Designing Incentives for Responsive Consensus Protocols

Mahimna Kelkar
Columbia University

Ertem Nusret Tas
a16z Crypto Research

Maryam Bahrani
Ritual

Tim Roughgarden
Columbia University

Abstract

Modern consensus protocols often aspire to be *responsive*—that is, to confirm transactions in time proportional to the actual network delays as opposed to a (typically much larger) worst-case bound on network delays. Responsiveness can yield substantial practical improvements in both protocol latency and throughput.

In blockchain settings, however, block proposers commonly have economic incentives (most notably MEV) to delay their blocks, a phenomenon repeatedly observed in practice, e.g., in the block production supply chains for Ethereum and Solana. Such incentives clash with responsiveness in existing protocol designs.

This paper develops a rigorous framework for designing *incentive-compatible* responsive consensus protocols. We first consider the canonical case of single-leader protocols and establish feasibility results characterizing the conditions on the distribution of network latency under which responsiveness can be incentivized through a suitable reward function. We further quantify the amount of stake required from the leader to deter dishonest delays.

We then show strong positive results for protocols with *multiple leaders*, demonstrating that the multi-leader approach is fundamentally superior to single-leader designs for resolving the tension between responsiveness and incentive-compatibility: by forcing leaders to compete for rewards, much of the burden otherwise placed on the incentive mechanism is alleviated. Notably, this results in simpler reward mechanisms with no stake requirements that also remain feasible in parameter regimes which are provably impossible under a single leader.

1 Introduction

Modern consensus protocols, including those underlying blockchain systems, increasingly aim to provide the property of *(optimistic) responsiveness* [6, 21, 31, 41]. Roughly speaking, a protocol is *responsive* if it can confirm transactions in time proportional to the *actual* network delay instead of a worst-case bound. Typically, these protocols are responsive under optimistic or "happy path" conditions (e.g., when the leader does not equivocate). In real network settings, the actual delay is typically much smaller on average than the worst-case bound—this means that by being able to adapt to network conditions, responsive protocols can yield practical improvements in both protocol latency and throughput. It is not surprising that responsive designs underpin layer-1 blockchains like Aptos, Monad, and Sui, and an upcoming upgrade on Solana.

While responsiveness has received significant attention in the consensus literature, minimal work has been done to understand the *incentives* behind it. This is despite common acknowledgment that for blockchain systems in particular, significant monetary incentives can lead to *timing games* [11, 35] where block proposals are delayed on purpose, resulting in a potential loss of responsiveness.

The core goal of this work is to formally understand the *incentives* underlying the property of responsiveness.

Timing games. In common consensus designs, a *leader* (or proposer) node collects transactions arriving since the last block to include in its own proposal. As a result, waiting longer to make its proposal can allow the leader to increase its revenue by including e.g., more or higher-paying transactions.

A similar, and arguably more prominent effect, arises in the context of maximal extractable value (MEV) [7, 12]—a quantity that captures profit resulting from a proposer's absolute power to decide on the transactions and their relative ordering within its block. MEV includes sources like arbitrage [42], frontrunning [12], and sandwich attacks [32], and in recent years, has increasingly become the major source of revenue for block proposers [29, 39]. MEV has been shown, both in theory and practice, to increase with the time elapsed since the last block [17, 29]. Moreover, this increase is typically superlinear [17, 26, 39], resulting in strong incentives to delay proposals, even at the risk of protocol timeouts and penalties. Unlike transaction fees which are visible within the block, the exact value of MEV is difficult for the protocol to determine, making it more challenging to prevent proposal delays due to MEV incentives.

Unsurprisingly, *timing games*, i.e., strategically delaying block proposals for monetary gain, have been observed on blockchains like Ethereum and Solana [11, 29]. Timing games are undesirable for protocol performance and stability—they result in higher latency and lower throughput for users, and more instability from the risk of missed blocks. Moreover, for blockchains that use responsive designs, the effect is more profound since it not only invalidates responsiveness, but may lead to slower block times than if the protocol had a fixed block time in the first place.

Proposed mitigations. Several techniques, many ad-hoc in nature, have been proposed to deter timing games. As a general theme, these involve some mechanism to penalize the leader when a delayed proposal is detected by the protocol—a detailed comparison is given in Section 1.2. The most concrete design among these penalty-based approaches was put forth and analyzed in very recent work by Alpturer et al. [2]. However, we find that implementing their design requires strong assumptions and significant knowledge on part of the mechanism designer. This includes knowing how the MEV increases with time, the exact distribution of peer-to-peer network latency among consensus nodes, and assuming limited malicious behavior among parties tasked with detecting delayed proposals. As a result, major hurdles exist for practical deployment.

Our approach: multiple competing leaders. We employ a different approach to deter timing games—*multiple competing leaders*. The intuition is simple: instead of designing a mechanism to detect and penalize a single leader's delayed proposal, we make multiple leaders *compete* with each other for rewards. Then, if one leader

delays its proposal to capture more MEV, it may lose out to another who does not delay and therefore has its proposal confirmed. The incentives faced by leaders in this design resemble those faced by players in the Prisoner's Dilemma, but with the game-theoretic forces pushing toward the optimal outcome ("propose early") rather than the worse outcome ("delay") from the system's perspective.

While multi-leader designs have been considered previously for other goals like performance, censorship-resistance, and privacy, our work shows that they also have significant advantages in this new context of responsiveness. In particular, we design mechanisms to deter timing games even when MEV is unbounded and with minimal to no knowledge of the network latency. Our mechanism is also simpler—it requires no sophisticated detection of delays, no payments, and no stake held from leaders. We emphasize that this is in stark contrast with single-leader designs, for which we show that these components are necessary.

1.1 Our Results

In this work, we develop a framework for analyzing the incentives underlying responsiveness in consensus protocols. Broadly, we design *responsive* mechanisms, whose goal is to incentivize leaders to take the honest action of not delaying their proposals. While we focus on responsive consensus for exposition, we note that our results apply more broadly to deter delays, allowing us to capture even traditionally non-responsive protocols such as Ethereum.

1.1.1 Two-Choice Model (Section 3). We begin with a simplified model that still captures the qualitative nature of our general results. As a starting point, consider a single leader for a given protocol round and suppose that it only has two choices for its action—it can either propose "early" (i.e., with no delay) or "late" (i.e., delay to capture more MEV). Our goal as the mechanism designer is to incentivize early proposal despite exogenous incentives (like MEV) to delay. Later, we also motivate our two-choice model by showing that it captures common practical scenarios (Section 4.2).

Formally, a leader's utility $\mathbb{U}$ is the sum of three components—observable transaction fees (T), unobservable MEV (M), and reward/penalty (R) from the mechanism. R is the only lever available to the mechanism designer to incentivize responsiveness (i.e., early proposal) by ensuring that $\mathbb{U}(\text{early}) > \mathbb{U}(\text{late})$. We give our main results in this model below:

Fee revenue does not affect responsiveness. We find that transaction fees do not affect responsiveness. Intuitively, since fees are *observable*, they can always be burned. As a consequence, early proposal can be incentivized even when the leader makes arbitrarily larger fee revenue from proposing late.

Responsiveness in the presence of MEV requires *detection*. On the other hand, even a slight presence of MEV, since it is unobservable, creates problems for responsiveness. In a setting where the protocol cannot detect the leader's chosen action, Theorem 3.5 shows that even with arbitrarily small MEV, no mechanism can incentivize responsiveness. This already elucidates a significant barrier in the context of asynchronous consensus protocols that cannot rely on network-based detection of late proposals.

Even when the leader's action can be fully detected, Theorem 3.4 shows that a responsive reward mechanism needs a budget of at least $\mathsf{M}(\text{late}) - \mathsf{M}(\text{early})$, i.e., the increase in MEV for late proposal. We model $\mathsf{M}(\mathbf{a})$ for action $\mathbf{a}$ as arising from an unknown distribution (but restricted to a specific class), but state it here as a concrete function for clarity.

Leveraging multiple leaders (Section 3.4). We show how a multi-leader design allows us to not only circumvent the previous impossibility, but also do so with zero budget! The core idea is to make leaders *race* with each other—only the earliest leader's proposal receives any utility from the block. In practice, this will be the leader whose proposal gets e.g., enough attestations earliest and is confirmed by the consensus protocol.

Within our simple two-choice model, suppose that a leader who plays early will always have its proposal chosen ahead of another leader who plays late. Then, under a suitable MEV bound, we show that for all leaders, proposing early is the *strictly dominant strategy*—i.e., the strongest-possible justification for choosing a given strategy. Notably, our mechanism needs no budget for rewarding or penalizing leaders, and makes no assumptions on being able to detect leader actions.

We take advantage of a prisoner's dilemma structure in our favor—while leaders would prefer to coordinate to the "all late" state where they would receive higher MEV, the incentives encourage each leader to deviate and propose early.

Concretely, we show that our mechanism works as long as the ratio $\mathsf{M}(\text{late})/\mathsf{M}(\text{early})$ is bounded by the number of leaders K. This also means that it is always possible to make early the strictly dominant strategy simply by increasing the number of leaders.

