Design and Analysis of Side Contracts Attacking Ethereum EIP-1559 from Theory, Simulation, and Empirical View

Zerui Cheng

Princeton University, Princeton NJ 08544, USA

Abstract. EIP-1559, which is currently adopted on the second most popular blockchain, Ethereum, is a new transaction fee mechanism for blockchains. Through the innovative burnt fee mechanism, both empirical and theoretical analysis have shown that, EIP-1559 can help improve user experience by both avoiding first price auction in bidding and decreasing expected waiting time. However, it remains questionable whether or not the new transaction fee mechanism can defend against possible side contracts or misbehavior from miners or users. In this work, we propose a possible side contract against EIP-1559, analyze its possibility and impact on the whole blockchain ecosystem in theory, simulate the side contract using history data to see whether or not it is profitable, and use empirical analysis to find evidence of whether or not there has been similar misbehavior in the wild. Our work sheds light on an inherent weakness of EIP-1559 which is well worth noticing, especially after "the Merge" in the near future where Ethereum adopts "Proof-of-Stake" as its consensus algorithm.

Keywords: Side contracts · EIP-1559 · Ethereum.

1 Introduction

Dating back to Bitcoin [6], blockchains have received plausible attention from all over the world as a secure, decentralized, reliable and tamper-proof ledger. And Ethereum [2], introduced in 2014, is widely acknowledged as one of the best-received blockchains worldwide. Smart contracts [9], as a fundamental building block of various tokens and NFTs, are enabled on Ethereum, which results in the popularity of Ethereum but also leads to congestion in the network. Thus a transaction fee mechanism (TFM) is essential for solving the potential auction and competition as well as improving users' experience in case of busy network.

Regarding TFM in most blockchains, the most intuitive way, where users send a bid along with their transactions and miners pick up the ones with highest bid to confirm, is applied. In this case, users simply pay what they bid upon inclusion in a block, and as we know, it will end up in first price auction among miners. However, no dominant bidding strategy exists for users in this case, and thus users are likely to underbid or overbid and have a terrible experience. To resolve the issue, EIP-1559 protocol [1] is applied in Ethereum. In this protocol, there exists a system configuration $base\_fee$, which fluctuates according to the market and network for better flow control. And instead of a single

number, users should bid a 2-dimension vector (cap, tip) for their transactions, where cap is the maximum total price they can tolerate per gas, and tip is the maximum price per gas they're willing to award to the miner. For an included transaction, the user needs to pay $min(cap, base\_fee + tip)$ per gas used, while $base\_fee$ is burnt by the system per gas, and only the remaining part is rewarded to the miner. After the activation of EIP-1559 after London hard fork [7], there is rising concern in both academic world and industrial world on the new TFM.

In this work, regarding the fluctuation of base fee and the fact that miners and users don't play a zero-sum game any more, we propose a new category of attack through cooperation between miners and users, and investigate into it from theory, simulation and empirical analysis. Our main contribution can be summarized as follows.

- We identify an inherent weakness of EIP-1559 owing to the fluctuating base fee, where miners and users have the possibility of cooperating to form a side contract to lower down the base fee, and thus increase their own profit in theory.
- We simulate the side contract based on history data, and show that the side contract can actually benefit both miners and users in practice.
- We also do an empirical analysis to identify whether or not there has already been such misbehavior in the wild. The results suggest that some miners tend to deviate from honest mining but have less on-chain profit than honest behavior, which implies that they may get extra profit from possible side contracts.

2 Related Work

Regarding Transaction Fee Mechanism Design (TFM), there have been some notable results recently. Chung et al. [3] give a rigorous definition on what properties a good TFM should achieve and use the criteria to inspect Bitcoin, successfully proving in theory that, under the assumption that block capacity is finite which is always true in practice, UIC (User Incentive Compatible), MIC (Miner Incentive Compatible) and 1-SCP (Side Contract Proof Between 1 Miner and 1 User) can't be satisfied at the same time under traditional definition of blocks in Bitcoin. Furthermore, they put forward a TFM allowing blocks to carry unconfirmed blocks, which is a variant of Second Price Auction, and show that it can satisfy all three properties at the same time. It is a groundbreaking work on TFM in blockchain, which is also inspiring on how we can analyze EIP-1559 in theory.

