

Timing Games in Responsive Consensus Protocols

Kaya Alpturer*
Princeton University

Kushal Babel†
Category Labs

Aditya Saraf‡
Cornell University

Abstract

Optimistic responsiveness—the ability of a consensus protocol to operate at the speed of the network—is widely used in consensus protocol design to optimize latency and throughput. However, blockchain applications incentivize validators to play *timing games* by strategically delaying their proposals, since increased block time correlates with greater rewards. Consequently, it may appear that responsiveness (even under optimistic conditions) is impossible in blockchain protocols. In this work, we develop a model of timing games in responsive consensus protocols and find a prisoner’s dilemma structure, where cooperation (proposing promptly) is in the validators’ best interest, but individual incentives encourage validators to delay proposals selfishly. To attain desirable equilibria, we introduce dynamic block rewards that decrease with round time to explicitly incentivize faster proposals. Delays are measured through a *voting* mechanism, where other validators vote on the current leader’s round time. By carefully setting the protocol parameters, the voting mechanism allows validators to coordinate and reach the cooperative equilibrium, benefiting all through a higher rate-of-reward. Thus, instead of responsiveness being an unattainable property due to timing games, we show that responsiveness itself can promote faster block proposals. One consequence of moving from a static to dynamic block reward is that validator utilities become more sensitive to latency, worsening the gap between the best- and worst-connected validators. Our analysis shows, however, that this effect is minor in both theoretical latency models and simulations based on real-world networks.

1 Introduction

A major focus in consensus protocol design for Proof-of-Stake blockchains has been designing protocols that are *optimistically responsive* [YMR⁺19]. In the happy path, these protocols proceed at the speed of the network, whereas non-responsive protocols always wait for the worst case network delay. Responsiveness helps the protocol adapt to network conditions, by making rounds move faster than in non-responsive protocols like Ethereum, which have a fixed round time. Adaptive round times benefit from lower latency and higher throughput. However, responsive protocol design is complicated by the fact that the blockchain setting has novel incentive structures that highly depend on the timing of the validators’ actions. One of the major sources of reward for validators is maximal extractable value (MEV) [BDKJ23, DGK⁺20], the additional value the validators are able to extract by building their blocks strategically. Examples of MEV include the profit that validators make from arbitrage between exchanges (centralized and decentralized), frontrunning or backrunning user transactions, transaction fees, and so on. As MEV tends to increase with proposal time [OKV⁺23], a selfish leader will delay their block proposal as long as possible, even risking a

*E-mail: kalpturer@princeton.edu. Work done while interning at Category Labs.

†E-mail: babel@cs.cornell.edu.

‡E-mail: as2777@cornell.edu. Work done while interning at Category Labs.

timeout. This phenomenon has been called a *timing game* [SSST+23], and has been observed on blockchains such as Ethereum [SSST+23] and recently on Solana [art25, Tea25].

Timing games are undesirable as they increase latency to users, and increase the likelihood of missed blocks, reducing throughput even in non-responsive protocols. Timing games pose an even greater risk in blockchains relying on responsive protocols, where slower round times would invalidate the responsiveness offered by the underlying consensus protocol. A natural question then is the following: **Is optimistic responsiveness possible in blockchain protocols when there are strong incentives to delay blocks due to MEV?**

In this paper, we propose explicitly rewarding timely block production to shift incentives against timing games. If faster block proposals are assigned greater block reward, the reward can counter-balance MEV.

Implementing block rewards that depend on proposal duration is not trivial. One challenge is that there is no reliable measure of time available to the protocol. For instance, relying on block timestamps alone is not possible as they can be strategically manipulated [YSZ23]. To measure proposal durations accurately, we introduce a *timeliness vote*. All validators report their observed durations, and the reported durations are aggregated to form a proxy to the real duration. However, strategic validators may misreport their observed time, either unilaterally or as part of a coalition.

Our results suggest that responsiveness can induce honest voting, which disincentivizes timing games. This is because timing games often take the form of a prisoner’s dilemma among validators. Validators can *cooperate* by proposing early, which increases the speed of the system, and hence the rate of block reward for everyone. Alternatively, validators can defect by proposing late, greedily gaining some MEV while slowing down the system. Since each validator is a small part of the network, proposing late is a dominant strategy in the absence of coordination (i.e., the standard prisoner’s dilemma result). Honest voting, along with time-decreasing block reward, allows validators to coordinate to reach the early-proposing equilibrium rather than the late-proposing equilibrium. While it is clear that MEV allows validators to profit from users in the single-block horizon, our work highlights the surprising fact that MEV can hurt validators in the long-run by slowing down the rate of block reward.

Finally, we consider a notion of *fairness* that models disparities in rewards to validators due to heterogeneous latencies. Being better connected to other validators is helpful under both dynamic and static block rewards. Unfortunately, this fact encourages geographical collocation, which compromises decentralization. We show that, while fairness does degrade with time-decreasing block rewards, the gap is small, in a variety of theoretical and empirical settings.

1.1 Overview of results

We develop a **model of timing games in optimistically responsive BFT protocols**. Diverging from prior work, we explicitly incorporate the randomness of peer-to-peer latencies. Modeling randomness ensures that validators cannot distinguish intentional delays from genuinely higher latency.

Early-proposing equilibria. We show that if the block reward is set appropriately and individual entities have bounded stake, early-proposing and honest-reporting of timeliness is a Nash equilibrium (Theorem 4.1). We then argue that this equilibrium is robust in several ways. We first consider “whales”, i.e., single validators who control many units of stake, and show that they won’t deviate as long as they are sufficiently small (e.g., controlling less than $1/3$ of the total stake under standard parameter choices). We also argue that rational coalitions are unstable. Assuming a low-latency condition, we prove that a certain natural family of “large” coalitions will revert back to the desirable

equilibrium (Theorem 5.4). We also argue that smaller coalitions are unstable because members would be incentivized to join many small coalitions, to the detriment of others.

Fairness. We analyze fairness with respect to latencies. We first theoretically study two models of latency: one where nodes are arranged on a line and one where they are clustered into two groups, and show that fairness gets marginally worse under dynamic block rewards compared to static block rewards (Theorems C.1 to C.3). We then show similar results in simulations of real-world latency data. Our results suggest that there is some fairness cost to timing game mitigation.

Alternative approaches. An alternative incentive knob to decreasing block reward is to decrease leader election probabilities as a function of their round duration. However, it’s more difficult to incentivize honest voting in that setting, as a decrease in others’ election probability directly increases one’s own.

1.2 Related work

Timing incentives. Timing games and incentives for delaying one’s block in order to propose a more rewarding block have been both empirically observed [art25, SSST⁺23, Tea25] and theoretically analyzed [SSST⁺23]. Besides modeling responsive protocols, our model diverges from the approach of [SSST⁺23] by modeling peer-to-peer latencies. The increase in MEV with time has been empirically measured by [OKV⁺23, WZQ⁺23]. Solutions to timing games have been explored as well [SS22, SXNC25]. The blog post [SS22] proposes rewarding timeliness using a proxy measurement of same-slot committee votes in Ethereum. [SXNC25] proposes measuring transaction arrival rates and correcting timing incentives by redistributing execution layer fees. Our setting, however, requires analyzing incentives explicitly, as the effect of deviations directly affects the protocol’s operation.

Time-rotating leaders. A recent proposal by [KGMS25] involves modifying consensus so that leaders are elected for a certain period of time, rather than a round (or number of rounds). Here, even with a constant block reward, leaders may prefer to propose as fast as possible, as that allows them to propose more often. This is a promising approach for ensuring timely inclusion of *non-MEV* transactions, such as direct transfers.

However, for the transactions that generate MEV, such as decentralized finance (DeFi) transactions, the leader would likely simply hold onto those transactions and only submit them in their final block. Doing so would allow them to better optimize MEV in the last period. For a specific example, the arbitrage profit between a centralized exchange and a decentralized exchange over a certain time range T grows proportional to $T^{3/2}$ [MMR24]. Thus, the leader may prefer to wait until their final block to include DeFi transactions.

Responsive Consensus. There is a vast literature on fault-tolerant consensus protocol design that utilizes responsiveness but does not consider timing incentives. Responsiveness was first introduced in [ADLS94], and optimistic responsiveness with the widely used fast-vs-slow path framework was introduced in [PS18]. Numerous BFT consensus protocols were designed focusing on blockchain settings with responsiveness in mind [GKKS⁺22, GSPD⁺23, JB25, YMR⁺19]. The tendency of leaders to slow down has been acknowledged in [KGMS25], and punishing misses by changing the leader election schedule has been explored by [CGK⁺22, TSKK24]. However, as far as we are aware, these works do not consider timing games from a game-theoretic perspective and the proposed solutions are not analyzed in a model where the validators are rational.

Blockchain protocol incentives. In addition to timing games, there is a growing body of work investigating incentive incompatibilities in blockchain protocols—including selfish mining in Bitcoin [ES14, SSZ16], Proof-of-Stake attacks [AW24, BCNPW19], and timestamp manipulation [YSZ23, YTZ22]. Moreover, MEV has been observed to interfere with performance and scaling [FM25]. Lastly, our results show that when block rewards are too low, timing games are inevitable. A similar observation was made for Bitcoin’s block reward in relation to forking attacks in [CKWN16].

1.3 Roadmap

In Section 2, we give a brief summary of the optimistically responsive consensus protocols that we consider and the solution concepts we use. In Section 3, we formalize timing games in responsive protocols with time-dependent block reward. In Section 4, we present our equilibria analysis. Section 6 analyzes the fairness for both early and late-proposing scenarios in several models of network latencies. Lastly we give guidance on parameter settings in Section 7 and conclude with a discussion of future work in Section 8.

Proofs are deferred to Sections A to C. Section D has simulations for the fairness models. Section E discusses utility modeling. Section F explores an alternative approach of modifying leader election probabilities as the incentive knob instead of block reward. For a glossary of notation, see Section G.

2 Preliminaries

Optimistically responsive Consensus. We adopt the rotating leader framework of optimistically responsive consensus protocols (such as [YMR⁺19, JB25]) where validators vote on a proposed block by the leader, and the next leader aggregates a certain threshold number of votes to form a “quorum certificate”. In the happy path where the quorum is able to be formed, the protocol makes progress as soon as messages arrive, rather than waiting for the worst case time bound. If the nodes do not hear from the leader within a certain duration, a *timeout* is triggered to proceed with the next leader. The responsive paradigm is in contrast to fixed-duration rounds. For instance, Ethereum blocks are 12 seconds regardless of how fast the proposal and vote messages are propagated. The same is true for Tendermint [Buc16] as well. As we analyze the game from a rational model, we will mostly focus on the common case of happy-path operation of the protocol. For leader election, we assume a leader is randomly sampled each round, as is the case for Proof-of-Stake blockchain protocols.

Subgame perfect Bayes-Nash equilibria. We are modeling the blockchain validators as strategic participants in an infinitely repeated game. Therefore, the solution concept we use is subgame perfect Bayes-Nash equilibria. A strategy profile $\vec{\sigma}$, which specifies the voting and proposal strategy of each validator, is a subgame perfect Bayes-Nash equilibrium if, in any possible subgame, the validator is best-responding with respect to its expected utility by utilizing $\vec{\sigma}$.

Coalition resistance. We will also consider coalitions best-responding to other validators in a strategy profile. While we discuss the types of coalition we consider in Section 3, the key condition we require is coalition-resistance: no coalition (satisfying some condition such as size) can deviate from honest behaviour (early proposing and honest voting) and strictly get more utility in expectation.

3 Model

Our model abstracts away many complexities of running consensus among a distributed set of validators in order to capture a wide range of protocol implementations. One key complexity that we do retain is the existence of pairwise latencies among validators and we are concerned with the time messages take from one validator to another. Whether the delivery is through a gossip network or direct communication is abstracted away. We first explain the model basics, then present the full game, and finally discuss agent utilities in more detail.

Setup. In our model, we consider the game induced by introducing a block reward that depends on timeliness of proposals. There are n players $\{1, \dots, n\} = [n]$ participating in the protocol.¹ A leader is elected uniformly randomly each round and, since the protocol is responsive, may propose a block as soon as they have received the previous block and a quorum of c validators' votes (the quorum threshold c is $\lfloor \frac{2n}{3} \rfloor + 1$ for most protocols). A strategic leader may choose to intentionally delay its block by some delay Δ . Validators measure the *round duration* locally as the time elapsed on their clock between receiving the previous leader's proposal and the current leader's block proposal. τ represents the timeout duration after which honest validators time out on the current leader's proposal. In other words, if the round duration exceeds τ for an honest validator, it will timeout. Validators move to the next round after a quorum (also c -sized for simplicity) has either timed out, or confirmed the current leader's proposal. Finally, $L_{i \rightarrow j}$ represents the latency distribution of messages from i to j .²

Timeliness measurement. We propose a novel timeliness measurement that does not exist in current consensus implementations. Upon receiving a proposal, each validator votes whether to approve the block, and additionally includes the observed proposal duration (their vote is implicitly infinite if not reported). The measurement time is then defined as the m^{th} smallest vote, i.e., the m^{th} order statistic of the votes. We assume votes are aggregated honestly.³

Optimal-delaying oracle, bounded latencies, and mostly honest timeouts. We assume that the validators are able to delay maximally without missing their block. We model this with the existence of an oracle that sees all the latency realizations and informs the leader of the largest delay that ensures their block doesn't timeout. This assumption does not map to practice, but it only makes our results stronger, as we will show that validators prefer to propose early despite this "timeout protection". Note that this oracle depends on the voting strategy of validators, as some validators may timeout dishonestly.