1.1.2 General Model (Sections 4 and 5). We extend our results to a more general model where the leader can choose to delay its block by any $\Delta \geq 0$. Our goal is to design mechanisms to incentivize honest behavior, i.e., $\Delta = 0$, or at least $\Delta \leq \delta$ for some small δ.

A budget lower bound for single-leader mechanisms (Section 4.3). We prove a general lower bound (Theorem 4.3) on the budget required for incentivizing responsiveness with a single leader. This lower bound increases with larger MEV and decreases as a function of how well the protocol can detect the leader's chosen delay (as measured by the total variation distance (TVD) between the observation distributions induced by a chosen delay of 0 vs. Δ).

Concrete responsive mechanisms (Section 4.4). We construct several concrete mechanisms to incentivize responsiveness under different assumptions about the distribution Ψ of network latency. First, we show that if only the mean of Ψ is known, no responsive mechanism with finite budget exists. Next, assuming bounded variance of Ψ, we construct a concrete 0-responsive mechanism (i.e., incentivizing $\Delta = 0$). Finally, when the exact distribution of Ψ is known, we construct a 0-responsive mechanism whose budget matches our lower bound—this shows its tightness.

As part of our mechanism, we use techniques that network observers (e.g., other consensus nodes) can use to detect the leader's behavior. We show how to make this detection robust to any minority fraction of observers colluding with the leader.

Multi-leader designs (Section 5). We design a responsive multi-leader mechanism as follows: first, each leader i proposes its block with a chosen delay $\Delta_i \geq 0$. Only the leader whose proposal is confirmed by consensus the earliest (e.g., after some network latency $\psi_i \sim \Psi$) obtains any utility. We find that, under appropriate

MEV bounds similar to the two-choice model, this simple approach incentivizes all leaders to play the honest strategy of $\Delta_i = 0$. While there is no strictly dominant strategy in this game, we show that all leaders choosing $\Delta_i = 0$ is the unique pure Nash equilibrium. Further, unlike with single-leader designs, our mechanism requires no sophisticated detection techniques, and works with zero budget.

1.2 Related Work

Responsive consensus. Responsiveness has been explored since the early days of consensus design [5], leading to a plethora of BFT protocols that target it as a first-order property [30, 31, 41]. Most modern consensus protocols are responsive, including those underlying popular blockchains—like Solana's Alpenglow [25] and Sui's Mysticeti [6]. Despite this, minimal work exists on understanding the incentives behind responsiveness.

Incentives and blockchain. Since blockchain applications often involves significant monetary aspects, aligning incentives with honest behavior is a major part of practical protocol design. This has been an active area of research for several years, with analysis on e.g., selfish-mining in Bitcoin [15, 33], manipulation of randomness beacons and block timestamps [3, 40], centralization [27], and rational BFT [4, 9, 20] to name a few. In the past few years, the emergence of MEV as a major source of revenue has introduced new tensions in aligning validator incentives to the honest action required for protocol security and performance [16, 29]. Timing games are a prime example of this, and we design mechanisms to deter them in this work.

Mitigations for timing games. A common theme among proposed mitigation techniques is to detect a block proposal's "timeliness," that is, by how long it was delayed by the leader. Typically, this is done through a quorum of committee members reporting the time they received the leader's block [2, 34, 38]. Reports are then aggregated by the protocol into a single timeliness measure. [34] and [38] propose to decrease the leader's share of transaction fees (for the current block, and potentially also in the future), but neither makes formal incentive-based arguments. Moreover, as we show, MEV rather than observable transaction fees is ultimately responsible for timing games.

A different strategy, albeit without formal analysis, considered by [22] is to elect leaders for a fixed *time* interval rather than a fixed number of blocks. The intuition is to incentivize them to propose as many blocks as possible during their allotted window. However, in practice, protocol-level block rewards are usually smaller than MEV revenue (e.g., in 59% of Ethereum blocks [29]). Further, as identified by [2], such a mitigation strategy does not work for prominent MEV sources like CEX-DEX arbitrage which scale super-linearly with time [26, 39]. In other words, there are still strong incentives for a leader to delay its block proposals.

Comparison with [2]. Recent work by Alpturer et al. [2] is the one closest to ours. [2] uses a timeliness measure of the leader's block to set a *dynamic* reward that decays with time. Abstractly, the block reward R is a function of the timeliness measure that decreases faster than the growth of MEV (with time). Their work formally analyzes incentives to show that the leader will play honestly since any delay hurts its revenue.

Our core design of using multiple competing leaders is fundamentally different and offers substantial advantages over the single-leader design in [2]. By making leaders compete with each other, we forgo any requirement for detection through timeliness, as well as the need to dynamically set block rewards. This not only makes our mechanism simpler but also means that we require no additional budget to reward or penalize leaders. In contrast, we show that for *any* single-leader design such as [2], it is necessary for the mechanism's budget to be larger than the MEV.

Further, even within a single-leader setting, we provide new analysis with more general results. Some examples of this are: (i) we model MEV as coming from an unknown distribution (within a given class) instead of a fixed known function; (ii) we do not restrict the leader's delay to a known $\Delta_{\max}$, allowing us to cleanly capture partially synchronous and asynchronous consensus; (iii) we provide a generic lower-bound on the budget required to incentivize responsiveness; (iv) we do not require complete knowledge of the network latency distribution, and instead construct responsive mechanisms as a function of the available distributional information (e.g., only bounds on the mean and variance).

2 Preliminaries

Notation. For $K \in \mathbb{N}$, let $[K] := \{1, \ldots, K\}$. For a random variable X, $\widetilde{X}$ denotes its expectation whenever it is well-defined, and $\mathrm{supp}(X)$ its support. Vectors are denoted as $\vec{x} = (x_1, \ldots, x_K)$, and for $i \in [K]$, $\vec{x}_{-i}$ denotes the vector obtained by removing the i-th coordinate, whereas $\vec{0}$ denotes the all-zero vector of the appropriate dimension. Unless stated otherwise, expectations are taken over all exogenous randomness and any randomization in the protocol or in players' strategies.

For probability distributions $\mathcal{D}_1$ and $\mathcal{D}_2$ with pdfs f_1 and f_2, we use the following definition of total variation distance:

$$\mathrm{TVD}(\mathcal{D}_1, \mathcal{D}_2) = \int_{f_1 > f_2} f_1(x) - f_2(x)\, dx.$$

Game-theoretic notions. Actions are denoted by $\mathbf{a}$. Strategy $\mathbf{s}_i$ of a player i is a distribution over the support of its actions. Pure strategies are identified with the actions they place unit mass on. A strategy profile $\vec{\mathbf{s}}$ is a vector of strategies $(\mathbf{s}_1, \ldots, \mathbf{s}_n)$ of n players. Utility of a player i given a strategy profile $\vec{\mathbf{s}}$ is denoted by $U_i(\vec{\mathbf{s}})$. A strategy profile $\vec{\mathbf{s}}$ is a *Nash equilibrium* if $\mathbb{E}\left[U_i(\mathbf{s}_i, \vec{\mathbf{s}}_{-i})\right] \geq \mathbb{E}\left[U_i(\mathbf{s}'_i, \vec{\mathbf{s}}_{-i})\right]$ for every $i \in [n]$ and every alternative strategy $\mathbf{s}'_i$. A strategy $\mathbf{s}_i$ is *strictly dominant* if $\mathbb{E}\left[U_i(\mathbf{s}_i, \vec{\mathbf{s}}_{-i})\right] > \mathbb{E}\left[U_i(\mathbf{s}'_i, \vec{\mathbf{s}}_{-i})\right]$ for every $\mathbf{s}'_i \neq \mathbf{s}_i$ and every $\vec{\mathbf{s}}_{-i}$; *weak dominance* is defined by replacing $>$ with $\geq$ and requiring strict inequality for some $\vec{\mathbf{s}}_{-i}$. We say that $\mathbf{s}_i$ is (strictly) *dominated* by $\mathbf{s}'_i$ when $\mathbf{s}'_i$ (strictly) dominates $\mathbf{s}_i$. For strategy profiles, we say that $\vec{\mathbf{s}}$ *dominates* $\vec{\mathbf{s}}'$ if $\mathbb{E}\left[U_i(\vec{\mathbf{s}})\right] \geq \mathbb{E}\left[U_i(\vec{\mathbf{s}}')\right]$ for all i, with strict inequality for at least one i.

3 The Basic Two-Choice Model

We start by considering a leader-based setting where a single leader (or *block proposer*) is in charge of proposing the current block. Here, the leader will only have two choices for its action: (1) It can propose the block *early* (the honest action); or (2) It can propose the block *late* to potentially obtain additional utility. Concretely,

the leader takes an action $\mathbf{a} \in \{\text{early}, \text{late}\}$. Broadly, our goal (as the mechanism designer) is to incentivize early proposal.

Our purpose with this simple two-choice model is to develop intuition for our general results. Looking ahead, we find realistic cases in the general model that effectively reduce to the two-choice model. Further, we emphasize that even this simple model serves as a useful proxy for guiding practical protocol incentive design.

