

On the Optimality of EIP-1559 for Patient Bidders

(Draft – Comments Welcome)

Moshe Babaioff* Noam Nisan†

November 4, 2024

Abstract

The “EIP-1599 algorithm” is used by the Ethereum blockchain to assemble transactions into blocks. While prior work has studied it under the assumption that bidders are “impatient”, we analyze it under the assumption that bidders are “patient”, which better corresponds to the fact that unscheduled transactions remain in the mempool and can be scheduled at a later time. We show that with “patient” bidders, this algorithm produces schedules of near-optimal welfare, provided it is given a mild resource augmentation (that does not increase with the time horizon). We prove some generalizations of the basic theorem, establish lower bounds that rule out several candidate improvements and extensions, and propose several questions for future work.

1 Introduction

1.1 Background on Blockchains

Blockchains like Bitcoin [Nak09] or Ethereum [But14] operate by repeatedly aggregating transactions into blocks. As a first approximation, each transaction requires some “size” and has some value for its user, and each block has some size limit, where the notion of “size” is defined by the blockchain.¹ Each block is assembled by a different “operator”² which the blockchain protocol designer tries to motivate to maximize the total value subject to the block’s capacity constraints. A key difficulty here is that both users and operators are selfish and strategic, so the mechanism that specifies which transactions are accepted to the block and how much do their users pay and to whom must take such strategic behavior into account. A large literature has studied such mechanisms. While Bitcoin’s “pay your bid” mechanisms is relatively simple (but not incentive compatible for the users), Ethereum’s mechanism [BCD⁺19], known as EIP-1559, is more sophisticated and, following [Rou21], a large literature (e.g. [LMR⁺21, CS23, GTW24, AGHH23]) studies its strategic properties.

In this paper we undertake an algorithmic analysis of this protocol and the family of protocols that vary its various parameters. Specifically, we ask to what extent does it indeed succeed in maximizing the social welfare given the size constraints. Our main result is essentially positive: for “patient bidders”, the social welfare achieved by EIP-1559, with a slight tweak of parameters, is close to optimal. Before specifying more precisely what “close to optimal” means, we need to be more specific about our time-model of transactions.

Blocks are created one at a time³ and transactions arrive continuously through time. Much analysis of blockchains assumes that a transaction that arrives at a certain point in time *must* be scheduled in the immediately next block or loses all value for its user. This is called the model of “impatient users”. In reality, however, blockchains operate under an opposite assumption and submitted transactions remain in the “mempool” until accepted, and thus may be scheduled at any future block. The model that assumes that transactions can be scheduled at any future time without losing their value, which is closer to how blockchains actually operate, is termed a model with “patient users”, and has only received

*Hebrew University of Jerusalem.

†Hebrew University of Jerusalem and Starkware.

¹E.g. “gas” for Ethereum or (virtual-)bytes for Bitcoin.

²In bitcoin these are the “miners” and in Ethereum these are the “validators”.

³For Ethereum a block every 12 seconds and for Bitcoin a block every 10 minutes on the average.

modest attention [Nis23, HLM21, PS24, GY24]. A more general model may consider “partially patient” users who may have some time-sensitivity, e.g., model the value of a transaction as decreasing with its execution time [HLM21, PS24, GY24]. Unfortunately, current mainstream blockchain transaction fees mechanisms do not allow the user to express any time preferences. We show the near-optimally of EIP-1559 for fully patient users in the worst case, i.e. without any stochastic or distributional assumptions on the demand.⁴

1.2 The EIP-1559 algorithm and its parameters

The EIP-1559 algorithm assembles a sequence of blocks whose *average size* is at a *target* of a quantity B per block.⁵ I.e. the original constraint of maximum block size is relaxed, and is replaced by effectively limiting the *average* block size to be B , while relaxing the constraint of *maximum* block size to being only slightly larger than B .⁶ This relaxation is justified by the specific constraints of the Ethereum blockchain⁷ and may be viewed as a type of resource augmentation which was not technically studied in the literature of bin-packing and knapsack algorithms, and may be of independent interest.⁸

Here is how it operates: for every block t it choose a price p_t (according to an algorithm that we will explain shortly) and accepts only transactions whose value per unit is at least the price p_t . A difficulty arises if there are too many transactions with such high value per unit, as the EIP-1559 algorithm does not allow any block to be of total size that is larger than c times the target block size, with the parameter choice $c = 2$.⁹ In this case only a subset of the transactions that fit into the maximum block size is chosen, with the subset determined by an additional “tip” bid that is strategically chosen by each user. In this paper we will not undertake a strategic analysis, and we will analyze the protocol under the assumption that the subset of transactions is being chosen by an *adversary*¹⁰, making only the following mild assumption: the adversary must pick a maximal (by inclusion) set of transactions that fit into the allowed *maximum* block size.¹¹

The basic idea of choosing the price p_t for block t is rather simple: if the previous block size was above the target block size B , then we need to increase the price, and if it was below the target then we decrease the price. The increase or decrease is multiplicative according to the following formula: $p_{t+1} = p_t \cdot (1 + \eta(Q_t - B)/B)$, where Q_t denotes the total size of block t , and η is a small constant picked in Ethereum to be $\eta = 1/8$. This formula should be viewed as an approximation to the theoretically more elegant price-adjustment formula¹² $p_{t+1} = p_t \cdot e^{\eta(Q_t - B)/B}$ which we will analyze in this paper. Effectively, there is also a minimum price¹³ p_{min} , and if the previous formula drops below it then this minimum price is taken instead. The initial price p_1 was chosen arbitrarily¹⁴. In the rest of this paper whenever we talk about “the EIP-1559 algorithm”, we actually consider the algorithm with any setting of these five parameters: the target block size B , the maximum block size $c \cdot B$, the price adjustment parameter η , the minimum price p_{min} and the initial price p_1 (with $B > 0, c > 1, \eta > 0, p_{min} > 0$, and $p_1 \geq p_{min}$).

This algorithm produces blocks whose size may vary in the range 0 to $c \cdot B$. The *target* block size B is reached only in an average sense, but in a rather strong sense of averaging that we wish to formalize as follows.

⁴It is not difficult to observe that if one looks, instead, at *impatient users*, then any algorithm that *relies on historical prices* to set inclusion criteria for the current block (as does EIP-1559) may catastrophically fail to achieve any decent social welfare under adversarial conditions. We will also show that no *online algorithm* can achieve near-optimally for *partially patient* users, even if they are “very patient”.

⁵Ethereum’s choice is $B = 15M$ gas.

⁶Maximum size of only $2B$, a factor 2 increase that is independent of the parameters of the input (such as the range of per-unit values and the time horizon).

⁷I.e. it was deemed that the “syncing time” of “full nodes” (i.e., reading the entire history of the block chain) is more significant than the maximum block size limitation, and for it, only the average block size is important.

⁸This relaxation may be related in motivation to open-ended bin packing (e.g. [Eps22]), but not technically so.

⁹I.e. the maximum allowed block size in Ethereum is $30M$ gas.

¹⁰This naturally only strengthens our main result.

¹¹This is a mild assumption as any rational operator that is myopic is never better off by excluding transactions that can simply be added to the block.

¹²When a similar adjustment mechanism was later adopted for “blob gas” [BFL⁺22] then a better approximation to the desired exponential form was taken.

¹³This happens in the Ethereum implementation because all arithmetic is done in integers denominated in Wei, i.e. 10^{-18} ETH, and the price cannot become 0.

¹⁴The first Ethereum EIP-1559 block was given a price of 1 GWei, i.e. 10^{-9} ETH.

Definition 1.1. A schedule is said to have average block size limit B with slackness $\Delta : \mathbb{N} \rightarrow \mathbb{N}$ if for any number of times steps k , the total size of transactions during any k consecutive blocks is at most $(k + \Delta(k)) \cdot B$. When Δ is a constant function (for constant Z it holds that $\Delta(k) = Z$), we abuse notation and simply write $\Delta = Z$.

The basic property of the EIP-1559 algorithm is that it indeed produces blocks whose average size is the target size B .¹⁵ While this general fact is well known, we formalize it in our notion of “average block size limit B with slackness Δ ” as follows:

Proposition 1.2. Any schedule produced by an EIP-1559 algorithm with parameters $(B, c, \eta, p_{min}, p_1)$ has average block size limit B with slackness $\Delta = \frac{1}{\eta} \cdot \ln \frac{H}{L} + (c - 1)$, where H and L are upper and lower bounds on the per-unit value of any input transaction, and on the constants p_1 and p_{min} .

For completeness we provide a proof, in our model and notation, in Proposition 2.11.

1.3 Our Main result

Our main result states that the EIP-1599 algorithm is near-optimal for patient bidders, under some mild relaxations. Specifically, its allocation over slightly larger time horizon has social welfare that is close to that of the optimal, clairvoyantly chosen, schedule, that has maximum block size B .

Maximizing the welfare is a challenging problem. Even for a single block it is a knapsack problem, and thus NP-hard. On top of the computational issue, there are two additional hurdles that the EIP-1559 algorithm has to overcome. First, the scheduling problem that the EIP-1559 algorithm faces is an online problem, and so the algorithm is an online algorithm. While the offline (clairvoyant) optimal benchmark can take into account future transactions that may “fit with” some existing transactions, online algorithms cannot do so. Nicely, as we will show, the relaxation of the maximum block size constraint to an average block size constraint suffices to compensate for this online disadvantage, provided that the maximum block size limit is at least twice that of the benchmark, i.e. $c \geq 2$.

Secondly, the EIP-1559 algorithm is not only an online algorithm, but is further constrained as it is a “pricing-based” algorithm. It has to make decisions that are based on prices (so it cannot pick any allocation), and moreover, the price at each block is based only on on-chain information from previous blocks, without any knowledge of the pending transactions in the mempool.

Such algorithms must “search” for the correct prices, which may be difficult under adversarial conditions. As we show, the EIP-1599 algorithm manages to do so with a loss that can be overcome by extending the time horizon by the number of blocks that is logarithmic in the per-unit value range. This single logarithmic loss is global over the whole range of blocks, even though prices may keep fluctuating.

Let us use ALG to denote the schedule produced by an EIP-1559 algorithm, and for any schedule S let $SW(S, [1, t])$ denote the total social welfare of S up to time t .

Theorem 1. Fix an EIP-1559 algorithm with parameters $(B, c, \eta, p_{min}, p_1)$ for some $c > 2$. Consider any (adversarially chosen) sequence \mathcal{C} of input transactions where each transaction $i \in \mathcal{C}$ has per-unit value v_i satisfying $e^\eta \cdot p_{min} \leq v_i \leq v_{max}$ for some $v_{max} \geq p_1$. Let OPT_B be an arbitrary schedule with maximum block size B . Then for any time T it holds that

$$SW(ALG, [1, T + \Gamma]) \geq (1 - \delta) \cdot SW(OPT_B, [1, T])$$

for $\delta = 1 - e^{-\eta} \leq \eta$ and any integer $\Gamma \geq \frac{1}{\eta} \cdot \ln \left(\frac{v_{max}}{p_{min}} \right) \cdot \max\{1, \frac{1}{c-2}\} + c$.

We thus see that the EIP-1559 algorithm can indeed compete with the optimal schedule, but after suffering two types of losses: first, there is a loss of a small $\delta \leq \eta$ fraction of the social welfare, where as the adjustment rate η approaches 0 this loss vanishes. Second, the algorithm is given Γ additional rounds, where Γ is inversely proportional to η . Conceptually, if T was given in advance we could balance between these two losses (for any fixed v_{max}, p_{min} and c) by choosing $\eta = O(1/\sqrt{T})$, which

¹⁵Our analysis is for the theoretically clean version with $p_{t+1} = p_t \cdot e^{(\eta(Q_t - B)/B)}$. As shown in [LRMP23], the variant with $p_{t+1} = p_t \cdot (1 + \eta(Q_t - B)/B)$ results in a slightly larger average block size.

would lose $O(1/\sqrt{T})$ fraction of welfare every step as well as require an additional $O(1/\sqrt{T})$ fraction of steps.¹⁶

We remark that the EIP-1559 algorithm is not aware of the time horizon (T), and in practice it essentially can continue running indefinitely. The guarantee of the theorem holds concurrently for every time horizon T (with v_{max} being the maximal per-unit value up to T), even with respect to the time horizon dependent optimal allocation $OPT(T)$.