As this maximum delay varies in each round due to the latency realizations, let Δ_r^* be the maximum delay in round r . We drop the subscript when it is clear from context. To ensure that all validators can participate, we require $\Delta^* \geq 0$. In other words, if a leader does not intentionally delay, their proposal will not timeout. This is equivalent to assuming (1) a bounded latency condition, (2) a sufficiently long timeout, and (3) a limited number of validators will dishonestly time out. Condition (3) essentially requires that c validators will not *timeout* dishonestly, though they may

¹Each agent is a uniform unit of stake. We will consider the effect of heterogeneous stake distributions and collusion via coalitions.

²While we do not differentiate between block propagation latency and voting latency, they are often different in real life protocols, and it is straightforward to introduce separate distributions for both in our results.

³One possible implementation would be to require that timing votes are directly submitted on-chain by validators, and aggregated over a sufficiently long window spanning a number of leaders. This allows one to sidestep any strategic aggregation issues that would arise if the next leader is responsible for aggregation.

vote dishonestly otherwise. We know that $\Delta_r^* \leq \tau$, since latencies are nonnegative. Thus, τ should be thought of as the maximum possible delay over all latency realizations. While there may exist an incentive for large coalitions to censor other validators (i.e., 2/3 of the network can censor out the other 1/3), or benefit their coalition members by extending the timeout, these issues already exist without timing games, and our work does not aim to address them. Instead, we focus on incentives regarding the timeliness vote.

3.1 Formal game

We now present the formal model of the game. We present a general round, and then note the special differences that apply to the first round. For most of our variables, we denote the variable's value in round R by a subscript.




Figure 1: Timeline: t_R represents the time at which the leader of round R broadcasts their block. The leader of round $R - 1$ is j and the leader of round R is i .

timing game

1. Let R be the current round, and suppose i is the current leader and j led the previous round. Let t_{R-1} represent the time when j broadcasts its block. Let s_{R-1} denote the duration between when j broadcasts their block and when i receives the block and its quorum.
2. The leader i learns their maximum delay Δ_R^* and chooses a delay $\Delta \in [0, \Delta_R^*]$.
3. Let $k \neq i$ be an arbitrary validator.
 - (a) The latency $l_{i \rightarrow k} \sim L_{i \rightarrow k}$ is sampled.
 - (b) k receives the block at time $t_R^k = t_{R-1} + s_{R-1} + \Delta + l_{i \rightarrow k}$.
 - (c) k votes according to their voting strategy, which is a function of the block times they have seen. So $v_R^k = \sigma_k^V(\mathbf{t}_{1:R}^k) \in \mathbb{R}^+$. k then publishes its vote where σ_k^V is the voting strategy of validator k .
4. Let v^* refer to the aggregated timeliness vote (m^{th} smallest vote in the (multi)set $\{v_R^k\}_k$). i receives utility $M(\Delta + s_{R-1}) + B(v^*)$, where M denotes the MEV and B the block reward.
5. The next leader is randomly sampled, and the game repeats from 1.

The first round of this game is slightly different. The game starts at time $t = 0$, so $s_0 = 0$, i.e., the leader i can propose as soon as the game starts. As a result, validator k receives the block at time $t_1^k = \Delta + l_{i \rightarrow k}$. Otherwise, the first round matches the others.

Assumptions. We highlight some additional assumptions for our model.

- (a) *Clock synchronization.* We do not deal with clock drift in this model. We assume that all validators agree on the start time of the first round and that their clocks tick at the same rate.
- (b) *Hidden latency realizations.* While we assume the latency distributions are public, the realized latencies are unknown. So, validators don't know exactly when the previous leader proposed, as they can't distinguish the low-latency with delayed-proposal case from the high-latency with early-proposal case. They also don't know exactly when other validators receive a block.
- (c) *Efficient block building.* Everyone proposes every round and extracts all potential MEV (given by function $M(\cdot)$) at the time of proposal.⁴

3.2 Utilities and coalitions

Rewards. There are several sources of reward that drive the utilities of validators in the system. The main source of timing games is the MEV, which increases for blocks that are built later in time due to an increase in arbitrage opportunities, more time to grind through ordering combinations for the given transactions, and a larger mempool (available user transactions). The protocol, on the other hand, assigns reward for proposing a block to the proposer and voter rewards for block validation. We model MEV as an increasing function of the time between block proposals. We ignore voting rewards as they are constant.

- *MEV reward:* $M(x) : \mathbb{R} \rightarrow \mathbb{R}$ is a positive increasing function. In the fairness analysis, we will assume that it is of the form $\mu x + \mu_0$ for convenience.
- *Block reward:* B is chosen by the protocol designer and can be a function of information available to the protocol such as the timeliness vote. We design a linearly decreasing block reward that reaches 0 when the timeliness vote is equal to τ . Thus, $B(x) = b_0 - bx$ for $b \geq 0$, where $b_0 = b\tau$.




Figure 2: Reward functions.

Long-run time-averaged utility. Each player i has utility equal to their reward per unit of time, in the long-run. For each round r , let Δ_r denote the delay chosen by the leader of that round leader_r , and let the random variable S_{r-1} denote the duration between the leader of round $r-1$ proposing their block to the time when the current leader can start proposing, which is a function of the latency distribution. The duration $\Delta_r + S_{r-1}$ determines the MEV reward that the leader

⁴There may be incentives for validators to coordinate across the reigns of different leaders and exclude MEV-generating transactions for several rounds before finally claiming and sharing the MEV. Note that this strategy already exists even with static block rewards and is not empirically observed, suggesting practical coordination barriers.

receives, as well as the duration that round lasts. Let v_r^* denote the timeliness vote in round r . Then, player i 's utility is:⁵

$$\mathbb{E} \left[\lim_{R \rightarrow \infty} \frac{\sum_{r=0}^R \mathbb{I}[i = \text{leader}_r] (M(\Delta_r + S_{r-1}) + B(v_r^*))}{\sum_{r=1}^R \Delta_r + S_{r-1}} \right]$$

Coalitions. We consider two types of coalitions: (1) a single entity (“whale”) controlling multiple units of stake; (2) separate agents cooperating in a group.

- *Whale:* A whale $C \subseteq [n]$ is controlled by a single entity, which is only concerned with the joint utility of all nodes in C . So, a strategy profile σ is α -whale proof if, for any $C \subseteq [n]$ such that $|C|/n \leq \alpha$, there exists no deviation that improves the joint utility of C . Note that for all $i, j \in C$, $L_{i,j} = 0$. In other words, whales’ stake is assumed to be co-located.
- *Rational coalition:* Let $C \subseteq [n]$ denote a rational coalition among separate agents. Here, we require that any suggested deviation is itself a Bayes-Nash equilibrium, and we consider each agent’s utility. In other words, a strategy profile σ is α -rational coalition proof if, for any deviation σ' , one of the following two conditions holds.
 - a Some $i \in C$ has higher utility when everyone plays σ as opposed to the coalition playing σ' while the rest play σ .
 - b Some $i \in C$ prefers to deviate from σ' , given that the rest of the coalition plays σ' and the non-coalition validators play σ .

4 Early-proposing truthful equilibria

Let σ^* be the truthful strategy profile where each leader proposes without delay ($\Delta = 0$) and voters honestly report their observed round duration. In this section, we investigate when σ^* is an equilibrium. Full results are deferred to [Section A](#).

4.1 Honest equilibria

We first define $S^{j,i}$, the random variable that represents the time from when current leader i can start proposing, given that j led the previous round. For some validator k , the time it takes for them to receive j 's block and send a vote to i is distributed according to $L_{j \rightarrow k} + L_{k \rightarrow i}$. For a quorum size of c , let $Q_c^{j \rightarrow i}$ denote the c^{th} order statistic on $\{L_{j \rightarrow k} + L_{k \rightarrow i}\}_k$. This is the exact time when i has enough votes to form a quorum. Since i also needs the block to build on it, $S^{j,i} = \max(Q_c^{j \rightarrow i}, L_{j \rightarrow i})$, where the max is necessary because it's possible for i to receive the quorum before they have seen j 's block.

If i delays by $\Delta \geq 0$, they extract $M(\Delta + S^{j,i})$ due to MEV. For the duration votes, note that validator k 's honest report, given that i delays by Δ , is distributed according to $V^k = \Delta + S^{j,i} + L_{i \rightarrow k} - L_{j \rightarrow k}$. Aggregating at least m of these votes, the block reward of leader i for a delay of Δ is distributed according to $B(\Delta + S^{j,i} + E_m^{i-j})$ where E_m^{i-j} represents the m^{th} order statistic of

⁵An astute reader may notice that we are not modeling inflation, which is relevant because players, through their voting, have some control over how much block reward is printed. However, inflation does not qualitatively change any of our results. We defer a longer discussion to [Section E](#).

$\{L_{i \rightarrow k} - L_{j \rightarrow k}\}_k$. Intuitively, E_m^{i-j} represents the error in measuring the leader's duration by the timeliness vote due to latency in the observations.

The honest proposal strategy is to set $\Delta = 0$, so leader i would receive total reward $M(S^{j,i}) + B(S^{j,i} + E_m^{i-j})$. Lastly, observe that $\Delta^* = \tau - S^{j,i} - E_c^{i-j}$ for leader i preceded by j .⁶ Figure 3 illustrates the timeline of voting honestly.




Figure 3: Round timeline for honest voting. i is the current leader, j is the previous leader, and k is some validator observing i 's duration. $l_{i \rightarrow k}^r$ represents the realized latency in round r between i and k .

Note that some votes may be negative; due to latency, validator k may receive the block from i before the block from j . If the m^{th} smallest vote was negative, then the utility under σ^* would not be well-defined, as B is defined on the interval $[0, \tau]$. We first identify when this is not the case.

Lemma 4.1. *If $m > n - c$, then the m^{th} smallest vote is always at least the round delay Δ in the early-proposing truthful strategy profile σ^* .*

We assume $m > n - c$ from now on. This ensures that the aggregated timeliness vote will always come from a validator who contributed to the quorum. We show that, under a certain design of the block reward function, σ^* (the early-proposing truthful strategy profile) is a subgame-perfect Bayes-Nash equilibrium. The key condition is the following.

Definition 4.1 (time-decreasing). *We call the reward functions $\langle M, B \rangle$ time-decreasing if $M(x) + B(x)$ is a decreasing quantity in x .*

The condition guarantees that the per-round reward is maximized by proposing at each validator's earliest possible time. Now, we show that σ^* is an equilibrium. Intuitively, if rewards are time-decreasing, delaying only hurts i 's utility. For voting, note that i can only vote on other leader's proposals, and thus i 's voting strategy will not affect their own utility.

Theorem 4.1. *Suppose $\langle M, B \rangle$ are time-decreasing and $m > n - c$. Then, the early-proposing truthful strategy profile σ^* is a subgame-perfect Bayes-Nash equilibrium, where the time-averaged utility of i is:*

$$U_i(\sigma^*) = \frac{\sum_j \left(\mathbb{E}[M(S^{j,i})] + \mathbb{E}[B(S^{j,i} + E_m^{i-j})] \right)}{\sum_{j,k} \mathbb{E}[S^{j,k}]}$$

5 Coalition resistance

In this section, we argue that σ^* , the honest-voting and early-proposing equilibrium, will be resistant even to coalition deviations. Our results depend on the size of the coalition, which we characterize below with $k = m + c - n$:

⁶Note that we now use the c^{th} order statistic as the quorum threshold is c as Δ^* is determined by c validators not timing out.

Coalitions	Small	Medium	Large
Coalition size	$< k$	$[k, m)$	$\geq m$
Example ($c = m = 2n/3$)	$< n/3$	$[n/3, 2n/3)$	$\geq 2n/3$

Table 1: Coalition size categories, where $k = m + c - n$.

Suppose that coalitions vote 0 for the duration of coalition leaders’ proposals.⁷ Consider a coalition leader that proposes with delay Δ in some round. For this proposal, the m^{th} smallest vote comes from an honest, quorum-participating voter (and thus the vote is at least Δ). For a large coalition, the aggregate always comes from a coalition member. Medium coalitions are less predictable. The aggregate will always come from an honest voter, but they may not have participated in the quorum (in fact, their vote could be arbitrarily small). We have three arguments for coalition resistance.

- (A) Small coalitions will never compromise responsiveness:

Theorem 5.1. *For small whales and rational coalitions, the optimal strategy is to always vote 0 on each other’s proposals, and propose at $\Delta = 0$.*

- (B) We argue *informally* that small and medium rational coalitions are unstable. Suppose multiple, overlapping coalitions could form, as is true in practice. If validator i acquires m “free” votes, then they would propose late, as they are guaranteed the maximum block reward regardless. Further, other validators may want to enter into a 0-voting pact with i , as doing so will increase their utility by decreasing their own latency penalty. They would only reject this pact if they believed that i would have m free votes, and thus delay. Estimating this probability would be challenging in practice, rendering small and medium coalition formation highly unpredictable. Taken together, we believe that small or medium rational coalitions will not form in practice, as coalition members will have an incentive to join additional coalitions.

Note that this argument assumes that rational coalitions can neither (i) detect and penalize coalition-defectors, nor (ii) identify surreptitious participation in other coalitions. In practice, small coalitions could detect if one of their members proposes much later than they should (due to their secret membership in other coalitions). However, that member can instead make a smaller delay which could be difficult to distinguish from latency fluctuations. This delay would still lead to them benefiting at the cost of the rest of the network, and thus the same qualitative issues arise. Coalition capabilities and formation dynamics are complex [KGP⁺25, RV15, Ray07, She79]; we defer the meta-game of coalition formation to future work.

- (C) Suppose that small rational coalitions do grow into large coalitions. Since the optimal voting strategy for small coalitions is to vote 0 on coalition leaders, we consider large coalitions that follow this same strategy. We also introduce a condition called z -low-latency, which (roughly) requires that at most k validators receive aggregated timeliness votes greater than $z\tau$ in the honest equilibrium (where $z \in [0, 1]$). We believe that this is a natural condition, as it would be satisfied if the timeout is set high enough to account for large latency fluctuations. Using this, we show the following theorem.