This section investigates the conditions under which it is possible to incentivize the leader to propose early. Further, even in this simple model, we demonstrate our powerful observation of the benefit of having multiple *competing leaders*—we show feasibility even under conditions where it is impossible with a single leader. As alluded to earlier, the intuition here is that we can make multiple leaders compete with each other instead of needing to design a mechanism that itself has to compete with a single leader.

3.1 Model Preliminaries

Sources of utility. We find it useful to separate the different sources of utility received by the leader as a result of its action $\mathbf{a}$. Informally, we model the total utility $\mathbb{U}^{\text{leader}}$ as the sum

$$\mathbb{U}^{\text{leader}}(\mathbf{a}) = \mathsf{T}(\mathbf{a}) + \mathsf{M}(\mathbf{a}) + \mathsf{R}$$

where the components are as follows:

- *Observable exogenous utility* $\mathsf{T}(\mathbf{a})$. The canonical example of this is *transaction fees*. In general, $\mathsf{T}(\mathbf{a})$ models any utility arising from factors outside the system (i.e., exogenous) that are still observable within the system (e.g., as part of the block description). The function T will be chosen by nature.
- *Unknown exogenous utility* $\mathsf{M}(\mathbf{a})$. M captures utility from exogenous sources that are unknown within the system. Typical examples of this are miner/maximal extractable value (MEV) [7, 12], side payments, or other external factors that may not appear within the block itself. M is also chosen by nature. In this paper, we often use MEV as a proxy when discussing the utility M.
- *Endogenous reward mechanism* R. The function $\mathsf{R}(\cdot)$ models the reward given by the system for producing a block and is the only component of the leader's utility that can be set by the mechanism designer. More details on how R can be chosen and what it can depend on are given later in this subsection.

Distributions for T and M. We denote the distributions from which the functions T and M are sampled, by $\mathcal{T}$ and $\mathcal{M}$ respectively. We further use the notation $\mathsf{T} := (t_0, t_1) \sim \mathcal{T}$ where $t_0 = \mathsf{T}(\text{early})$ and $t_1 = \mathsf{T}(\text{late})$; we use (m_0, m_1) similarly as the notation for $\mathsf{M} := (\mathsf{M}(\text{early}), \mathsf{M}(\text{late}))$. $\widetilde{\mathsf{T}}$ and $\widetilde{\mathsf{M}}$ denote their respective expected values.

To capture the intuition that proposing late will result in more fees and MEV in isolation, we model $t_1 \geq t_0$ and $m_1 \geq m_0$. Formally, let $\mathcal{G}_+$ be the set of tuples $(x_0, x_1) \in \mathbb{R}^2$ with $0 \leq x_0 \leq x_1$[1]. Then, $\mathcal{T}$ and $\mathcal{M}$ can be defined as distributions over $\mathcal{G}_+$.

An important point is that in general, these distributions will not be known to the mechanism designer or the leader. This captures the typical folklore understanding that modeling the exact distribution of MEV is both challenging and fickle. Instead, looking ahead, for our model, we will define a setting Γ as a set of possible "worlds" (e.g., set of tuples $(\mathcal{T}, \mathcal{M})$) that will be known to all parties. Nature will choose any world among these.

[1]To prevent the degenerate case where the leader has no additional utility for proposing late, and therefore can only be weakly incentivized to propose early, we will only consider tuples $(t_0, m_0) \neq (t_1, m_1)$.

Detectability of the leader's action. We model the possibility that the mechanism can detect whether the leader played early or late and execute R accordingly. We call these *detectable* settings, where the leader's action $\mathbf{a}$ can be fully detected—here, R is allowed to depend on $\mathbf{a}$. On the other hand, for *undetectable* settings, R will get no information about $\mathbf{a}$.

For simplicity, we only consider these two types here and leave open exactly how detection takes place within the protocol. Later, in Section 4, for our general model, we will consider a variety of partial notions of detectability and further discuss how detection can be done in practice based on factors such as network delays.

Setting. Formally, we define a setting Γ as $(\mathcal{D}, \mathsf{d})$ where $\mathcal{D}$ is a set of tuples of distributions $(\mathcal{T}, \mathcal{M})$ and $\mathsf{d} \in \{\text{detectable}, \text{undetectable}\}$. In essence, Γ defines what distributions the utility could come from, as well as whether the leader's action can be detected. Intuitively, each $w = (\mathcal{T}, \mathcal{M}, \mathsf{d})$ denotes a "possible world." We often drop d from the tuple since it is the same for all worlds in Γ.

The setting Γ will be provided to the mechanism designer and the leader, and nature will choose one possible world that is consistent with Γ. Since the designer does not know the exact utility distributions, given Γ, its task will be to set a reward function R that incentivizes early proposal for *all* worlds consistent with Γ.

Reward function parameters. In general, we allow R to depend only on parameters observable in the system—i.e., R can take as input the realized value of the observable exogenous utility $t = \mathsf{T}(\mathbf{a})$. To simplify the formalism, we allow the designer to provide two functions $\mathsf{R}_{\text{early}}(t)$ and $\mathsf{R}_{\text{late}}(t)$. These can be distinct only in detectable settings; for undetectable settings, they must be the same since R cannot observe the action $\mathbf{a}$.

We permit R to be negative subject to budget constraints. This can model the mechanism holding a deposit from the leader, which could be slashed. We provide more details in Section 3.2.

Remark 3.1 (Two-choice model as a proxy for practice). In addition to the two-choice model capturing the qualitative nature of our general model results, it may itself be sufficient to capture some practical settings. For example, we might only care about preventing the leader from delaying by too much in a way that affects e.g., protocol stability or liveness. Further, in practice, it is often easy to distinguish "small" delays (i.e., playing early) from "large" ones (i.e., playing late). Our two-choice model captures this application well. We expand on this "threshold detection" setting in Section 4.2.

3.2 Formal Game Description

We are now ready to describe our game, which proceeds as follows: First, nature provides the setting Γ to the mechanism designer and the leader and chooses any $(\mathcal{T}, \mathcal{M})$ consistent with Γ. Next, given Γ but not $(\mathcal{T}, \mathcal{M})$, the designer chooses the endogenous reward $\mathsf{R} = (\mathsf{R}_{\text{early}}, \mathsf{R}_{\text{late}})$. If Γ is undetectable, then $\mathsf{R}_{\text{early}} = \mathsf{R}_{\text{late}}$ must hold. Now, given Γ and R, the leader chooses its action $\mathbf{a} \in \{\text{early}, \text{late}\}$.

Finally, the exogenous utilities $\mathsf{T} = (t_0, t_1)$ and $\mathsf{M} = (m_0, m_1)$ are sampled from $\mathcal{T}$ and $\mathcal{M}$ respectively, and the leader receives utility

$$\mathbb{U}^{\text{leader}}_{(\mathsf{T},\mathsf{M})}(\mathbf{a}) = \begin{cases} t_0 + m_0 + \mathsf{R}_{\text{early}}(t_0) & \text{if } \mathbf{a} = \text{early} \\ t_1 + m_1 + \mathsf{R}_{\text{late}}(t_1) & \text{if } \mathbf{a} = \text{late} \end{cases}$$

In turn, the expected utility for an action $\mathbf{a}$ in a given world $w = (\mathcal{T}, \mathcal{M})$ can be given by

$$\mathbb{EU}^{\text{leader}}_{w}(\mathbf{a}) = \mathbb{E}_{\mathsf{T}\sim\mathcal{T};\, \mathsf{M}\sim\mathcal{M}} \left[\mathbb{U}^{\text{leader}}_{(\mathsf{T},\mathsf{M})}(\mathbf{a})\right].$$

Responsive reward. We now define a *responsive* reward R as one that incentivizes the leader to propose early. We consider two variations—a standard notion that considers the expected utility, and a *point-wise* or *always* responsive version that needs to hold over all (T, M) in the support.

Definition 3.2 (Responsive R)**.** A reward function R is *responsive* w.r.t. setting Γ if for all $w = (\mathcal{T}, \mathcal{M})$ consistent with Γ, it holds that:

- $\mathbb{EU}^{\text{leader}}_{w}(\text{early}) > 0$;
- $\mathbb{EU}^{\text{leader}}_{w}(\text{early}) > \mathbb{EU}^{\text{leader}}_{w}(\text{late})$.

The first condition ensures that the leader wants to propose a block instead of staying silent. The second condition incentivizes early proposal over late proposal.

We also consider a stronger notion—*always* or *point-wise* responsiveness where the conditions need to hold for all realizations of exogenous utility. Concretely, R is always responsive if for all worlds $(\mathcal{T}, \mathcal{M})$ and all (T, M) in the support, it holds that (i) $\mathbb{U}^{\text{leader}}_{(\mathsf{T},\mathsf{M})}(\text{early}) > 0$; and (ii) $\mathbb{U}^{\text{leader}}_{(\mathsf{T},\mathsf{M})}(\text{early}) > \mathbb{U}^{\text{leader}}_{(\mathsf{T},\mathsf{M})}(\text{late})$. Whenever possible, we will try to achieve this within our results.