Regarding EIP-1559 protocol [1], Tim Roughgarden [8] gives a theoretical result that EIP-1559 can achieve 3 plausible properties , under the assumptions that the block capacity is infinite (or in other words, there is no network congestion) and both miners and users are myopic, building a theoretical foundation for similar analysis. However, both assumptions are unlikely to hold, since there is always severe congestion in Ethereum and miners are usually regarded to care more about long-term profit in common practice. Nonetheless it is still the foundation of theoretical analysis on EIP-1559. Apart from theory, Liu et al. [5] give a comprehensive overview on the impact of EIP-1559 in practice from empirical analysis, and the positive effect of EIP-1559 protocol gets confirmed in this work. Also, Lerner et al. [4] put forward several possible attacks

to EIP-1559 TFM, but neither any rigorous theoretical analysis nor any empirical or practical experiments have followed up. Nowadays, although EIP-1559 has triggered much interest in research, both academic and industrial communities are still in a lack of analysis on whether or not it is a secure and robust TFM. Either any successful attack or a theoretical proof on plausible properties with reasonable assumptions on EIP-1559 can help a lot.

3 Side Contract Design Rationale

3.1 Revisiting Base Fee in EIP-1559

In the original Ethereum transaction system, all gas fee is transferred to miners, and the total supply of ethers is increasing due to newly produced ethers. However, in EIP-1559, some gas fee is burnt in order to enhance user experience by avoiding underbidding or overbidding, in which a dominant bidding strategy exists for users according to Tim Roughgarden's paper [8]. Moreover, another positive effect of EIP-1559 is to decrease the supply of circulating ethers, suppress the inflation, and result in a higher value of ETH token.

Base fee is the basic amount of ethers for a transaction. Instead of being transferred to the corresponding miner who includes the transaction like what is done in traditional settings, it should be burnt by the system, instead. And the base fee fluctuation process is subject to the following rule.

Suppose r_c is the base fee of the current block, r_p is the base fee of its predecessor block, and s_p and s_t are size of predecessor block and target block size, correspondingly. The rule of base fee adjustment is:

$$r_c = r_p \cdot (1 + \frac{1}{8}\frac{s_p - s_t}{s_t})$$

And we can see that the base fee can fluctuate fiercely in a short period of time, and they can also be easily and deliberately manipulated by miners.

Also, from users' view, a user should specify *tip* and *cap* for a transaction, where the miner pays $min(cap, tip + base\_fee)$ and $min(tip, cap - base\_fee)$ will go to the miner who includes the transaction while the other part is burnt by the system.

In this way, as analyzed in Tim Roughgarden's paper [8], a user has a dominant strategy in a bidding process, which is to set *tip* to be negligible and *cap* to be the interior value of the transaction. It's a great improvement compared to first price auction, which can significantly reduce overbidding and underbidding.

On top of that, however, it's not hard for us to observe that, a decrease in base fee will benefit both miners and users. For miners, a decrease in base fee means less burnt fee and more transactions to include in a block, which results in higher transaction fee. And for users, if she bids according to the dominant strategy mentioned above, then she will pay $min(cap, base\_fee + tip)$ which will also likely decrease with $base_fee$, and that's because $base\_fee < cap$ should hold for the transaction to be included, from which we can derive $base\_fee + tip \leq cap$ if tip is a negligible amount.

The observation above gives us the motivation for designing a possible side contract. And the three facts to motivate us to design a possible side contract can be summarized as follows.

1. **With EIP-1559 protocol, miners and users don't play zero-sum game any more.** The base fee, a significant part of total transaction fee, is burnt by the system. Thus a reduction in base fee can benefit both miners and users.
2. **Some users aren't time-sensitive, and they are willing to have longer confirmation time while pay less.** In Tim Roughgarden's paper, it has been shown that, for users who want to get their transactions confirmed ASAP, they have a dominant bidding strategy by staying aligned with protocol. However, some users want to save more by waiting. In particular,for honest users, they can wait for valley of natural base fee fluctuation, but for rational users, it is a motivation for deviation from being honest.
3. **Through coalition of miners and users, the base fee can be decreased.** A simple appetizer is that, if a miner mines two consecutive blocks of size s_t where s_t is the target block size, then the transactions included in the blocks will yield the current base fee. However, if the first block is set empty intentionally and the second block has size $2s_t$ instead, all the users can enjoy $\frac{7}{8}$ current base fee by the aforementioned formula, while the miner won't suffer from any lost. Thus it sheds light on a possibility that miners and users can form a coalition to lower down the base fee.

3.2 A Paradigm for Side Contracts

In this part, we shed light on a simple side contract in our design. The side contract is set up by a miner, and users will decide whether or not to join it by themselves. Once sufficient number of transactions are collected, the miner can start the attack.

Contract Setup In the contract, miners can specify a bid b_0 for users' total paid transaction fee along with a bid b_1 for the transaction's on-chain max-fee bid. For users accepting the contract, they need to do the following.

