miner will do better than the myopic miner if she holds more than 26% of the total hashing power.

Figure 3, shows the minimum required computational power (in percentage) for different values of the base fee $b_t \geq \frac{4}{5}F(1 - \frac{T}{n})$, when the transactions are uniformly distributed on $(0, 10)$, and the target size is set to 15.

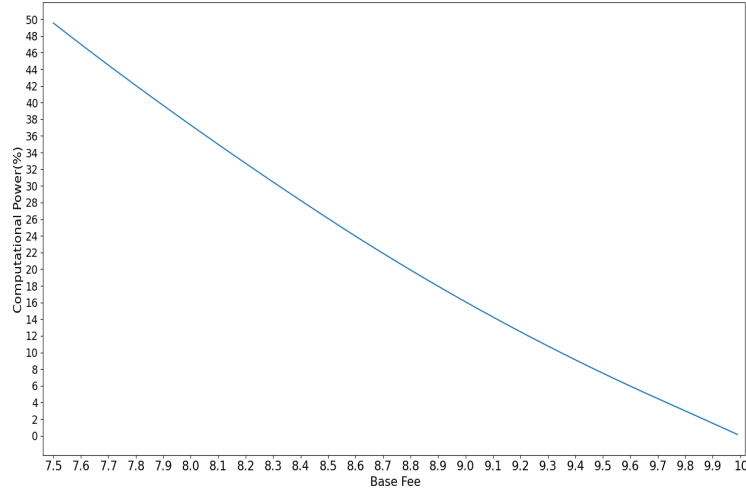


Fig. 3: The minimum computational power required to benefit from 2-block foresight compared to being myopic, for given levels of the base fee, with $T = 15$, $F = 10$, and $n = 100$.

To sum up, for a 2-block farsighted miner with α percent hashing power we will use the following strategy for our simulations:

- Network mines first block (happens with probability $1 - \alpha$):
 - First block: No action.
 - Second block:
 - * Network mines (with probability $1 - \alpha$): No action.
 - * Miner mines (with probability α): play myopically.
- The miner mines the first block (happens with probability α):
 - First block. Play a mixed strategy:
 - * with probability $1 - \alpha$: play myopically.
 - * with probability α : play as a 2-block farsighted monopolist.
 - Second block:
 - * Network mines (with probability $1 - \alpha$): No action.
 - * Miner mines (with probability α): Play as a 2-block farsighted monopolist (coinciding with myopic play).

Note that this strategy is optimal (i.e., maximizing expected revenue) for the miner: always play myopically in the second block since this is the last. Given this, in case the miner gets to mine the first block she will be a monopolist miner with 2-block foresight with probability α , and this is the only block she will get to mine with probability $1 - \alpha$. In the latter case, she should optimally fill up the first block as much as possible (i.e., play myopically).

5 Simulation Results

Clearly, analytical results for higher levels of k -block foresight are increasingly complex. We therefore turn to simulation studies. Specifically, we consider the case where the transactions at every time interval flow according to a *Poisson Distribution*, with constant arrival rate λ . Thus, the number of new transactions created at time t , denoted n_t , is given by:

$$n_t \sim P(n_t = x) = \frac{e^{-\lambda} \lambda^x}{x!}, \quad \text{and} \quad (x = 0, 1, 2, \dots) \quad (6)$$

For the following simulations we assume the arrival rate $\lambda = 100$. Transactions are sequentially added to the mempool. As such, *users* do not act strategically. That is, the users do not submit their transactions based on the current base fee, but following the Poisson process. Moreover, we assume the users' bids are i.i.d. from a fixed uniform distribution on $[0, 10]$, and that all the transactions are of the same size. Also, we set the target size of the block to $T = 15$.

For the execution, first we use the Poisson distribution to generate a data set for a sequence of 1000 blocks (specifically, the full data set contains 100,467 transactions). This data is used in all the subsequent simulations. Then we repeat 50 times a random pick of the miner who gets to mine blocks in the sequence of 1000 blocks. Every value regarding, base fee, miner revenue, and block distribution, is then averaged over those 50 repeated runs. We also assume that the rest

of the network acts myopically. As such, a farsighted miner with varying degrees of hashing (i.e, market) power is competing with the rest of the network that acts myopically. Throughout the sequences of 1000 blocks the base fee follows the EIP-rule (1) according to the chosen block sizes.

5.1 Myopic miner

We first simulate a myopic miner, i.e., a miner that at every block fills the block as much as possible according to the number of eligible transactions in the mempool and the available block size. As discussed previously, given that transactions are distributed uniformly on $[0, 10]$, and if at every round 100 new transactions are added to the mempool, the base fee must stabilize around $100 \frac{10-b}{10} = 15$, i.e., 8.5. In the simulations, since the number of new transactions are different at every round and the transactions are accumulated in the mempool, the base fee stabilizes at 8.9 with a variance of 0.14. Figure 4 shows the evolution of the base fee for myopic miners.

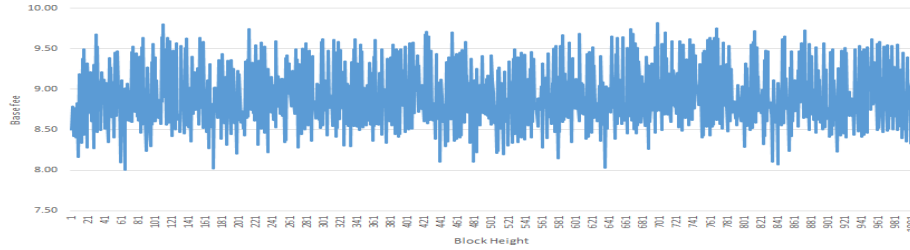


Fig. 4: The average and variance of the base fee for myopic miners.