Theorem 5.2. *If the system has z -low-latency with $z \leq \frac{n-c}{3n-c}$, then $|C| - k$ validators within the coalition would prefer defecting to the early-proposing truthful strategy.*

⁷We prove that this is optimal for small coalitions, and assume it for large coalitions due to its intuitive appeal and because large coalitions may grow from small coalitions.

Thus, large, low latency, 0-voting coalitions will not form, because a large fraction of the coalition would prefer to return to the honest equilibrium.

We defer full results for arguments (A) and (C) to [Section B](#).

Our coalition resistance results are summarized in [Table 2](#). While we do not have arguments against medium or large whales, such a whale must control a very large fraction of the total stake, and is thus unlikely to exist or violate the assumptions of the underlying BFT protocol itself.

	Whale	Rational
Small	(A)	(A), (B)
Medium	—	(B)
Large	—	(C)

Table 2: Cells show which arguments apply to which sizes and types of coalitions.

Here, we flesh out two of the arguments against coalition resistance that we sketched. To review, small coalitions have less than $k = m + c - n$ members and large coalitions have more than m members. We show the following.

1. For small whales and rational coalitions, the optimal strategy is to always vote 0 on each other’s proposals, and propose with $\Delta = 0$.
2. For large coalitions that vote 0 on each other’s proposals, if the average latencies are low relative to the timeout, then large coalitions would not form.

5.1 Small coalitions

Notice that the reported duration may be less than Δ if the coalition is too large. For $m = c = 2n/3$, suppose that $n/3$ validators got the block so late that they receive the next block before the current one (since $2n/3$ is needed for the quorum, the chain can advance without these validators). Then, if $n/3$ reports 0, the reported duration is going to be 0 as well, even though the leader took non-zero duration to propose the block. Formally, we have the following result.

Lemma 5.1. *Let $k = m + c - n$. If at least $n - k$ validators vote honestly, then for any round with delay Δ , the m^{th} smallest vote will be at least Δ . If fewer than $n - k$ validators vote honestly, then the m^{th} smallest vote may be less than Δ .*

This lemma is important because in small coalitions, delaying is always punished by an honest voter, while in larger coalitions that may not be the case. Now, we turn to whales, and show that small whales do not benefit from deviating, as they are already co-located.

Proposition 5.1. *Consider a whale of size at most k . The optimal strategy for the whale is the early-proposing honest strategy, σ^* .*

We now consider rational coalitions, and show that while small rational coalitions can benefit from deviating (in their voting strategy), they still propose early.

Theorem 5.3. *Consider a rational coalition of size at most k . The optimal strategy for the coalition members is to always vote 0 on each other’s proposals (despite any latency they may observe), and propose at $\Delta = 0$.*

5.2 0-voting large coalitions

What if everyone in the coalition had m free votes? In other words, suppose the coalition is large ($|C| \geq m$), and they vote 0 for coalition leaders.⁸ Is this large coalition profitable?

To rule this out, we identify when the timing game behaves similarly to a prisoner’s dilemma. Consider two different strategy profiles: one in which everyone proposes early and one in which everyone proposes late but votes 0. Our approach, which uses voting to ensure that the validators can all commit to proposing early, is robust to large coalitions only if the validators would prefer the early-proposal equilibrium to the late-proposal equilibrium. In terms of the prisoners dilemma, validators preferring the late-proposal equilibrium is equivalent to agents preferring the “all defect” equilibrium.

Theorem 5.4 characterizes when enough validators prefer the early-proposing equilibrium to make any large, late-proposing coalition infeasible. This is because enough validators would want to leave that the large coalition would become small (and thus their optimal strategy would be early-proposing). This matches a prisoner’s dilemma; each agent would individually prefer to propose late (defect) but, many agents would rather everyone propose early (cooperate) than late.

To introduce the theorem, we first define a “low-latency” setting. To do so, let $\mathbb{E}[\bar{S}^i] = \frac{1}{n} \sum_j \mathbb{E}[S^{j,i}]$ denote the expected starting time for validator i and let $\mathbb{E}[\bar{E}_x^i] = \frac{1}{n} \sum_j \mathbb{E}[E_x^{i-j}]$ denote the expected x^{th} smallest latency differential for validator i .

Definition 5.1 (z -low latency). *For $z \in [0, 1]$, a set of validators and their latency distributions satisfies z -low-latency if there exist at most k nodes i such that $\mathbb{E}[\bar{S}^i + \bar{E}_{\min(c,m)}^i] > z\tau$.*

Note that $\mathbb{E}[\bar{S}^i + \bar{E}_{\min(c,m)}^i]$ is the average reported round duration when validator i leads. Thus, the condition requires that $n - k$ validators are well-connected, in that their average reported round duration is a z fraction of the timeout.

Theorem 5.4. *Suppose a large coalition C votes 0 for leaders from the coalition and votes honestly for other leaders, while the rest of the validators use truthful strategy σ^* . If the system has z -low-latency with $z \leq \frac{n-c}{3n-c}$, then $|C| - k$ validators within the coalition would prefer defecting to the early-proposing truthful strategy.*

Under the conditions in the theorem, we see that 0-voting large coalitions that propose late become unsustainable. $|C| - k$ can “defect” from the coalition, and if they do so, then the coalition becomes small ($|C| \leq k$), and thus will not compromise responsiveness (as well as facing practical barriers to adoption due to a multiplicity of coalitions).

To ensure that the conditions of the theorem hold in practice, there are multiple options. Note that the “low-latency” condition requires that latencies are bounded away from $z\tau$. Thus, one can increase τ to make this condition hold. Intuitively, this makes the late-proposing equilibrium worse by decreasing the speed of the system. However, increasing τ would require higher block rewards. This is because $B(\tau) = 0$, and to ensure $\langle M, B \rangle$ is time-decreasing, the slope of B must not change; thus, $B(0)$ must increase.

Another option is to uniformly increase the block reward function (i.e., shift it up). One reason that validators may prefer the late-proposing equilibrium is that they always get rewarded as if their rounds had zero duration, as opposed to the reward of $B(\mathbb{E}[S^i] + \mathbb{E}[\bar{E}_m^i])$ in the honest equilibrium. Increasing the block reward function will thus decrease the multiplicative advantage that occurs from gaining b_0 in the late-proposing equilibrium.

⁸We consider the restricted set of large coalitions that always vote 0 for each other because this voting strategy is optimal for small coalitions, and thus may arise naturally if a coalition starts small and grows. There is also an incentive for the coalition to grow, as members’ utilities rise as the size approaches k .

6 Fairness

Prior works (e.g., [PS17]) defined a blockchain as “fair” (among validators) when the total welfare is distributed proportionally to stake. One of the main relevant obstacles to fairness is latency—with time-varying block reward, the utility of a validator decreases as it experiences higher latency to other validators. Note that latency-based unfairness already exists with static block rewards. Validators with lower latency can propose closer to the timeout while still getting their blocks approved, resulting in more MEV than validators with higher latency. Throughout our analysis of fairness, we assume that validators are in the early-proposing and honest-voting equilibrium with dynamic block rewards, and in the late-proposing equilibrium with static block rewards.

We will capture *unfairness* as the utility ratio between the player with highest utility and the lowest utility. At best, this ratio is 1 and the equilibrium is completely fair. As we will see in this section, the ratio is often less than 1.

Definition 6.1 (latency advantage). *Let $U_i(\vec{\sigma})$ be the utility of player i at some fixed strategy profile $\vec{\sigma}$.*

$$\text{ADVANTAGE}(\vec{\sigma}) = \frac{\max_i U_i(\vec{\sigma})}{\min_i U_i(\vec{\sigma})} - 1$$

Closed-form results. We first analyze fairness in two simple network topologies with deterministic latencies to gain intuition about fairness under dynamic block rewards. In the *line model*, n validators are spaced evenly on a line, and the latency between two validators is proportional to their distance. In the *cluster model*, n validators are placed into two clusters, and the latency between two validators is low if they are in the same cluster and high otherwise. The following informal theorem summarizes results from [Theorems C.1 to C.3](#).

Theorem 6.1 (informal). *In the line model and the cluster model, with deterministic latencies and reasonable parameters, fairness is worse under dynamic block rewards than static block rewards.*

We defer full results to [Section C](#). The theorem suggests that shifting the proposal times earlier has a fairness cost, although the cost remains small in the scenarios we analyze. For instance, if we plug in reasonable parameters inspired by current Ethereum settings, the unfairness (i.e., the percentage increase between the lowest utility player and the highest utility player) remains within a couple percentage points of the static reward setting, in both the line and cluster models (see [Fig. 5](#)).

Simulation details. We also simulate both the dynamic and static block reward settings in network topologies inspired by real-world latencies. On a world latency model, where pairwise latencies are sampled from a lognormal, we observe that the unfairness under dynamic rewards is only slightly worse than under static rewards (see [Fig. 4](#)). We defer full details to [Section D](#).

7 Guidance for practice

We now discuss some details that may be relevant for implementing our work.

Empirical feasibility of time-decreasing rewards. Due to our requirements on total reward, observe that the maximum block reward needs to be greater than the accumulated MEV at timeout.

Observation 7.1. For MEV $M(x) = \mu x + \mu_0$ and block reward $B(x) = b_0 - bx$ where $B(\tau) = 0$, $\langle M, B \rangle$ is time-decreasing ([Definition 4.1](#)), if and only if $b_0 > \mu\tau$.




Figure 4: Simulation of 100 nodes placed in cities around the world, where the latency between i and j is log-normal with mean as average ping latencies.

We take the example of Ethereum to illustrate that b_0 can indeed be set sufficiently high in practice. Prior empirical studies on Ethereum have looked at the marginal value of time in the block building auction⁹. The increase in MEV is estimated as $\mu = 6.5 \times 10^{-6}$ ETH/ms in [SSST⁺23] and $\mu = 5.71 \times 10^{-6}$ ETH/ms in [OKV⁺23]. For a block reward curve of the form $b_0 - \frac{b_0}{12s}x$, b_0 needs to be greater than $\mu \cdot 12000$ ms per Observation 7.1. For $\mu = 6.5 \times 10^{-6}$ ETH/ms, this boils down to $b_0 > 0.078$ ETH. Fortunately, the total consensus reward per block in Ethereum is much greater than 0.078 ETH [Edg23, OKV⁺23].

Setting m . Our results only require that $m > n - c$. This is a natural condition, as consensus protocols assume that at most $n - c$ nodes can be faulty, implying that the timeliness vote should not come from a faulty node. Moreover, since $c > n/2$ for all reasonable BFT protocols, this implies that m is not required to be higher than the quorum size c of the BFT protocol. The coalition size condition is also quite natural. In general, m should be set sufficiently high for resistance to coalitions (including whales), but not so high that correct nodes in BFT protocols are insufficient to form the timeliness vote. If fewer than m validators submit a timing vote, the block reward would be zero, erasing any incentive to propose early. One natural setting is $m = 2(n - c)$, which ensures that $k = n - c$, and so the small coalition size accounts for all byzantine faulty nodes in BFT consensus models. Concretely, if $c = 2/3n$, then a natural setting is $m = 2/3n$, resulting in coalitions of all $n/3$ byzantine nodes to be characterized as small and not affecting the protocol’s responsiveness (Theorem 5.1).

Target duration. While responsiveness is desirable, in practice, decentralized protocols require that the round times are not too fast and slow nodes are able to keep up. Observe that the fastest c nodes can drive the protocol, leaving the rest behind. In practice, blockchain protocols attempt to have a *target duration* (minimum round time), e.g., 0.3 seconds in Solana [Tea25]. Our block reward function can be simply modified as follows to ensure that proposing before the target duration is never beneficial: the block reward starts at 0 for time 0, and increases until the target duration, and then decreases from there. This ensures that leaders will not want to propose before the target

⁹On Ethereum, most validators auction off the right to build their block to sophisticated block builders that are able to extract more MEV. Assuming that the auction is competitive, this provides a close lower bound for estimating MEV.

duration, but also not want to delay their proposal beyond the target duration.

8 Conclusion

We modeled the timing game in responsive BFT protocols and argued for measuring proposal timeliness via a vote. Our results indicate that with a carefully chosen block reward curve and timeout, responsiveness itself can be used to align timing incentives.

For future work, there are several avenues to explore. The timeliness measurement requires a voting protocol that needs to be designed and implemented. The design space for this protocol is rich. One option is to make the next leader aggregate timeliness votes and post (potentially with a delay). While this would be easy to implement, it increases the scope of strategic manipulations by allowing the aggregator to filter. Another promising approach is to aggregate and post timeliness votes on-chain.

Alternative incentive knobs to make early-proposing desirable should also be explored. Our discussion in [Section F](#) could serve as a first step in considering more advanced leader-election schemes that punish late-proposing. Moreover, as these approaches are not necessarily contradictory, further analysis of hybrid approaches could be explored.

Our coalition resistance results could also be extended and formalized further. For example, one could define a meta-game, where validators can join or leave coalitions over time. This would allow for formalizing our argument against small and medium rational coalitions and investigating further coalition dynamics.

Acknowledgments

We would like to thank Michael Setrin and Matthew Weinberg for useful comments on earlier drafts.