1. Initiate a transaction with max-fee bid b_1 and tip 0 and broadcast the transaction on chain.
2. Privately send $b_0 - b_1$ ethers to the miner and specify the transaction.

Additionally, the miners can specify a guarantee for transaction inclusion time and refund the users if the time constraint isn't met, but it is just optional to attract more miners to join the coalition without risk. Also, it is worth noting that, miners are free to set up the contracts, and thus a miner can set up multiple such contracts with different parameters.

User Decision Process From users' view, they usually have a utility function for making decisions. And a possible utility function for rational users is: $U(b) = V(WT(b)) -$

b, where $WT(b)$ is the predicted waiting time for bid b, $V(t)$ is the value of a transaction (from users' view) with latency t, and the users select the bid which maximizes their utilization.

Based on simulation on historic data (the details are hereby omitted), a possible model for $WT(b)$ and $V(t)$ for honest users is

$$WT(b) = \frac{k}{b} \qquad V(T) = Val_{max} e^{-\lambda T}$$

In the second expression, Val_{max} is the value of the transaction if it's immediately included, and λ is a parameter determined by the user's eagerness.

For the side contract, the utility value is fixed (determined by the contract). The value should be larger than the maximum value of utility function to incentivize users to collude with the miners.

Attack Launching Process Then, the malicious miner can launch the attack once sufficient colluding transactions are collected, where "sufficient" means transactions can fill in a full block of size $2s_t$.

There is hardly any explicit criteria for deciding "enough" due to undulating base fee, but, since malicious miners can freely switch from attacking to honest mining, they can mine honestly if colluding transactions aren't enough, and thus the attacking process is actually dynamic and miners won't have concern on inability to mine a full block when needed.

To launch the attack, the colluding miner always follows "empty-full" alternation pattern to produce blocks. Note that, they just follow this way themselves, regardless of whether or not the two blocks are consecutive in the global view of the whole blockchain (i.e. If the last block mined by her is empty, then the following one is full, and vice versa).

For a colluding transaction, once the base fee for a full block drops below b_1 , it gets included in the corresponding block.

4 Theoretical Analysis on the Side Contracts

4.1 A Practical and Explicit Side Contract Instance

Based on the aforementioned paradigm, we can get an explicit instance for side contracts based on the following assumptions for further analysis.

1. All such side contracts yield $b_0 = b_1 + \epsilon$ (i.e. The miners almost produce "empty-full" alternating blocks voluntarily , it's reasonable because there can be multiple such malicious miners competing in attracting users to join.).
2. We assume any contract with arbitrary b_0 exists. (i.e. Users can select their ideal price for the transaction.)
3. A colluding transaction with bid b_0 gets included in the first full block mined by the contract-initiating miner with base fee not exceeding b_0. (i.e. it's an optimistic assumption for users which discards possible congestion.)

Also, we make the following assumptions for honest transactions.

– An honest transaction with bid b (i.e. max fee cap) gets included in the first block which is mined by anyone but the contract-initiating miner and has base fee strictly less than b. (Note that it's an optimistic assumption for users which discards possible congestion.)

4.2 Two-Party Assumption

To further model the miners' and the users' behavior, we also adopt the "two-party assumption" as follows.

First, we can make the following observation. For colluding transactions, since the tip is 0, they won't be included by honest miners. While for honest transactions, since the colluding miner doesn't receive the privately-paid compensation, they won't be included by colluding miners, either.

Thus we can establish a model where honest miners and honest users form one party, while colluding miners and users form another independent one. The two parties have independent mempool, whereas the only thing connecting them is the base fee.

And we'll base all the following theoretical computation on the concrete instance mentioned above.

4.3 Theoretical Analysis on Why the Side Contract can be Profit-Winning for Both Parties

In this part, we rigorously analyze how much on earth the side contract can benefit the coalition. Suppose the hash rate of the malicious party is x.

A Myopic View Then let's analyze the net profit of the coalition. In the coalition, the miner should publish an empty block first, where the duration for publishing one block is $\frac{1}{x}$ block time for the colluding miner. Then, by our assumption, for the next block mined by the same miner, the expected latency is $\frac{1}{x}$ block time, and in expectation the base fee for the block will be $\frac{7}{8}b$ where b is the original base fee, since the "two-party assumption" guarantees that any change in coalition market won't affect the honest market. For transactions, since all the colluding transactions form an independent pool by our assumption, they can't be grabbed by honest miners, and therefore there won't be any loss for miners in the sense of transaction fee.