The average revenue per block is 7.94, so ex-ante, a myopic miner with hashing power α has an expected average revenue of $\alpha \times 7.94$ for every block when the rest of the network are myopic as well.

5.2 2-block farsighted miner

We consider a 2-block farsighted miner with varying degrees of hashing power. Specifically, we focus on simulations where the miner has hashing power $\alpha \in \{0.1, 0.15, 0.2\}$. In comparison, the largest mining pool in the Ethereum network has approximately 27% of the total hashing power³.

Base fee: Figure 5, shows the base fee a 2-block farsighted miner with different computational powers. The summary of the results are presented in Table 1. Note that, as the computational power increases, the number of empty blocks increases, therefore the average base fee decreases and the variance increases.

³ <https://miningpoolstats.stream/ethereum>

	Average Variance	
$\alpha = 10\%$	8.9	0.15
$\alpha = 15\%$	8.88	0.18
$\alpha = 20\%$	8.82	0.20

Table 1: The average and variance of the base fee for miners with 2-block foresight.

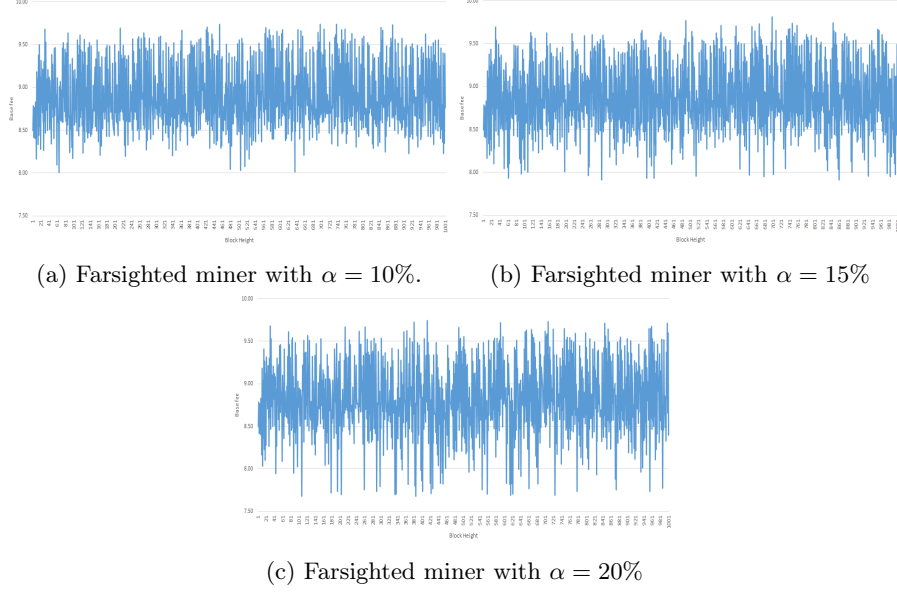


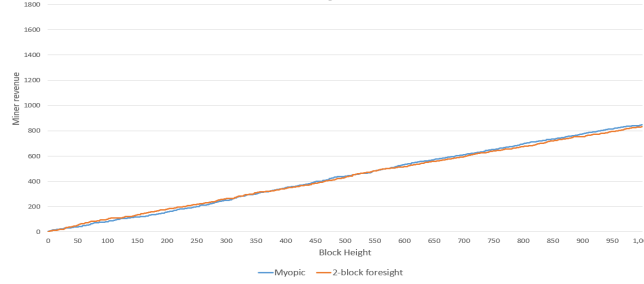
Fig. 5: Evolution of the base fee for different computational power of 2-block foresight miners.

Reward Distribution: Next, we consider the aggregate reward of a farsighted miner with $\alpha = 10\%$, 15% and 20% hashing power and compare the results with the case that the miner acts myopically. The results are shown in Figure 6.

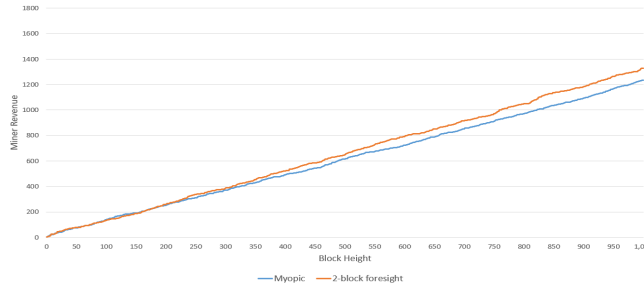
For a miner with $\alpha = 10\%$, the expected revenue of the being farsighted is similar to that of being myopic. This is due to the fact that for a miner with $\alpha = 10\%$, Equation 5 and Figure 3, implies that the base fee must be larger than 9.35 before it pays off (in expectation) to produce an empty block for the farsighted miner. However, with respect to Figure 5, the base fee is often less than 9.35 and in those instances the miner only have a 10% chance to mine the block. Therefore, the revenue of a farsighted miner is basically similar to that of a myopic miner in this case.

For a 2-block farsighted miner, with $\alpha = 20\%$, Equation 5 and Figure 3, implies that the base fee must be larger than 8.79 in order to benefit from producing an empty block, which happens more frequently. Furthermore, as

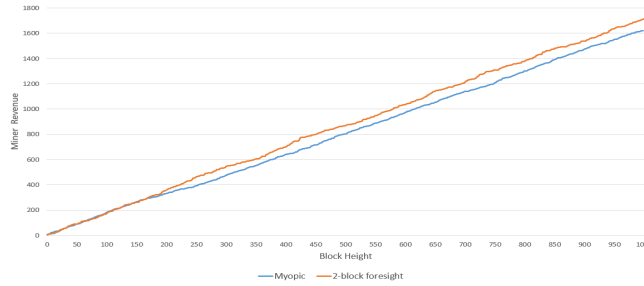
the miner has 20% of the computational power, the miner gets to produce more consecutive blocks, which increases the farsighted miner's expected revenue. This corresponds to the finding in Section 4.2.