Before looking more carefully at generalizations and limitations of this theorem, it is useful to compare the result to two similar-in-spirit papers that study more difficult models and consequently obtain much weaker results for their models. The first is the analysis of [FMPS21] of price-based mechanisms in the model of impatient buyers. In this model they only obtain a constant factor approximation and only under a stochastic model. The second is the analysis of [ADM24] of a generalization of EIP-1559 to multiple resources. They obtained near-optimality (a regret bound somewhat similar to ours) for the case of impatient bidders, but only relative to the weak benchmark of a single fixed price.

We actually prove a significant generalization of this basic theorem simultaneously along several dimensions (see Theorem 2 for the exact formal result):

1. **Maximum Block Size vs. Maximum Transaction Size:** The limitation on the maximum block size $c > 2$ is required only when the largest transaction size q_{max} can be as large as the full target block size B . If we have a better bound on the maximum transaction size, then the limitation on c can be relaxed to $c > 1 + q_{max}/B$, and in the bound on Γ the expression $1/(c-2)$ is replaced by $1/(c-1-q_{max}/B)$ (with a slight change to the additive term c as well). Notice that when we have any bound on the maximum transaction size $q_{max}/B < 1$ then the actual choice made in Ethereum of $c = 2$ suffices for obtaining a good approximation and losing only $\eta = 1/8$ fraction of social welfare when allowed a small number of extra blocks.
2. **A Stronger Benchmark with Average Block Size Constraints:** As mentioned above, the EIP-1559 algorithm’s schedule has *average block size* B , a constraint that is more relaxed than having *maximum block size* B which is the benchmark in the theorem. We extend the theorem to show that the algorithm also competes with the benchmark of schedules with *average block size* B with any slackness Δ , where the bound on the constant Γ is now increased by Δ . In this generalization, the EIP-1559 algorithm with maximum block size $c > 1 + q_{max}/B$ competes with any schedule with average block size B , without requiring any upper bound on the benchmark’s maximum block size.
3. **A Stronger Benchmark with Fractional Allocations:** One of the things we get “for free” when we relaxed the maximum block size restriction to an average block size restriction, is that EIP-1559 turns out to also compete with any *fractional* allocation that is allowed to arbitrarily split transactions among blocks or even only fractionally allocate transactions (even though EIP-1559 does not do that, obviously).

1.4 Lower Bounds

Let us now look deeper at the losses as stated in the main theorem. Not only will we show that EIP-1559 indeed may suffer such losses, but we will also prove that such losses must be suffered by two classes of algorithms of which EIP-1599 is a member: a loss for the general class of online algorithms and an additional loss for a subclass that we term “price-based algorithms” that operate by setting a price for each block that only depends on on-chain information from previous blocks (see Definition 2.9 for a formal definition).

1.4.1 Maximum Block size

The theorem requires that $c > 2$, i.e. that the EIP-1559 algorithm is allowed *maximum* block size that is more than twice the original block size. We prove such a blowup in the *maximum* allowed block size is required by *any online algorithm* that obtains a welfare guarantee like ours.

¹⁶This is somewhat in the vein of regret bounds that balance between losing a small factor (that is proportional to the learning rate) every step and suffering an additional additive loss that is inversely proportional to the learning rate. Unlike in classic regret bounds the theorem does not combine both losses into a single $O(1/\sqrt{T})$ regret measure.

The lower bound is proven even if Γ and Δ can grow with T , as long as they grow as $o(T)$. Additionally, they hold even if the welfare approximation needs to be guaranteed only for one, known value T (that is large enough).

Proposition 1.3. *Any online algorithm that produces schedules with maximum block size $c \cdot B$, for some $c < 2$ and average block size limit B with slackness $\Delta(T) = o(T)$, even with an extension $\Gamma(T) = o(T)$, must lose at least $\delta_c = \min(1/8, (2-c)/3)$ fraction of social welfare relative to schedules with worst-case block size limit B (even when the values of all transactions are in $\{1, 2\}$).*

Of special interest is the case $c = 2$ which is the actual parameter used in Ethereum. While we do present an online algorithm (that can use information about non-scheduled transactions) with $c = 2$ that loses no social welfare, we show (Section 4.3) that for any horizon T there is a possible input sequence where EIP-1559 with Ethereum’s choice of parameter, $c = 2$, does *not* compete with the optimum. We do note however that this is due to very large transactions and, as mentioned above, if we limit the maximum transaction size by $q_{max}/B < 1$, then a positive result emerges also for $c = 2$.

1.4.2 Dependence on Range of per-unit Values

The theorem suffers a double loss: beyond allowing the algorithm Γ more steps as well as having Δ slackness, we lose a fraction δ of welfare. Furthermore, both the number of extra steps Γ and the slackness parameter Δ increase as a function of the range of possible per-unit values of transactions. While we exhibit an *online algorithm* with zero loss of welfare and $\Delta = 1, \Gamma = 1$, we prove that the dependence of Γ or Δ on the range of per-unit values is unavoidable for *price-based algorithms with good welfare guarantee*.

Proposition 1.4. *Fix any price-based algorithm with average block size B , slackness Δ , and extension Γ . There exists an instance with per-unit values in the range $[1, H]$ for which the algorithm does not obtain a $H^{1/4^{(\Gamma+\Delta)}}$ fraction of the optimal social welfare of a schedule with worst-case block size limit B (with no slackness or extension). Thus, if a price-based algorithm guarantees a constant fraction of this optimum, then $\Gamma + \Delta$ must grow as $\Omega(\log \log H)$.*

Note that there is an exponential gap between our lower bound (which grows at rate $\Omega(\log \log H)$) and the upper bound (that is only singly-logarithmic in the per-unit value range). While we do show that the EIP-1559 algorithm may indeed require such a singly-logarithmic additional number of blocks in order to compete with the optimum, closing the gap for general price-based algorithms remains open. Especially intriguing is the possibility of designing a useful “variant” of EIP-1559 with $\Delta + \Gamma$ that only grows in a rate double-logarithmic in the per-unit value range.

1.5 Model Extensions

Given our basic positive result of EIP-1559’s near-optimality, it is tempting to generalize this to natural extensions of the basic model. Specifically, we look at two tempting extensions: to “partially patient” bidders and to “multi-dimensional fees”. In both cases we prove the impossibility of an extension, and show that every online algorithm must lose at least some fixed constant fraction $\delta_0 > 0$ of the optimal social welfare. We leave open the question of whether one can recover at least some, smaller, constant fraction of optimal social welfare.

1.5.1 Partially Patient Bidders

All our results so far assumed “patient bidders”, i.e. where a transaction’s value for its user remains the same over time. The opposite assumption of “impatient bidders” assumes that a transaction *must* be scheduled in the immediate block or it loses all value for its user. A more general model would capture some sensitivity of the value of a transaction to the time of its execution, where the value decreases with time. The most common model for such time-dependent value [HLM21, PS24, GY24] would have a discount rate: A transaction with value v and *discount factor* ρ (where $0 \leq \rho < 1$) has value for its user of $v \cdot (1 - \rho)^{t_e - t_a}$, where t_a is its arrival time and t_e its execution time.

Thus the fully patient model corresponds to discount factor $\rho = 0$, while the impatient model corresponds to $\rho \rightarrow 1$. As the EIP-1559 algorithm does not allow its input to specify any discount factor, it is not hard to observe that if we just run EIP-1559 on partially patient bidders then, since

it does not distinguish between “new” transactions and old ones that already lost most of their value, it cannot produce highly efficient results. But perhaps if we just let the algorithm take into account, at every block, the *current* value of the transaction, then we regain efficiency? I.e., suppose that at block t , a pending transaction i that has already arrived is considered for allocation if and only if $v_i \cdot (1 - \rho_i)^{t-t_i} \geq p_t$? Unfortunately, we get a negative answer and, in fact, a lower bound for any online algorithm. Significantly, this lower bound holds for arbitrarily low discount factors.

Proposition 1.5. *Fix an online algorithm with average block size B , slackness Δ , and extension Γ , where $\Delta(T) + \Gamma(T) = o(T)$. For every minimum discount rate $\rho_{min} > 0$ there exist a time horizon T , and an input sequence where all bidders have discount rates $\rho_i \leq \rho_{min}$ and values in $\{1, 2\}$, for which the algorithm loses at least a fraction $\delta_0 = 1/20$ of welfare relative to the optimal schedule with worst-case block size limit B .*

This impossibility result applies to scenarios where different transactions may have different discount rates. We do conjecture that the “modified” EIP-1559 is near optimal if the discount rate is global and shared by all transactions.¹⁷ We do show however, that if we replace the model of “discount factor” with a model where transactions have a “patience level” where they can be scheduled within p steps of their arrival time without losing any value, but afterwards lose all value, then again, online algorithms must lose a constant fraction of welfare, even if all transactions have the same patience level.

1.5.2 Multi-dimensional Fees

In our basic model each transaction has a “single-dimensional” size q_i . More generally one may consider a model where there are m different resources, each block has a size limit B_j for every resource j and each transaction i uses the amount q_{ij} of each resource j (see e.g. [DECA23, ADM24]). The introduction of the “blob” resource to Ethereum [BFL⁺22] is a step in this direction. The natural generalization of the slackness condition to the multi-dimensional case would require that every T consecutive blocks use at most $(T + \Delta) \cdot B_j$ amount of each resource j . While one may hope to extend the near-optimality result to such a multi-dimensional model, it turns out any online algorithm must lose some constant fraction of social welfare in the multi-dimensional case.

Proposition 1.6. *If there are at least three resources then any online algorithm with average block size B_j , slackness $\Delta(T) = o(T)$ for each resource j and given an extension $\Gamma(T) = o(T)$, must lose at least a constant fraction $\delta_0 = 1/6$ of social welfare on some input sequence.*

1.6 Strategic Points of View

Our analysis in this paper is purely algorithmic, so our results can be viewed as proving near-optimally in terms of the *declared* social welfare: the welfare with respect to the values given as input by the bidders to the algorithm (rather than the true values). Regarding the operators (miners, validators), our analysis assumed adversarial behavior, subject to the weak constraint that they must schedule a maximal by inclusion subset of admissible transactions. While we do not formally make any strategic analysis, we do want to shortly mention some strategic points of view.

From the point of view of the bidders, the known incentive compatibility of EIP-1559 is only for impatient bidders¹⁸, while patient bidders may in fact shade their bids profitably. So can we expect truthful bidding from patient bidders? Well, practically, bidders may simply act myopically – in our case, truthfully – due to the complexity of predicting the future, and their bounded rationality. In our case, such behavior is also justified by risk aversion: Bidders that are sufficiently risk-averse with respect to the future will bid truthfully, taking the first available slot at a profit rather than risk never being scheduled due to possible future increase in demand. An appealing strategic analysis of patient bidders’ behavior that includes a reasonable model of future uncertainty is of course of much interest, and is left as a future challenge.¹⁹

¹⁷Likely suffering an additional loss that behaves like $(1 - \rho)^{O(\Delta + \Gamma)}$ when all transactions have discount rate ρ .

¹⁸Even for impatient bidders, the incentive compatibility of EIP-1559 does not hold in the case of a rapid demand increase as then tips kick in and winners pay the tips they offer.

¹⁹It is interesting to note that if in equilibrium the declared values are monotone in the true values – as may be expected in many models – then optimizing for the former is closely related to optimizing for the latter. In that case, our algorithmic results may carry over to the strategic setting in the model in question. A precise statement to such a result would obviously depend on many details of the model.

From the point of view of the operators that assemble the blocks, our analysis undertook an adversarial point of view. Thus, it should hold under a very general class of strategic behaviors of the operators, including any reasonable model of myopic operators as studied so far in the literature. The only required assumption is that scheduling a set of transactions is not less beneficial to the operator than scheduling only a subset of them.