Thus the expected combined revenue for miners and users in every $\frac{2}{x}$ blocks is $\frac{1}{8}b2s_t$ due to decreased base fee, averaging $\frac{bxs_t}{8}$ gwei per block time. Take base fee 50 gwei per gas, target block size $15M$ gas, and block time 13 seconds in today's Ethereum, and the net combined revenue is around $0.0072x$ ethers per second, and that's $11.88x$ dollars per second on today's ether price. Take $x = 0.1$ as an example, it means that the coalition can earn more than 1 dollar per second, and it will lead to a considerable long-term stable profit if the coalition continues.

From Non-Myopic View Another reasoning for why the coalition can form is that, apart from short-term interest-driven agents, rational agents who care about their long-term revenue should also be willing to join the coalition.

Note that, through the aforementioned "empty-full" alternation, the base fee will be decreased to $\frac{7}{8} * \frac{9}{8} = \frac{63}{64} < 1$ times original base fee in every $\frac{2}{x}$ blocks. As a result, it will have a long-term effect on the stability of base fee because the decrease is actually irreversible, since the quantity of transactions included in the honest case and the colluding case are actually identical. So, if the attack continues, the base fee will gradually decrease to a negligible amount. Take

Since both miners and users appreciate a low base fee thanks to fewer burnt tokens, they will be willing to create and join the side contract to cooperate in lowering down the base fee.

What's even worse is that, although one may argue that the first point can be regarded as MEV (Miner Extractable Value) which is common in today's blockchain world which won't make a negative impact on the Ethereum ecosystem, the latter point indicates that the base fee can be permanently lowered down to a very low level. And it essentially declares the death sentence to EIP-1559, and that's because, with negligible base fee, EIP-1559 will degenerate into the original first price auction on tip.

In conclusion, not only miners and users can get illegal income from the collusion, but this kind of side contracts can harm the Ethereum ecosystem permanently, as well. Thus, from the theoretical analysis, there's a strong need that we investigate into and put forward possible approaches to prevent these harmful side contracts.

5 Simulation of Side Contracts

5.1 Simulation Configuration

To see whether or not users can actually gain more benefit by joining the side contract, we adopt the following assumption to control variables.

- **For the coalition:** We suppose there exists exactly one coalition of certain strength where strength is defined as $min($portion of colluding users, the miner's mining power$)$, i.e. the expected portion of blocks produced by the coalition. The rationale of the coalition is exactly the same as the aforementioned concrete instance. And users will individually decide whether or not to join it.
- **For users' decision:** To better control the variables, we suppose each user has a threshold of the maximum gas price they can afford for a transaction, and the strategy set is $\{honest, collude\}$, where honest means that the user will wait for the fluctuation of base fee until it's lower than the users' expectation, while collude means that the user joins the coalition. To better show the possibility of such side contracts, we also suppose that all users are myopic, which means that they will always choose the option that minimizes their waiting time.
- **Network condition:** Here we optimally suppose that there's no congestion in the network, which implies that a transaction can be included as long as the base fee is lower than the maximum possible gas price that the corresponding user can afford.

Then, to show whether or not the miners will join the coalition is equivalent to showing the relationship between bid and waiting time under the two selections.

To explore the relationship between user bid and waiting time of Ethereum after the application of EIP-1559 in real world, we exploit the mempool data in history and the block information for further analysis. We sample two time intervals on Ethereum main chain after London hard fork for our simulation.

- One has the height interval $[14771453, 14777776]$.
  It is the set of all blocks on May 14, 2022 (EDT), which is a normal day, and we use I1 to denote this interval.
- The other has the height interval $[14687398, 14691144]$.
  This is the set of blocks from the afternoon of Apr 30 to the morning of May 1 in 2022 (EDT), during which there is a surge in needs, and the gas price was as high as several thousand gwei for a long time. And it's denoted as I2.

The two sampled time intervals can altogether give us a comprehensive view on the actual relationship between bid and waiting time under all circumstances.

5.2 Bid → Waiting-time Relation for Included Transactions in Reality

With the help of the data of all transactions entering the mempool on May 14, we analyze the set of transactions that satisfy the following two properties.

- The transaction enters the mempool at anytime on May 14.
- The transaction is included in a block in the height interval $[14774524, 14777776]$ (the set of blocks produced in the later 12 hours on May 14).

In this way, we consider all transactions included in the later 12 hours on May 14, only except those abnormal ones which have super long latency of more than a half day. And the set of transactions is symbolic enough to represent general transactions on a normal day.

A Simple Appetizer: From bid value to waiting-time We collect the bid gas price, where the gas price is defined as cap for Type-2 transactions, and waiting time of all these transactions, and the scatter diagram is shown in Figure 1.