References

- [ADLS94] Hagit Attiya, Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Bounds on the time to reach agreement in the presence of timing uncertainty. *J. ACM*, 41(1):122–152, January 1994. [3](#)
- [art25] art3mis.cloud. <https://slowpoke.art3mis.cloud/>, 2025. Accessed: 2025-09. [2](#), [3](#)
- [AW24] Kaya Alpturer and S Matthew Weinberg. Optimal RANDAO manipulation in ethereum. In *6th Conference on Advances in Financial Technologies (AFT 2024)*, pages 10–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024. [4](#)
- [BCNPW19] Jonah Brown-Cohen, Arvind Narayanan, Alexandros Psomas, and S Matthew Weinberg. Formal barriers to longest-chain proof-of-stake protocols. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 459–473, 2019. [4](#)
- [BDKJ23] Kushal Babel, Philip Daian, Mahimna Kelkar, and Ari Juels. Clockwork finance: Automated analysis of economic security in smart contracts. In *2023 IEEE Symposium on Security and Privacy (SP)*, pages 2499–2516. IEEE, 2023. [1](#)
- [Buc16] Ethan Buchman. *Tendermint: Byzantine fault tolerance in the age of blockchains*. PhD thesis, University of Guelph, 2016. [4](#)

- [CGK⁺22] Shir Cohen, Rati Gelashvili, Lefteris Kokoris Kogias, Zekun Li, Dahlia Malkhi, Alberto Sonnino, and Alexander Spiegelman. Be aware of your leaders. In *Financial Cryptography and Data Security: 26th International Conference, FC 2022, Grenada, May 2–6, 2022, Revised Selected Papers*, page 279–295, Berlin, Heidelberg, 2022. Springer-Verlag. 3
- [CKWN16] Miles Carlsten, Harry Kalodner, S. Matthew Weinberg, and Arvind Narayanan. On the instability of bitcoin without the block reward. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS ’16*, page 154–167, New York, NY, USA, 2016. Association for Computing Machinery. 4
- [DGK⁺20] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *2020 IEEE symposium on security and privacy (SP)*, pages 910–927. IEEE, 2020. 1
- [Edg23] Ben Edgington. *Upgrading Ethereum*. Capella edition, 2023. 14
- [ES14] Ittay Eyal and Emin Gün Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *International Conference on Financial Cryptography and Data Security*, pages 436–454. Springer, 2014. 4
- [FM25] Flashbots and Robert Miller. Mev and the limits of scaling. <https://writings.flashbots.net/mev-and-the-limits-of-scaling>, 2025. Accessed: 2025-06. 4
- [GKKS⁺22] Rati Gelashvili, Lefteris Kokoris-Kogias, Alberto Sonnino, Alexander Spiegelman, and Zhuolun Xiang. Jolteon and ditto: Network-adaptive efficient consensus with asynchronous fallback. In Ittay Eyal and Juan Garay, editors, *Financial Cryptography and Data Security*, pages 296–315, Cham, 2022. Springer International Publishing. 3
- [GSPD⁺23] Neil Giridharan, Florian Suri-Payer, Matthew Ding, Heidi Howard, Ittai Abraham, and Natacha Crooks. Beegees: Stayin’ alive in chained bft. In *Proceedings of the 2023 ACM Symposium on Principles of Distributed Computing, PODC ’23*, page 233–243, New York, NY, USA, 2023. Association for Computing Machinery. 3
- [JB25] Mohammad Mussadiq Jalalzai and Kushal Babel. Monadbft: Fast, responsive, fork-resistant streamlined consensus, 2025. 3, 4
- [KGMS25] Dakai Kang, Suyash Gupta, Dahlia Malkhi, and Mohammad Sadoghi. Hotstuff-1: Linear consensus with one-phase speculation, 2025. 3
- [KGP⁺25] Mahimna Kelkar, Aadityan Ganesh, Aditi Partap, Joseph Bonneau, and S. Matthew Weinberg. Breaking omertà: On threshold cryptography, smart collusion, and whistleblowing. Cryptology ePrint Archive, Paper 2025/1582, 2025. 10
- [MMR24] Jason Milionis, Ciamac C Moallemi, and Tim Roughgarden. Automated market making and arbitrage profits in the presence of fees. In *International Conference on Financial Cryptography and Data Security*, pages 159–171. Springer, 2024. 3
- [OKV⁺23] Burak Öz, Benjamin Kraner, Nicolò Vallarano, Bingle Stegmann Kruger, Florian Matthes, and Claudio Juan Tessone. Time moves faster when there is nothing you anticipate: The role of time in mev rewards. In *Proceedings of the 2023 Workshop on*

- Decentralized Finance and Security*, DeFi '23, page 1–8, New York, NY, USA, 2023. Association for Computing Machinery. 1, 3, 14
- [PS17] Rafael Pass and Elaine Shi. Fruitchains: A fair blockchain. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, PODC '17, page 315–324, New York, NY, USA, 2017. Association for Computing Machinery. 13
- [PS18] Rafael Pass and Elaine Shi. Thunderella: Blockchains with optimistic instant confirmation. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018*, pages 3–33, Cham, 2018. Springer International Publishing. 3
- [Ray07] Debraj Ray. *A game-theoretic perspective on coalition formation*. Oxford University Press, 2007. 10
- [RV15] Debraj Ray and Rajiv Vohra. Chapter 5 - coalition formation. volume 4 of *Handbook of Game Theory with Economic Applications*, pages 239–326. Elsevier, 2015. 10
- [She79] Prakash P Shenoy. On coalition formation: a game-theoretical approach. *International journal of game theory*, 8(3):133–164, 1979. 10
- [SS22] Caspar Schwarz-Schilling. Retroactive proposer rewards. <https://notes.ethereum.org/@casparschwa/S1vcyXZL9>, 2022. Accessed: 2025-06. 3
- [SSST⁺23] Caspar Schwarz-Schilling, Fahad Saleh, Thomas Thiery, Jennifer Pan, Nihar Shah, and Barnabé Monnot. Time Is Money: Strategic Timing Games in Proof-Of-Stake Protocols. In Joseph Bonneau and S. Matthew Weinberg, editors, *5th Conference on Advances in Financial Technologies (AFT 2023)*, volume 282 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 30:1–30:17, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. 2, 3, 14
- [SSZ16] Ayelet Sapirshtein, Yonatan Sompolsky, and Aviv Zohar. Optimal selfish mining strategies in bitcoin. In *International conference on financial cryptography and data security*, pages 515–532. Springer, 2016. 4
- [SXNC25] Weijie Sun, Zihuan Xu, Wangze Ni, and Lei Chen. Intime: Towards performance predictability in byzantine fault tolerant proof-of-stake consensus. In *Proc. ACM Manag. Data*, volume 3, New York, NY, USA, feb 2025. Association for Computing Machinery. 3
- [Tea25] ChorusOne Team. Timing games on solana: Validator incentives, network impacts, and agave’s hidden inefficiencies. <https://x.com/ChorusOne/status/1955615856580301252>, 2025. Accessed: 2025-09. 2, 3, 14
- [TKSKK24] Giorgos Tsimos, Anastasios Kichidis, Alberto Sonnino, and Lefteris Kokoris-Kogias. Hammerhead: Leader reputation for dynamic scheduling. In *2024 IEEE 44th International Conference on Distributed Computing Systems (ICDCS)*, pages 1377–1387, 2024. 3
- [WZQ⁺23] Anton Wahrstätter, Liyi Zhou, Kaihua Qin, Davor Svetinovic, and Arthur Gervais. Time to bribe: Measuring block construction market, 2023. 3

- [YMR⁺19] Maofan Yin, Dahlia Malkhi, Michael K. Reiter, Guy Golan Gueta, and Ittai Abraham. Hotstuff: Bft consensus with linearity and responsiveness. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, PODC '19, page 347–356, New York, NY, USA, 2019. Association for Computing Machinery. [1](#), [3](#), [4](#)
- [YSZ23] Aviv Yaish, Gilad Stern, and Aviv Zohar. Uncle maker: (time)stamping out the competition in ethereum. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, CCS '23, page 135–149, New York, NY, USA, 2023. Association for Computing Machinery. [2](#), [4](#)
- [YTZ22] Aviv Yaish, Saar Tochner, and Aviv Zohar. Blockchain stretching & squeezing: Manipulating time for your best interest. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 65–88, 2022. [4](#)

A Equilibria Results (Section 4)

Lemma 4.1. *If $m > n - c$, then the m^{th} smallest vote is always at least the round delay Δ in the early-proposing truthful strategy profile σ^* .*

Proof. Let i be an arbitrary current leader and j the previous leader in round r . If a validator k contributed to the quorum in i 's proposal, they must have started their round- r clock before i receives the quorum. Thus, if i delays by Δ , k 's honest vote would be above Δ . Since $m > n - c$, the m^{th} smallest vote will come from one of the above quorum-contributing validators, which means that the aggregated vote is at least Δ . $\square$

Theorem 4.1. *Suppose $\langle M, B \rangle$ are time-decreasing and $m > n - c$. Then, the early-proposing truthful strategy profile σ^* is a subgame-perfect Bayes-Nash equilibrium, where the time-averaged utility of i is:*

$$U_i(\sigma^*) = \frac{\sum_j \left(\mathbb{E}[M(S^{j,i})] + \mathbb{E}[B(S^{j,i} + E_m^{i-j})] \right)}{\sum_{j,k} \mathbb{E}[S^{j,k}]}$$

Proof. Consider an arbitrary subgame: i is the leader, and j is the previous leader. We will show that i never wants to propose a delay $\Delta \in (0, \Delta^*]$ or vote dishonestly.

First, note that i 's own utility does not depend on their voting strategy, as their votes only determine the block reward when other validators are the leader. So, we focus on proposal delays. When i is proposing their block, they have observed the time they received the block and the time they received the quorum. Since all other validators follow σ^* , i knows that j did not intentionally delay their block. Thus, by subtracting these times, they can learn $q_c^{j \rightarrow i} - l_{j \rightarrow i}$, where $q_c^{j \rightarrow i} \sim Q_c^{j \rightarrow i}$ and $l_{j \rightarrow i} \sim L_{j \rightarrow i}$. Depending on the latency distributions, this could reveal some information about when validators have received block j and started their clock for i 's round. Let $\tilde{E}_m^{i-j}$ represent the distribution of the m^{th} latency difference, conditional on i 's observations. Similarly, let $\tilde{S}^{j,i}$ represent the distribution of i 's start time, conditional on their observations.

From Lemma 4.1, we know that $\Delta + \tilde{S}^{j,i} + \tilde{E}_m^{i-j} \geq \Delta > 0$. From the optimal-delaying oracle, we also know that $\Delta + \tilde{S}^{j,i} + \tilde{E}_m^{i-j} \leq \tau$. Thus, $B(\Delta + \tilde{S}^{j,i} + \tilde{E}_m^{i-j})$ is well defined. Now, i 's expected reward for the current round for delay Δ , given their observations, is:

$$\begin{aligned} \mathbb{E}[\text{REWARD}_i^{(r)} \mid j, r-1] &= \mathbb{E} \left[M(\Delta + \tilde{S}^{j,i}) + B(\Delta + \tilde{S} + \tilde{E}_m^{i-j}) \right] \\ &= \mathbb{E}[M(\Delta + \tilde{S}^{j,i}) + B(\Delta + \tilde{S}^{j,i})] - b\tilde{E}_m^{i-j} \end{aligned}$$

Where b is the block reward slope.

Since $\langle M, B \rangle$ are time-decreasing, $\Delta = 0$ maximizes the expectation, regardless of the realization $s^{j,i} \sim S^{j,i}$, and thus maximizes the overall reward, regardless of i 's information. Thus, σ^* is a subgame-perfect Bayes-Nash equilibrium.

We apply the Renewal-Reward theorem to compute the time-averaged expected utility of i in σ^* . That allows us to first compute the expected reward of i in a single round, and then the expected duration of a round, and then divide. For the reward, note that the leaders are elected uniformly

at random. Thus, consider the expected reward of i in a given round:

$$\begin{aligned}
\mathbb{E}[\text{REWARD}_i] &= \Pr[i = \text{leader}] \sum_j \Pr[j = \text{leader}_{\text{prev}}] \mathbb{E}[\text{REWARD}_i \mid j = \text{leader}_{\text{prev}}] \\
&= \frac{1}{n} \sum_j \frac{1}{n} (\mathbb{E}[M(S^{j,i})] + \mathbb{E}[B(S^{j,i} + E_m^{i-j})]) \\
&= \frac{1}{n^2} \sum_j \mathbb{E}[M(S^{j,i})] + B(\mathbb{E}[S^{j,i}]) - b(\mathbb{E}[E_m^{i-j}])
\end{aligned}$$

Now, note that since $\Delta = 0$, the expected duration of a round where i leads and j led the previous round is simply $\mathbb{E}[S^{j,i}]$. Thus, the expected duration is $\frac{1}{n^2} \sum_{i,j} \mathbb{E}[S^{j,i}]$. Dividing the expected reward by the expected round duration yields the desired time-averaged utility. $\square$

B Coalition Resistance Results (Section 5)

Lemma 5.1. *Let $k = m + c - n$. If at least $n - k$ validators vote honestly, then for any round with delay Δ , the m^{th} smallest vote will be at least Δ . If fewer than $n - k$ validators vote honestly, then the m^{th} smallest vote may be less than Δ .*

Proof. We count the leader's delay Δ from the earliest time they could have proposed, which is after they receive the quorum. At that point, at least c validators have received the previous block, and will therefore observe an duration of at least Δ . Since k of these validators may dishonest, only $c - k = n - m$ are guaranteed to report a duration of at least Δ . Thus, the $(n - m)^{\text{th}}$ largest vote (which is the m^{th} smallest vote) is at least Δ .

If fewer validators voted honestly, then note that all non-quorum validators may receive the previous block after the current leader's block (and would thus observe a negative round time), and dishonest validators can similarly report a negative round time. In this case, the m^{th} smallest vote would be less than Δ (in fact, it would be negative). $\square$

Proposition 5.1. *Consider a whale of size at most k . The optimal strategy for the whale is the early-proposing honest strategy, σ^* .*

Proof. Recall that the time-averaged utility under σ^* for node i is:

$$\frac{\sum_j \mathbb{E}[M(S^{j,i})] + B(\mathbb{E}[S^{j,i}]) - b\mathbb{E}[E_m^{i-j}]}{\sum_{i,j} \mathbb{E}[S^{j,i}]}$$

Notice that the denominator (the duration) cannot decrease, as in σ^* everyone proposes as early as possible. So, the only way for the whale to benefit is to increase the numerator, which is their reward. We split into cases, based on whether the whale deviates by delaying their proposal. First, suppose that they don't delay, and propose at $\Delta = 0$. Then, they could only benefit by lowering their latency penalty ($\sum_j b\mathbb{E}[E_m^{i-j}]$) through dishonest voting. However, since a whale has no latency to its own nodes, the whale already votes 0 for their rounds under the honest strategy. So, the whale can't lower their latency penalty.