2 Model and Notations

2.1 Blocks, Schedules, and Social Welfare

We consider transactions that are processed in blocks. At each time $t \in \{1, 2, \dots\}$ one block is processed. We consider a set \mathcal{C} of transactions that arrive online over time, at each time $t \in \mathbb{N}$ a finite set of transactions arrives. Each *transaction* $i \in \mathcal{C}$ has an arrival time $t_i \in \mathbb{N}$, size $q_i > 0$, and value per-unit size of $v_i > 0$ (thus its total value if executed is $q_i \cdot v_i$). Transactions are patient, i.e. may be scheduled at time t_i or at any time afterwards, without loss in value. We use $q_{max}^T = \max_{i \in \mathcal{C}} q_i$ to denote the maximal size of any transaction arriving up to time T , and use $v_{min}^T = \min_{i \in \mathcal{C}} v_i$ and $v_{max}^T = \max_{i \in \mathcal{C}} v_i$ to denote the minimal and maximal per-unit value of any transaction arriving up to time T , respectively. When T is clear from the context we drop it from these notations and write q_{max} , v_{min} and v_{max} (instead of q_{max}^T , v_{min}^T and v_{max}^T).

Definition 2.1. *Given a set of transactions \mathcal{C} (where finitely many arrive at each time step) a schedule S is an assignment of transactions to times, such that a transaction cannot be assigned before it arrives). Formally, for each $i \in \mathcal{C}$ and $t \in \mathbb{N}$, a schedule $S = \{x_{i,t}\}_{i \in \mathcal{C}, t \in \mathbb{N}}$ where $x_{i,t} \in \{0, 1\}$ specifies, for each $i \in \mathcal{C}$, $t \in \mathbb{N}$, whether transaction i was scheduled at time t , under the constraints that for every transaction $i \in \mathcal{C}$:*

- the transaction is not scheduled before it arrives: $t < t_i$ implies that $x_{i,t} = 0$.
- the transaction is assigned no more than one slot: $\sum_{t \in \mathbb{N}} x_{i,t} \leq 1$.

Definition 2.2. *A fractional schedule $S = \{x_{i,t}\}_{i \in \mathcal{C}, t \in \mathbb{N}}$ is a schedule satisfying the exact same two constraints, but relaxing the integrality requirement on $x_{i,t}$, allowing for any fraction $x_{i,t} \in [0, 1]$.*

Definition 2.3. *We use $Q_t(S) = \sum_{i \in \mathcal{C}} x_{i,t} \cdot q_i$ to denote the total capacity used by the schedule S at time t . When the schedule S is clear from the context we may omit it in the notation and denote $Q_t = Q_t(S)$. We say that a schedule has maximum block size B if for all $t \geq 1$ we have that $Q_t \leq B$.*

Definition 2.4. *We denote the set of transaction in \mathcal{C} with per-unit value at least θ by $\mathcal{C}(\theta)$. For a schedule $S = \{x_{i,t}\}_{i \in \mathcal{C}, t \in \mathbb{N}}$ we use $Q_t(S, \theta) = \sum_{i \in \mathcal{C}(\theta)} x_{i,t} \cdot q_i$ to denote the total capacity of transactions with per-unit value at least θ that are schedule at time t . We use $Q_{[1,T]}(S, \theta) = \sum_{t=1}^T Q_t(S, \theta)$ to denote the total capacity of transactions with per-unit value at least θ that are schedule at times $1, 2, \dots, T$.*

Note that for brevity we have omitted the implied underlying set of transactions \mathcal{C} from the notations of $Q_t(S, \theta)$ and $Q_t(S)$. We do the same in the following definition of social welfare.

Definition 2.5. *The social welfare of a schedule S of transactions in \mathcal{C} up to time T , is the sum of the values of the transactions that were scheduled up to T : $SW(S, [1, T]) = \sum_{i \in \mathcal{C}, t \in [1, T]} x_{i,t} \cdot q_i \cdot v_i$.*

The social welfare of a schedule may be easily expressed using the quantity parameters:

Lemma 2.6. *For any T and any schedule S it holds that*

$$SW(S, [1, T]) = \int_0^\infty Q_{[1, T]}(S, \theta) d\theta$$

Proof. By definition $Q_{[1, T]}(S, \theta) = \sum_{i \in \mathcal{C}(\theta), t \in [1, T]} x_{i,t} \cdot q_i = \sum_{i \in \mathcal{C}, t \in [1, T]} x_{i,t} \cdot q_i \cdot \mathbb{1}_{v_i \geq \theta}$ (notice that the summation over i in the first summation is only over transactions with per-unit value of at least θ , while in the second summation it is over all transactions). Integrating by θ we have $\int_0^\infty Q_{[1, T]}(S, \theta) d\theta = \sum_{i \in \mathcal{C}, t \in [1, T]} x_{i,t} \cdot q_i \cdot \int_0^\infty \mathbb{1}_{v_i \geq \theta}(\theta) d\theta = \sum_{i \in \mathcal{C}, t \in [1, T]} x_{i,t} \cdot q_i \cdot v_i$. Now notice that this is exactly the definition of $SW(S, [1, T])$. \square

Definition 2.7. A schedule is said to have average block size limit B with slackness $\Delta : \mathbb{N} \rightarrow \mathbb{N}$ if for any number of times steps k , the total size of transactions during any k consecutive blocks is at most $(k + \Delta(k)) \cdot B$. I.e. for any $t_0 \leq t_1$ we have that $\sum_{t=t_0}^{t_1} Q_t(S) \leq (t_1 - t_0 + 1 + \Delta(t_1 - t_0 + 1)) \cdot B$. When Δ is a constant function (for constant Z it holds that $\Delta(k) = Z$), we abuse notation and simply write $\Delta = Z$.

2.2 Online Scheduling Algorithms

In this paper we consider *online* scheduling algorithms. While (offline) scheduling algorithms map all information about the transactions to a schedule, online algorithms are constrained to assemble transactions into a sequence of blocks using only information about transactions arriving so far.

Definition 2.8. A scheduling algorithm is called *online* if its allocation at any time t is determined only as a function of the information known by time t (inclusive), i.e. on information from every transaction i with $t_i \leq t$.

We say that a scheduling algorithm has *maximum block size* B if on every input it produces a schedule with maximum block size B .

It is perhaps useful to present the EIP-1559 algorithm by first defining a sub-class of online algorithms, termed “price-based”²⁰, which the EIP-1559 algorithm is a member of.

Definition 2.9. A scheduling algorithm is called *price-based*, if at every time t it sets a per-unit price p_t and maximum size B_t based only on the parameters of the transactions that were executed at previous times (i.e. every transaction i such that $x_{it'} = 1$ for $t' < t$). The transactions that are scheduled at time t are chosen as follows: look at all yet unscheduled transactions i that have already arrived by time t ($t_i \leq t$) and are willing to pay p_t per unit ($v_i \geq p_t$). From these, an adversary chooses a maximal (by inclusion) subset of transactions of total size at most B_t .

Thus price-based scheduling algorithms are only allowed to use “on chain” information (i.e. from transactions that were scheduled before time t) and their allocation decisions must be simply by price per unit. This is in contrast to the more general class of *online* algorithms that may take decisions based on all information known by decision time, and use that information in an arbitrary way.

We next define a family of “EIP-1559 algorithms” - these are price-based algorithms that are parameterized by several parameter: 1) target block size $B > 0$, 2) maximum block size $c \cdot B$ for some constant $c > 1$, 3) a price-update parameter $\eta > 0$, 4) minimum price $p_{min} \geq 0$, and 5) an initial price p_1 for the first round, satisfying $p_1 \geq p_{min}$.

Definition 2.10. The EIP-1559 algorithm with parameters $(B, c, \eta, p_{min}, p_1)$ such that $B > 0, c > 1, \eta > 0, p_1 \geq p_{min} > 0$, is a price-based algorithm where $B_t = c \cdot B$ for every time t (so it has maximal block size $c \cdot B$), and where the per-unit price p_t is computed iteratively as follows: the price at time $t = 1$ is p_1 , and for any $t > 1$ it holds that $p_{t+1} = \max \{p_{min}, p_t \cdot e^{\eta \cdot (Q_t - B)/B}\}$, where Q_t is the total size of the transactions scheduled at time t .

Observe that the price goes up when $Q_t > B$, and it goes down when $Q_t < B$. The step size η controls the rate in which the price is updated.

The actual EIP-1559 algorithm with Ethereum’s choice of parameters is such an algorithm with target block size 15 Mega gas, maximum block size twice as large ($c = 2$), price-update parameter $\eta = 1/8$, $p_1 = 10^{-9}$ ETH and $p_{min} = 10^{-18}$ ETH.²¹ We also remark that Ethereum’s EIP-1559 algorithm approximates the price update multiplier $e^{\eta \cdot (Q_t - B)/B}$ by $1 + \eta \cdot (Q_t - B)/B$ (using the first term of the Taylor expansion, $e^x \approx 1 + x$). Finally, when there is more demand at the current block’s price than the maximum block size, then in Ethereum’s EIP-1559 algorithm, the transactions taken are chosen according to their offered “tip” rather than adverserially. From an algorithmic (non-strategic) point of view, as we take here, since the tips may be unrelated to the actual values of the transactions, they may indeed result in any adversarial schedule.

It is generally known that the EIP-1559 algorithm yields blocks whose *average* size approaches the target size B . We present the formal statement and proof in our setting and notation.

²⁰This class is conceptually similar to the class considered in [FMPS21] but we take an adversarial abstraction of the tip mechanism rather than their randomized choice. This whole class shares EIP-1559’s incentive properties in the impatient case.

²¹The minimum price is implicit since all price calculations are done in integer multiples of Wei ($= 10^{-18}$ ETH) and the first block’s price, when moving to EIP-1559 was set to one Gwei ($= 10^{-9}$ ETH).

Proposition 2.11. *The EIP-1559 algorithm with parameters $(B, c, \eta, p_{min}, p_1)$ has average block size limit B with slackness $\Delta = \frac{1}{\eta} \cdot \ln \frac{v_{max}}{p_{min}} + (c - 1)$.*

Proof. Fix any $k \in \mathbb{N}$, and consider a sequence of k consecutive blocks $[i, j]$ of length $k \geq 1$ (i.e. $j - i = k - 1$). If $\sum_{t=i}^j Q_t = 0$ the claim trivially holds. Otherwise, let $i' \geq i$ be the lowest index such that $Q_{i'} > 0$. Note that $i' \leq j$ and that $\sum_{t=i}^j Q_t = \sum_{t=i'}^j Q_t$.

Since the prices are updated multiplicatively, we have that

$$\frac{p_{j+1}}{p_{i'}} = \prod_{t=i'}^j \frac{p_{t+1}}{p_t} = \prod_{t=i'}^j \max \left\{ \frac{p_{min}}{p_t}, e^{\eta \cdot (Q_t - B)/B} \right\} \geq e^{\eta \cdot (\sum_{t=i'}^j Q_t - k \cdot B)/B}$$

It holds that $p_{i'} \geq p_{min}$. Additionally, $p_{j+1} \leq v_{max} \cdot e^{\eta \cdot (c-1)}$ by the following lemma and the fact that $Q_{i'} > 0$ for $i' \leq j < j + 1$:

Claim 2.12. *If for time t it holds that $Q_t > 0$ then for any $k > t$ it holds that $p_k \leq v_{max} \cdot e^{\eta \cdot (c-1)}$.*

Proof. We prove the claim by induction. Assume the claim was true up to time $k - 1 \geq t$, so $p_{k-1} \leq v_{max} \cdot e^{\eta \cdot (c-1)}$. If $p_k \leq p_{k-1}$ then $p_k \leq p_{k-1} \leq v_{max} \cdot e^{\eta \cdot (c-1)}$. Otherwise $p_k > p_{k-1}$, implying that demand at price p_{k-1} was positive, so $p_{k-1} \leq v_{max}$. In that case, the price can only rise above v_{max} by a factor of at most $e^{\eta \cdot (Q_{k-1} - B)/B} \leq e^{\eta \cdot (c \cdot B - B)/B} = e^{\eta \cdot (c-1)}$ since $Q_{k-1} \leq c \cdot B$. \square

We conclude that $\frac{p_{j+1}}{p_{i'}} \leq \frac{v_{max} \cdot e^{\eta \cdot (c-1)}}{p_{min}}$. Combining the two inequalities we get

$$\frac{v_{max} \cdot e^{\eta \cdot (c-1)}}{p_{min}} \geq \frac{p_{j+1}}{p_{i'}} \geq e^{\eta \cdot (\sum_{t=i'}^j Q_t - (j - i' + 1) \cdot B)/B}$$

taking natural logs

$$\ln \frac{v_{max}}{p_{min}} + \eta \cdot (c - 1) \geq \frac{\eta}{B} \cdot \left(\sum_{t=i'}^j Q_t - (j - i' + 1) \cdot B \right)$$

thus

$$B \left((j - i' + 1) + \frac{1}{\eta} \cdot \ln \frac{v_{max}}{p_{min}} + (c - 1) \right) \geq \sum_{t=i'}^j Q_t = \sum_{t=i}^j Q_t$$

As $i' \geq i$ we have $j - i' + 1 \leq j - i + 1 = k$ and thus $B \cdot k \geq B \cdot (j - i' + 1)$. Thus, the EIP-1559 algorithm with parameters $(B, c, \eta, p_{min}, p_1)$ has average block size limit B with slackness $\Delta = \frac{1}{\eta} \cdot \ln \frac{v_{max}}{p_{min}} + (c - 1)$. \square

3 EIP-1559 has High Welfare for Patient Bidders

In this section we prove our main result, showing that when bidders are patient, the EIP-1559 algorithm produces schedules of near-optimal welfare, provided it is given a mild resource augmentation (that does not increase with the time horizon). As a warm up we first analyze a Greedy Online Algorithm (Section 3.1), illustrating some of the techniques we use in proving our main result. We then present our main result (Section 3.2), its proof is based on a central lemma which we prove in Section 3.3.