An always responsive R is particularly useful when the leader may have more knowledge about the world than the mechanism. For example, such an R incentivizes early proposal even when the leader knows the exact exogenous realization (T, M) but the mechanism does not.

Budget. Notice that adding a constant c to a reward R will not affect responsiveness. It is therefore useful to compare different responsive R's in terms of their *budget*. Intuitively, the budget of R, denoted by $\text{budget}(\mathsf{R})$, is the maximum amount the mechanism will need to pay out of pocket. An equivalent way to think about this when R is negative is as the "stake" held from the leader which could be slashed as needed[2].

An easy bound is $\text{budget}(\mathsf{R}) \leq |r_{\text{neg}}| + r_{\text{pos}}$ where r_{neg} is the most negative value and r_{pos} is the largest positive value of R. Note however that the real budget required may be lower since a portion of any negative R could be taken from the other utility parts—as an example, if $\mathsf{R}(t) = -t$, this arguably does not consume any "budget" since it only deducts e.g., the observable fees that are already received. Observe that a tight bound on the budget will depend on the setting Γ. One way to formalize this is to "shift" R by adding the smallest constant such that the resultant utility is always non-negative. Then, the budget of R will be the maximum reward given out. We formalize this intuition below:

[2]For repeated games in practice, it is easier to hold a single stake to deter against dishonest behavior than to reward honest behavior every round.

Definition 3.3 (Budget)**.** Let $c \geq 0$ be the minimum value such that R^* defined by $\mathsf{R}^*_{\mathbf{a}}(t) = \mathsf{R}_{\mathbf{a}}(t) + c$ results in $\mathbb{EU}^{\text{leader}}_{w}(\mathbf{a}) \geq 0$ for all $\mathbf{a}, w$. Then, $\text{budget}(\mathsf{R}) = \sup_{\mathbf{a},t} \mathsf{R}^*_{\mathbf{a}}(t)$.

When considering always responsiveness, the above needs to hold for all realizations, i.e., $\mathbb{U}^{\text{leader}}_{(\mathsf{T},\mathsf{M})}(\mathbf{a}) \geq 0$ for all (T, M).

When the setting Γ does not constrain $\mathcal{T}$ in any way (in particular, it allows transaction fees to be zero in some world), $\text{budget}(\mathsf{R})$ is exactly equal to the previous easy bound.

3.3 Feasibility Results for a Single Leader

Detectable settings. For our first result, we look at settings where the leader's action can be detected. We show feasibility here but *only when* there is a known upper bound on $\mathsf{M}(\text{late}) - \mathsf{M}(\text{early})$. The intuition for this is simple—if the leader receives arbitrarily large unknown MEV from proposing late, there is no finite penalty that the mechanism can apply to prevent the leader from doing so.

Concretely, given a setting Γ, define m^+_Γ to be the supremum of $\mathsf{M}(\text{late}) - \mathsf{M}(\text{early})$ over all $(\mathcal{T}, \mathcal{M})$ consistent with Γ and all M in the support of $\mathcal{M}$. Given a finite m^+_Γ, the idea to achieving an always responsive R is to set a heavy penalty for late proposal that is larger than any additional MEV received. Concretely, we can set $\mathsf{R}_{\text{early}}(t) = -t + \epsilon$ and $\mathsf{R}_{\text{late}}(t) = -t - m^+_\Gamma$ for any small $\epsilon > 0$. It is not hard to see that for such an R, due to the high penalty in R_{late}, a leader will always be worse off choosing late. Note that there is no restriction needed on $\mathcal{T}$ since its realization is observable.

Theorem 3.4 below characterizes *all* reward functions that are always responsive in the detectable setting. Further, it shows that the minimal budget of such an R is m^+_Γ.

Theorem 3.4. *Consider a detectable setting Γ with finite m^+_Γ. Then, $\mathsf{R} = (\mathsf{R}_{\text{early}}, \mathsf{R}_{\text{late}})$ is always responsive iff. the following hold:*

- $\mathsf{R}_{\text{early}}(t) > -t \quad$ *for all t;*
- $\mathsf{R}_{\text{late}}(t) < \inf_{x \leq t}(\mathsf{R}_{\text{early}}(x) + x) - t - m^+_\Gamma \quad$ *for all t.*

Over all such responsive R*, the infimum of the required budget is m^+_Γ. As a corollary, if m^+_Γ is infinite (i.e., MEV from proposing* late *is unbounded), no* R *can be always responsive.*

The full proof is deferred to Appendix A.1. An analogous result on the budget also holds for standard responsiveness (over expected utility). In particular, the minimum budget required here is $\widetilde{m}^+_\Gamma$, i.e., the supremum over Γ of the expected difference $\widetilde{\mathsf{M}}(\text{late}) - \widetilde{\mathsf{M}}(\text{early})$.

Undetectable settings. Unsurprisingly, incentivizing early proposal is a lot more challenging in settings where the leader's action cannot be detected at all. Still, in a world without any unknown utility (i.e., $\mathsf{M}(\mathbf{a}) = 0$), we construct responsive R (Theorem A.1 in Appendix A.1). The intuition here is that the known utility itself can be used as a proxy for early/late detection—a larger $t = \mathsf{T}(\mathbf{a})$ signals a higher likelihood that late was chosen. In turn, we will be able to incentivize early proposal by ensuring that R decreases in t.

When there is no MEV, the budget requirement on R can be arbitrarily small—concretely, for any $\epsilon > 0$, there exists a feasible R with $\text{budget}(\mathsf{R}) = \epsilon$.

On the other hand, in the presence of any MEV (i.e., $\mathcal{M} \neq 0$), we show in Theorem 3.5 below that it is impossible to incentivize proposing early. Notably, this holds even given upper-bounds on all

distributions. We use the notation $\mathcal{M} < \epsilon$ to denote that $m_0, m_1 < \epsilon$ for any (m_0, m_1) in its support. $\mathcal{T} < \epsilon$ is defined similarly.

Theorem 3.5. *Given any $\epsilon > 0$, there exists an undetectable setting Γ where every $(\mathcal{T}, \mathcal{M})$ consistent with it satisfies $\mathcal{T} < \epsilon$ and $\mathcal{M} < \epsilon$ and there is no responsive reward (even in expectation) w.r.t. Γ.*

Proof. We define the undetectable setting Γ that intuitively contains all worlds with point-mass distributions less than ϵ as follows: First, consider any $\mathsf{T} = (t_0, t_1)$ and $\mathsf{M} = (m_0, m_1)$ with $0 < t_0 < t_1 < \epsilon$ and $0 < m_0 < m_1 < \epsilon$. Next, define the world $w_{(\mathsf{T},\mathsf{M})}$ in which the distributions $\mathcal{T}$ and $\mathcal{M}$ are point distributions at (T, M). The setting Γ contains all such worlds $w_{(\mathsf{T},\mathsf{M})}$. The vision for considering such point-wise distributions is to effectively reduce the problem to showing that there is no always responsive R.

Suppose now for contradiction that some R is responsive w.r.t. Γ. Since the distributions are point masses, for any $w_{(\mathsf{T},\mathsf{M})} \in \Gamma$, and action $\mathbf{a}$, we have

$$\mathbb{E}\mathbb{U}_w^{\text{leader}}(\mathbf{a}) = \mathbb{U}_{(\mathsf{T},\mathsf{M})}^{\text{leader}}(\mathbf{a}).$$

The responsiveness condition therefore collapses to the point-wise condition below:

$$t_0 + m_0 + \mathsf{R}(t_0) > t_1 + m_1 + \mathsf{R}(t_1).$$

Now, define $g(t) = \mathsf{R}(t) + t$ and rewrite the above inequality as

$$g(t_0) > (m_1 - m_0) + g(t_1).$$

Note that this needs to hold for all worlds i.e., all (t_0, t_1, m_0, m_1). Further, it is easy to see that g must be strictly decreasing since the inequality needs to hold for arbitrarily small $(m_1 - m_0)$. Such a g can be discontinuous only at countably many points.

Now, choose arbitrary $t_0' < \epsilon/2$ where g is continuous. This means that we can find positive τ, ν smaller than $\epsilon/2$ such that

$$g(t_0') - g(t_0' + \tau) < \nu.$$

Now, set $t_1' = t_0' + \tau$, and any (m_0', m_1') with $\nu < m_1' - m_0' < \epsilon$. Consider $\mathsf{T}' = (t_0', t_1')$ and $\mathsf{M}' = (m_0', m_1')$. Observe that,

$$\begin{aligned}\mathbb{U}_{(\mathsf{T}',\mathsf{M}')}^{\text{leader}}(\text{late}) - \mathbb{U}_{(\mathsf{T}',\mathsf{M}')}^{\text{leader}}(\text{early}) &= m_1' + g(t_0' + \tau) - m_0' - g(t_0') \\ &= (m_1' - m_0') + (g(t_0' + \tau) - g(t_0')) \\ &> \nu - \nu = 0.\end{aligned}$$

In other words, R does not incentivize playing early in the world $w_{(\mathsf{T}',\mathsf{M}')}$ which contradicts the assumption that R is responsive. □

Fee burning. From the results above, observe that any responsive R (in both the detectable and undetectable settings) includes a $-t$ term. Consequently, as a corollary, it is not hard to see that a responsive R with minimum budget essentially "burns" all observable utility (e.g., transaction fees). This is perhaps not surprising and is in fact a core reason why the same R can work without any restriction on the distribution $\mathcal{T}$ (including when it is unbounded).