Now, suppose the whale proposed at $\Delta > 0$. From 5.1, we know that the m^{th} smallest vote will be at least Δ . Since $\langle M, B \rangle$ are time-decreasing, a proposal of $\Delta = 0$ would maximize the whale's utility. $\square$

Theorem 5.3. *Consider a rational coalition of size at most k . The optimal strategy for the coalition members is to always vote 0 on each other's proposals (despite any latency they may observe), and propose at $\Delta = 0$.*

Proof. The previous proof illustrates why delaying one's proposal is suboptimal, no matter how the k other coalition members vote. Now, given that everyone is still proposing at $\Delta = 0$, voting 0 clearly reduces the latency penalty as much as possible.

In particular, if everyone votes at 0, the time-averaged utility of coalition member i is almost the same as in the honest case, except for the latency penalty. That penalty was proportional to $\sum_j b\mathbb{E}[E_m^{i-j}]$ before, but now it is $\sum_j b\mathbb{E}[\tilde{E}_{m-k}^{i-j}]$, where $\tilde{E}_{m-k}^{i-j}$ is the $(m-k)^{\text{th}}$ order statistic of $\{L_{i \rightarrow k} - L_{j \rightarrow k}\}_{k \notin C}$, where C is the coalition. In other words, the latency penalty now depends on the $(m-k)^{\text{th}}$ smallest latency among the $n-k$ honest players, instead of depending on the m^{th} smallest latency among all honest players. In particular, this latency penalty in the coalition can never be larger than the latency penalty in the honest strategy. Further, from i 's perspective, the most valuable coalition is the one which features the nodes furthest from i (i.e., those nodes which have highest latency to i). $\square$

Theorem 5.4. *Suppose a large coalition C votes 0 for leaders from the coalition and votes honestly for other leaders, while the rest of the validators use truthful strategy σ^* . If the system has z -low-latency with $z \leq \frac{n-c}{3n-c}$, then $|C| - k$ validators within the coalition would prefer defecting to the early-proposing truthful strategy.*

Proof. We start by bounding the utility of coalition members. Observe that if a coalition leader i delayed by Δ , their reward for that round would be $M(S^{j,i} + \Delta) + b_0$. Due to the optimal-delaying oracle, coalition leaders will know the time Δ^* that is the latest time they can propose and not timeout, and their utility is clearly maximized at Δ^* .

Now, note that Δ^* averaged over the previous leader can be derived from $\bar{S}^i + \bar{E}_c^i + \Delta^* = \tau$ where $\bar{E}_c^i = \frac{1}{n} \sum_{k \in [n]} E_c^{i-k}$. This means that the c^{th} slowest validator sees i 's block exactly when their local clock is τ . Thus, $\mathbb{E}[\Delta^* + S^i] = \tau - \mathbb{E}[\bar{E}_c^i]$, which is the average round duration of coalition leader i . This means that their utility per-round is $M(\tau - \mathbb{E}[\bar{E}_c^i]) + b_0 \leq M(\tau) + b_0$. Thus, the time-averaged utility of coalition member i is at most:

$$\frac{\frac{1}{n}(M(\tau) + b_0)}{\frac{|C|}{n} \sum_{i \in C} \tau - \mathbb{E}[\bar{E}_c^i] + \frac{n-|C|}{n} \sum_{i \notin C} \mathbb{E}[\bar{S}^i]}$$

We now bound the increase in utility.

Claim B.1. *The average round duration is at least $\frac{|C|-k}{zn}$ times higher for the coalition than in the honest equilibrium.*

Proof. From z -low-latency, we know that at least $|C| - k$ validators in the coalition have $\mathbb{E}[S^i] \leq z\tau - \mathbb{E}[\bar{E}_c^i]$. These validators will increase their round duration from at most $z\tau - \mathbb{E}[\bar{E}_c^i]$ to $\tau - \mathbb{E}[\bar{E}_c^i]$, which is an multiplicative increase of more than z . Since no validator takes less time under the coalition, the total round duration increases by a multiplicative factor of at least $\frac{|C|-k}{nz}$. $\square$

Claim B.2. *For any coalition C , there exists at least $|C| - k$ members whose reward when leading a round in the coalition is at most $\frac{2}{1-z}$ times as high as in the honest equilibrium.*

Proof. In the honest equilibrium, the reward for validator i when leading is $\mathbb{E}[M(S^i)] + B(\mathbb{E}[S^i] + \mathbb{E}[\bar{E}_m^i])$, which is highest for nodes with lowest $\mathbb{E}[S^i] + \mathbb{E}[\bar{E}_m^i]$. From z -low-latency, we know that at least $|C| - k$ validators in the coalition have $\mathbb{E}[S^i + \bar{E}_m^i] \leq z\tau$. For these validators, their reward in the honest equilibrium is at least $B(\mathbb{E}[S^i] + \mathbb{E}[\bar{E}_m^i]) \geq B(z\tau)$. Since B is linear and $B(\tau) = 0$, $B(z\tau) = (1 - z)b_0$.

In the coalition, all validators get reward $b_0 + M(\tau)$. Since $\langle M, B \rangle$ is time-decreasing, we know that $M(\tau) + B(\tau) \leq M(0) + b_0$. Letting $B(\tau) = M(0) = 0$, we get $M(\tau) \leq b_0$. Thus, the validators get reward at most $2b_0$ in the coalition. Dividing these terms results in a multiplicative gain of at most $\frac{2}{1-z}$ for $|C| - k$ validators in the coalition. $\square$

Putting the two claims together, we get that at least $|C| - k$ members in the coalition would prefer to leave the coalition if $\frac{|C|-k}{nz} \geq \frac{2}{1-z}$, which simplifies to $z \leq \frac{|C|-k}{2n+|C|-k}$. The right hand side is minimized when $|C| = m$. Since $m - k = n - c$, if $z \leq \frac{n-c}{3n-c}$, $|C| - k$ members of the coalition prefer the honest equilibrium to the coalition. $\square$

C Fairness Results (Section 6)

In this section, we analyze fairness for two concrete latency models under dynamic or static block rewards. We refer to the former as the “early-proposing equilibrium” and the latter as the “late-proposing equilibrium”.

We analyze two latency models inspired by concerns of geographical collocation. The first model is one where nodes are relatively evenly dispersed, and the second is one with two clusters of nodes, with $2/3$ and $1/3$ of the nodes in the two clusters. Our results indicate that fairness slightly degrades when moving to early-proposing equilibrium. At the end of the section, we provide an empirical analysis of latency and fairness experienced in a real-life distributed system and explore broadcast protocols as a promising avenue for mitigating unfairness. Our theoretical results are derived via deterministic latencies and consider the following two models:

- *Line model:* Nodes are evenly dispersed on a number line. Latency is directly proportional to distance. Intuitively, we want to measure the difference in utility between the center node and the endpoint node.
- *Cluster model:* Nodes are clustered into two clusters. Within a cluster, latency is very low and across the clusters latency is very high. Intuitively, we will compare the gap between a node in one cluster and a node in the other cluster.

In [Section D](#) we also simulate these models with lognormal latency distributions and get similar results to the theoretical ones.

Reward slope and timeout slack. Since early-equilibria requires the slope of block reward to be greater than the MEV slope (Observation 7.1), we will assume that $b > \mu$ throughout this section. Therefore, we will phrase our results in terms of k such that $b = k\mu$ for $k > 1$. Another intuitive parameter we use in our results is the timeout slack, how many times larger the timeout is as compared to the slowest node’s start time in order to allow for everyone to participate. Intuitively, when consensus parameters are set, the unhappy path timeout gets set based on how early the slowest node can propose their blocks. Let d represent this slack in the timeout.

Late equilibria. Here, we assume a fixed block reward of b_0 . We assume validators will as late as possible while ensuring that every block still makes it in. Hence we propose at $\Delta^* = \tau - S^{j,i} - E_c^{i-j}$ where i is the current leader preceded by leader j . Note that this is even better than proposing late in real life since there is no risk of missing blocks. We call this the late proposing strategy profile $\vec{\sigma}_{\text{late}}$.

Throughout the section, we theoretically analyze deterministic latency models, where links have fixed latency both on the line and the cluster model. Later on, we simulate both models with log-normal latency distributions.¹⁰

We observe that there is a cost to early equilibria in terms of fairness. For parameters¹¹ such that $k = 1.5$, and $m = c = 2/3$, we see that the fairness to utilizing early-equilibria compared to late (see Fig. 5) is up to 8% worse in terms of advantage.

C.1 Nodes on a line

We begin by considering deterministic latency. Validators are evenly spaced along a number line, with validator i at number i . The latency between i and j is just $|j - i|$. We aim to measure the fairness index, which occurs between validator $n/2$ (the most centrally located) and validator 0. To do so, we assume that block reward and MEV are linear: $B(x) = b_0 - bx$ and $M(x) = \mu x + \mu_0$. We also assume that $m = c = 2n/3$. We call this the *deterministic line* model.

We now compute utility ratios to measure fairness in the deterministic line model. First observe that the ratio of their utilities in the honest equilibrium depends on $\bar{S}^i$ and $\bar{E}_m^i$ for $i \in \{0, n/2\}$.

Lemma C.1. *For the early-proposing strategy profile σ^* , if $b' > m$, then the fairness index in the deterministic line model is given by the following:*

$$\text{ADVANTAGE}(\vec{\sigma}^*) = \frac{(b_0 + \mu_0) - (b - \mu)\mathbb{E}[\bar{S}^{n/2}] - b\mathbb{E}[\bar{E}_m^{n/2}]}{(b_0 + \mu_0) - (b - \mu)\mathbb{E}[\bar{S}^0] - b\mathbb{E}[\bar{E}_m^0]} - 1$$

Proof. Observe that the time contribution of average round duration cancels out in the ratio of two utilities. Hence, consider the reward earned in expectation by each player. Observe that player i becomes leader right after all players j with probability $1/n^2$, in which case it proposes at $S^{j,i}$. Hence, using $\bar{S}^i = \frac{1}{n} \sum_{j \in [n]} S^{j,i}$ and $\bar{E}^i = \frac{1}{n} \sum_{j \in [n]} E_m^{i-j}$, we have the following:

$$\frac{U_i(\sigma^*)}{U_j(\sigma^*)} = \frac{b_0 - b(\mathbb{E}[\bar{S}^i + \bar{E}_m^i]) + \mu(\mathbb{E}[\bar{S}^i]) + \mu_0}{b_0 - b(\mathbb{E}[\bar{S}^j + \bar{E}_m^j]) + \mu(\mathbb{E}[\bar{S}^j]) + \mu_0}$$

Since the player on the edge of the line and the middle of the line have the two extreme latency distributions and $b' > m$, player $n/2$ earns the most and player 0 earns the least. We get the desired conclusion after applying linearity of expectation. $\square$

Now, the expectations range only over the leader election randomness, as latencies are otherwise deterministic.

¹⁰For all our results, we assume that n is divisible by 2 and 3 to simplify the expressions.

¹¹For example, for high performance chains, 500 ms block times with a 2500 ms timeout and a quorum threshold at $2/3$ is very reasonable. The parameter k depends on calibration of rewards. For a discussion of reward parameters see Section 7.

Lemma C.2. *In the deterministic line model, assume without loss of generality that $i \leq j$ for $i, j \in [n]$. Then,*

$$\mathbb{E}[Q_c^{i \rightarrow j}] = \mathbb{E}[Q_c^{j \rightarrow i}] = \begin{cases} j - i & j - i + 1 \geq 2n/3 \\ 2n/3 - 1 + \kappa_{ij} & j - i + 1 < 2n/3 \wedge j - i + 1 + 2w_{ij} \geq 2n/3 \\ 4n/3 - 2w_{ij} - j + i - 2 & \text{otherwise} \end{cases}$$

where $w_{ij} = \min(n - j - 1, i)$ and $\kappa_{ij} = \mathbb{I}[2n/3 - (j - i + 1) \bmod 2 = 1]$.

Proof. We need to compute $Q_c^{j \rightarrow i}$ for all j and for $i = 1, n/2$. So, we need the $(2/3)^{\text{th}}$ order statistic of $\{L_{j \rightarrow k} + L_{k \rightarrow i}\}_k = \{|j - k| + |k - i|\}_k$. Since each path is symmetric, $\mathbb{E}[Q_c^{j \rightarrow i}] = \mathbb{E}[Q_c^{i \rightarrow j}]$ so without loss of generality assume that $j \geq i$. Then, there are three cases to consider for k :

1. For all $k \in [i, j]$, $L_{j \rightarrow k} + L_{k \rightarrow i} = j - i$.
2. For all $k < i$: $L_{j \rightarrow k} + L_{k \rightarrow i} = j + i - 2k > j - i$.
3. For all $k > j$, $L_{j \rightarrow k} + L_{k \rightarrow i} = 2k - (j + i) > j - i$.