Process the data to retrieve the median and 90-percent quantile for every bid interval of length 1 gwei, and results are shown in Figure 2.

Be More Practical: What about considering the ratio of the bid value to the base fee at the proposal time? From the system's perspective, the bid itself wouldn't be determinant in the inclusion of a transaction, because the introduction of base fee in EIP-1559 also plays an important role in determining a transaction's inclusion status. And from users' perspective, they can always have a good idea of the current base fee and bid accordingly. So it's practical to use the ratio of bid and the base fee at the transaction-proposal time instead of the bid value itself to analyze the bid-waiting time mapping.

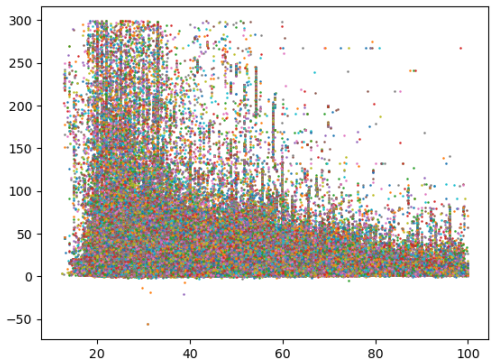

Fig. 1: Overall mapping from bid value to waiting time

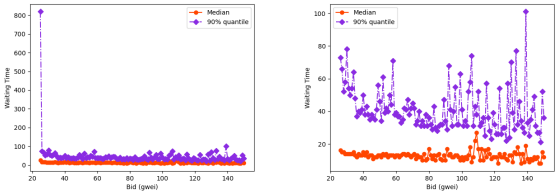

(a) Global View (b) Partial View

Fig. 2: Relation Between Bid Value and Waiting Time.

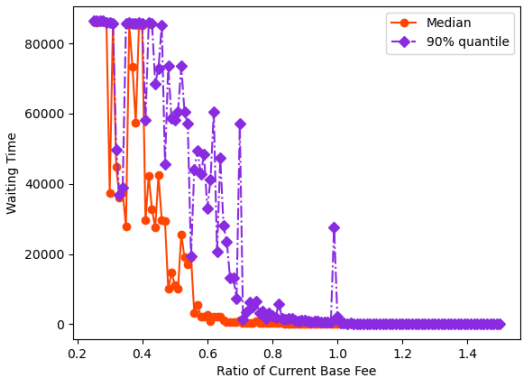

Fig. 3: Relation Between Bid-Base Ratio and Waiting Time (Global View).

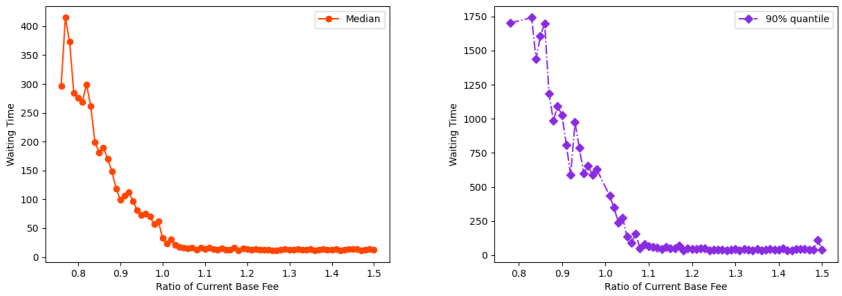

(a) Median (b) 90% Quantile

Fig. 4: Relation Between Bid-Base Ratio and Waiting Time (Partial View).

Process the data to retrieve the median and 90-percent quantile for every ratio interval of length 0.01, and we can get the following, as shown in Figure 3 4.

It can be seen that the waiting time is highly relevant to the bid ratio to base fee rather than the net value of the bid itself, so, in the following, we'll lay more emphasis on the ratio rather than the value.

5.3 Bid → Waiting-time Prediction for Hypothetical Transactions by Honest Waiting

In this part, we further investigate into what we should expect to happen when we insert a transaction of a certain bid at a certain time.

From the diagrams in the previous subsection, we can find out that, if we bid more than the current base fee, it's very likely that our transaction can be included within a satisfactory short time. So here we lay emphasis on the contrary, what if we think the current base fee is too high to accept and bid a value smaller than the current base fee?

Here we measure the waiting time for a transaction with bid b by the number of blocks between its appearance in mempool and the first block with base fee not exceeding b. For a given ratio r, we iterate through all time slots t in the sampled interval, and

calculate the waiting time $\{\Delta_t\}_t$. The following diagrams plot two symbolic quantiles (i.e. 50% and 90%) of $\{\Delta_t\}_t$, and can help us understand the change in waiting time if we bid $r \cdot b$ for ratio r and current base fee b at a certain time in the sampled interval.