3.1 Warm up: a Greedy Online Algorithm

Before diving into the main theorem, we can get some intuition by proving that a simple online greedy algorithm with slightly relaxed constraints, competes with the optimum. Specifically, we show that a simple online greedy algorithm with $c = 2$, slackness of 1 and extension of 1, has welfare at least as high as the optimum. This algorithm can be viewed as setting prices at each time, but these prices are allowed to depend on all input arriving so far. To some extent it is possible to view the EIP-1559 algorithm as attempting to mimic this online algorithm by gradually approximating its prices. This gradual approximation will take extra steps, require bounds on the per-unit valuations, and incur some losses which the simple greedy algorithm does not suffer from (as, unlike the EIP-1559 algorithm, it has full access to all parameters of the transactions that has already arrived).

Consider any time horizon T (that need not be known by the algorithm). For any set of arriving transactions, we compare the welfare of the algorithm to the welfare of the optimal schedule over T steps, where the optimal schedule is limited to have maximal block size B at each time step (and thus we can assume that there are only transactions of size at most B). At each time $t = 1, 2, \dots, T$ our algorithm will greedily schedule pending transactions with the highest per-unit value, until the first point in which total size of the transactions scheduled by time t is at least $t \cdot B$ (the case there are not enough pending transactions is an edge case that we may ignore, by pretending to have an unlimited amount of 0 value transactions). As transactions are of size at most B , we have that $t \cdot B \leq \sum_{i=1}^t q_i < (t+1) \cdot B$, where q_t is the total size of transactions scheduled for block t .

Let us denote by p_t the value (per unit) of the last (lowest value per unit) transaction scheduled at time t .

We first observe that for any sequence of Z blocks we have: $(Z-1) \cdot B \leq \sum_{i=t}^{t+Z-1} q_i < (Z+1) \cdot B$, and thus this algorithm has average block size B with slackness 1. This also implies that it has maximum block size $2 \cdot B$.

We next observe that when given an extension of one extra block, this greedy online algorithm always achieves social welfare at least as high as the optimum schedule of maximal block size B .

We denote the schedule produced by this algorithm by ALG and the optimal (welfare-maximizing up to time T) fractional schedule of maximal block size B by OPT .

Proposition 3.1. *For any T it holds that $SW(ALG, [1, T+1]) \geq SW(OPT, [1, T])$.*

Proof. In order to analyze its performance, we make the following claim. Fix an arbitrary positive number θ and denote by $Q_{[1,t]}(ALG, \theta)$ the total size of transactions whose value (per unit) is at least θ that were scheduled by the algorithm up to (and including) time t . Similarly, denote by $Q_{[1,t]}(OPT, \theta)$ the total size of transactions whose value (per unit) is at least θ that were scheduled by the optimal (even fractional) schedule that has maximum block size B up to (and including) time t .

Lemma 3.2. *For every t and θ we have that $Q_{[1,t+1]}(ALG, \theta) \geq Q_{[1,t]}(OPT, \theta)$.*

Proof. The proof of this lemma proceeds by splitting the algorithm's run into two phases. Let α_θ be the last (largest) time t such that $p_t < \theta$ (If this never happens then we take $\alpha_\theta = 0$). By the end of the first phase, i.e. at time α_θ , there are *no* pending transactions whose value per unit is at least θ (since these should have been scheduled before the transaction whose value per unit is p_{α_θ}). Thus, during the first phase, the total size of transactions with value at least θ that ALG schedules is at least as high as the total scheduled by OPT, that is, $Q_{[1,\alpha_\theta]}(ALG, \theta) \geq Q_{[1,\alpha_\theta]}(OPT, \theta)$. If $\alpha_\theta \geq T$, the claim follows. Otherwise, $\alpha_\theta < T$. Now we look at the second phase, i.e. during times steps $\alpha_\theta + 1, \dots, T, T+1$. For any time t in this time range it holds that $p_t \geq \theta$, and thus every transaction scheduled by the algorithm in that range has value per unit of at least θ . Since, as mentioned above, in every Z consecutive blocks the algorithm schedules transactions of size at least $(Z-1) \cdot B$, we have that during the second stage (including time $T+1$) our algorithm scheduled transactions of size at least $((T+1) - \alpha_\theta - 1) \cdot B$, while OPT can schedule transactions of size at most $(T - \alpha_\theta) \cdot B$ during the range $\alpha_\theta + 1, \dots, T-1, T$ (i.e without the extra step at time $(T+1)$). It follows that in the second phase, the total size of transactions with value at least θ that ALG schedules (with the extra time step) is at least as high as the total scheduled by OPT, that is, $Q_{[\alpha_\theta+1, T+1]}(ALG, \theta) \geq (T - \alpha_\theta) \cdot B \geq Q_{[\alpha_\theta+1, T]}(OPT, \theta)$. \square

The proposition can be directly deduced by from this lemma using Lemma 2.6 that states that the social welfare achieved by ALG for time 1 till $t+1$ can be expressed as $\int_0^\infty Q_{[1,t+1]}(ALG, \theta) d\theta$, and similarly, the social welfare achieved by OPT for time 1 till t can be expressed as $\int_0^\infty Q_{[1,t]}(OPT, \theta) d\theta$. \square

Remark: This analysis did not assume any bound on the maximum transaction size beyond being bounded by the target block size $q_{max} \leq B$ (since we are competing with schedules with maximum block size B). If we have a bound on the maximum transaction size then the maximum block size of this algorithm is in fact sharper: $c = 1 + q_{max}/B$.

3.2 Our Main Result

Our main result is that EIP-1559 algorithms have the attractive property that for patient bidders (with some bounds on the per-unit valuations), for any time horizon, they give close to optimal welfare with respect to the benchmark of all schedules that are restricted to maximum block size B , with slightly less steps.

This may be viewed as surprising since the algorithm is both restricted to online decisions, and furthermore, is restricted to be price-based, and yet competes with all offline clairvoyant schedules. On the other hand, it is not limited by a worse-case block size limit B , but rather block size limit B holds only on the average with some slackness Δ (as shown in Proposition 2.11). Additionally, the algorithm gets an extension of additional Γ time steps, for some appropriately chosen Γ (that does *not* grow with the time T).

We use ALG to denote the schedule produced by an EIP-1559 algorithm. Recall that for schedule S we use $SW(S, [1, t])$ the total social welfare of S up to time t .

Theorem 1. *Fix an EIP-1559 algorithm with parameters $(B, c, \eta, p_{min}, p_1)$ for some $c > 2$. Consider any (adversarially chosen) sequence \mathcal{C} of input transactions where each transaction $i \in \mathcal{C}$ has per-unit value v_i satisfying $e^\eta \cdot p_{min} \leq v_i \leq v_{max}$ for some $v_{max} \geq p_1$. Let OPT_B be an arbitrary schedule with maximum block size B . Then for any time T it holds that*

$$SW(ALG, [1, T + \Gamma]) \geq (1 - \delta) \cdot SW(OPT_B, [1, T])$$

for $\delta = 1 - e^{-\eta} \leq \eta$ and any integer $\Gamma \geq \frac{1}{\eta} \cdot \ln \left(\frac{v_{max}}{p_{min}} \right) \cdot \max\{1, \frac{1}{c-2}\} + c$.

Theorem 1 is a corollary of a more general and stronger result that we prove. First, instead of the benchmark being only integral allocations, we can take the benchmark to be all *fractional* allocations (in which a transaction with per-unit value v_i that is allocated a fraction $x_i \in [0, 1]$ contributes $v_i \cdot q_i \cdot x_i$ to the value of the schedule). Second, we present tighter results for the case that $q_{max} < B$, replacing the assumption that $c > 2$ with the assumption that $c > 1 + q_{max}/B$. Third, instead of the benchmark have worst-case block size limit B , we allow it to only have *average block size limit B with slackness Δ'* . To allow for this extra slackness we will allow the algorithm an additional Δ' time steps (effectively increasing Γ by Δ'). Finally, we will allow the initial price p_1 to be arbitrary. When it is larger than v_{max} it will imply that Γ grows as $\ln \frac{p_1}{p_{min}}$.

Theorem 2. *Fix an EIP-1559 algorithm with parameters $(B, c, \eta, p_{min}, p_1)$. Consider any (adversarially chosen) sequence \mathcal{C} of input transactions where each transaction $i \in \mathcal{C}$ has size at most q_{max} and has per-unit value v_i satisfying $e^\eta \cdot p_{min} \leq v_i \leq v_{max}$. Assume that $c > 1 + \frac{q_{max}}{B}$. Let $OPT_{B, \Delta'}$ be an arbitrary fractional schedule with average block size limit B and slackness Δ' . Then for any time T it holds that*

$$SW(ALG, [1, T + \Gamma]) \geq (1 - \delta) \cdot SW(OPT_{B, \Delta'}, [1, T])$$

for $\delta = 1 - e^{-\eta} \leq \eta$ and for any integer $\Gamma \geq \max\left\{\frac{1}{\eta} \ln \left(\frac{p_1}{p_{min}} \right), \frac{1}{\eta(c'-1)} \cdot \ln \left(\frac{v_{max}}{p_{min}} \right) + c - 1 + \frac{c-2}{c'-1}\right\} + \Delta'$, where $c' = c - \frac{q_{max}}{B}$.

We note that we do not assume any knowledge of v_{max} by the algorithm (it is only used for analysis).

To see that Theorem 1 follows from Theorem 2, observe that when $q_{max} = B$ we have $c > 1 + \frac{q_{max}}{B} = 2$. Additionally, the benchmark of integral schedules with maximal block size B is clearly weaker than the benchmark of fractional schedule with average block size limit B and slackness $\Delta' \geq 0$. Finally, as $c = c' + 1$, the assumption that $v_{max} \geq p_1$ implies that $\frac{1}{\eta} \ln \left(\frac{p_1}{p_{min}} \right) \leq \frac{1}{\eta} \ln \left(\frac{v_{max}}{p_{min}} \right)$ and thus Γ must be at least $c + \max\left\{1, \frac{1}{c-2}\right\} \cdot \frac{1}{\eta} \cdot \ln \left(\frac{v_{max}}{p_{min}} \right)$ as

$$c + \max\left\{1, \frac{1}{c-2}\right\} \cdot \frac{1}{\eta} \cdot \ln \left(\frac{v_{max}}{p_{min}} \right) = c - 1 + \frac{c-2}{(c-1)-1} + \max\left\{1, \frac{1}{c-2}\right\} \cdot \frac{1}{\eta} \cdot \ln \left(\frac{v_{max}}{p_{min}} \right) \geq \frac{1}{\eta} \ln \left(\frac{p_1}{p_{min}} \right)$$

We remark that if $p_{min} > v_{min} \cdot e^{-\eta}$ then the result still hold, but with respect to schedules that are only allowed to schedule transactions with value per unit at least $p_{min} \cdot e^\eta$. That is, there is an

additional additive loss in welfare that is bounded by the welfare of any schedule of transactions of value per unit less than $p_{min} \cdot e^\eta$.