For the $(2/3)^{\text{rd}}$ fastest node k , note first that the nodes in $[i, j]$ are fastest. After those nodes are exhausted, the nodes $i - 1$ and $j + 1$ have the same latency: $j - i + 2$. So, then we pair up these nodes while they exist. After that point (i.e., we hit $k = 1$ or n on the sides), we select single nodes on the remaining side to get to the $2/3$ fastest node. Thus, $Q_c^{j \rightarrow i}$ can be broken into three cases:

- (a) If $j - i + 1 \geq 2n/3$, then all $2n/3$ nodes are contained in $[i, j]$ so $Q_c^{j \rightarrow i} = j - i$.
- (b) If $j - i + 1 < 2n/3$ but $j - i + 1 + 2 \min(n - j - 1, i) \geq 2n/3$, then we can pair nodes from both sides of the line to form $2n/3$ votes so $Q_c^{j \rightarrow i} = j - i + 2 \lceil (\frac{2n}{3} - (j - i + 1)) / 2 \rceil$. The second term is the offset from the paired nodes before i and after j . This simplifies to $2n/3 - 1 + \kappa_{ij}$.
- (c) Otherwise, $Q_c^{j \rightarrow i} = (j - i) + 2 \min(n - j - 1, i) + 2 (\frac{2n}{3} - (2 \min(n - j - 1, i) + j - i + 1))$. Here, the $\min(n - j - 1, i)$ term comes from pairing nodes outside $[i, j]$, and the last term comes from the remaining nodes needed on one side. These terms are multiplied by two, since the offset increases by 2 each time. This simplifies to $Q_c^{j \rightarrow i} = 4n/3 - 2 \min(n - j - 1, i) - j + i - 2$.

□

Lemma C.3. *In the deterministic line model, $\mathbb{E}[\bar{S}^{n/2}] = \frac{25n}{36} - \frac{2}{3}$ and $\mathbb{E}[\bar{S}^0] = \frac{17n}{18} - \frac{7}{6}$.*

Proof. Recall that $S^{j,i} = \max(Q_c^{j \rightarrow i}, L_{j \rightarrow i})$. In this model, the quorum time always takes longer than the second term (due to these deterministic latencies obeying the triangle inequality). By Lemma C.2, we consider three cases for both $\bar{S}^0$ and $\bar{S}^{n/2}$.

For $\bar{S}^0$ observe that the middle case vanishes, as there are no nodes before the first. Further, $\min(n - j - 1, i) = 0$, as $i = 0$. Putting this all together:

$$\begin{aligned} \bar{S}^0 &= \frac{1}{n} \sum_{j \in [n]} S^{j,0} = \frac{1}{n} \sum_{j \in [n]} Q_{2/3}^{j \rightarrow 0} \\ &= \frac{1}{n} \left(\sum_{j=0}^{2n/3-2} \left(\frac{4n}{3} - j - 2 \right) + \sum_{j=2n/3-1}^{n-1} j \right) && \text{Lemma C.2} \\ &= \frac{1}{n} \left(\frac{(n-1)(2n-3)}{18} + \frac{(n+3)(5n-6)}{18} \right) \\ &= \frac{17n}{18} - \frac{7}{6} \end{aligned}$$

Now, we consider $\bar{S}^{n/2}$. Note that our earlier derivations assumed that $j \geq i$. However, $S^{j,i} = S^{i,j}$ since $L_{j \rightarrow i} = L_{i \rightarrow j}$. So, we simply compute $S^{n/2,j}$ for $j < n/2$. Note that in this case, $j - n/2 + 1$ is always less than $2n/3$, so the first case for $Q_c^{j \rightarrow n/2}$ vanishes. Further, whether there's enough nodes for the “pairing” strategy to work simply depends on if $j \leq n/6$ or $j > 5n/6$; otherwise, we use the second case. Putting it all together:

$$\begin{aligned} n\bar{S}^{n/2} &= \sum_{j \in [n]} S^{j,n/2} \\ &= \sum_{j=0}^{n/2-1} S^{n/2,j} + \sum_{j=n/2}^n S^{j,n/2} \\ &= \sum_{j=0}^{n/2-1} Q_{2/3}^{n/2 \rightarrow j} + \sum_{j=n/2}^{n-1} Q_{2/3}^{j \rightarrow n/2} \end{aligned}$$

Now, we compute these terms using [Lemma C.2](#). The four segments we consider are $[0, n/6 - 2]$, $[n/6 - 1, n/2 - 1]$, $[n/2, 5n/6 - 1]$, $[5n/6, n - 1]$.

First consider $[0, n/6 - 2], [n/6 - 1, n/2 - 1]$. Observe that for $Q_{2/3}^{n/2 \rightarrow j}$, $\min(n/2 - 1, j) = j$. Moreover, for $j \in [0, n/6 - 2]$ case 3 applies and for $j \in [n/6 - 1, n/2 - 1]$ case 2 applies of [Lemma C.2](#) always applies. Therefore,

$$\begin{aligned} \sum_{j=0}^{n/2-1} Q_{2/3}^{n/2 \rightarrow j} &= \sum_{j=0}^{n/6-2} \left(\frac{4n}{3} - j + \frac{n}{2} - 2 \right) + \sum_{j=n/6-1}^{n/2-1} \left(\frac{2n}{3} - 1 + \kappa_{ij} \right) \\ &= \left(\sum_{j=0}^{n/6-2} \left(\frac{5n}{6} - j + \frac{n}{2} - 2 \right) \right) + \left(\frac{n}{3} + 1 \right) \left(\frac{2n}{3} - 1 \right) + \frac{n}{6} \\ &= \left(\frac{n^2}{8} - \frac{11n}{18} + 1 \right) + \left(\frac{2n^2}{9} + \frac{n}{2} - 1 \right) \end{aligned}$$

where we use the fact that there are $n/6$ integers that make $\kappa_{(n/2)j} = 1$ in the second segment for n divisible by 6.

Similarly, for $[n/2, 5n/6 - 1]$ case 2 applies and for $[5n/6, n - 1]$ case 3. We use the fact that in $[5n/6, n - 1]$, $\min(n - 1 - j, n/2) = n - 1 - j$ and in $[n/2, 5n/6 - 1]$ there are $n/6$ that make $\kappa_{(n/2)j} = 1$.

$$\begin{aligned} \sum_{j=n/2}^{n-1} Q_{2/3}^{j \rightarrow n/2} &= \sum_{j=n/2}^{5n/6-1} \left(\frac{2n}{3} - 1 + \kappa_{ij} \right) + \sum_{j=5n/6}^{n-1} \left(j - \frac{n}{6} \right) \\ &= \left(\frac{n}{3} \right) \left(\frac{2n}{3} - 1 \right) + \frac{n}{6} + \left(\sum_{j=5n/6}^{n-1} \left(j - \frac{n}{6} \right) \right) \\ &= \left(\frac{2n^2}{9} - \frac{n}{6} \right) + \left(\frac{n^2}{8} - \frac{n}{12} \right) \end{aligned}$$

Lastly, adding all and dividing by n , we get the following:

$$\bar{S}^{n/2} = \frac{1}{n} \left(\sum_{j=0}^{n/2-1} Q_{2/3}^{n/2 \rightarrow j} + \sum_{j=n/2}^{n-1} Q_{2/3}^{j \rightarrow n/2} \right) = \frac{25n}{36} - \frac{2}{3}$$

□

Lemma C.4. $\mathbb{E}[\overline{E}_m^0] = \frac{7n-15}{18}$ and $\mathbb{E}[\overline{E}_m^{n/2}] = \frac{n}{36} - \frac{2}{3}$.

Proof. Recall that $E_m^{i-j} = f(m, \{L_{i \rightarrow k} - L_{j \rightarrow k}\}_k) = f(m, \{|i-k| - |j-k|\})$, i.e., the m th smallest number in the set. Now, suppose that $i < j$. Then, there are three cases:

1. If $k < i$, then $i - k - (k - j) = i - j$.
2. If $k \in [i, j]$, then $k - i - j + k = 2k - (i + j)$.
3. If $k > j$, then $k - i - k + j = j - i$.

Note that the values goes up monotonically with k . So, $E_m^{i-j} = L_{i \rightarrow m} - L_{j \rightarrow m}$, for $i < j$. A similar argument shows that when $j > i$, the values go down monotonically with k , and thus $E_m^{i-j} = L_{i \rightarrow (n-m)} - L_{j \rightarrow (n-m)}$. Note that $E_m^{i-i} = 0$ and the m^{th} index when indexing from 0 is $2n/3 - 1$.

We first compute $\overline{E}_m^0$:

$$\begin{aligned}
n\overline{E}_m^0 &= \sum_{j=0}^{n-1} E_m^{0-j} \\
&= \sum_{j=1}^{2n/3-1} \left(\left(\frac{2n}{3} - 1 \right) - \left(\frac{2n}{3} - 1 - j \right) \right) + \sum_{j=2n/3}^{n-1} \left(\left(\frac{2n}{3} - 1 \right) - \left(j - \frac{2n}{3} + 1 \right) \right) \\
&= \sum_{j=1}^{2n/3-1} j + \sum_{j=2n/3}^{n-1} \left(\frac{4n}{3} - j - 2 \right) \\
&= \frac{2n^2 - 3n}{9} + \frac{n^2 - 3n}{6} \\
&= \frac{n(7n - 15)}{18}
\end{aligned}$$

Thus, $\overline{E}_m^1 = \frac{7n-15}{18}$, as desired.

Now, we compute $\overline{E}_m^{n/2}$:

$$\begin{aligned}
n\overline{E}_m^{n/2} &= \sum_{j=1}^{n/2-1} E_m^{n/2-j} + \sum_{j=n/2+1}^n E_m^{n/2-j} \\
&= \sum_{j=1}^{n/2-1} \left(\left| \frac{n}{2} - \frac{n}{3} \right| - \left| j - \frac{n}{3} \right| \right) + \sum_{j=n/2+1}^n \left(\left| \frac{n}{2} - \left(\frac{2n}{3} - 1 \right) \right| - \left| j - \left(\frac{2n}{3} - 1 \right) \right| \right) \\
&= \frac{n^2}{36} - \frac{2n}{3}
\end{aligned}$$

So, $\overline{E}_m^{n/2} = \frac{n}{36} - \frac{2}{3}$, as desired. □

We now combine these results to prove the following theorem.

Theorem C.1 (line, σ^*). *In the line model with deterministic latencies, the advantage is:*

$$\text{ADVANTAGE}(\sigma^*) = \frac{d + \frac{\mu_0}{kn\mu} - \frac{k-1}{k} \left(\frac{25}{36} - \frac{2}{3n} \right) - \frac{1}{36} - \frac{2}{3n}}{d + \frac{\mu_0}{kn\mu} - \frac{k-1}{k} \left(\frac{17}{18} - \frac{7}{6n} \right) - \frac{7}{18} - \frac{5}{6n}}$$

Proof. Recall that the ratio of utilities is:

$$\frac{U_{n/2}}{U_0} = \frac{(b_0 + \mu_0) - (b - \mu)\mathbb{E}[\bar{S}^{n/2}] - b\mathbb{E}[\bar{E}_m^{n/2}]}{(b_0 + \mu_0) - (b - \mu)\mathbb{E}[\bar{S}^1] - b\mathbb{E}[\bar{E}_m^1]}$$

Plugging in the values from [Lemmas C.3](#) and [C.4](#):

$$\frac{U_{n/2}}{U_0} = \frac{(b_0 + \mu_0) - (b - \mu) \left(\frac{25n}{36} - \frac{2}{3} \right) - b \left(\frac{n}{36} - \frac{2}{3} \right)}{(b_0 + \mu_0) - (b - \mu) \left(\frac{17n}{18} - \frac{7}{6} \right) - b \left(\frac{7n}{18} - \frac{5}{6} \right)}$$

Now, note that the longest start time, $S^{0,n-1} = n - 1 \approx n$. Suppose the timeout is a multiple of the longest start time, so that $\tau = dn$.

Now, since $B(\tau) = 0$, we know that $b_0 - b\tau = 0$. Since $b > \mu$, let $b = k\mu$ for $k > 1$. Then, we get that $b_0 - k\mu\tau = 0$, or $b_0 = k\mu\tau = k\mu dn$. Now, the utility ratio simplifies to:

$$\begin{aligned} \frac{U_{n/2}}{U_0} &= \frac{k\mu dn + \mu_0 - (k-1)\mu \left(\frac{25n}{36} - \frac{2}{3} \right) - k\mu \left(\frac{n}{36} - \frac{2}{3} \right)}{k\mu dn + \mu_0 - (k-1)\mu \left(\frac{17n}{18} - \frac{7}{6} \right) - k\mu \left(\frac{7n}{18} - \frac{5}{6} \right)} \\ &= \frac{d + \frac{\mu_0}{kn\mu} - \frac{k-1}{k} \left(\frac{25}{36} - \frac{2}{3n} \right) - \frac{1}{36} - \frac{2}{3n}}{d + \frac{\mu_0}{kn\mu} - \frac{k-1}{k} \left(\frac{17}{18} - \frac{7}{6n} \right) - \frac{7}{18} - \frac{5}{6n}} \end{aligned}$$

As n increases, this quickly converges to:

$$\frac{U_{n/2}}{U_0} = \frac{d - \frac{25}{36} \cdot \frac{k-1}{k} - \frac{1}{36}}{d - \frac{17}{18} \cdot \frac{k-1}{k} - \frac{7}{18}}$$

□

C.2 Nodes in two clusters

We now consider a different latency model. Suppose the validators are clustered into two clusters: X and Y , where $|X| \geq |Y|$. Latencies within a cluster are small: ε . Latencies across the cluster are larger ($l > \varepsilon$). We call this the *two cluster model*. What is the ratio of utilities of a node in one cluster to the other?

As before, we assume that latencies are deterministic, and that block reward and MEV are linear and $c = m = 2n/3$.

We compute the utility ratios in the two cluster model. For notational convenience, let $\bar{X}$ denote the relative fraction $\bar{X} = \frac{|X|}{n}$ and $\bar{Y}$ denote $\bar{Y} = \frac{|Y|}{n}$.

Lemma C.5. *In the two cluster model, for $x \in X$ and $y \in Y$ the following holds:*

$$\mathbb{E}[\bar{S}^y] = \bar{X}(\varepsilon + l) + 2\bar{Y}l \quad \text{and} \quad \mathbb{E}[\bar{S}^x] = \begin{cases} 2\bar{X}\varepsilon + \bar{Y}(\varepsilon + l) & |X| \geq c \\ 2\bar{X}l + \bar{Y}(\varepsilon + l) & \text{otherwise} \end{cases}$$

Proof. Consider $L_{j \rightarrow k} + L_{k \rightarrow i}$. There are three cases:

1. If j and i are in different clusters, then this is $\varepsilon + l$ (one hop within the cluster, the other hop between clusters).
2. If j, i , and k are in the same cluster, then this is 2ε .
3. If j and i are in the same cluster while k is in the other cluster, this is $2l$.