3.4 Leveraging Multiple Leaders

So far, in a sense, our goal has been to design a mechanism to "police" the leader's behavior of delaying its block. As is evident, this requires non-trivial knowledge of exogenous leader utility, and the ability to detect the leader's action.

A powerful insight is that in multi-leader protocols, we can effectively make the leaders compete with each other for their reward. As an illustration, rewarding only the first received proposal incentivizes the leaders to not delay their proposal—otherwise, they risk losing to other leaders.

In this section, we design responsive rewards for multi-leader settings, even in cases that are infeasible with a single leader. As an added benefit, we show significant other advantages of using multiple leaders—on simplicity of the mechanism, budget requirements, and robustness.

Multiple leader preliminaries. Since early proposal can be incentivized in detectable settings even with a single leader, here we will focus on *undetectable* settings to demonstrate the advantage of using multiple leaders.

First, we observe that there are several ways that consensus protocols may use multiple leaders. As one example, each leader may propose its own block with only one of them being selected. Another possibility is that multiple leaders propose but the final block is jointly decided by all proposers (e.g., as the union of all transactions). For the purpose of this section, we will focus on protocols where one block proposal is fully selected. In our general multi-leader model in Section 5, we discuss other approaches.

Multi-leader model. Our model extends the basic one from Section 3.2 to support multiple leaders; we only detail the changes here. First, we model K competing leaders that will each propose a block.

Actions and strategies. Let $\mathbf{a}_i \in \{\text{early}, \text{late}\}$ denote the action taken by the i-th leader; denote the combined action vector by $\vec{\mathbf{a}} = (\mathbf{a}_1, \ldots, \mathbf{a}_K)$. We also allow leaders to randomize between the two choices. In particular, define a *strategy* $\mathbf{s} = \mathbf{s}^{(\phi)}$ as a tuple $(\phi, 1-\phi)$ where $\mathbf{a} = \text{early}$ is played with probability ϕ and $\mathbf{a} = \text{late}$ is played with probability $1 - \phi$.

We use $\mathbf{s}_i$ to denote leader i's strategy and $\vec{\mathbf{s}} = (\mathbf{s}_1, \ldots, \mathbf{s}_K)$ to denote the combined strategy profile. $\mathbf{s}^{(\mathbf{a})}$ denotes the pure strategy of (a single leader) playing $\mathbf{a}$ and $\vec{\mathbf{s}}^{(\mathbf{a})}$ denotes the corresponding pure strategy profile where all leaders play $\mathbf{a}$.

Probability of selection. Given an action vector $\vec{\mathbf{a}}$, we use $p_i(\vec{\mathbf{a}})$ to denote the probability that i's proposal is selected (e.g., by being the first proposal to receive enough votes). Note that p_i may depend on factors specific to the chosen consensus protocol as well as network latencies. Without loss of generality, we have $\sum_i p_i(\vec{\mathbf{a}}) = 1$. For our multi-leader setting, we expand Γ to also include $\vec{p} = (p_1, \ldots, p_K)$.

Responsive reward. R now represents the reward given to the leader whose proposal is selected; other leaders receive no utility. The goal of a responsive R in the multi-leader setting is to incentivize all leaders to play early, i.e., the conditions in Definition 3.2 need to hold individually for all leaders.

A nicer way to capture this goal is through standard game-theoretic equilibrium notions. Concretely, a minimal requirement will be for the pure strategy $\mathbf{s}^{(\text{early})}$ to be a Nash equilibrium—i.e., no leader is incentivized to individually deviate from this state. A stronger goal is for early proposal to be the *dominant* strategy—that is, regardless of what others do, it is always best for a leader i to play $\mathbf{a}_i = \text{early}$. Looking ahead, our result in Theorem 3.7 will show a reward R that makes early proposal the dominant strategy.

Leader competition. To provide a flavor of what kind of results we can expect through leader competition, we define a simple competitive setting below within our two-choice model.

Definition 3.6 (Simple Competitive Setting)**.** We say that a multi-leader setting Γ is (simple) *competitive* if the following holds:

- $p_i(\vec{\mathbf{a}}) = 0$ if $\mathbf{a}_i = \text{late}$ and $\exists j$ with $\mathbf{a}_j = \text{early}$.
- $p_i(\vec{\mathbf{a}}) = 1/E$ if $\mathbf{a}_i = \text{early}$ and $\left|\{j \mid \mathbf{a}_j = \text{early}\}\right| = E$.
- $p_j(\vec{\mathbf{a}}) = 1/K$ if $\mathbf{a}_j = \text{late}$ for all $j \in \{1, \ldots, K\}$.

In the simple competitive setting, a leader who plays late will never be chosen ahead of another leader playing early. Leaders playing early will be chosen with equal probability. In effect, this captures the simplest framing of leader competing to be "first."

We will define a broader notion of leader competition within our general multi-leader model in Section 5. Still, we emphasize that the insights from the above simple competitive setting carry over to more general settings and already illustrate the potency of a multi-leader approach.

Also note that a competitive setting can still be undetectable—the mechanism has no information about what actions were chosen by the leaders *including the winner*. It only learns whose proposal was picked as a result of the consensus protocol. Note that even if i was chosen, this does not guarantee that i played early—it could have been because all leaders played late.

Theorem 3.7 below shows that as long as the *ratio* between $\mathsf{M}(\text{late})$ and $\mathsf{M}(\text{early})$ is not too large, there exists a reward R that makes early the dominant strategy. In addition, we essentially find a *sharp phase transition*—there is a threshold below which, we achieve the best-possible game-theoretic notion for incentivizing early, and above which, playing early is arguably worse than playing late.

Following that, Section 3.4.1 provides a discussion on how the result illustrates advantages of multi-leader protocol designs.

Theorem 3.7. *Consider an undetectable and simple competitive* Γ *with* $K \geq 2$ *leaders. Define* $\widetilde{\mathsf{M}}_\Gamma$ *as the maximum ratio of the MEV for proposing* late *to the MEV for* early *over all worlds in* Γ*. That is,*

$$\widetilde{\mathsf{M}}_\Gamma = \sup_{\mathcal{M} \in \Gamma} \frac{\mathbb{E}_{\mathsf{M} \in \mathcal{M}}[\mathsf{M}(\text{late})]}{\mathbb{E}_{\mathsf{M} \in \mathcal{M}}[\mathsf{M}(\text{early})]}.$$

There exists a reward function R *with* $\text{budget}(\mathsf{R}) = 0$ *such that:*

- *If* $\widetilde{\mathsf{M}}_\Gamma < K$*, then for all leaders, playing* early *is a strictly dominant strategy.*
- *If* $\widetilde{\mathsf{M}}_\Gamma > K$ *while* $\vec{\mathbf{s}}^{(\text{early})}$ *(everyone plays* early*) is still a Nash equilibrium, it is dominated by* $\vec{\mathbf{s}}^{(\text{late})}$ *(everyone plays* late*).*

As a corollary, for a fixed Γ*, there exists a large enough* K^* *number of leaders, such that for each leader, playing* early *is again a strictly dominant strategy.*

Proof. Consider a reward $\mathsf{R}(t) = -t$ that is given only to the chosen leader. We will show that R incentivizes early proposal.

Our first goal is to show that when $\widetilde{\mathsf{M}}_\Gamma < K$, playing early is the strictly dominant strategy[3]. For this, first recall from standard game-theory results, that it is sufficient to only show dominance w.r.t. any pure strategy profile—if early is dominant w.r.t. late, then due to linearity of expectation, it is also dominant w.r.t. any mixing.

[3]The boundary $\widetilde{\mathsf{M}}_\Gamma = K$ results in playing early being at least weakly dominant.

Formally, consider a particular world with $\mathcal{M}$ in Γ. Consider $\vec{\mathbf{s}} = (\mathbf{s}_1, \ldots, \mathbf{s}_K)$ where all players play some pure strategy and i plays late. That is, $\mathbf{s}_{j \neq i} \in \{\mathbf{s}^{(\text{early})}, \mathbf{s}^{(\text{late})}\}$ and $\mathbf{s}_i = \mathbf{s}^{(\text{late})}$. We will show that it is strictly better for i to switch to $\mathbf{s}^{(\text{early})}$. Let $\vec{\mathbf{s}'}$ denote the strategy profile identical to $\vec{\mathbf{s}}$ except that i now plays $\mathbf{s}'_i = \mathbf{s}^{(\text{early})}$. Let $U_i^{\vec{\mathbf{s}}}$ and $U_i^{\vec{\mathbf{s}'}}$ denote the utility of i in $\vec{\mathbf{s}}$ and $\vec{\mathbf{s}'}$ respectively.