The cornerstone of the proof of the main result is the next lemma which shows that for any θ , the EIP-1559 algorithm schedules large enough size of transactions of per-unit value θ , as long as it get a long enough extension. The theorem will be easily deduced from this bound using Lemma 2.6. For the next lemma, recall that $Q_t(S, \theta)$ denotes the total size of transactions with per-unit value of at least θ that were scheduled by S at time t .

Lemma 3.3. *Fix an EIP-1559 algorithm with parameters $(B, c, \eta, p_{min}, p_1)$. Consider any (adversarially chosen) sequence of input transactions, and assume that each transaction up to time T has a per-unit value at least v_{min} , and that $v_{min} \geq e^\eta \cdot p_{min}$. Assume also that $c > 1 + \frac{q_{max}}{B}$. Let E denote the schedule of the EIP-1559 algorithm on the input. For any fractional schedule S with average block size limit B and slackness Δ' , for any $\theta \geq 0$ it holds that*

$$Q_{[1, T]}(S, \theta) = \sum_{t=1}^T Q_t(S, \theta) \leq \sum_{t=1}^{T+\Gamma} Q_t(E, \theta \cdot e^{-\eta}) = Q_{[1, T+\Gamma]}(E, \theta \cdot e^{-\eta}) \quad (1)$$

for any integer $\Gamma \geq \max \left\{ \frac{1}{\eta} \ln \left(\frac{p_1}{p_{min}} \right), \frac{1}{\eta(c'-1)} \cdot \ln \left(\frac{v_{max}}{p_{min}} \right) + c - 1 + \frac{c-2}{c'-1} \right\} + \Delta'$, where $c' = c - \frac{q_{max}}{B}$.

The theorem easily follows the main lemma by integrating Eq. (1) over θ , and then applying Lemma 2.6. By integrating Eq. (1) we conclude that

$$\int_0^\infty Q_{[1, T]}(S, \theta) d\theta \leq \int_0^\infty Q_{[1, T+\Gamma]}(E, \theta \cdot e^{-\eta}) d\theta \quad (2)$$

By Lemma 2.6 the LHS of Eq. (2) is exactly $SW(S, [1, T])$. To evaluate the RHS of Eq. (2), we make a linear change of variable (with our notation reusing θ) $\int_0^\infty Q_{[1, T+\Gamma]}(E, \theta \cdot e^{-\eta}) d\theta = e^\eta \int_0^\infty Q_{[1, T+\Gamma]}(E, \theta) d\theta$, and we use Lemma 2.6 again to conclude that the RHS is exactly $e^\eta \cdot SW(E, [1, T + \Gamma])$, as needed.

3.3 Proof of The Main Lemma

Before giving a formal proof of Lemma 3.3, let us give the intuition drawing parallels with the online algorithm above. Similarly to the previous analysis we will look at the last point in time α_θ (to be simply called α below) where the price was less than θ and the quantity scheduled is at most $Q_t \leq c \cdot B - q_{max}$, meaning that every transaction with per-unit value at least the price was scheduled. As previously, up to that point, the EIP-1559 algorithm scheduled everything with per-unit value of at least θ , so it competes with any schedule S (and also with a fractional one). From this point on there are always enough pending transactions that are willing to pay at least θ per-unit (quantity larger than $c \cdot B - q_{max}$), so EIP-1559 prices will increase rapidly until they first reach θ (at a time to be called β below) from which point they can never decrease much below θ (price will always be at least $\theta \cdot e^{-\eta}$, the minimum price when price decreases from θ in a single step). So after this point in time, the EIP-1559 algorithm fills its quota with transactions of value per unit of at least $\theta \cdot e^{-\eta}$, and since its average block size limit is the same as that of S , it competes with it successfully if having enough scheduling time. To compensate for the time it took for the price to rise up to θ (between α and β), we allow the algorithm an extended time of extra Γ steps to schedule transactions. Additionally, we allow an extra Δ' steps to the algorithm, in order for it to compete also with schedules that have *average* block size limitations with slackness Δ' . We suffer a loss of a $1 - e^{-\eta}$ factor due to the possible single price decrease below θ .

We next present the proof of Lemma 3.3.

Proof. We prove the claim for $\theta \in [v_{min}, v_{max}]$ since the claim is trivial for $\theta > v_{max}$, and immediate for $\theta < v_{min}$ from the claim for $\theta = v_{min}$. Recall that $c' = c - \frac{q_{max}}{B}$. As $c' > 1$ it holds that $c \cdot B - q_{max} = c' \cdot B > B$. We denote $\theta' = \theta \cdot e^{-\eta}$ and observe that $\theta > \theta' > \theta \cdot (1 - \eta)$ and that $\theta' = e^{-\eta} \cdot \theta \geq e^{-\eta} \cdot v_{min} \geq p_{min}$ as $\theta \geq v_{min} \geq e^\eta \cdot p_{min}$. We say that *demand at price θ was met at time t* if every transaction i such that $v_i \geq \theta$ and $t_i \leq t$ has been scheduled by the algorithm at or before time t . Observe that if at time t it holds that $p_t \leq \theta$ and $Q_t \leq c \cdot B - q_{max} = c' \cdot B$, then every transaction i such that $v_i \geq \theta$ and $t_i \leq t$ must have been scheduled by the algorithm (as any such unscheduled transaction can fit the block t as $Q_t + q_i \leq c \cdot B - q_{max} + q_i \leq c \cdot B$). In such a case

we say that *demand at price θ was necessarily met at time t* (as all demand at price θ was scheduled for sure by time t).

Trivially, demand at any price θ was necessarily met at time $t = 0$ (as no demand arrived yet). Let $\alpha = \alpha_\theta \in [0, T + \Gamma]$ denote the latest time $t \in [0, T + \Gamma]$ such that demand at price θ was necessarily met at time t . We assume for convenience that $p_0 = p_{min} \leq \theta$ and $Q_0 = 0$, so for $t = \alpha$ (even when $\alpha = 0$) it holds that $p_\alpha \leq \theta$ and $Q_\alpha \leq c' \cdot B$.

To simplify notation we use $S_t(\theta) = Q_t(S, \theta)$ and $E_t(\theta) = Q_t(E, \theta)$. Let $Z_t(\theta)$ denote the total size of all transactions that arrived by time t with value at least θ , i.e. $i \in Z_t(\theta)$ if $v_i \geq \theta$ and $t_i \leq t$. It holds that

$$\sum_{t=1}^{\alpha} S_t(\theta) \leq Z_\alpha(\theta) = \sum_{t=1}^{\alpha} E_t(\theta) \leq \sum_{t=1}^{\alpha} E_t(\theta')$$

where the rightmost inequality is due to the fact that E_t is a non-increasing function of θ and since $\theta \geq \theta'$.

Thus, if $\alpha \geq T$ then our claim follows as

$$\sum_{t=1}^T S_t(\theta) \leq \sum_{t=1}^{\alpha} S_t(\theta) \leq \sum_{t=1}^{\alpha} E_t(\theta') \leq \sum_{t=1}^{T+\Gamma} E_t(\theta')$$

So we are left with the case that $0 \leq \alpha < T$. In that case we have

$$\sum_{t=1}^T S_t(\theta) = \sum_{t=1}^{\alpha} S_t(\theta) + \sum_{t=\alpha+1}^T S_t(\theta) \leq Z_\alpha(\theta) + \sum_{t=\alpha+1}^T S_t(\theta)$$

As demand at price θ was necessarily met at α , we know that $\sum_{t=1}^{\alpha} E_t(\theta') \geq \sum_{t=1}^{\alpha} E_t(\theta) = Z_\alpha(\theta)$ so

$$\sum_{t=1}^{T+\Gamma} E_t(\theta') = \sum_{t=1}^{\alpha} E_t(\theta') + \sum_{t=\alpha+1}^{T+\Gamma} E_t(\theta') \geq Z_\alpha(\theta) + \sum_{t=\alpha+1}^{T+\Gamma} E_t(\theta')$$

We conclude that to complete the proof we only need to show that $\sum_{t=\alpha+1}^T S_t(\theta) \leq \sum_{t=\alpha+1}^{T+\Gamma} E_t(\theta')$, knowing that $0 \leq \alpha < T$ and that demand at price θ was not necessarily met at any time $t > \alpha$.

Thus, in the rest of the proof we assume that demand at price θ was not necessarily met at any time $t \in [\alpha + 1, T + \Gamma]$, and move to lower bound the total capacity scheduled by the EIP-1559 algorithm between time $\alpha + 1$ and $T + \Gamma$. As demand at price θ was not necessarily met at any time $t \in [\alpha + 1, T + \Gamma]$, then for t in this range of times, whenever $p_t \leq \theta$ then $Q_t > c' \cdot B \geq B$. We conclude that if $p_t \leq \theta$ then $p_{t+1} = p_t \cdot e^{\eta \cdot (Q_t - B)/B} > p_t \cdot e^{\eta \cdot (c' - 1)}$. Thus, if $p_\alpha < \theta$ then price will start rising till at some time $t \geq \alpha + 1$ it will hold that $p_t \geq \theta \geq \theta'$.

Additionally, we claim that if for some $t > \alpha$ it holds that $p_t \geq \theta'$, then $p_{t+1} \geq \theta'$ (implying that for any $t' > t$ it also holds that $p_{t'} \geq \theta'$). Indeed, if $\theta > p_t \geq \theta'$ then $p_{t+1} = p_t \cdot e^{\eta \cdot (Q_t - B)/B} > p_t \cdot e^{\eta \cdot (c' - 1)} \geq p_t \geq \theta'$. On the other hand, when $\theta \leq p_t$ then $p_{t+1} = p_t \cdot e^{\eta \cdot (Q_t - B)/B} \geq p_t \cdot e^{-\eta} \geq \theta \cdot e^{-\eta} \geq \theta'$. We conclude that once $p_t \geq \theta'$ it holds $p_{t'} \geq \theta'$ for any $t' > t$.

Now, we use $\beta \in [\alpha + 1, T + \Gamma]$ to denote the first time the price was at least θ' (the existence of such a time is implied by the arguments below since Γ is large enough). So, if $\beta > 1$ then it holds that $p_\beta \geq \theta'$ and $p_{\beta-1} < \theta'$ (note that for any time $t \in [\beta, T + \Gamma]$ it holds that $p_t \geq \theta'$).

To complete the proof we only need to prove the following lemma:

Lemma 3.4. *It holds that*

$$\sum_{t=\alpha+1}^{T+\Gamma} E_t(\theta') \geq \sum_{t=\beta}^{T+\Gamma} E_t(\theta') = \sum_{t=\beta}^{T+\Gamma} Q_t \geq B \cdot (T - \alpha + \Delta') \geq \sum_{t=\alpha+1}^T S_t(\theta)$$

Proof. Observe that the leftmost inequality is trivial as $\beta \geq \alpha + 1$. Let us now consider the rightmost inequality. It holds due to the average block size limitation constraint of B with slackness Δ' , which state that $\sum_{t=\alpha+1}^T S_t(\theta) \leq (T - \alpha + \Delta') \cdot B$. The equality $\sum_{t=\beta}^{T+\Gamma} E_t(\theta') = \sum_{t=\beta}^{T+\Gamma} Q_t$ holds as for any time $t \in [\beta, T + \Gamma]$ it holds that $p_t \geq \theta'$, and thus all capacity is used to schedule transaction with value per unit at least θ' .