Now, if j and i are in different clusters, then $Q_c^{j \rightarrow i} = \varepsilon + l$. Otherwise, if they are in the same cluster, it depends on whether that cluster has size at least c . If so, then $Q_c^{j \rightarrow i} = 2\varepsilon$; otherwise, $Q_c^{j \rightarrow i} = 2l$.

Now, let x be an arbitrary node in X . If $|X| \geq c$:

$$n\bar{S}^x = \sum_j S^{j,x} = \sum_{j \in X} 2\varepsilon + \sum_{j \in Y} \varepsilon + l = |X|2\varepsilon + |Y|(\varepsilon + l)$$

Dividing by n yields the desired result. If $|X| < c$, then $\bar{S}^x = 2\bar{X}l + 2\bar{Y}(\varepsilon + l)$. Since $|X| \geq |Y|$ and $c > 0.5$, an arbitrary node $y \in Y$ only has one case: $\mathbb{E}[\bar{S}^y] = \bar{X}(\varepsilon + l) + 2\bar{Y}l$. $\square$

Lemma C.6. *In the two cluster model, for $x \in X$ and $y \in Y$, the following holds:*

$$\mathbb{E}[\bar{E}_m^y] = \bar{X}(l - \varepsilon) \quad \text{and} \quad \mathbb{E}[\bar{E}_m^x] = \begin{cases} \bar{Y}(\varepsilon - l) & |X| \geq m \\ \bar{Y}(l - \varepsilon) & \text{otherwise} \end{cases}$$

Proof. Consider $\{L_{i \rightarrow k} - L_{j \rightarrow k}\}_k$. There are three cases:

1. If i, j are in the same cluster, then this is 0, regardless of which cluster k is in.
2. If i, j are in different clusters, and k is in i 's cluster, then this is $\varepsilon - l$.
3. If i, j are in different clusters, and k is in j 's cluster, then this is $l - \varepsilon$.

So, $E_m^{i-j} = 0$ if i and j are in the same cluster. Otherwise, if i 's cluster's size is at least m , then $E_m^{i-j} = \varepsilon - l$; if the cluster is smaller, then $E_m^{i-j} = l - \varepsilon$.

Now, let $x \in X$ be arbitrary. If $|X| \geq m$:

$$n\bar{E}_m^x = \sum_j E_m^{x-j} = \sum_{j \in X} 0 + \sum_{j \in Y} \varepsilon - l = |Y|(\varepsilon - l)$$

If $|X| < m$, then $\bar{E}_m^x = \bar{Y}(l - \varepsilon)$.

Since $|X| \geq |Y|$ and $m > 0.5$, an arbitrary node $y \in Y$ only has one case: $\bar{E}_m^y = \bar{X}(l - \varepsilon)$. $\square$

Theorem C.2 (cluster, σ^*). *In the cluster model with deterministic latencies with $|X| > c = m$, the advantage is:*

$$\text{ADVANTAGE}(\sigma_{\text{early}}) = \frac{2dl + \frac{\mu_0}{k\mu} - \frac{k-1}{k}(2\bar{X}\varepsilon + (1 - \bar{X})(\varepsilon + l)) - (1 - \bar{X})(\varepsilon - l)}{2dl + \frac{\mu_0}{k\mu} - \frac{k-1}{k}(\bar{X}(\varepsilon + l) + 2(1 - \bar{X})l) - \bar{X}(l - \varepsilon)}$$

Proof. First, note that the slowest start time is simply $2l$; thus the timeout $\tau = 2ld$. Since $B(\tau) = 0$, we get that $b_0 - b \cdot 2dl = 0$ or $b_0 = 2bdl$. Further, let $b = k\mu$, so $b_0 = 2k\mu dl$.

Recall that

$$\begin{aligned}
\frac{U_x}{U_y} &= \frac{(b_0 + \mu_0) - (b - \mu)\mathbb{E}[\bar{S}^x] - b\mathbb{E}[\bar{E}_m^x]}{(b_0 + \mu_0) - (b - \mu)\mathbb{E}[\bar{S}^y] - b\mathbb{E}[\bar{E}_m^y]} \\
&= \frac{2k\mu dl + \mu_0 - (k-1)\mu\mathbb{E}[\bar{S}^x] - k\mu\mathbb{E}[\bar{E}_m^x]}{2k\mu dl + \mu_0 - (k-1)\mu\mathbb{E}[\bar{S}^y] - k\mu\mathbb{E}[\bar{E}_m^y]} \\
&= \frac{2dl + \frac{\mu_0}{k\mu} - \frac{k-1}{k}\mathbb{E}[\bar{S}^x] - \mathbb{E}[\bar{E}_m^x]}{2dl + \frac{\mu_0}{k\mu} - \frac{k-1}{k}\mathbb{E}[\bar{S}^y] - \mathbb{E}[\bar{E}_m^y]}
\end{aligned}$$

We know apply [Lemmas C.5](#) and [C.6](#) using the fact that $\bar{Y} = 1 - \bar{X}$. The denominator is always $2dl + \frac{\mu_0}{k\mu} - \frac{k-1}{k}(\bar{X}(\varepsilon + l) + 2(1 - \bar{X})l) - \bar{X}(l - \varepsilon)$. For the numerator, it depends on whether the relative magnitudes of $|X|$, c , and m . Using the assumption that $|X| > c, m$, the ratio becomes:

$$\frac{U_x}{U_y} = \frac{2dl + \frac{\mu_0}{k\mu} - \frac{k-1}{k}(2\bar{X}\varepsilon + (1 - \bar{X})(\varepsilon + l)) - (1 - \bar{X})(\varepsilon - l)}{2dl + \frac{\mu_0}{k\mu} - \frac{k-1}{k}(\bar{X}(\varepsilon + l) + 2(1 - \bar{X})l) - \bar{X}(l - \varepsilon)}$$

□

C.3 Delayed equilibria as a baseline

The following lemma will be useful for evaluating the utility ratios in the late-proposing profile.

Lemma C.7. *For the late strategy profile $\sigma_{late}^{\vec{\sigma}}$,*

$$\frac{U_i(\sigma_{late}^{\vec{\sigma}})}{U_j(\sigma_{late}^{\vec{\sigma}})} = \frac{(b_0 + \mu_0) + \mu\tau - \mu\mathbb{E}[\bar{E}_c^i]}{(b_0 + \mu_0) + \mu\tau - \mu\mathbb{E}[\bar{E}_c^j]} \quad \text{where} \quad \bar{E}_c^i = \frac{1}{n} \sum_{k \in [n]} E_c^{i-k}$$

Proof. First consider the utility ratio of two validators i and j . Since we are in the late-proposing case, every validator will propose at Δ^* . Since i is elected leader preceded by some j with $1/n^2$ probability, the expected reward is $\mathbb{E}[\frac{1}{n^2} \sum_j M(\tau - E_c^{i-j}) + \frac{1}{n} b_0] = \frac{1}{n}(b_0 + \mu_0 + \mu\tau) - \mu \frac{1}{n^2} \sum_j \mathbb{E}[E_c^{i-j}]$. Hence,

$$\frac{U_i(\sigma_{late}^{\vec{\sigma}})}{U_j(\sigma_{late}^{\vec{\sigma}})} = \frac{(b_0 + \mu_0) + \mu\tau - \mu\mathbb{E}[\bar{E}_c^i]}{(b_0 + \mu_0) + \mu\tau - \mu\mathbb{E}[\bar{E}_c^j]}$$

□

Late proposing on a line. We now look at the line model.

Lemma C.8. *For the late strategy profile $\sigma_{late}^{\vec{\sigma}}$, the advantage in the deterministic line model is given by the following:*

$$\text{ADVANTAGE}(\sigma_{late}^{\vec{\sigma}}) = \frac{\frac{b_0 + \mu_0}{\mu n} + d - (\frac{1}{36} - \frac{2}{3n})}{\frac{b_0 + \mu_0}{\mu n} + d - (\frac{7}{18} - \frac{5}{6n})}$$

Proof. In the line model, since latencies are proportional to distance, the worst case ratio is then the ratio between the player in the middle and the player on the edge. Note that the longest start

time is $S^{0,n-1} = n - 1 \approx n$. Hence, $\tau = dn$. We also note that as $m = c$ we can use [Lemma C.4](#). Now, using [Lemma C.7](#),

$$\begin{aligned} \frac{(b_0 + \mu_0) + \mu\tau - \mu\left(\frac{n}{36} - \frac{2}{3}\right)}{(b_0 + \mu_0) + \mu\tau - \mu\left(\frac{7n}{18} - \frac{5}{6}\right)} &= \frac{\frac{b_0 + \mu_0}{n} + \mu d - \mu\left(\frac{1}{36} - \frac{2}{3n}\right)}{\frac{b_0 + \mu_0}{n} + \mu d - \mu\left(\frac{7}{18} - \frac{5}{6n}\right)} \\ &= \frac{\frac{b_0 + \mu_0}{\mu n} + d - \left(\frac{1}{36} - \frac{2}{3n}\right)}{\frac{b_0 + \mu_0}{\mu n} + d - \left(\frac{7}{18} - \frac{5}{6n}\right)} \end{aligned}$$

□

Late proposing in clusters. Now consider two clusters.

Lemma C.9. *For the late strategy profile $\sigma_{late}^{\rightarrow}$ with $|X| > c$, the advantage in the deterministic cluster model is given by the following:*

$$ADVANTAGE(\sigma_{late}^{\rightarrow}) = \frac{(b_0 + \mu_0) + 2dl\mu - \mu((1 - \bar{X})(\varepsilon - l))}{(b_0 + \mu_0) + 2dl\mu - \mu(\bar{X}(l - \varepsilon))}$$

Proof. In the cluster model, we know that $|X| > |Y|$ and therefore $x \in X$ will have utility that always dominates $y \in Y$. Note that the slowest start time is also $2l$ so $\tau = 2dl$ and we calculated $\mathbb{E}[\bar{E}_c^i]$ in [Lemma C.6](#). Using [Lemma C.7](#),

$$\frac{U_x}{U_y} = \frac{(b_0 + \mu_0) + 2dl\mu - \mu\mathbb{E}[\bar{E}_c^x]}{(b_0 + \mu_0) + 2dl\mu - \mu\mathbb{E}[\bar{E}_c^y]} = \frac{(b_0 + \mu_0) + 2dl\mu - \mu\mathbb{E}[\bar{E}_c^x]}{(b_0 + \mu_0) + 2dl\mu - \mu\bar{X}(l - \varepsilon)}$$

The value of $\mathbb{E}[\bar{E}_c^x]$ depends on the size of X . Since $|X| > c$, advantage is:

$$\frac{U_x}{U_y} = \frac{(b_0 + \mu_0) + 2dl\mu - \mu(\bar{Y}(\varepsilon - l))}{(b_0 + \mu_0) + 2dl\mu - \mu(\bar{X}(l - \varepsilon))}$$

□

Theorem C.3 (line and cluster, σ_{late}). *In the deterministic line model, the unfairness can be quantified as,*

$$ADVANTAGE(\sigma_{late}^{\rightarrow}) = \frac{\frac{b_0 + \mu_0}{\mu n} + d - \left(\frac{1}{36} - \frac{2}{3n}\right)}{\frac{b_0 + \mu_0}{\mu n} + d - \left(\frac{7}{18} - \frac{5}{6n}\right)}$$

and in the deterministic cluster model with $|X| > c$, the unfairness can be quantified as,

$$ADVANTAGE(\sigma_{late}^{\rightarrow}) = \frac{(b_0 + \mu_0) + 2dl\mu - \mu((1 - \bar{X})(\varepsilon - l))}{(b_0 + \mu_0) + 2dl\mu - \mu(\bar{X}(l - \varepsilon))}$$

Proof. The result follows directly from [Lemmas C.8](#) and [C.9](#).

□




Figure 5: Advantage comparison between early and late strategy profiles. Plotted expressions are from Theorems C.1 to C.3. The parameters are set such that $\mu_0 = 0.005$, $b_0 = 0.038$, $d = 5$, $\mu = 6 \times 10^{-6}$, $k = 1.5$, $\varepsilon = 0.1$, $\bar{X} = 0.67$.

D Simulation details

World latency simulation. For the world latency model, we use a lognormal where the underlying normal is distributed with mean 1 and std 0.5 within cities. Between cities, we fix a std of 0.8 and the log of the table below as the mean.¹²

Table 3: List of cities for the world map simulation.

Amsterdam	Dublin	Frankfurt	Los Angeles	New York
Paris	Seoul	Sydney	Vienna	Warsaw
Zurich				

Table 4: Inter-city latency (ms)

	Amsterdam	Dublin	Frankfurt	Los Angeles	New York	Paris
Amsterdam	0.00 ms	18.90 ms	8.76 ms	151.91 ms	77.25 ms	35.70 ms
Dublin	18.93 ms	0.00 ms	22.57 ms	148.11 ms	79.36 ms	17.04 ms
Frankfurt	8.79 ms	22.44 ms	0.00 ms	157.84 ms	83.33 ms	15.54 ms
Los Angeles	151.80 ms	148.00 ms	157.76 ms	0.00 ms	69.00 ms	144.77 ms
New York	77.23 ms	79.84 ms	83.57 ms	68.95 ms	0.00 ms	72.74 ms
Paris	35.49 ms	17.69 ms	15.34 ms	144.80 ms	72.78 ms	0.00 ms
Seoul	216.87 ms	248.61 ms	222.59 ms	134.22 ms	202.45 ms	310.96 ms
Sydney	252.22 ms	258.32 ms	267.89 ms	160.50 ms	200.35 ms	242.46 ms
Vienna	18.63 ms	33.62 ms	24.06 ms	168.07 ms	94.42 ms	38.41 ms
Warsaw	22.94 ms	43.37 ms	22.77 ms	171.28 ms	95.17 ms	33.04 ms
Zurich	16.29 ms	35.95 ms	7.00 ms	159.11 ms	104.17 ms	18.16 ms

¹²Taken from <https://wondernetwork.com/pings/> on 13-08-2025.