(a) Farsighted miner with $\alpha = 10\%$.



(b) Farsighted miner with $\alpha = 15\%$



(c) Farsighted miner with $\alpha = 20\%$

Fig. 6: Cumulative reward for 2-block foresight miners.

Block Distribution and throughput: Figure 7, shows the distribution of the blocks for myopic and farsighted miners. Note that, as the computational power of the farsighted miner increases then the number of empty blocks increases as well. However, considering the total throughput of the network, i.e., the number of transactions included in the blocks, for a period of 1000 blocks, increases. The throughput of the network (i.e., the total transactions included

in blocks) is summarized in Table 2. Therefore, the miners by being farsighted tend to increase the throughput of the network. This is due to the fact that, as shown in Table 1, the base fee decreases when there are strategic miners that produce empty blocks. As the base fee goes down, more transactions become eligible to be included in a block, which in turn increases the throughput of the system (in the extreme case, if the base fee is always close to zero, then every block is almost full).

	Average
Myopic	15342
$\alpha = 10\%$	15363
$\alpha = 15\%$	15433
$\alpha = 20\%$	15468

Table 2: The total throughput for myopic and 2-block foresight miners.

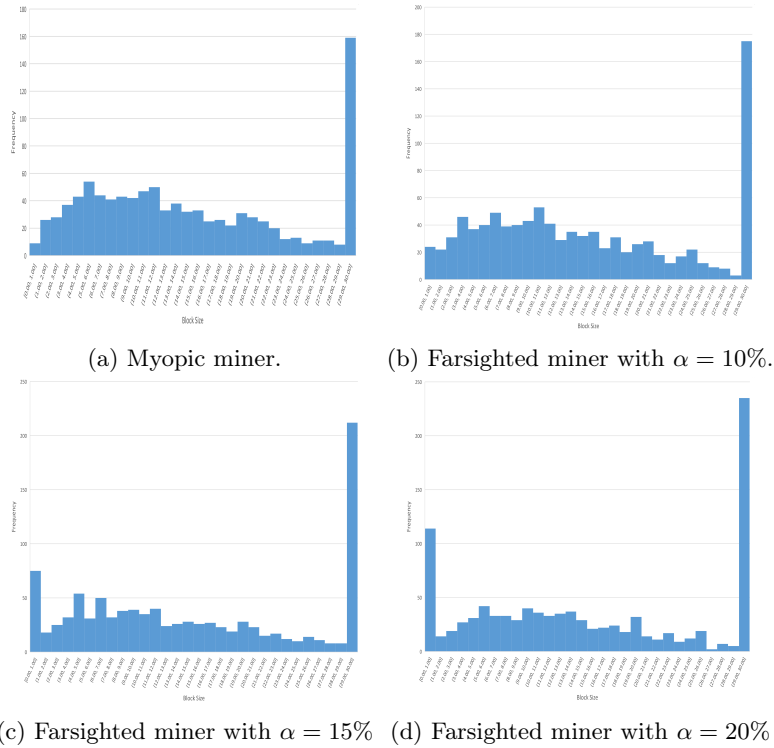


Fig. 7: Distribution of 1000 the blocks for myopic miner and 2-block farsighted miners.

5.3 3-block foresight

In this section, we consider a miner with 3-block foresight. Roughly speaking, a 3-block farsighted miner considers the expected revenue of the next three blocks

and tries to strategize based on the first block and the second block (if she gets the chance to mine these blocks). Similar to the case of 2-block foresight, the decision tree of a 3-block farsighted miner with computational power α is shown in Figure 8.

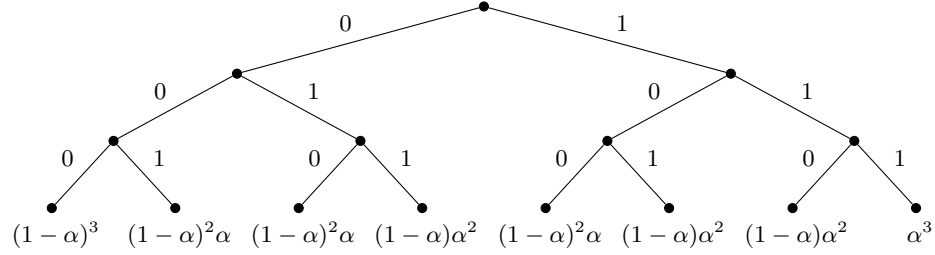
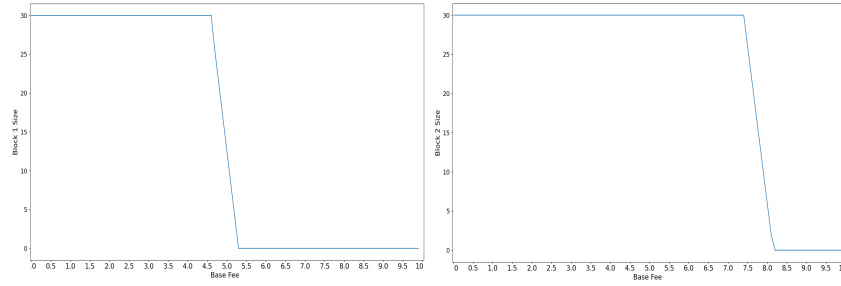


Fig. 8: The decision tree for a 3-block farsighted miner with α computational power. The probability of each mining sequence is denoted on each leaf.

First, we consider the optimal strategy for the monopolist miner, i.e., $\alpha = 100\%$. Using simulation, the optimal sizes of the first and the second block as a function of the initial base fee, is shown in Figure 9.



(a) Optimal size of the first block. (b) Optimal size of the second block.