Suppose that E players play early in $\vec{\mathbf{s}}$. If $E \neq 0$, then i has 0 utility in $\vec{\mathbf{s}}$ since it plays late, i.e., $U_i^{\vec{\mathbf{s}}} = 0$. In this case, switching to playing early already gives it positive utility.

On the other hand, if $E = 0$, this means that all leaders are playing late. In this case,

$$U_i^{\vec{\mathbf{s}}} = \frac{1}{K} \cdot \mathbb{E}_{\mathsf{M} \in \mathcal{M}}[\mathsf{M}(\text{late})].$$

Now, if i switches to early as in $\vec{\mathbf{s}'}$, it will be the only one playing early. This means that i will receive utility

$$U_i^{\vec{\mathbf{s}'}} = \mathbb{E}_{\mathsf{M} \in \mathcal{M}}[\mathsf{M}(\text{early})].$$

Observe that i will prefer playing early exactly when $U_i^{\vec{\mathbf{s}'}} > U_i^{\vec{\mathbf{s}}}$. In other words, when $\widetilde{\mathsf{M}}_\Gamma < K$ holds, playing early is strictly dominant for i, and consequently for all leaders.

As a corollary, it is also easy to see that increasing the number of leaders to above the ratio $\widetilde{\mathsf{M}}_\Gamma$ will always allow us to recover the strict dominance of playing early.

Now, for the second part of the result, it is easy to see that $\vec{\mathbf{s}}^{(\text{early})}$ is always a Nash equilibrium regardless of $\widetilde{\mathsf{M}}_\Gamma$. This is because playing early has strictly positive utility while individually switching to late from $\vec{\mathbf{s}}^{(\text{early})}$ will always result in 0 utility.

When $\widetilde{\mathsf{M}}_\Gamma > K$, there is another equilibrium where everyone plays late that dominates the pure early equilibrium. The proof for this is a direct consequence of the $E = 0$ case from the first part. In particular, when all leaders play late in $\vec{\mathbf{s}}$ (i.e., $E = 0$), a leader does not want to individually switch to early (i.e., resulting in $\vec{\mathbf{s}'}$) exactly when $U_i^{\vec{\mathbf{s}}} > U_i^{\vec{\mathbf{s}'}}$. Over all worlds in Γ, the profile $\vec{\mathbf{s}}^{(\text{late})}$ is an equilibrium exactly when $\widetilde{\mathsf{M}}_\Gamma > K$. Note further in this case, the equilibrium $\vec{\mathbf{s}}^{(\text{late})}$ dominates $\vec{\mathbf{s}}^{(\text{early})}$. □

3.4.1 Discussion. Theorem 3.7 shows that leader competition allows incentivizing early proposal in undetectable settings even in the presence of MEV despite this not being feasible with a single leader. We now discuss several other advantages of the undetectable multi-leader setting to even the detectable single-leader setting.

Budget comparison. Observe that the MEV doesn't even need to be bounded in absolute terms—the only requirement is on the *ratio* between the MEV for proposing late vs early. Further, the budget requirement of the reward mechanism is zero! This itself is significant—in the single leader setting, even in the best-case scenario where the leader action can be detected, the budget required is proportional to the bound on the MEV.

Robustness to changes in MEV. In the single-leader setting, since the MEV bound is explicitly embedded into the mechanism R, handling any changes to the distributions over repeated runs will require explicit updates to the mechanism's penalty. In contrast, the mechanism in the multi-leader setting can remain unchanged.

Advantage with knowledgeable leaders. So far, we model the world chosen by nature to be hidden from both the leader(s) and the

mechanism. But what if the leader is more knowledgeable about the available MEV than the mechanism? In the extreme case, suppose that the leader knows not just the exact world, but also the realizations $\mathsf{M}(\text{early})$ and $\mathsf{M}(\text{late})$ before it make its decision. These are still hidden from the mechanism as before.

Observe that in the single-leader case, even with full detection, the mechanism's budget will need to depend on the *maximum* value of $\mathsf{M}(\text{late})-\mathsf{M}(\text{early})$ over all realizations instead of the expectation $\widetilde{\mathsf{M}}(\text{late}) - \widetilde{\mathsf{M}}(\text{early})$. This is because the leader's knowledge of the world forces a requirement on the mechanism to be point-wise (i.e., always) responsive. Note that this may result in a significantly larger budget requirement. In contrast, in the multi-leader setting, the same mechanism from Theorem 3.7 with 0 budget will suffice.

4 General Model with a Single Leader

We formalize our general model in this section. Here, the leader can now choose any arbitrary delay $\Delta \geq 0$ when proposing its block.

We formalize the model with a single leader here, and will later incorporate multiple leaders in Section 5.

Leader utility. As before, we model three components in the utility of the leader: the observable (T) and unknown (M) exogenous utilities chosen by nature, and the endogenous reward (R) set by the mechanism designer.

We model the exogenous utility to be non-decreasing when a larger Δ is chosen—this captures the intuition that in isolation, the leader is always incentivized to delay its proposal. To maintain generality, we can define T and M as being sampled from distributions $\mathcal{T}$ and $\mathcal{M}$ over the set $\mathcal{G}_+ = \{g : \mathbb{R}^{\geq 0} \to \mathbb{R}^{\geq 0} \mid g \text{ is non-decreasing}\}$.

Further, we use $\widetilde{\mathsf{T}}(\Delta) = \mathbb{E}[\mathsf{T}(\Delta)]$ and $\widetilde{\mathsf{M}}(\Delta) = \mathbb{E}[\mathsf{M}(\Delta)]$ to denote the expected exogenous utilities for a particular $(\mathcal{T}, \mathcal{M})$.

Note that similar to our two-slot model, these distributions will typically be *unknown* to the mechanism designer. We emphasize the practical nature of such a general model considering the challenge and inherent fickleness in estimating the distribution of MEV.

Now, if the leader chooses delay Δ, its expected utility can be computed as:

$$\mathbb{EU}^{\text{leader}}(\Delta) = p(\Delta) \cdot \mathbb{E}\left[\mathsf{T}(\Delta) + \mathsf{M}(\Delta) + \mathsf{R}(\mathsf{T}(\Delta), d)\right] \qquad (1)$$

where $\mathsf{T} \sim \mathcal{T}, \mathsf{M} \sim \mathcal{M}$, $p(\Delta)$ is the probability that the block is accepted by the protocol, and d denotes the "perceived" value of Δ (which generalizes our detectability notion). We discuss these parameters in more detail next.

Probability of block acceptance. For a delay $\Delta \geq 0$ chosen by the leader, we use $p(\Delta)$ to denote the probability that the leader's block is accepted. This models the intuition that if the leader waits for too long, the consensus protocol will move forward without waiting for the leader. Here, as opposed to simply modeling a maximum delay $\Delta_{\max}$, we allow for any $\Delta \geq 0$ and consider a more general notion using the probability $p(\Delta)$ for block acceptance. We emphasize that this generalization also enables us to capture responsiveness in partially synchronous and asynchronous protocols, for which a $\Delta_{\max}$ assumption often cannot be made. For our results, we will consider natural settings where e.g., p decreases with larger Δ.

Generalized detectability. Recall that in our two-choice model, detectability models whether the mechanism can detect a leader's action as early or late. We generalize this notion of "detecting Δ" as follows: Consider a parametric family of distributions $\{\chi_\Delta\}$ indexed by $\Delta \in \mathbb{R}^{\geq 0}$. Here, each χ_Δ is a probability distribution over $\mathbb{R}^{\geq 0}$, such that when the leader chooses delay Δ, the mechanism will observe some $d \sim \chi_\Delta$ instead as a "detection" of Δ. Typically, this will be a "noised" representation of Δ due to e.g., inherent randomness in network latency or other external factors.

The mechanism can now use this observation d to set its reward R. For instance, a natural approach is to use d to decide whether the leader proposed early (i.e., $\Delta = 0$) or not, and penalize appropriately. Abstractly, this will be the mechanism's "pull" to prevent a leader from simply choosing a larger Δ. Note that $\{\chi_\Delta\}$ will typically not be known in full to the designer.

We keep the description of χ_Δ generic in the model here to allow for more general feasibility results. Later, in Section 4.4, we show practical examples of how this $d \sim \chi_\Delta$ can be provided to the reward mechanism by the consensus protocol itself. Looking ahead, one example could be an observer (e.g., a consensus node) providing the time it received the leader's proposal to the mechanism. In effect, we discuss guidelines on consensus protocol design so that the reward mechanism can be provided an observation from a suitable χ_Δ for which responsiveness can be incentivized.

Reward function. We model the reward function $\mathsf{R}(\cdot,\cdot)$ as taking inputs (i) the realization of the known exogenous utility, i.e., $\mathsf{T}(\Delta)$, and (ii) An observation $d \sim \chi_\Delta$ dependent on the leader's chosen Δ.

As before, without loss of generality, we allow R to be negative (subject to budget constraints). Its budget, denoted by $\text{budget}(\mathsf{R})$ can be defined exactly as Definition 3.3 from our two-choice model.