So to complete the proof we only need to show that $\sum_{t=\beta}^{T+\Gamma} Q_t \geq B \cdot (T - \alpha + \Delta')$. For any time $t \in [\beta, T + \Gamma]$ it holds that $p_t \geq \theta' = e^{-\eta} \cdot \theta \geq e^{-\eta} \cdot v_{min} \geq p_{min}$ (and thus also $p_{t+1} \geq p_{min}$) so $p_{t+1} = p_t \cdot e^{\eta \cdot \frac{Q_t - B}{B}}$.

Thus

$$\frac{p_{T+\Gamma+1}}{p_\beta} = \prod_{t=\beta}^{T+\Gamma} \frac{p_{t+1}}{p_t} = \prod_{t=\beta}^{T+\Gamma} e^{\eta \cdot \frac{Q_t - B}{B}} = e^{\eta \cdot (\sum_{t=\beta}^{T+\Gamma} \frac{Q_t - B}{B})}$$

So

$$\frac{B}{\eta} \ln \left(\frac{p_{T+\Gamma+1}}{p_\beta} \right) = \left(\sum_{t=\beta}^{T+\Gamma} Q_t \right) - B \cdot (T + \Gamma - \beta + 1)$$

or equivalently,

$$\sum_{t=\beta}^{T+\Gamma} Q_t = B \cdot \left(T + \frac{1}{\eta} \ln \left(\frac{p_{T+\Gamma+1}}{p_\beta} \right) + \Gamma - (\beta - 1) \right) \quad (3)$$

So proving that $\sum_{t=\beta}^{T+\Gamma} Q_t \geq B \cdot (T - \alpha + \Delta')$ is immediate from the inequality

$$\frac{1}{\eta} \ln \left(\frac{p_{T+\Gamma+1}}{p_\beta} \right) + \Gamma - (\beta - 1) \geq \Delta' - \alpha$$

which we prove in Lemma 3.5. □

Lemma 3.5. *It holds that*

$$\Gamma \geq \frac{1}{\eta} \ln \left(\frac{p_\beta}{p_{T+\Gamma+1}} \right) + \Delta' + \beta - (\alpha + 1)$$

Proof. Let us first consider the case that $\beta = 1$ and $\alpha = 0$. In this case $p_\beta = p_1$ and $p_{T+\Gamma+1} \geq p_{min}$. So

$$\Gamma \geq \Delta' + \frac{1}{\eta} \ln \left(\frac{p_1}{p_{min}} \right) \geq \Delta' + \frac{1}{\eta} \ln \left(\frac{p_\beta}{p_{T+\Gamma+1}} \right)$$

where the left inequality holds by the definition of Γ , and the claim follows. Thus, in the rest of the proof we assume that $\beta > 1$.

Recall that $\beta \in [\alpha + 1, T + \Gamma]$ is the first time such $p_\beta \geq \theta'$. Thus, as $\beta > 1$ it holds that $p_{\beta-1} < \theta'$ (note that we have shown that for any time $t \in [\beta, T + \Gamma]$ it holds that $p_t \geq \theta'$). Thus for $\beta > 1$ it holds that $p_{\beta-1} < \theta'$, and as $p_{\alpha+1} \geq p_{min}$ we get

$$\frac{\theta'}{p_{min}} \geq \frac{p_{\beta-1}}{p_{\alpha+1}} \geq e^{\eta \cdot (\sum_{t=\alpha+1}^{\beta-1} (Q_t - B)) / B} \geq e^{\eta \cdot (\sum_{t=\alpha+1}^{\beta-1} (c' - 1))} = e^{\eta \cdot (\beta - \alpha - 1)(c' - 1)}$$

Thus

$$\ln \left(\frac{\theta'}{p_{min}} \right) \geq \eta \cdot (\beta - \alpha - 1)(c' - 1)$$

and as $\eta > 0$ and $c' > 1$ we get

$$\frac{1}{\eta(c' - 1)} \cdot \ln \left(\frac{\theta'}{p_{min}} \right) \geq \beta - \alpha - 1 \quad (4)$$

As $\eta > 0$ and $c > c' > 1$ (so $c - 2 > -1$) it holds that $v_{max} \cdot e^{\eta \cdot (c-2)} > v_{max} \cdot e^{-\eta} \geq \theta \cdot e^{-\eta} \geq \theta'$, by Eq. (4)

$$\frac{1}{\eta(c' - 1)} \cdot \ln \left(\frac{v_{max} \cdot e^{\eta \cdot (c-2)}}{p_{min}} \right) \geq \frac{1}{\eta(c' - 1)} \cdot \ln \left(\frac{\theta'}{p_{min}} \right) \geq \beta - (\alpha + 1) \quad (5)$$

Since $\beta > 1$ is the first time the price is at least θ' , the price was rising from time $\beta - 1$, and $p_{\beta-1} < \theta' = \theta \cdot e^{-\eta}$. In this case p_β is at most $\theta' \cdot e^{\eta \cdot (c-1)}$. As once the price get to θ' it never drops below it, we have $p_{T+\Gamma+1} \geq \theta'$. We conclude that $\frac{p_\beta}{p_{T+\Gamma+1}} \leq \frac{\theta' \cdot e^{\eta \cdot (c-1)}}{\theta'} = e^{\eta \cdot (c-1)}$ and therefore

$$\frac{1}{\eta} \ln \left(\frac{p_\beta}{p_{T+\Gamma+1}} \right) \leq \frac{1}{\eta} \ln(e^{\eta \cdot (c-1)}) = c - 1$$

Combining this with Eq. (5) we get

$$\frac{1}{\eta} \ln \left(\frac{p_\beta}{p_{T+\Gamma+1}} \right) + \Delta' + \beta - (\alpha + 1) \leq c - 1 + \Delta' + \frac{1}{\eta(c' - 1)} \cdot \ln \left(\frac{v_{max} \cdot e^{\eta \cdot (c-2)}}{p_{min}} \right) \leq \Gamma$$

where the rightmost inequality follows from the assumption about Γ . □

□

4 Lower Bounds

In this section we show that certain limitations of our main theorem are inevitable. We do so by proving appropriate lower bounds. Our lower bounds apply more generally to classes of algorithms of which the EIP-1559 algorithm is a member, some to all online algorithms (see Section 4.2) and others just for the subclass of price-based algorithms (see Section 4.1). Where there are gaps between our upper bound for the EIP-1559 algorithm and our general lower bounds, we also show that the upper bound is in fact the tight one for the EIP-1559 algorithm.

Our lower bounds apply to algorithms that, like the EIP-1559 algorithm, are allowed both some slackness Δ and some extension Γ in order to compete with the optimum schedule. As these two notions of augmentation may be traded-off against each other to some extent²², we will generally prove lower bounds on their sum.

4.1 Augmentation of Price-Based Algorithms must depend on Value Range

In this section we prove that the dependence of Γ or Δ on the range of per-unit values is unavoidable for *price-based algorithms with good welfare guarantee*, proving Proposition 1.4.

Proposition 4.1. *Fix any price-based algorithm with average block size limit B , slackness Δ , and extension Γ . There exists an instance with values per unit in the range $[1, H]$ for which the algorithm does not obtain better than a $H^{1/4^{(\Gamma+\Delta)}}$ fraction of the optimal social welfare of a schedule with worst-case block size limit B (with no slackness or augmentation). Thus, if a price-based algorithm guarantees a constant fraction of this optimum, then $\Gamma + \Delta$ must grow as $\Omega(\log \log H)$.*

Proof. Given an algorithm with parameters Δ and Γ , and assume wlog that $B = 1$, we will choose some parameters T, r and design a set of input scenarios that for at least one of them the algorithm cannot achieve more than $(1 + (\Gamma + \Delta)/T)/r$ fraction of the social welfare. This, with our choices of parameters r and T , will give the desired upper bound.

We will look at our algorithm running for $R = T + \Gamma$ steps, where all transactions arrive at time 1. We will have $2^R + 1$ possible inputs scenarios, where in input scenario number m ($0 \leq m \leq 2^R$) we have, for every integer $0 \leq i \leq m$, exactly $(R + \Delta)$ transactions with value r^i (all arriving at time 1), each of size 1. We also specify how our adversary will respond to every price p in each input scenario m . If $p > r^m$ then there are no transactions that are willing to pay this price so no transactions will be scheduled. if $p \leq r^m$ then the adversary will always schedule only transactions with the lowest possible value (i.e. those with value r^i for the smallest i so that $r^i \geq p$). Note that there are enough transactions of this value to suffice for anything that can be scheduled by the end of time R .

Now let us look at a price-based algorithm running for R steps on one of these input scenarios. For any price p that is queried by the algorithm there are at most two possible answers by the adversary on any one of these scenarios: either no transaction is scheduled or as many transactions as allowed by the block's capacity, each of the value r^i for the i completely determined by the price, is scheduled. It follows that the complete set of possible behaviors of the algorithm on all possible scenarios can be described as a binary tree of depth R , containing at most $2^R - 1$ possible prices in all and splitting the possible transactions into the 2^R leaves of the tree. Since there are $2^R + 1$ possible scenarios, two of them must map into the same leaf and thus produce the exact same schedule. So let us assume that

²²All our lower bounds hold even for a single horizon T that is known to the algorithm in advance. In that setting, shifting between the two parameters is indeed possible. If there is a single horizon T that is known by the algorithm, the extension may be completely transferred to additional slackness by immediately scheduling at the last block all transactions that were intended for the extension period. Similarly, slackness can be shift to an extension when T is known, by postponing some scheduling to the extension time.

input scenarios $m < m'$ get the same behavior by the algorithm, which must thus never schedule any transactions with any value strictly greater than r^m as such do not exist in scenario m .

So now let us compare the social welfare of our algorithm on input scenario m' to that of the optimal schedule. The optimal schedule for T blocks with worst-case block size 1 always uses transactions with value $r^{m'}$ for a total welfare of $r^{m'} \cdot T$. Since our algorithm cannot use transactions with value greater than r^m and has slackness Δ and runs for Γ extra steps, its welfare is bounded by $r^m \cdot (T + \Gamma + \Delta)$. The fraction of welfare obtained in scenario m' is bounded from above by $(1 + (\Gamma + \Delta)/T)/r$.

Now for our choice of parameters. We let $T = \Gamma + \Delta$ which ensures that the nominator in the expression above is the constant 2, and thus, the fraction of social welfare (with respect to the benchmark of any schedule with worst-case block size limit 1, and no slackness or augmentation) that is achieved by the algorithm is bounded from above by $2/r$. Our choice of T also ensures that $R = T + \Gamma \leq 2T = 2(\Gamma + \Delta)$ and since for our choice of r we only needed that $r^{2R} \leq H$ we may choose $r = H^{1/2^{2T}} = H^{1/2^{2(\Gamma+\Delta)}} = H^{1/4^{(\Gamma+\Delta)}}$, which concludes the proof. \square

4.2 Online Algorithms need to Relax the Maximum Block Size

In this section we restate and prove Proposition 1.3, showing that every online algorithm with maximum block size $c \cdot B$ for $c < 2$, must lose a constant fraction of welfare.

Proposition 4.2. *Any online algorithm that produces schedules with maximum block size $c \cdot B$, for some $c < 2$ and average block size limit B with slackness $\Delta(T) = o(T)$, even with an extension $\Gamma(T) = o(T)$, must lose at least $\delta_0 = \min(1/8, (2-c)/3)$ fraction of social welfare relative to schedules with worst-case block size limit B (even when the values of all transactions are in $\{1, 2\}$).*

Remark: this proposition applies to the case where the largest transaction size can be as large as B . If we have an upper bound on q_{max}/B then the lower bound applies whenever $c < 1 + q_{max}/B$, matching the positive result in Section 3.1.

Proof. Here is the construction for the lower bound (assuming w.l.o.g. that $B = 1$): We take T to be large enough so that Γ and Δ that are $o(T)$ are as small as we want relative to T . In each of the first $T/2$ steps we get one “red” transaction with size 1 and value 1 and one “green” transaction with size $c/2 + \varepsilon$ and value 2. Notice that as $c > 1$ at most one of these transactions can fit into each block. Let us look at G which is the number of green transactions that were scheduled during the first $T/2$ steps.