Fig. 9: The optimal size of the first and second block for a miner with 3-block foresight.

Figure 10 compares the aggregate revenue of a monopolist myopic miner, a monopolist 2-block farsighted miner, and a monopolist 3-block farsighted miner.

We emphasize that the revenue of a (monopolist) 2-block farsighted miner is higher than that of a (monopolist) 3-block farsighted miner. At first glance this may seem counter-intuitive since one would expect a miner with higher levels of foresight to be able to do at least as good as a miner with lower levels of foresight. However, this line of reasoning is only valid if we consider per block revenue in the respective 2, and 3-block sequences. When we consider longer sequences of blocks another feature sets in. For instance, consider a series of, say, 6 blocks: here, a miner with 2-block foresight will at best be able to produce 3 empty blocks, whereas a miner with 3-block foresight at best will be able to produce 2 empty blocks. As it is a lowering of the base fee that creates the opportunity

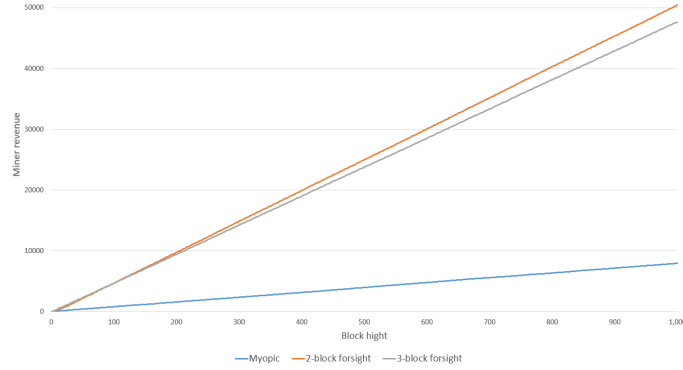


Fig. 10: The aggregate revenue of monopolist miners for myopic, 2-block foresight and 3-block foresight.

to increase revenue, a monopolist miner with 2-block foresight can decrease the base fee further than a miner with 3-block foresight, and thereby obtain a higher revenue over repeated sequences of 2 and 3-block constellations.

Next, consider a miner with 3-block foresight and α computational power. In the simulations we assume that the miner uses the following strategy:

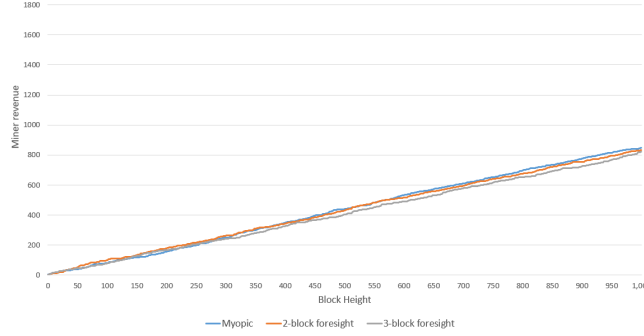
1. Network gets to mine the first block (with probability $1 - \alpha$):
 - (a) First block: No action.
 - (b) Onwards: Copy the strategy of a miner with 2-block foresight and α hashing power.
2. The miner mines the first block (with probability α):
 - First block. Play a mixed strategy:
 - (a) with probability $1 - \alpha$: play myopically.
 - (b) with probability α : play as a monopolist miner with 3-block foresight.
 - Onwards: Copy the strategy of a miner with 2-block foresight and α hashing power.

Note that playing this strategy we “pretend” that the mempool is empty when entering the second block. This is clearly not the case, so in principle the miner could have done better if she took the size of the mempool when entering block 2 into account. Yet, for realistic levels of hashing power, this mistake is arguably limited and we claim that the above strategy is near optimal in that case.

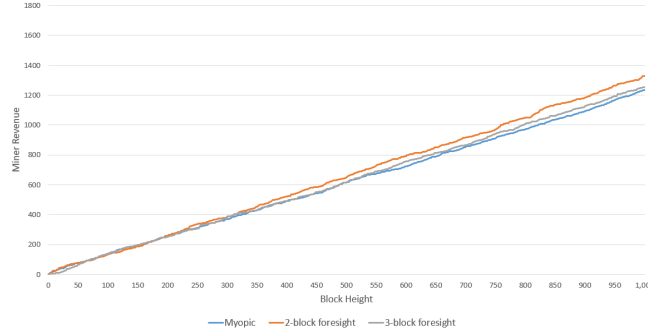
The results are shown in Figure 11(a)-(c) below. With only 10% hashing power the aggregate revenue of the myopic miner is slightly better than that of both a 2, and 3-block farsighted miner. With 15% and 20% hashing power both farsighted miners do better than the myopic miner, and the 2-block farsighted miner consistently do better than the 3-block farsighted miner. We emphasize that this may partly be due to the fact that the strategy of the 3-block farsighted miner is only near optimal so the 3-block farsighted miner could potentially have done slightly better if allowed to play an optimal strategy. Yet, as shown above,

when we simulate the optimal strategy for a 3-block farsighted monopolist miner she will still be worse off than a 2-block farsighted monopolist miner. So the results for lowers levels of hashing power seem in line with this.

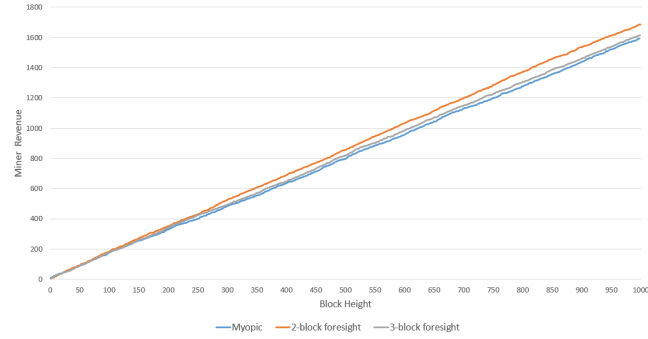
The system throughput (Table 3 below) is increasing in the levels of hashing power and is slightly higher for farsighted than myopic miners. Again, average base fee (Table 4 below) decreases in the levels of hashing power whereas the variance increases - similar to the case of the miner with 2-block foresight.