Game description. As before, the mechanism designer will not know the utility distributions. Instead, it will be provided a setting Γ, which effectively denotes which tuples $(\mathcal{T}, \mathcal{M}, \{\chi_\Delta\}, p(\cdot))$ are possible. Nature picks one of these "worlds" w (unknown to others).

Given Γ, the designer will choose its reward function $\mathsf{R}(\cdot,\cdot)$. The leader then learns Γ, R and chooses $\Delta \geq 0$ by which to delay its proposal. Finally, values are sampled from the distributions in the world w chosen by nature, and the leader receives total utility as described earlier. We will consider the leader's expected utility (Eq. (1)) in our analysis.

Responsive R**.** We now define the properties we want from R. Specifically, we define a δ-responsive R as one where the leader is incentivized to choose a delay no larger than δ. In other words, this effectively leads to responsiveness with at most δ slack.

Definition 4.1 (δ-responsive R)**.** A reward function R is δ-responsive w.r.t. Γ if there exists $\Delta^* \leq \delta$ such that for all $(\mathcal{T}, \mathcal{M}, \{\chi_\Delta\}, p(\cdot))$ consistent with Γ, the following hold:

- $\mathbb{EU}^{\text{leader}}(\Delta^*) > 0$;
- $\mathbb{EU}^{\text{leader}}(\Delta^*) > \mathbb{EU}^{\text{leader}}(\Delta')$ for all $\Delta' > \delta$.

As a non-triviality condition, we will only consider δ such that $p(\delta) > 0$—otherwise, especially when $p(\cdot)$ is decaying (see Section 4.1), the leader is already incentivized to never delay by more than δ since its block will never be accepted.

Subsection structure. Having described the general model, we now provide an outline for the rest of this section. First, in Section 4.1, we detail practical assumptions that make our analysis cleaner. Next, in Section 4.2, we find cases where the general model effectively reduces to the simpler two-choice model. Section 4.3

shows when responsiveness is feasible and proves a general tight lower bound on the budget required based on $\{\chi_\Delta\}$. Finally, Section 4.4 shows how to interpret this lower-bound in practice and how $d \sim \chi_\Delta$ itself can be provided to the mechanism.

4.1 Practical Assumptions on Model Parameters

While, we keep the model as general as possible, for our feasibility results, it will be helpful to make certain assumptions, motivated by practice, on the model parameters.

Decaying $p(\Delta)$. We assume that the probability of block acceptance $p(\Delta)$ is non-increasing in Δ and that $\lim_{\Delta\to\infty} p(\Delta) = 0$. This models the realistic situation that a leader's block is usually less likely to be accepted the longer it waits. Further, without loss of generality, we can set $p(0) = 1$, i.e., a leader who immediately proposes will always have its block accepted. Of course, $p(\cdot)$ may remain at 1 until some Δ' before it starts to decrease.

Importantly, we do not assume that our mechanism has knowledge of $p(\cdot)$. This is equivalent to considering Γ that includes worlds with all possible p consistent with the above properties.

Bounded MEV. While we allow specific realizations M of MEV to be unbounded, we will need to assume that the expected MEV taking block acceptance into account is bounded. In particular, for each world $w = (\mathcal{T}, \mathcal{M}, \{\chi_\Delta\}, p(\cdot))$, define the function:

$$m^*_w(\Delta) = p(\Delta) \cdot \widetilde{\mathsf{M}}(\Delta)$$

and consider its maximum over all worlds as:

$$m^*_\Gamma(\Delta) = \sup_{w\in\Gamma} m^*_w(\Delta).$$

We assume that $m^*_\Gamma(\Delta)$ is upper-bounded for all Δ. We use $M^*_\Gamma = \sup_\Delta m^*_\Gamma(\Delta)$ to denote the bound. This could be satisfied in several ways—either each M is bounded, or M is unbounded but the expectation $\widetilde{\mathsf{M}}$ is bounded, or both are unbounded but $p \cdot \widetilde{\mathsf{M}}$ is bounded. This is a necessary condition for incentivizing responsiveness in our single-leader model (see Lemma A.2 in Appendix A.2). Later, in Section 5, we show how multi-leader designs get around this.

We note that the function m^* captures a more general framing of MEV-boundedness and so results will typically relate to it instead of $\widetilde{\mathsf{M}}$. As a result, going forward, we often refer to m^* as the MEV.

4.2 Settings Reducing to the Two-Choice Model

We first identify settings that, either in their model or their results, effectively reduce to the two-choice model from Section 3.

Perfect detection. First, define a *perfectly detectable* setting Γ in the general model to be one where $\chi_\Delta = \Delta$ is a constant—in other words, the mechanism can perfectly observe the delay Δ chosen by the leader. A result analogous to Theorem 3.4 from the two-choice model will apply here. The theorem below shows exactly when an R with budget B exists. The intuition is simple: when the difference between the MEV from choosing a "large" delay and the MEV from choosing a "small" delay is less than B, an R with budget B can incentivize the leader to choose the smaller delay.

Theorem 4.2. *Consider a perfectly detectable setting Γ, and define*

$$b_\delta = \sup_{w\in\Gamma} \left\{ \sup_{\delta'>\delta}(m^*_w(\delta')) - \sup_{\delta^*\le\delta}(m^*_w(\delta^*)) \right\}.$$

There exists a δ-responsive reward R with budget $B \ge 0$ if $b_\delta < B$.

Threshold detection. Often, while the exact delay Δ chosen may be hard to determine, it is easy to distinguish between a "small" delay ($\Delta \le \delta$) and a "large" delay ($\Delta > \delta + \phi$) for some slack parameter ϕ. This is because, intuitively, in most practical networks, a message sent sufficiently after another will almost never arrive earlier. We show that such a "threshold" detection setting effectively reduces to the two-choice model.

As a concrete instantiation, for a fixed δ, ϕ, suppose that for $\Delta \le \delta$, the protocol always thinks that the leader acted honestly—that is, $\chi_\Delta = 0$ or "early". On the other hand, when $\Delta > \delta + \phi$, the protocol always detects dishonest behavior—that is, $\chi_\Delta = \delta + \phi = L$, or "late". In the intermediate range, if the leader chooses $\delta < \Delta \le \delta + \phi$, the protocol provides partial detection—one example is a linear decay represented by $\Pr_{\chi_\Delta}[d = 0] = 1 - \frac{\Delta-\delta}{\phi}$ and $\Pr_{\chi_\Delta}[d = L] = \frac{\Delta-\delta}{\phi}$.

Our first goal is to incentivize the leader to choose a small delay $\Delta \le \delta$, since this corresponds to the best possible detection in our threshold setting. The intuition for achieving this is exactly like our two-choice model—heavily penalize the leader when it is detected as "late." Concretely, for given fixed δ and ϕ, let

$$B = \sup_{w\in\Gamma} \sup_{\Delta>\delta} \frac{m^*_w(\Delta) - m^*_w(\delta)}{\Pr_{d\sim\chi_\Delta}[d = L]}$$

Intuitively, this denotes the maximum utility a leader can make by delaying by $\Delta > \delta$ compared to $\Delta \le \delta$. In other words, by designing R to penalize the leader by B when its proposal is detected as "late", we can incentivize δ-responsiveness. The full proof is deferred to Theorem A.3 (Appendix A.2).

Undetectable setting. An undetectable setting is one where there are no additional constraints on Γ. Consequently, a mechanism that hopes to incentivize responsiveness must also work when χ_Δ provides no information about Δ—that is $\chi_{\Delta_1} = \chi_{\Delta_2}$ for all Δ_1, Δ_2. As expected based on our results in the two-choice model, with any MEV, there is no hope of responsiveness since there is no way for the mechanism to distinguish between the leader choosing Δ_1 vs Δ_2. Since there is a strong impossibility result even in the two-choice model, and the framing here does not provide any additional insights—we choose to omit the formal results for brevity.

Our focus for the rest of the section will therefore be on *partially* detectable settings. In essence, the goal is to identify the kinds of detection that are sufficient to incentivize early proposal.

4.3 Partially Detectable Settings

We now investigate settings where the family $\{\chi_\Delta\}$, from which the detection d is sampled, provides partial information on the leader's actual delay Δ. Our broad goal is to find the conditions on $\{\chi_\Delta\}$ that will allow for designing a responsive mechanism R, and how they affect the budget requirements of R.

Our main result in this section is a *general lower bound* on the budget required for a responsive R that depends on properties of $\{\chi_\Delta\}$—namely, its total variation distance (TVD). Note that this automatically captures feasibility—when the budget lower bound is infinite, no responsive R exists.

As alluded to earlier, properties on $\{\chi_\Delta\}$ are directly influenced by what is known about the distribution of network latency. In Section 4.4, we discuss details on exactly how this plays out.

We also show that our lower bound is *tight*—when the exact distribution of network latency is known, there exists a responsive R with a budget that matches the lower bound.

4.3.1 A general budget lower-bound. We now describe our general lower bound on the budget of a responsive R. We make two simplifications that make the description cleaner but do not affect the nature of the result. First, we focus only on 0-responsiveness but note that our bound will naturally extend to δ-responsiveness.