The first two figures in 5 are for Interval 1 (normal case).

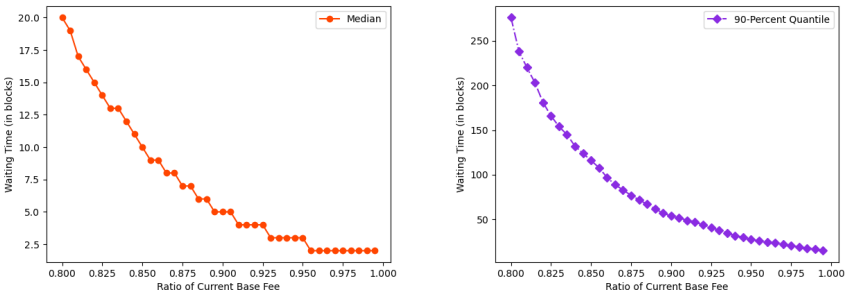

(a) Median (b) 90% Quantile

Fig. 5: Predicted Relation Between Bid-Base Ratio and Waiting Time (I1).

The second two figures in 6 are for Interval 2 (a surge in needs).

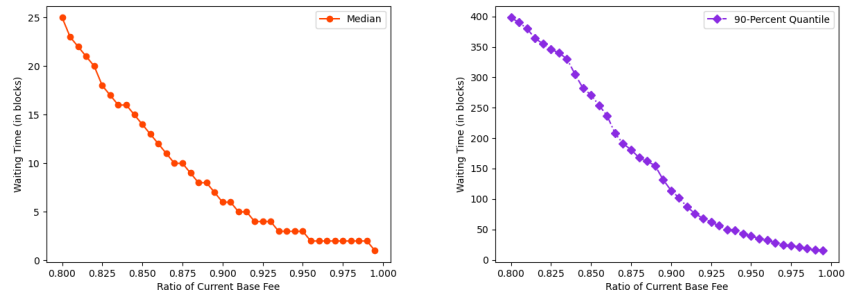

(a) Median (b) 90% Quantile

Fig. 6: Predicted Relation Between Bid-Base Ratio and Waiting Time (I2).

5.4 Possibility of Side Contracts

Consider the side-contract attack mentioned before, where miners produce blocks in an "empty-full-empty-full-..." manner and attract users to join the coalition to consistently reduce the base fee.

First, it's worth mentioning that, due to less burnt fee, the reduction of base fee is actually profitable for both miners and users in the long term. Thus, for non-myopic miners or users who are willing to sacrifice a bit of current interest for long-lasting benefit in low base fee, the side contract problem is inevitable.

However, in practice, most miners and users are myopic and won't voluntarily donate their short-term profit for an illusory re-gain in the future. And here we argue that, even if miners and users are myopic, there is still possibility that the collusion f. The detailed evidence is as follows.

In our simulation, we suppose there has already been a collusion with strength $x \in (0, 1)$ where the strength of a collusion is defined as the portion of blocks on the blockchain corresponding to the coalition, and can be calculated by $min($miner's strength, portion of colluding users$)$. Then, for different x and each time slot in Interval I1, we suppose there exists a user with highest possible gas price pb where $p \in (0, 1)$ and b is the current base fee. Then we investigate the waiting time of the transaction for both two possible selections, and the results are shown in Figure 7 and 8.