(a) Farsighted miner with $\alpha = 10\%$.



(b) Farsighted miner with $\alpha = 15\%$



(c) Farsighted miner with $\alpha = 20\%$

Fig. 11: Cumulative reward for 3-block foresight miners.

	Average
Myopic	15342
$\alpha = 10\%$	15356
$\alpha = 15\%$	15389
$\alpha = 20\%$	15423

Table 3: The throughput for miners with 3-block foresight.

	Average	Variance
$\alpha = 10\%$	8.9	0.15
$\alpha = 15\%$	8.89	0.16
$\alpha = 20\%$	8.84	0.18

Table 4: The average and variance of the base fee for miners with 3-block foresight.

Average base fee updating rule: Since the base fee only depends on the size of the previous block, then a farsighted miner can manipulate the base fee so that he can extract more value in the next block. Here, in line with Reijbergen et al. (2021), we propose the *average base fee updating rule* for the base fee. The simplest form, is to update the base fee based on the average of the previous two blocks. Formally,

$$b_t = \left(b_{t-1} \left(1 + \frac{1}{8} \cdot \frac{s_{t-1} - T}{T} \right) + b_{t-1} \right) \cdot \frac{1}{2} \quad (7)$$

Figure 12 compares the base fee evolution for the myopic miners with the EIP1559 and the average updating rule of Equation 7. Note that, as the expected demand at every block is the same then the average of the base fee for using both updating rule with myopic miners is similar. However, the variance of the base fee with the average updating rule is lower (to be more precise, the variance of the base fee by applying the average updating rule is 0.028 whereas the variance of EIP1559 is 0.14.)

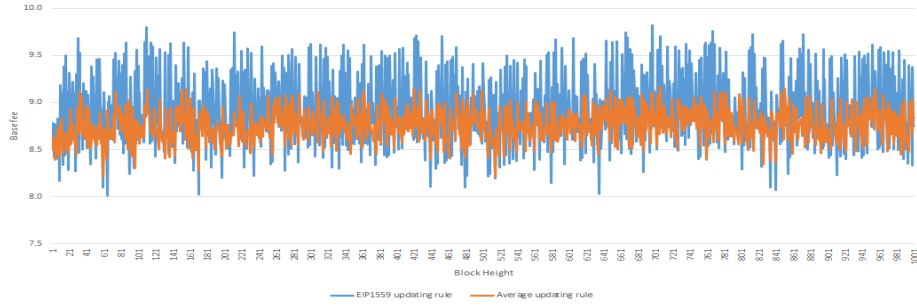


Fig. 12: Comparison of the base fee with the new updating rule for myopic miners.

The effect of applying the average updating rule on miners' revenue is shown in Figure 13. Given that the network reaches a steady state, the first part of Equation 7 is approximately equal to b_{t_1} , and hence $b_t = b_{t-1}$. Therefore, the revenue of myopic miner is roughly the same with both updating rules. However, in case of a farsighted miner producing an empty block it has less effect on the base fee. This implies that a miner needs more computational power to be able to manipulate the new updating rule for the base fee.

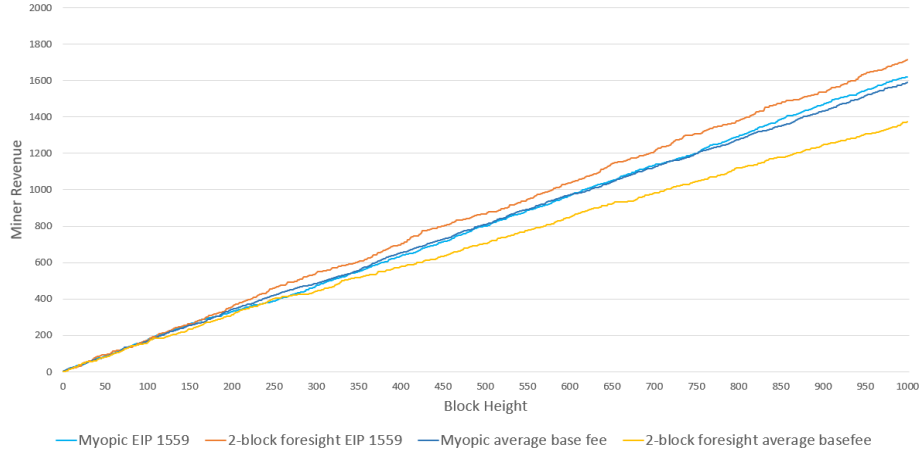


Fig. 13: Comparison of the revenue of myopic and 2-block foresight miners with $\alpha = 20\%$.

5.4 Robustness

The simulation results of the previous section used the uniform distribution to model the users' willingness to pay. To examine whether the main conclusions obtained above are robust to changes in this distributional assumption, we have also run the simulations using a normal distribution with average 10, and variance 5 to model the users' willingness to pay. The remaining parameters, i.e., the distribution of the new transactions and the number of blocks kept the same as in the previous section. The simulation results shows the previous conclusions still hold, i.e., miners with 2 and 3-block foresight are better off by strategically producing empty blocks. The intuition for this results is that the miners can carry the high value transactions into the next block, and hence lower the base fee which results in less burning and thereby higher miner revenue. Moreover, as before, the miner with 2-block foresight does better than a miner with 3-block foresight on aggregate revenue and system throughput also increases with farsighted miners.