Case I: $G \leq T/4$. In this case the total value achieved in the first $T/2$ steps is at most $2 \cdot G + 1 \cdot (1/2 - G) \leq 3/4$. In this case consider the case where for the next $T/2$ steps, in each step a new transaction of size 1 and value 2 enters. In this case the best that the mechanism can do is get value 2 for each of the remaining $T/2$ steps and the extension period $\Gamma = o(T)$ for a total social welfare of at most $7T/4 + o(T)$. In comparison the optimal schedule with block size 1 would schedule all the green transactions in the first half, getting a value of 2 every step for a total of $2T$, so the ratio of welfares in this case is bounded from above by $7/8 + o(1)$.

Case II: $G > T/4$. In this case consider the case where for the next $T/2$ steps we get a large number of small transactions of size γ with value γ (i.e. value per unit of 1) each. The total size of all blocks produced by our algorithm during all T steps and the extension Γ is at most $T + \Delta + \Gamma = T + o(T)$. The total welfare achieved by the algorithm can thus be bounded by $T + o(T)$ (obtained from red transactions and from small transactions) plus $T/2$ (the extra 1 value from every green transaction) minus $(1 - c/2 - \varepsilon) \cdot G$ which is the capacity completely lost by scheduling green transactions in the first half. Thus the total welfare is at most $T \cdot (3/2 + c/2 - 1 + \varepsilon + o(1))$. The optimal schedule with maximum block size 1 would keep all the green transactions for the second half, getting a social welfare of exactly $3T/2$. The ratio between these two welfare expressions is $(1 + c + o(1))/3$, so the fraction lost is at least $1 - (1 + c + o(1))/3 = (2 - c - o(1))/3$. \square

4.3 Lower Bounds Examples for the EIP-1559 Algorithm

Our lower bounds for general classes of algorithms presented above do not match the upper bounds of our main theorem for the EIP-1559 algorithm (Theorem 2). In this section we observe, by providing examples, that the upper bounds of our main theorem are in fact tight for the EIP-1559 algorithm.

4.3.1 $c = 2$ Does not Suffice for the EIP-1559 Algorithm

In Proposition 4.2 we showed that the maximum block size must be chosen to be at least $2 \cdot B$ in any online algorithm that attempts to compete with offline schedules whose worst case block size is B (when $q_{max} = B$), i.e., in our notation, the lower bound showed that we need $c \geq 2$. Proposition 3.1 showed that $c = 2$ indeed suffices for the online algorithm that we presented. However, our main theorem for $q_{max} = B$ requires the maximum block size to be strictly greater than twice the average block size limit, $c > 2$, and in fact the result also uses an extension that grows in rate that is inversely proportional to $c - 2$. So the question is whether having the maximum block size $2 \cdot B$ in the EIP-1559 algorithm, i.e. $c = 2$, suffices for near-optimally. This question is especially interesting since the actual choice of parameters in Ethereum is exactly $c = 2$, as the target block size is $B = 15M$ gas and the maximum block size is defined to be $2B = 30M$ gas.

Here is an example showing that the EIP-1559 algorithm with $c = 2$, even with an extension that grows as $\Gamma(T) = o(T)$, cannot give even half of the optimum social welfare.²³ For notional convenience let us assume without loss of generality that $p_{min} = 1$ (otherwise just scale all the values in the rest of the proof by p_{min}). Our example starts with steps where no transactions arrive, and so blocks are empty and the price decreases to $p_{min} = 1$. From this point on, every time step t , we get as input one “high” transaction of size exactly B and value per unit 10, and additional “low” transactions of total size $(1 + \varepsilon) \cdot B$ and value per unit 2, where $\varepsilon > 0$ is a parameter, to be specified below, that may be as small as we wish. (The size composition of the low transactions does not matter as long as their total size is $(1 + \varepsilon) \cdot B$.) In our example the low transactions offer a high tip, i.e. the adversary always schedules them before the high transaction.

Now let us see how the EIP-1559 algorithm deals with this input. As long as $p_t \leq 2$, both the high and the low transactions can be scheduled, and since the low transactions have priority they will be scheduled first, filling the block to size $(1 + \varepsilon) \cdot B$. At this point the high transaction cannot get scheduled, since scheduling it will just exceed the maximum allowed block size of $2 \cdot B$, so the block closes at size $(1 + \varepsilon) \cdot B$. Thus the next price will be $p_{t+1} = p^t \cdot e^{\varepsilon\eta}$. Notice that at this rate of price increase, it will take $\Omega(1/(\varepsilon\eta))$ blocks for the price to reach the value 2. As long as this is the case, the EIP-1559 algorithm’s welfare per block is less than a quarter of the optimal schedule’s welfare per block (since the optimal schedule can at the very least accept a high transaction of value 10 every block while the algorithm accepts transactions of value $2 \cdot (1 + \varepsilon)$ which for small enough ε is less than $2.5 = 10/4$). Now, if we take T to be half of the time it took for the price to rise to 2 (so, $T = \Omega(1/(\varepsilon\eta))$) we must have the extension also be at least $\Gamma \geq T$ in order to reach even half of the optimal social welfare (since the algorithm that accrues welfare at a rate of a quarter of the optimal, must run at least for twice as many blocks in order to reach half the social welfare). In order to show that no extension $\Gamma(T) = o(T)$ suffices to get even half of the social welfare on this example, just choose $\varepsilon = O(1/(\eta T))$ for T so that $\Gamma(T) \leq T$.²⁴

Thus, this lower bound shows that Ethereum’s choice of parameter $c = 2$ is problematic when there are transactions of size $q_{max} = B$ that contributes significantly to the welfare. As the general version of our main theorem showed, this problem disappears once transaction sizes are limited by some $q_{max} < B$.

4.3.2 EIP-1559 suffers a single logarithmic loss in the range size

While, as we showed in Proposition 3.1, online algorithms may compete with the optimum using constant slackness and extension ($\Delta = 1$ and $\Gamma = 1$), Proposition 4.1 proved a lower bound of $\Delta + \Gamma = \Omega(\log \log H)$ for price-based algorithms when the transaction per-unit values are in the range $[1, H]$. Our upper bound for the EIP-1559 algorithm in the main theorem is exponentially worse, with $\Gamma = O(\log H)$. Additionally, by Proposition 2.11 we know that the EIP-1559 algorithm has slackness $\Delta = O(\log H)$.

We show that these upper bounds are the correct ones for the EIP-1559 algorithm. Our example starts with steps where no transactions arrive, and so blocks are empty and the price decreases to

²³This example also shows that even when $c > 2$, as c approaches 2, the extension length must grow inversely proportional to $c - 2$. Simple tweaks in the example will show that even longer extensions cannot guarantee even a lower fraction of the welfare.

²⁴When $c > 2$ but is close to 2, one may take this example with $\varepsilon \approx c - 2$ and get a lower bound of $\Omega\left(\frac{1}{\eta(c-2)}\right)$ on the extension.

$p_{min} \leq 1$. At this point of time, a large amount of “high” transactions with value per unit of $v = H$ arrive. Additionally, at that time also arrive a large amount of “low” transactions with value of L per unit, where $H > L > 1$ are parameters. The exact size composition of low and high transactions does not matter, as long as there is enough demand in each of these values to always completely fill blocks to the maximum block size. For our example, the low transactions will offer high “tips”, i.e. the adversary will give them priority.

So how will the EIP-1559 algorithm react to this input? After the price decreases to $p_{min} \leq 1$, only low transactions are scheduled, filling blocks to the maximum block size, so the prices will increase as $p_{t+1} = p_t \cdot e^{(c-1)\eta}$, where c is the maximum block size multiplier. So after k steps we have $p_{t+k} = p_t \cdot e^{k \cdot (c-1)\eta}$. As long as $L \geq p_t$ the adversary will only schedule low transactions, while the optimum schedule will only have high transactions, getting a ratio of H/L in the social welfare. This can go on for at most $k = \Theta(\log L / ((c-1)\eta))$ steps (after which p_t that started at $p_{min} \leq 1$ can exceed L), which by choosing L to be an arbitrary constant factor lower than H (so $\log H = \Theta(\log L)$) gives us that extension $\Gamma \geq k = \Omega(\log H / ((c-1)\eta))$ is needed to before we get better than the constant factor approximation H/L to the welfare. Also, looking at these k steps of price increase, note that all blocks are completely full, i.e. their total size is $k \cdot c \cdot B$, which implies that the slackness is $\Delta \geq k \cdot (c-1) = \Omega(\log H / \eta)$.

5 Model Extensions

In this section we consider two model extensions: to partially patient bidders and to multiple resources. In both cases we present impossibilities, showing that losing a constant fraction of the welfare is inevitable even for general online algorithms (not only price-based algorithms or the EIP-1559 algorithm).

5.1 Partially Patient Bidders: Online Algorithms must Lose a Constant Factor

All our results so far assumed “patient bidders”, i.e. where a transaction’s value for its user remains the same over time. The opposite assumption of “impatient bidders” assumes that a transaction *must* be scheduled in the immediate block or it loses all value for its user. A more general model would capture some sensitivity of the value of a transaction on the time of its execution, where the value decreases with time. There are two simple models for such dependence on time:

Definition 5.1. *A transaction with value v and discount factor ρ (where $0 \leq \rho < 1$) has value for its user of $v \cdot (1 - \rho)^{t_e - t_a}$ if where t_a is its arrival time and t_e its execution time.*

Definition 5.2. *A transaction with value v and patience p has value of v for its user if it have been scheduled within p blocks of its arrival time and value 0 otherwise. I.e. if t_a is its arrival time and t_e its execution time then the user’s value is v if $t_a \leq t_e \leq t_a + p$ (and 0 otherwise).*

Thus the fully patient model corresponds to patience level $p \rightarrow \infty$ or discount factor $\rho = 0$, while the impatient model corresponds to $p = 0$ or $\rho \rightarrow 1$. Clearly the EIP-1559 algorithm does not allow its input to specify any patience level or discount factor. It is not hard to observe that if we just run EIP-1559 on partially patient bidders without taking the patience or discount factor into account then since it does not distinguish between “new” transactions and old ones that already lost most or all of their value, it cannot produce efficient results. But perhaps if we just let the algorithm at every block take into account the *current* value of the transaction, then we regain efficiency? I.e. for the discount factor model, at block t , a pending transaction i that has already arrived is considered for allocation if and only if $v_i \cdot (1 - \rho_i)^{t - t_i} \geq p_t$. Unfortunately, we get a negative answer and, in fact, lower bounds for any online algorithm. Significantly, these lower bounds hold for arbitrarily high patience (or low discount factor).

Proposition 5.3. *Fix an online algorithm with average block size limit B , slackness Δ , and extension Γ , where $\Delta(T) + \Gamma(T) = o(T)$. For every minimum discount rate $\rho_{min} > 0$ there exist a time horizon T , and an input sequence where all bidders have discount rates $\rho_i \leq \rho_{min}$ and values in $\{1, 2\}$ for which the algorithm loses at least a fraction $\delta_0 = 1/20$ of welfare relative to the optimal schedule with worst-case block size limit B .*

Proof. Fix an online algorithm with average block size limit 1, slackness $\Delta(T)$, and extension $\Gamma(T)$. Choose T that is larger than $\Gamma(T) + \Delta(T)$ by a large factor K , and also such that $(1 - \rho_{min})^{T/3} \leq 1/2$. Let $p = T/3$ and denote $\rho = \rho_{min}$. Now consider the following input scenario: At block 1, We have p “patient” transactions arriving, each with size $q_i = 1$, value $v_i = 2$, and discount rate $\rho_i = 0$. Also at each of the first p blocks arrives one “hasty” transactions, each with size $q_i = 1$, value $v_i = 1$, and discount rate $\rho_i = \rho$. So now our algorithm has to decide what to schedule during the first p steps: the high-value patient transactions or the low-value hasty transactions.

Case I: by the end of time p at least $p/2$ hasty transactions have been scheduled. In this case, for block $p+1$ our adversary will choose as input $2p$ more transactions each with size $q_i = 1$, value $v_i = 2$, and discount rate $\rho_i = 0$. Now by the end of time $T = 3p$ at most $T + \Delta(T) + \Gamma(T) = T \cdot (1 + O(1/K))$ transactions have been scheduled, at least $p/2$ of them of value 1 and the rest of value 2. So the total social welfare achieved is at most $2 \cdot (5/2 + O(1/K))p + 1 \cdot p/2 = (11/2 + O(1/K))p$. On the other hand, the optimal offline benchmark would only schedule the high-value patient transactions for a social welfare of $2 \cdot 3p = 6p$, so the ratio of welfares is $11/12 + O(1/K)$ which, by choosing K large enough, is less than $19/20$.