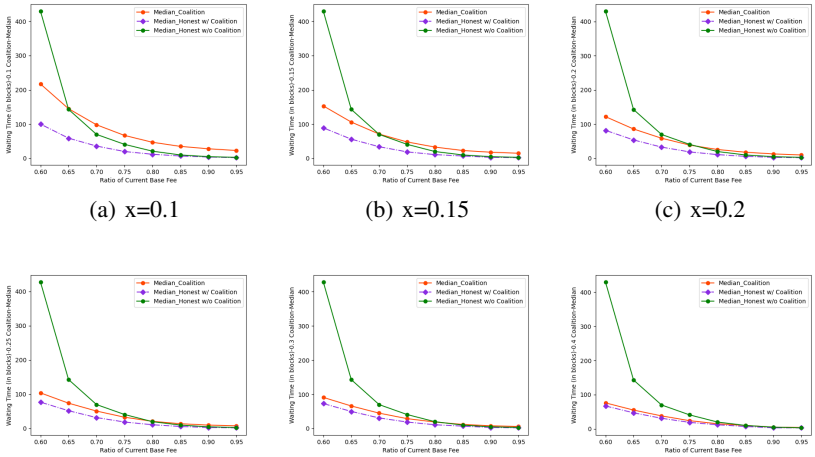

(a) x=0.1 (b) x=0.15 (c) x=0.2

(d) x=0.25 (e) x=0.3 (f) x=0.4

Fig. 7: Median of Waiting Time under Different Selections

From the diagrams above, we can figure out that, when compared with honest waiting, in most cases joining a coalition is more profitable for users due to less waiting time when paying the same amount of transaction fee, especially when $x > 0.2$ or $p < 0.7$. And it suggests that such a coalition is really a risk to the whole system, because users can actually be incentivized to join it, showing the high possibility of such side contracts even if users are myopic.

5.5 Instability of Side Contracts

Although we draw the conclusion that, even if users are myopic, such side contracts are likely to form if miners and users are self-interest driven, the diagrams 7 8 also suggests the instability of such side contracts.

From the diagrams, it is clear that, if a coalition does form, miners who choose honest waiting can also benefit from the decrease in base fee. And as a result, as long

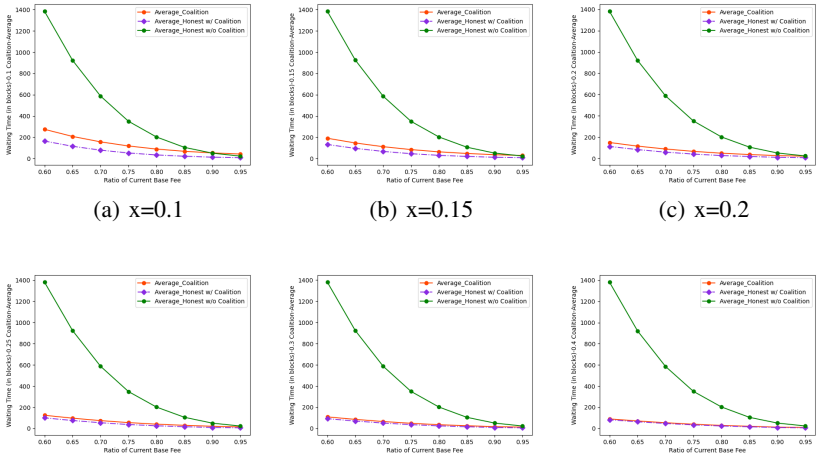

(a) x=0.1 (b) x=0.15 (c) x=0.2

(d) x=0.25 (e) x=0.3 (f) x=0.4

Fig. 8: Expected Waiting Time under Different Selections

as $x < 0.5$, they can even enjoy less waiting time than colluding miners. Thus, once the coalition forms and the side contract begins to work, it is more profitable for users to jump out of the contract and deviate to honest waiting. However, if all such users quit, then the coalition stops working, and it is more profitable to deviate to joining the coalition to resume it afterwards. Thus it is predictable that the size of coalition and whether or not it can work will be erratic.

Nevertheless, it doesn't rule out the risk of such side contracts to the system. On one hand, miners who initiate the contract can further incentivize the users not to deviate. On the other hand, even if such side contracts don't always exist, once it exists and works for a certain period of time, the impact on base fee will be irreversible, and the system is thus facing the risk of negligible base fee and degrading into first price auction if the contract lasts for a relatively long time.

6 Evidence of Possible Collusion in the Wild

To show the evidence of possible collusion that is already happening in the wild, we adopt the following two metrics to identify some suspicious behavior.

6.1 Evidence from Block Size of Each Miner

We sampled 100,000 consecutive blocks on Ethereum in early July of 2022 and investigate into every non-negligible miner (i.e. miners who occupy more than 0.1% of total computation power). We use "over-small" to denote the blocks that is less than 0.4 target size and "over-large" to denote the blocks that is larger than 1.6 target size. The results are shown in Table 1.

Miner Address	Hash Power	Over-small Ratio	Over-large Ratio
All blocks	100.0%	23.41%	27.96%
0xab3b229eb4bcff881275e7ea2f0fd24eeac8c83a	1.573%	45.96%	15.32%
0x0c7213bac2b9e7b99aba344243c9de84227911be	0.879%	29.24%	22.08%
0x1ca43b645886c98d7eb7d27ec16ea59f509cbe1a	0.369%	40.38%	17.89%
0x1ad91ee08f21be3de0ba2ba6918e714da6b45836	10.212%	30.13%	20.96%

Table 1: Results on Block Size

From the table, we can find out that, the four mentioned miners' block size has an obvious deviation from the general public, and they are suspicious for deliberately mining small blocks due to abnormal high rate of "over-small" blocks. However, the evidence isn't quite convincing because the difference in block size may be coincidence, and thus we conduct the following experiments.