6 Conclusion

To sum up, our simulation study illustrates that a 2-block farsighted miner/pool may gain from strategic behavior if their hashing power is sufficiently high (here more than approx. 15%). Strategic behavior leads to more empty blocks and variation in block sizes, but on average increases system throughput slightly. Somewhat counter intuitively, allowing miners to reason more than 2 blocks ahead does not add to the advantage of having strategic foresight. In particular, our result illustrate that for all levels of hashing power a miner with 2-block foresight do better in terms of aggregate revenue than a miner with 3-block foresight. We conjecture that for higher levels of foresight this remains true for realistic levels of hashing power (that is, between 15-40%).

Finally, changing the base fee up-dating rule slightly to average over the two past blocks reduces the variance of the base fee considerably. Thereby, it also significantly limits the expected gains from strategic behavior. These results seem robust to changes in the distributional assumptions on users' willingness to pay.

Bibliography

- Beiko, T. (2021). London mainnet announcement. <https://blog.ethereum.org/2021/07/15/london-mainnet-announcement/>.
- Buterin, V. (2014). Blockchain resource pricing. <https://ethresear.ch/uploads/default/original/2X/1/197884012ada193318b67c4b777441e4a1830f49.pdf>.
- Buterin, V. (2016). On inflation, transaction fees and cryptocurrency monetary policy. <https://blog.ethereum.org/2016/07/27/inflation-transaction-fees-cryptocurrency-monetary-policy/>.
- Buterin, V. (2018). On transaction fees, and the fallacy of market-based solutions. <https://blog.ethereum.org/2014/02/01/on-transaction-fees-and-the-fallacy-of-market-based-solutions/>.
- Buterin, V. (2021). EIP-1559 FAQ. <https://notes.ethereum.org/@vbuterin/eip-1559-faq>.
- Buterin, V. et al. (2014). A next-generation smart contract and decentralized application platform. *white paper 3*(37).
- Buterin, V., E. Conner, R. Dudley, M. Slipper, I. Norden, and A. Bakhta (2019). EIP-1559: Fee market change for ETH 1.0 chain. <https://eips.ethereum.org/EIPS/eip-1559>.
- Chung, H. and E. Shi (2021). Foundations of transaction fee mechanism design. *arXiv preprint arXiv:2111.03151*.
- Ferreira, M. V., D. J. Moroz, D. C. Parkes, and M. Stern (2021). Dynamic posted-price mechanisms for the blockchain transaction-fee market. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, pp. 86–99.
- Leonardos, S., B. Monnot, D. Reijnders, E. Skoulakis, and G. Piliouras (2021). Dynamical analysis of the EIP-1559 ethereum fee market. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, pp. 114–126.
- Liu, Y., Y. Lu, K. Nayak, F. Zhang, L. Zhang, and Y. Zhao (2022). Empirical analysis of EIP-1559: Transaction fees, waiting time, and consensus security. *arXiv preprint arXiv:2201.05574*.
- Marshall, A. W., I. Olkin, and B. C. Arnold (1979). *Inequalities: theory of majorization and its applications*, Volume 143. Springer.
- Maskin, E. et al. (2001). *Auctions and efficiency*. School of Social Science, Institute for Advanced Study.
- Monnot, B., Q. Hum, C. Koh, S. Leonardos, and G. Piliouras (2020). Ethereum’s transaction fee market reform of EIP 1559. In *WINE’20 Workshop on Blockchain*.
- Reijnders, D., S. Sridhar, B. Monnot, S. Leonardos, S. Skoulakis, and G. Piliouras (2021). Transaction fees on a honeymoon: Ethereum’s EIP-1559 one month later. In *2021 IEEE International Conference on Blockchain (Blockchain)*, pp. 196–204. IEEE.

Roughgarden, T. (2020). Transaction fee mechanism design for the ethereum blockchain: An economic analysis of EIP-1559. *arXiv preprint arXiv:2012.00854*.

Roughgarden, T. (2021). Transaction fee mechanism design. *ACM SIGecom Exchanges* 19(1), 52–55.

7 Appendix

8 Proof of Proposition 1

Proposition 1. Assuming that there are n transactions uniformly distributed on $[0, F]$, the expected number of transactions that are eligible for inclusion in block t is $n \frac{F-b_t}{F}$.

Proof. Let χ_i denote the random variable that shows the inclusion of a transaction in the block B_t . Formally,

$$\chi_i = \begin{cases} 1, & \text{if } i\text{'s fee cap satisfies } f_i \geq b_t \\ 0, & \text{otherwise} \end{cases}$$

Therefore,

$$\Pr(\chi_i = 1 | b_t) = \Pr(f_i > b_t) = 1 - \Pr(f_i \leq b_t) = 1 - F(b_t) = \frac{F - b_t}{F}$$

As the users are independent and identically distributed, then so is χ_i . Hence, every transaction inclusion in the block follows, i.e., the χ_i , follows a Bernoulli distribution with parameter $p = 1 - F(b_t)$ (this is well defined since $0 \leq F(b_t) \leq 1$). Using Wald's equation $\mathbf{E}[\sum_{i=1}^n \chi_i | b_t] = E[n]E[X_1] = n \times \frac{F-b_t}{F}$.

9 Proof of Proposition 2

Proposition 2. Let (B_t, B_{t+1}) be two consecutive blocks, such that $s_t \neq 0$ and $s_{t+1} \neq 2T$. The total expected miner payoff of (B_t, B_{t+1}) is smaller than that of $(\bar{B}_t, \bar{B}_{t+1})$ when $\bar{s}_t = 0$ and $\bar{s}_{t+1} \neq 0$ or $\bar{s}_t \neq 0$ and $\bar{s}_{t+1} = 2T$.