Case II: on the end of time p at most $p/2$ hasty transactions have been scheduled. In this case, in each of the next p steps we get a single transaction with size $q_i = 1$, value $v_i = 1$, and discount rate $\rho_i = \rho$ (and nothing more arrives in the last p steps). We can assume wlog that our online algorithm was “clever enough” to schedule all new transactions during the “middle” p steps as otherwise switching a scheduled hasty transaction or patient transaction with the new transaction that arrived at that time step will only increase the social welfare (since both the hasty transaction and the new transaction lose value at the same rate ρ but the new transaction has not lost value yet when it arrived while the hasty transactions have already lost some of their initial value. Patient transactions do not lose value at all.) It follows that at most Δ hasty transactions were scheduled during the “middle” p blocks and so the social welfare that our algorithm gets can be bounded from above by $(19/4 + o(1)) \cdot p$: $2p$ from all the new transactions (that are scheduled in the middle p steps) plus $2p$ from all the patient transactions plus at most $3p/4 + \Delta$ from the hasty transactions since at least half of them (minus Δ) are scheduled in the last p block after they have lost at least half their value (since each of them waited at least for the middle p steps and $(1 - \rho)^p \leq 1/2$). The optimal offline solution is to schedule the hasty transactions as they come in the first p steps, schedule the new transactions, as they come, in the middle p steps, and leave the patient transactions for the last p steps, getting the full value of $5p$. The ratio between these welfares is at most $19/20 + o(1)$. \square

Notice that the proof uses transactions with different discount rates. It is interesting to ask what can be done when all transactions share the same, global, discount rate. We conjecture that, in fact, the modified EIP-1559 (that takes into account current values) does produce nearly-optimal schedules with a loss that approaches the loss of the fully patient case, as the minimum discount rate ρ_{min} approaches 0.

For the other model, of bidders with given patience levels, we can show that having a global patience level does not suffice and exhibit an impossibility result for online algorithms where all transactions share a fixed, arbitrarily high, patience level $p_i = p$.

Proposition 5.4. *Fix an online algorithm with average block size limit B , slackness Δ , and extension Γ . For every global patience level p and every time horizon $T \geq 2p$, there exists an input sequence where all bidders have patience $p_i = p$ and values in $\{1, 2\}$ where the algorithm loses at least a fraction $\delta_0 = 1/10$ of welfare relative to the optimal schedule with worst-case block size limit B .*

Comment: If the patience of all bidders is just slightly larger than the time horizon, $p_i \geq T$, then we are back to the fully patient case and near-optimality is regained.

Proof. Let $B = 1$ and fix an online algorithm with average block size limit B and slackness Δ . Let $p = (\Gamma + \Delta)/\varepsilon$ for some small constant ε to be specified below. We will prove the lemma for $T = 2p$ and the same follows for any larger T by adding an arbitrarily long time prefix with no input transactions. At time 1 arrive p “green” transactions with value 1 and size 1. Additionally, every time step for the next $p - 1$ steps arrives a single “red” transaction with value 2 and size 1. Let us look at how many red transactions were scheduled by our algorithm by the end of time p .

Case I: at least $p/2$ red transactions were scheduled. Since the total number of scheduled transactions by this time is at most $p + \Delta$, this means that at least $p/2 - \Delta$ green transactions were not

scheduled and can never be scheduled any more since their deadline has passed. Now assume that no more transactions arrive in the next p steps, so our algorithm can schedule the rest of the red transactions but, again, not the unscheduled green ones. This gives an upper bound of $2p + p/2 + \Delta$ on the social welfare of our online algorithm. The Optimal clairvoyant algorithm could have scheduled only the green transactions during the first p steps and then the red transactions, one after another, after that for a total social welfare of $2(p - 1) + p$. The ratio between these amounts is $5/6 + O(\varepsilon)$.

Case II: less than $p/2$ red transactions were scheduled by the end of time p . Now assume that another p “new” transactions arrive at time $p+1$ each with value 2 and size 1. Since the total quantity scheduled in the first p steps is at most $p + \Delta$ and since at most $p/2 - 1$ of this quantity has a value per unit of 2 and the rest have value 1, then during the first p steps our online algorithm could have made at most $p + \Delta + p/2 - 1$ social welfare. In the second half, including the extra allowed Γ steps, an upper bound to the social welfare that it can make is $2(p + \Delta + \Gamma)$. Thus the total welfare of our online algorithm is bounded from above by $7p/2 + O(\varepsilon)$ welfare. The Optimal clairvoyant algorithm could have scheduled only the red transactions during the first p steps and then scheduled the new transactions in the next p steps for a total welfare of $4p$. The ratio between these amounts is $7/8 + O(\varepsilon)$. Choosing ε small enough that this amount is less than $1/10$ concludes the proof of the theorem. \square

5.2 Multidimensional Case: Online Algorithms Must Lose a Constant Factor

In our basic model each transaction has a “single-dimensional” size q_i . More generally one may consider a model where there are m different resources, each block has a size limit B_j for every resource j and each transaction i uses the amount q_{ij} of every resource j (see e.g. [DECA23, ADM24]). The natural generalization of the slackness condition to the multi-dimensional case would require that every T consecutive blocks use at most $(T + \Delta) \cdot B_j$ amount of each resource j . The introduction of the “blob” resource to Ethereum [BFL⁺22] is a step in this direction. While one may hope to extend the near-optimality result to such a multi-dimensional model, it turns out any online algorithm must lose some constant fraction of social welfare in the multi-dimensional case.

Theorem 5.5. *If there are at least three resources then any online algorithm with average block size B_j and slackness $\Delta(T) = o(T)$ must lose at least a constant fraction $\delta_0 = 1/6 - o(1)$ of social welfare on some input sequence, even when allowed $\Gamma(T) = o(T)$ extension.*

Proof. Consider the case of three resources, X, Y, and Z, and assume that the block size for each resource is $B_X = B_Y = B_Z = 1$. Let $t = T/2$ where T is the horizon over which our online algorithm aims to approximate OPT and fix at an online algorithm that uses each resource at total capacity of at most $(t + \Delta)$ during any consecutive t blocks.

At time 1 we get t demand for the bundle $\{X, Z\}$ and t demand for the bundle $\{Y, Z\}$ each at value of 1 per unit of the bundle.

During the first t steps there is only capacity of $t + \Delta$ of resource Z so at least one of the two types of demanded bundles is allocated at most at total capacity $(t + \Delta)/2$. Without loss of generality assume that the bundle $\{X, Z\}$ is allocated at most $(t + \Delta)/2$ during the first t steps. Now at time $t+1$ we get another demand of t for the singleton bundle $\{X\}$ again at value 1 per unit of X . (Of course, had the bundle $\{Y, Z\}$ been the one that was at least half-un-allocated then our adversary would have brought demand for the singleton bundle $\{Y\}$ instead.) The best that our online algorithm can do at this point is pack singleton bundles X together with any leftover of the original bundles $\{Y, Z\}$ into the remaining blocks, and leaving at least $t/2 - \Delta - \Gamma$ demand for the bundle $\{X, Z\}$ unsatisfied for a total value of at most $(5t + \Delta + \Gamma)/2$. OPT, on the other hand, could have fully allocated everything by allocating only the $\{X, Z\}$ bundles for the first t blocks and packing the demand for $\{Y, Z\}$ with $\{X\}$ in each of the last t blocks getting a total value of $3tB$. The ratio between these two levels of social welfare is $5/6 + o(1)$. \square

We do not know whether the case of two resources has a similar impossibility result, we leave this problem, and the problem of getting some constant-factor approximation with any number of resources, for future work.

Acknowledgments

We thank Shiri Ron for useful discussions during the early stages of this work. Noam Nisan thanks his colleagues in Starkware for many useful discussions. Moshe Babaioff’s research is supported in part by a Golda Meir Fellowship and the Israel Science Foundation (grant No. 301/24). Noam Nisan’s research was supported by a grant from the Israeli Science Foundation (ISF number 505/23).

References

- [ADM24] Guillermo Angeris, Theo Diamandis, and Ciamac Moallemi. Multidimensional Blockchain Fees are (Essentially) Optimal. *arXiv preprint arXiv:2402.08661*, 2024.
- [AGHH23] Sarah Azouvi, Guy Goren, Lioba Heimbach, and Alexander Hicks. Base Fee Manipulation In Ethereum’s EIP-1559 Transaction Fee Mechanism. *arXiv preprint arXiv:2304.11478*, 2023.
- [BCD⁺19] Vitalik Buterin, Eric Conner, Rick Dudley, Matthew Slipper, Ian Norden, and Abdelhamid Bakhta. EIP-1559: Fee market change for ETH 1.0 chain. <https://eips.ethereum.org/EIPS/eip-1559>, 2019. Ethereum Improvement Proposals [Online serial].
- [BFL⁺22] Vitalik Buterin, Dankrad Feist, Diederik Loerakker, George Kadianakis, Matt Garnett, Mofi Taiwo, and Ansgar Dietrichs. EIP-4844: Shard Blob Transactions. <https://eips.ethereum.org/EIPS/eip-4844>, 2022. Ethereum Improvement Proposals [Online serial].
- [But14] Vitalik Buterin. Ethereum: A Next-Generation Smart Contract and Decentralized Application Platform. <https://ethereum.org/en/whitepaper>, 2014.
- [CS23] Hao Chung and Elaine Shi. Foundations of transaction fee mechanism design. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3856–3899. SIAM, 2023.
- [DECA23] Theo Diamandis, Alex Evans, Tarun Chitra, and Guillermo Angeris. Designing Multidimensional Blockchain Fee Markets. In *5th Conference on Advances in Financial Technologies (AFT 2023)*, volume 282, pages 4:1–4:23, 2023.
- [Eps22] Leah Epstein. Open-end bin packing: new and old analysis approaches. *Discrete Applied Mathematics*, 321:220–239, 2022.
- [FMPS21] Matheus VX Ferreira, Daniel J Moroz, David C Parkes, and Mitchell Stern. Dynamic posted-price mechanisms for the blockchain transaction-fee market. In *Proceedings of the 3rd ACM conference on Advances in Financial Technologies*, pages 86–99, 2021.
- [GTW24] Aadityan Ganesh, Clayton Thomas, and S Matthew Weinberg. Revisiting the Primitives of Transaction Fee Mechanism Design. In *ACM EC*, 2024.
- [GY24] Yotam Gafni and Aviv Yaish. Competitive revenue extraction from time-discounted transactions in the semi-myopic regime. *arXiv preprint arXiv:2402.08549*, 2024.
- [HLM21] Gur Huberman, Jacob D Leshno, and Ciamac Moallemi. Monopoly without a monopolist: An economic analysis of the bitcoin payment system. *The Review of Economic Studies*, 88(6):3011–3040, 2021.
- [LMR⁺21] Stefanos Leonardos, Barnabé Monnot, Daniël Reijsbergen, Efstratios Skoulakis, and Georgios Piliouras. Dynamical analysis of the EIP-1559 Ethereum fee market. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, pages 114–126, 2021.
- [LRMP23] Stefanos Leonardos, Daniël Reijsbergen, Barnabé Monnot, and Georgios Piliouras. Optimality despite chaos in fee markets. In *International Conference on Financial Cryptography and Data Security*, pages 346–362. Springer, 2023.
- [Nak09] Satoshi Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System. <http://www.bitcoin.org/bitcoin.pdf>, 2009.

- [Nis23] Noam Nisan. Serial monopoly on blockchains. *arXiv preprint arXiv:2311.12731*, 2023.
- [PS24] Paolo Penna and Manvir Schneider. Serial monopoly on blockchains with quasi-patient users. *arXiv preprint arXiv:2405.17334*, 2024.
- [Rou21] Tim Roughgarden. Transaction Fee Mechanism Design. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, EC '21, page 792, 2021.