Table 5: Inter-city latency (ms)

	Seoul	Sydney	Vienna	Warsaw	Zurich	w_i
Amsterdam	216.45 ms	252.29 ms	18.93 ms	23.05 ms	16.31 ms	18
Dublin	248.55 ms	259.37 ms	33.80 ms	43.35 ms	35.98 ms	17
Frankfurt	222.78 ms	267.78 ms	24.02 ms	22.53 ms	7.01 ms	17
Los Angeles	134.15 ms	160.57 ms	168.40 ms	171.43 ms	159.03 ms	8
New York	202.48 ms	200.36 ms	95.40 ms	95.10 ms	104.10 ms	8
Paris	309.10 ms	242.42 ms	38.39 ms	33.00 ms	18.15 ms	11
Seoul	0.00 ms	137.45 ms	235.10 ms	237.83 ms	259.89 ms	5
Sydney	137.55 ms	0.00 ms	260.27 ms	275.89 ms	276.44 ms	6
Vienna	235.01 ms	260.58 ms	0.00 ms	27.63 ms	19.38 ms	7
Warsaw	237.68 ms	275.93 ms	27.68 ms	0.00 ms	23.27 ms	3
Zurich	259.77 ms	276.52 ms	19.42 ms	23.30 ms	0.00 ms	15

The weight indicates how many nodes are in that city out of 100.

E Inflation and griefing incentives

In our standard model, the block reward is *printed* by the protocol, and thus is a form of inflation. Thus, there is an incentive to lower others’ rewards, as that increases the value of one’s own rewards. We refer to this as a *griefing* incentive. Here, we argue that this incentive would not significantly change the results.

First, for the Nash equilibrium without coalitions, note that a validator can lower other validator’s rewards, on average, by always timing out (or voting infinity). This would lower the leader’s reward if and only if the dishonest validator was one of the m earlier validators to receive the leader’s proposal. However, with large n , any one validator would have a very small impact on the aggregate measurement – the leader would receive the $(m + 1)$ -smallest vote instead of the m -smallest vote. So, adding inflation to our model would result in an ε -BNE here, where ε would be very small.

Second, small coalitions would have a larger ability to affect others’ rewards. If they vote 0 on coalition leaders and vote ∞ for non-coalition leaders, they would lower the honest leaders’ reward from the m th smallest vote to the $(m + |C|)$ th smallest vote (among honest validators). This would shift rewards from non-coalition members to coalition members. However, the optimal proposal strategy, for both coalition and non-coalition leaders, is still to propose early – this follows directly from the time-decreasing property. So, responsiveness is still intact.

Third, large coalitions would be able to completely censor non-coalition leaders. While this is problematic, we note that the same problem occurs in existing system (i.e., with static block rewards). It would always be in a large coalition’s interest to deny rewards to non-coalition members, whether they are dynamic or static.

Finally, the fairness ratios would be unchanged, as the net inflation would devalue all profits by the same multiplicative factor.

F Alternative incentive knob: changing leader election

We now consider varying the other incentive knob: the leader schedule. We think that this alternative might be of practical interest for ease of implementation.

Incentive knobs. As we have discussed so far, once the timeliness of a protocol is measured, there are two aspects of the protocol that can be used by the mechanism designer to set the incentives. Both of these result in a rich incentive structure that we analyze separately in this paper.

- *Block reward:* The block reward can be a decreasing function of the proposal duration. Intuitively, if the initial reward is high enough compared to MEV gain, the incentives shift to early proposal.
- *Leader probabilities:* We explore designing a leader election schedule such that for each player, the later the proposal is, the next leader election probability is proportionally reduced. The result is a direct reduction in future rewards with the goal of incentivizing early proposals.

Leader weights. Each player now has a weight variable w_i . Intuitively, our goal is to reduce the weight of player i proportional to how late it was to propose in its previous proposal. Initially, the weights are equal to 1 so that each player is selected with probability $1/n$ (Note that our model normalizes by stake, and consider coalitions to account for stake distributions). After each proposal, we update the leader's weight as a function of its measured proposal duration using a weight function $\mathcal{W} : [0, \tau] \rightarrow [0, 1]$. Then, the election probability for each i is $w_i / \sum_j w_j$.

Stake transfer. We do not explicitly model the mechanisms that allow exiting and entering the system. Clearly, if a staked entity is reduced to 0 weight, it is not permanent as they are able to withdraw their stake and re-stake at a new keypair to start over. In practice, implementations could choose to start new stake at lower weight and have a delay on withdrawals in order to prevent incentives to propose late and avoid punishment by withdrawing.

Choosing the weight function. While the design space for selecting the weight function is quite large, we would like the weight function to not introduce strong incentives to run more validators. Intuitively, it shouldn't matter if you are running a single validator with more stake compared to lots of validators with the minimum amount of stake. Since our model considers each entity as a single unit of stake anyway, we capture this notion by designing the weight function to apply an *additive* effect on previous weight. Hence, no matter how large your group is as a coalition, for an action Δ , you get the same effect of $+\mathcal{W}(\Delta)$, effectively reducing the consideration to total weight.

The timing game with leader decay. The changes from the timing game in [Section 3](#) are in bold.

timing game with leader decay

1. Let R be the current round, and suppose i is the current leader and j led the previous round. Let t_{R-1} represent the time when round $R-1$ ended, i.e., when j broadcasted their block. Let s^{R-1} denote the time it takes between when j broadcasts their block and when i receives the block and its quorum. i learns their maximum delay Δ_R^* and chooses a delay $\Delta \in [0, \Delta_R^*]$.
2. Let k be an arbitrary validator.
 - (a) The latency $l_{i \rightarrow k} \sim L_{i \rightarrow k}$ is sampled.
 - (b) k receives the block at time $t_k^R = t_{R-1} + s^{R-1} + \Delta + l_{i \rightarrow k}$.
 - (c) k votes according to their voting strategy, which is a function of the block times they have seen. So $v_k^R = \sigma_k^V(\mathbf{t}_k^{1:R}) \in \mathbb{R}^+$. k then broadcasts^a their vote.^b
3. Let v_m^* refer to the m^{th} smallest vote in the (multi)set $\{v_k^R\}_k$. i receives utility $M(\Delta + s^{R-1}) + B(v_m^*)$, where M denotes the MEV and B the block reward.

4. **Update the weight of the leader by $w_i := w_i + \mathcal{W}(v_m^*)$.**

5. The next leader is randomly sampled **from the distribution $\{w_j / \sum_k w_k\}_j$** and the game repeats from 1.

The first round of this game is slightly different. **Initialize weights $w_i = 1$ for all $i \in [n]$.** The game starts at time $t = 0$, and so $s^0 = 0$, i.e., the leader i can propose as soon as the game starts. As a result, validator k receives the block at time $t_k^1 = \Delta + l_{i \rightarrow k}$. Otherwise, the first round matches the others.

^aDepending on the protocol, k may know the next leader and send their vote only to them, but this detail is not relevant in our model.


^bAs discussed earlier, in practice, one would decouple consensus from the duration voting, by only voting “yes”, “no”, or “timeout” in the critical path, and submitting duration votes on-chain for later aggregation. For the purposes of this model, we do not make this distinction.

Let σ^* denote the strategy profile, in which players report honestly, each leader i proposes with $\Delta = 0$. Similar to last time, we refer to σ^* as the *early-proposing truthful* strategy profile.


F.1 Equilibria under binary state

As an instructive example, consider a step function as leader weight that oscillates weights between high weight and low weight. Nodes never go below $1 - \rho$, and above 1. If the step function is simply $\pm\rho$ we get a two state Markov process to analyze. Formally, let $\mathcal{T}$ be a target time such that, intuitively, most quorums can be formed before $\mathcal{T}$. Let ρ be weight function parameter such that if people propose before $\mathcal{T}$, we apply $+\rho$, and anything later than $\mathcal{T}$, we apply $-\rho$. We call this the *binary weight* model.

We will begin by analyzing the early-proposing truthful strategy σ^* . A key probability due to latencies is the probability that the reported delay is above $\mathcal{T}$, hence resulting in a penalty. For player i , we will represent this probability with $p_i := \Pr[S^{j,i} + E_m^{i-j} > \mathcal{T}]$.



(a) Step function $\mathcal{W}(\Delta) = \rho \operatorname{sgn}(\mathcal{T} - \Delta)$.



(b) Possible weights of player i .

Figure 6: Binary weight model.

We first observe that once we fix the strategy to σ^* , the stationary probabilities of any player i having weight w_i equal to $1 - \rho$ or 1 is independent of the latency distributions of other players. For analysis, consider the following equivalent formulation and the induced Markov chain M^{σ^*} with two states, 1 and $1 - \rho$. In each state, first uniformly randomly pick some player. If this player is i , toss a coin with probability w_i . If i survives this coin, toss a last coin with probability p_i to determine whether to transition to $1 - \rho$ or 1. On all other cases, stay in the same state, essentially skipping i .

Lemma F.1. *In the binary weight model with players adopting σ^* , the stationary distribution of weights in M^{σ^*} for any player i is the following:*

$$\nu_{1-\rho} = \frac{p_i}{1 - \rho(1 - p_i)} \quad \text{and} \quad \nu_1 = \frac{(1 - \rho)(1 - p_i)}{1 - \rho(1 - p_i)}$$

Proof. Observe that the probability of going from 1 to $1 - \rho$ requires i to be selected with probability proportional to its weight 1 and having high latency. Similarly, from $1 - \rho$ to 1, i gets elected with probability proportional to $1 - \rho$ and has low latency. Therefore, we can compute the stationary probabilities by:

$$\nu_{1-\rho} = \frac{p_i}{p_i + (1 - \rho)(1 - p_i)} = \frac{p_i}{1 - \rho(1 - p_i)}$$

and

$$\nu_1 = \frac{(1 - \rho)(1 - p_i)}{p_i + (1 - \rho)(1 - p_i)} = \frac{(1 - \rho)(1 - p_i)}{1 - \rho(1 - p_i)}$$

Note that we omit the factor of $1/n$ as it cancels out. $\square$

Unfortunately, however, there exists an explicit incentive to deviate from σ^* . Consider the scenario in which player j is leaders and during aggregation the latencies were borderline—slightly aggregating to below $\mathcal{T}$ with i having the power to push it above $\mathcal{T}$. In this case, if player i lies, j 's weight decreases, increasing i 's probability of winning and hence utility.

Low latency. We do however recover equilibria if i never has the potential to cast a deciding vote. Consider the scenario in which even the $(m + 1)^{\text{th}}$ vote is always guaranteed to be less than $\mathcal{T}$. If no reported duration is ever above $\mathcal{T}$, then the players always have weight 1.

Corollary F.1. *If $\Pr[S^{j,i} + E_{m+1}^{i-j} > \mathcal{T}] = 0$ for all i and j and $m \geq n - c$, then each player i has the following utility:*

$$U_i(\sigma^*) = \frac{\sum_j \mathbb{E}[M(S^{j,i})] + \mathbb{E}\left[B\left(S^{j,i} + E_m^{i-j}\right)\right]}{\sum_{i,j} \mathbb{E}[S^{j,i}]}$$

Moreover, if $\langle M, B \rangle$ are time-decreasing, σ^* is a subgame-perfect Bayes-Nash equilibria.

Proof. Observe that if the assumption holds, Lemma F.1 implies that all players stay at weight 1, even if some player i votes dishonestly, as at best it can influence the $(m + 1)^{\text{th}}$ vote but not the m^{th} . Hence, we recover the case in Theorem 4.1. $\square$

This section highlights the fact that the equilibria structure after introducing decaying leader election probabilities depends highly on the latency distributions and how the decay mechanism is designed.

G List of symbols

n	number of players (or units of stake)
$[n]$	set of all players $\{0, 1, \dots, n - 1\}$
Δ	proposal delay
$M(x)$	MEV reward function $\mu x + \mu_0$
$B(x)$	block reward function $b_0 - bx$

μ_0	MEV intercept
μ	MEV slope
b_0	block reward intercept
b	block reward slope
τ	timeout duration
Δ^*	maximum delay
$\vec{\sigma}$	strategy profile
$U_i(\vec{\sigma})$	time averaged utility of player i with strategy profile σ
$L_{i \rightarrow j}$	latency from i to j
m	timeliness measurement threshold
c	quorum threshold
Δ	strategic proposal delay of the leader (in addition to latencies)
$Q_c^{j \rightarrow i}$	quorum time of leader i when previous leader is j (c^{th} order statistic of $\{L_{j \rightarrow k} + L_{k \rightarrow i}\}_k$)
$\overline{Q}_c^i$	average quorum time of leader i ($\frac{1}{n} \sum_{k \in [n]} Q_c^{i \rightarrow k}$)
$S^{j,i}$	earliest proposal duration (starting time) of leader i when previous leader is j (the maximum of quorum duration and block duration $\max\{Q_c^{j \rightarrow i}, L_{j \rightarrow i}\}$)
$\overline{S}^i$	average starting time of leader i ($\frac{1}{n} \sum_{j \in [n]} S^{j,i}$)
E_x^{i-j}	honest measurement error of leader i when previous leader is j (x^{th} order statistic of $\{L_{i \rightarrow k} - L_{j \rightarrow k}\}_k$)
$\overline{E}_x^i$	average honest measurement error of leader i ($\frac{1}{n} \sum_{j \in [n]} E_x^{i-j}$)