Proof. Let R denote the payoff of the two blocks (B_t, B_{t+1}) , and \bar{R} denote the payoff of the two blocks $(\bar{B}_t, \bar{B}_{t+1})$. That is,

$$\begin{aligned} R &= \left(\sum_{\tau \in B_t} f_\tau + \sum_{\tau \in B_{t+1}} f_\tau \right) - (s_t b_t + s_{t+1} b_{t+1}) \\ \bar{R} &= \left(\sum_{\tau \in \bar{B}_t} f_\tau + \sum_{\tau \in \bar{B}_{t+1}} f_\tau \right) - (\bar{s}_t \bar{b}_t + \bar{s}_{t+1} \bar{b}_{t+1}). \end{aligned}$$

Note that in the above equations, $b_t = \bar{b}_t$, since the base fee of the block at height t only depends on the size of the block at height $t - 1$. Furthermore, $b_{t+1} = \frac{b_t}{8} (7 + \frac{s_t}{T})$ and $\bar{b}_{t+1} = \frac{b_t}{8} (7 + \frac{\bar{s}_t}{T})$.

Based on the size of the first block, i.e., s_t , we consider different cases and in each case we show that the miner can increase her payoff by decreasing the size of the first block.

Case 1. $s_t \leq T$, and $s_{t+1} \leq T$.

Since $s_t \leq T$ implies that $b_{t+1} \leq b_t$, all transactions in B_t are eligible for inclusion in B_{t+1} . Furthermore, as $s_{t+1} \leq T$, B_{t+1} has sufficient space to include all the transactions in B_t . Therefore, setting $\bar{B}_t = \emptyset$ and $\bar{B}_{t+1} = B_t \cup B_{t+1}$, implies $\left(\sum_{\tau \in B_t} f_\tau + \sum_{\tau \in \bar{B}_{t+1}} f_\tau \right) = \left(\sum_{\tau \in \bar{B}_t} f_\tau + \sum_{\tau \in \bar{B}_{t+1}} f_\tau \right)$. Consequently, $\bar{R} > R$ if and only if $\bar{s}_t b_t + \bar{s}_{t+1} \bar{b}_{t+1} < s_t b_t + s_{t+1} b_{t+1}$. In particular, we have $\bar{s}_t = 0$ and $\bar{s}_{t+1} = s_t + s_{t+1}$. Moreover, because $\bar{s}_t = 0$, we get $\bar{b}_{t+1} = \frac{7}{8} b_t$. Thus, we have

$$\begin{aligned} \bar{s}_t b_t + \bar{s}_{t+1} \bar{b}_{t+1} &= \bar{s}_{t+1} \bar{b}_{t+1} = (s_t + s_{t+1}) \frac{7}{8} b_t \\ &< s_t b_t + s_{t+1} \left(\frac{b_t}{8} \left(7 + \frac{s_t}{T} \right) \right) = s_t b_t + s_{t+1} b_{t+1} \end{aligned}$$

To conclude, in case $s_t \leq T$, and $s_{t+1} \leq T$, the payoff of the miner is increased if instead we set $\bar{s}_t = 0$ and $\bar{s}_{t+1} = s_t + s_{t+1}$.

Case 2. $s_t \leq T$, and $s_{t+1} > T$.

Since $s_t \leq T$ implies that $b_{t+1} \leq b_t$, all the transactions in B_t are eligible for inclusion in B_{t+1} . Let $k = 2T - s_{t+1}$ be the remaining space in B_{t+1} , and let Γ be any subset of k transactions in B_t . Setting $\bar{B}_t = B_t \setminus \Gamma$ and $\bar{B}_{t+1} = B_t \cup \Gamma$, implies $\left(\sum_{\tau \in \bar{B}_t} f_\tau + \sum_{\tau \in \bar{B}_{t+1}} f_\tau \right) = \left(\sum_{\tau \in \bar{B}_t} f_\tau + \sum_{\tau \in \bar{B}_{t+1}} f_\tau \right)$. Therefore, again $\bar{R} > R$ if and only if $\bar{s}_t b_t + \bar{s}_{t+1} \bar{b}_{t+1} < s_t b_t + s_{t+1} b_{t+1}$. Now, $\bar{s}_t = s_t - k$ and $\bar{s}_{t+1} = s_{t+1} + k$. Moreover, note that, $\bar{b}_{t+1} = \frac{b_t}{8} (7 + \frac{\bar{s}_t}{T}) = \frac{b_t}{8} (7 + \frac{s_t - k}{T}) = \frac{b_t}{8} (7 + \frac{s_t}{T}) - k \frac{b_t}{8T} = b_{t+1} - k \frac{b_t}{8T}$. Therefore,

$$\bar{s}_t \bar{b}_t + \bar{s}_{t+1} \bar{b}_{t+1} = (s_t - k) b_t + (s_{t+1} + k) (b_{t+1} - k \frac{b_t}{8T}) < s_t b_t + s_{t+1} b_{t+1}$$

(since $b_{t+1} \leq b_t$). To conclude, the miner's payoff increases whenever we set $\bar{s}_t = s_t - k$ and $\bar{s}_{t+1} = s_{t+1} + k$.

Case 3. $s_t > T$, and $s_{t+1} < T$.

Since, $s_t > T$ then $b_{t+1} > b_t$. Let $k = 2T - s_{t+1}$, as $s_{t+1} < T$, then $k > T$. Let Γ be any subset of k transactions in B_t . As $k > T$, then removing all the set of Γ transactions from B_t , implies that $\bar{b}_{t+1} < b_t$, so these transactions can be included in the \bar{B}_{t+1} . Setting $\bar{B}_t = B_t \setminus \Gamma$ and $\bar{B}_{t+1} = B_t \cup \Gamma$, and a similar argument as that of Case 2 applies.