6.2 Contrast Between Honest Mining and Miners' Actual Behavior

For 10,000 consecutive blocks in mid August of 2022, we further simulate a simplified honest behavior for each block as follows.

- **Transaction Pool:** All transactions that are included in the following 2,000 blocks whose timestamp recorded by our full node is prior to the block's publication time.
- **Strategy:** We adopt a simple greedy algorithm to select transactions. First sort all transactions in the pool by their profit per gas (computed by $min(tip, cap - base)$), and then scan over them in descending order of profit per gas, and include the transaction if it is feasible.
- **Simplification:** For our light-weighted simulation, we don't bother to check the validity of each transaction using complicated criteria. We just record each transaction's nonce, and regard a transaction feasible as long as it yields the smallest nonce that hasn't appeared on the blockchain yet.

Note that it's neither the optimal strategy nor a valid strategy because we simplify the transaction validity check phase and apply a simple greedy algorithm, and there does exist some blocks that earn more reward than our simulation, but it does set a baseline. The difference of transaction fee earning between our simulation and reality is shown in the Table 2.

Although there exists some system error in the simulation due to the simplification, we can still figure out that some miners tend to give up some on-chain reward and deliberately produce small blocks. Since we know miners are self-interest driven and rational, a possible reason to account for this phenomenon is that, they are also engaged in some side contracts (not necessarily our aforementioned side contract, but can be something similar) and receive compensation for their on-chain loss from some other side channels. And it provides us the evidence that some miners have already deviated from honest mining and increase their profit from possible side contracts.

Miner Address	Hash Power	Avg Tx-Fee Deviation	Avg Tx-Fee Deviation Rate
All blocks	100.0%	0.07235 ETH	60.16%
0xab3b229eb4bcff881275e7ea2f0fd24eeac8c83a	2.512%	0.09712 ETH	85.36%
0x0c7213bac2b9e7b99aba344243c9de84227911be	0.725%	0.08744 ETH	70.28%
0x1ca43b645886c98d7eb7d27ec16ea59f509cbe1a	0.475%	0.08029 ETH	68.88%
0x1ad91ee08f21be3de0ba2ba6918e714da6b45836	10.463%	0.08346 ETH	65.12%

Table 2: Results on Transaction Fee Deviation

7 Conclusion

In this work, we first identify an inherent weakness of EIP-1559 transaction fee mechanism caused by fluctuation in base fee, and then put forward a side contract attack towards the weakness. Furthermore, we use both theoretical analysis and simulation to show that the side contract can be profitable to both miners and users, and thus it's a potential risk to the system, although the attack may not be long-lasting. In the end, we use empirical analysis to show that the attack or similar ones may already happen in reality. Moreover, after "the Merge" where Ethereum changes into "Proof-of-Stake" in the near future, the miners can know when they can produce the blocks in advance, and the success rate of such attacks can be even higher, casting a potential security risk of the current Ethereum system.

References

1. Buterin, V., Conner, E., Dudley, R., Slipper, M., Norden, I., Bakhta, A.: Eip-1559: Fee market change for eth 1.0 chain. Published online on GitHub (2019)
2. Buterin, V., et al.: A next-generation smart contract and decentralized application platform. white paper **3**(37) (2014)
3. Chung, H., Shi, E.: Foundations of transaction fee mechanism design. arXiv preprint arXiv:2111.03151 (2021)
4. Lerner, S.D.: Flaws in ethereum's eip-1559. Medium Blog URL: https://medium.com/iovlabs-innovation-stories/flaws-in-ethereums-eip-1559-c0f91838ce23 (2022)
5. Liu, Y., Lu, Y., Nayak, K., Zhang, F., Zhang, L., Zhao, Y.: Empirical analysis of eip-1559: Transaction fees, waiting time, and consensus security. arXiv preprint arXiv:2201.05574 (2022)
6. Nakamoto, S.: Bitcoin: A peer-to-peer electronic cash system. Decentralized Business Review p. 21260 (2008)
7. Pramod: Ethereum 2.0 london hard fork: A brief snapshot. Published online as a blog post on Medium forum (2021)
8. Roughgarden, T.: Transaction fee mechanism design for the ethereum blockchain: An economic analysis of eip-1559. arXiv preprint arXiv:2012.00854 (2020)
9. Szabo, N.: Formalizing and securing relationships on public networks. First monday (1997)