Case 4. $s_t > T$, and $s_{t+1} > T$. Let $k = 2T - s_{t+1}$, as $s_{t+1} > T$, then $k < T$. Next, we show that there are a set of T transactions in B_t , that are eligible for inclusion in B_{t+1} . Note that as $s_t > T$, then $b_t < b^*$, where $b^* = F(1 - \frac{T}{n})$ is the stable base fee. So, there are T transaction in B_t , with a fee cap larger than the stable base fee, that is $Tops = \{\tau \in B_t \mid f_\tau \geq F(1 - \frac{T}{n})\}$. As, $s_{t+1} > T$, then $b_{t+1} < b^*$, therefore all the transactions in $Tops$ are eligible for inclusion in block B_{t+1} . Note that, reducing the size of the first block results in a lower base fee, hence $\bar{b}_{t+1} < b_{t+1}$. Therefore the set of transactions in $Tops$ are eligible for inclusion in B_{t+1} . As $k < T$, then Γ is the subset of any subset of k transactions in $Tops$. Setting, $\bar{B}_t = B_t \setminus \Gamma$ and $\bar{B}_{t+1} = B_t \cup \Gamma$, a similar argument as in Case 2 applies.

10 Proof of Proposition 3

Proposition 3. The optimal size of the first block for a 2-block farsighted miner is:

$$s_t = \max\{0, \min\{2T, (\frac{5}{4}b_t + \frac{TF}{n} - F)(\frac{-2n}{F})\}\} \quad (8)$$

Proof. The steady state base fee (leaving exactly T eligible transactions) is given by $b^* = F(1 - \frac{T}{n})$. We consider two cases where the base fee is above and below b^* and show that the optimal size follows Equation 4.

- Case 1. $b_t > b^*$. In this case the, the number of eligible transactions for two consecutive blocks is less than $2T$. Thus, by Proposition 2 the optimal size of the first block is $s_t = 0$ which is also what we get from (4) when inserting $b_t \geq F(1 - \frac{T}{n})$.
- Case 2. $b_t < b^*$. In this case there are sufficient eligible transactions so that the miner can fill the second block to $2T$. In the first round the miner chooses s_t which is bounded by either by the maximum size of the block $2T$ or by the maximum number of eligible transactions. At time t , n new transactions with values being uniformly distributed on $[0, F]$ is added to the pool. By Proposition 1, the miner is limited by $\min\{2T, n\frac{F-b_t}{F}\}$. Hence the miner chooses all the transactions τ with $f_\tau \geq F(1 - \frac{s_t}{n})$, such transactions are uniformly distributed on $[F(1 - \frac{s_t}{n}), F]$. Therefore the gross revenue of the miner is $\frac{s_t F}{2}(2 - \frac{s_t}{n})$. As for each transaction, b_t is burnt, the net revenue of the miner from the first block is

$$\frac{s_t F}{2}(2 - \frac{s_t}{n}) - s_t b_t. \quad (9)$$

At time $t + 1$, n new transactions are added to the mempool, with values that are uniformly distributed on $[0, F]$. Therefore the miner gets to create a block of size s_t transactions from the new transactions which, similar to the previous case, yields a revenue of $\frac{s_t F}{2}(2 - \frac{s_t}{n}) - s_t b_{t+1}$. The remaining space in the second block, i.e., $\bar{S} = 2T - s_t$, must be filled up from the transactions in the mempool. Let Λ denote the transactions

remaining in the mempool. Note that, $\Lambda = 2n - 2s_t$ and the values of these transactions are uniformly distributed on $[0, F(1 - \frac{s_t}{n})]$. Out of the Λ transactions, the miner chooses the top \bar{S} transactions. That is, the miner includes transactions with $f \geq F(1 - \frac{s_t}{n})(1 - \frac{\bar{S}}{\Lambda})$. These transactions have values that are uniformly distributed on $[F(1 - \frac{s_t}{n})(1 - \frac{\bar{S}}{\Lambda}), F(1 - \frac{s_t}{n})]$. Therefore, the average revenue of these transactions is $\frac{\bar{S}}{2}F(1 - \frac{s_t}{n})(2 - \frac{\bar{S}}{\Lambda})$. Note that for each transaction, b_{t+1} is burnt. Hence, the revenue of the second block is $\frac{\bar{S}}{2}F(1 - \frac{s_t}{n})(2 - \frac{\bar{S}}{\Lambda}) + \frac{s_t F}{2}(2 - \frac{s_t}{n}) - 2Tb_{t+1}$. Note that, $b_{t+1} = b_t(1 + \frac{1}{8}\frac{s_t - T}{T})$. Therefore, $2Tb_{t+1} = \frac{b_t}{4}(7T + s_t)$. All in all, the revenue of the second block is

$$\frac{\bar{S}}{2}F\left(1 - \frac{s_t}{n}\right)\left(2 - \frac{\bar{S}}{\Lambda}\right) + \frac{s_t F}{2}\left(2 - \frac{s_t}{n}\right) - \frac{b_t}{4}(7T + s_t) \quad (10)$$

Putting Equations 9 and 10 together the total revenue of the miner from two consecutive blocks is:

$$\begin{aligned} R(s_t) &= \frac{1}{2}(2T - s_t)\left(F - \frac{s_t}{n}\right)\left(2 - \frac{2T - s_t}{2(n - s_t)}\right) \\ &\quad + Fs_t\left(2 - \frac{s_t}{n}\right) - s_t b_t - \frac{7T}{4}b_t - \frac{s_t}{4}b_t \end{aligned}$$

The first order condition implies:

$$\frac{\partial R}{\partial s_t} = -\frac{F}{n}(2T + 2n - 3s_t) + F\left(2 - \frac{2s_t}{n}\right) - \frac{5}{4}b_t = 0$$

Therefore, the optimal size of the first block is $s_t = (\frac{5}{4}b_t + \frac{TF}{n} - F)(\frac{-2n}{F})$.