Second, we consider settings Γ that have no constraints on $\mathcal{T}$ (i.e., exogenous fees). In other words, the mechanism R we design will need to work for any $\mathcal{T}$. While this restriction is not necessary since we make it more challenging for the mechanism designer, it results in a cleaner formulation and budget bounds. Doing so also allows to consider only those R that "burn" all transaction fees.

We are now ready to state our budget lower bound (in Theorem 4.3). Note that the even in the best case, budget(R) cannot be smaller than the MEV gained by delaying dishonestly—this will be an important contrast to the multi-leader case where responsiveness can be incentivized even without any budget requirements!

Theorem 4.3. *Consider any setting* Γ *(with no restriction on* $\mathcal{T}$*) and suppose that* R *is* 0*-responsive w.r.t.* Γ*. Then, it holds that*

$$\text{budget}(\mathsf{R}) \geq \sup_{w \in \Gamma} \sup_{\Delta > 0} \frac{m^*_w(\Delta) - m^*_w(0)}{\text{TVD}(\chi^w_\Delta, \chi^w_0)}$$

Proof. Suppose that R is 0-responsive. Since Γ has no restrictions on $\mathcal{T}$, we can consider $R(t, d) = -t - P(d) + c$ where WLG, $P(d) \geq 0$ is the penalty and $c \geq 0$ is an additive constant (note that to achieve minimal budget, c will be arbitrarily small).

Now, for any world w and $\Delta > 0$, directly from the second responsiveness condition,

$$\begin{aligned} m^*_w(\Delta) - m^*_w(0) &\leq \mathbb{E}_{y \sim \chi^w_\Delta}[P(y)] - \mathbb{E}_{y \sim \chi^w_0}[P(y)] \\ &= \int P(y) \left[f^w_\Delta(y) - f^w_0(y) \right] dy \end{aligned}$$

where f^w_D denotes the p.d.f. for χ^w_D. Define $h(y) = f^w_\Delta(y) - f^w_0(y)$. For an upper-bound, it is sufficient to restrict to only the positive part of the integral, i.e., in the region H^+ where $h(y) > 0$. In this case, since $P(y) \leq B$, we get:

$$\int P(y) \cdot h(y)\, dy \leq \int_{H^+} B \cdot h(y)\, dy$$

Observe that the definition of total variation distance directly gives:

$$\text{TVD}(\chi^w_\Delta, \chi^w_0) = \int_{H^+} h(y) dy$$

since $h(y)$ is the difference between the density functions of χ^w_Δ and χ^w_0. Note that H^+ is non-empty when $\chi^w_\Delta \neq \chi^w_0$ which is without loss of generality assuming Γ is not undetectable.

Putting it together, we have

$$\frac{m^*_w(\Delta) - m^*_w(0)}{\text{TVD}(\chi^w_\Delta, \chi^w_0)} \leq B.$$

Since R needs to be 0-responsive over all of Γ, taking the supremum over all $\Delta > 0$ and all worlds w gives us the desired result. □

4.4 Practical instantiations of $d \sim \chi_\Delta$

We now discuss how to interpret our lower-bound result in practice and also show concrete mechanisms to achieve responsiveness. In particular, we show what kinds of $\{\chi_\Delta\}$ will arise in practice given specific information about the distribution of network latency, and how this affects both the feasibility of responsiveness and the budget required by our mechanism. We further show that our lower bound is tight when the exact network distribution is known.

The core idea is to have observers (e.g., other consensus nodes) report on when the leader's proposal was received so that the mechanism can appropriately apply penalties. Section 4.4.1 considers a single honest observer, while Section 4.4.2 shows how to simulate an honest observer using multiple maliciously colluding or rational utility-seeking observers.

4.4.1 A single observer. Suppose that an observer in the system (e.g., another consensus node) receives the leader's block proposal and reports on the time. Assume for now that the observer is honest or at least weakly rational (i.e., reports honestly if it receives no benefit from dishonest reporting). Intuitively, our mechanism can use this report for its "detection" of the leader's delay Δ.

Formally, let Ψ denote the *unknown* distribution of network latency. While this could be different for different leaders, we can consider only the one corresponding to the current leader. Then, the observer receives the block at time $o = \Delta + \psi$ where $\psi \sim \Psi$.

$\mathbb{E}[\psi]$ is known. When the mean of Ψ is known, the observer effectively learns an unbiased estimator of Δ as $o - \mathbb{E}[\psi]$. This will serve as the observation d used by the mechanism. Note here that $d \sim \Delta + \varepsilon$ where ε is a "noise term" with mean 0. Note that ε is allowed to be negative here.

We find that only knowing the mean of the network latency is not sufficient to incentivize responsiveness. Theorem 4.4 states a stronger version of this impossibility result.

Theorem 4.4. *Consider any* $\tau(\Delta) : \mathbb{R}^{\geq 0} \to \mathbb{R}^{\geq 0}$ *and let* Γ *be any setting that includes all* $\{\chi_\Delta\}$ *satisfying* $\tau(\Delta) = \mathbb{E}[\chi_\Delta]$ *for all* Δ*. Then, there is no reward* R *with finite budget that is* 0*-responsive.*

Known mean and bounded variance. Suppose now that for the network latency distribution Ψ, both its mean and an upper-bound σ^2 on its variance are known. Again, subtracting $\mathbb{E}[\psi]$ from the observed time results in an unbiased estimator d for Δ but now one with bounded variance. In other words, $d \sim \Delta + \varepsilon$ where the noise ε is such that $\mathbb{E}[\varepsilon] = 0$ and $\text{Var}(\varepsilon) \leq \sigma^2$.

In such a case, we construct an explicit 0-responsive mechanism R given extremely mild conditions on the MEV. Concretely, we require the derivative of MEV at 0, i.e., $\frac{\partial}{\partial \Delta} m^*_w(\Delta)|_{\Delta=0}$ to be finite.

Theorem 4.5. *Let* Γ *be any setting such that each world* w *with* $\{\chi^w_\Delta\}$ *is such that* $\chi^w_\Delta = \Delta + \varepsilon$ *where* $\varepsilon \in \mathbb{R}$ *satisfies* $\mathbb{E}[\varepsilon] = 0$ *and* $\text{Var}(\varepsilon) \leq \sigma^2$*. Further,* $\frac{\partial}{\partial \Delta} m^*_w(\Delta)|_{\Delta=0}$ *is finite.*

Then, there exists a 0*-responsive* R *with budget* $B = \max\{16\sigma\gamma, 8M^*_\Gamma\}$ *where* $M^*_\Gamma = \sup_\Delta m^*_\Gamma(\Delta)$ *and*

$$\gamma = \sup_{w \in \Gamma} \sup_{\Delta > 0} \frac{m^*_w(\Delta) - m^*_w(0)}{\Delta}.$$

The proof is deferred to Section A.2.5.

Perfect knowledge of the distribution. If the network latency distribution is perfectly known, we achieve can responsiveness with budget that matches the TV distance lower bound from Theorem 4.3. Let f_Ψ be the known density of Ψ. The idea now is to penalize the leader if it is more likely that the observation d could have come from $\Delta > 0$ compared to $\Delta = 0$.

For simplicity, consider Γ with a single world with $\chi_\Delta = \Delta + \Psi$. Define B^* as the minimal budget from our lower bound and suppose that it is achieved uniquely at Δ^*. Let $\mathbf{I}_d$ denote the indicator variable that $f_\Psi(d - \Delta^*) > f_\Psi(d)$ and consider the mechanism

$$R(t, d) = -t - \mathbf{I}_d \cdot B^*$$

Theorem A.4 (Appendix A.2) shows that this is 0-responsive.

Practical example. As a concrete example to illustrate practical budget requirements, consider based on [35], that the function 6.5×10^{-3} ETH/s upper-bounds the real MEV growth. Suppose that the network latency for a leader's proposal to reach a globally-distributed set of consensus nodes is exponential with mean 2 seconds. The leader can delay by at most 6 seconds.

We compute the budget lower-bound for incentivizing responsiveness as 0.041 ETH which is only slightly larger than the MEV gained in 6 seconds ($M^* = 0.039$ ETH). We achieve this bound exactly when the full distribution is known. When only the mean is known, no finite budget is sufficient. If we know a bound on $\sigma \leq 3$ sec without any other information about the distribution, we can still design a responsive mechanism with budget roughly 0.31 ETH, i.e., 8 times larger than the lower-bound.

For context, we can compare these numbers to existing rewards. Ethereum mints about 1700 ETH per day as consensus-layer rewards distributed to validators [14]—this amounts to 0.23 ETH per block. In other words, rewards would not have to substantially increases to accommodate even conservative responsive mechanisms.

4.4.2 Multiple observers and voting. We can remove the honesty assumption on a single observer by aggregating together timestamps from multiple observers. For simplicity, suppose that all consensus nodes apart from the leader act as observers. Let there be ℓ of them.

Now, if the leader chooses delay Δ, each observer i receives the proposal at time $o_i = \Delta + \psi_i$ where the ψ_i are sampled i.i.d. from the network latency distribution Ψ. Let μ and σ^2 denote its (potentially unknown) mean and variance respectively.

Basic detection process under honesty. A simple aggregation technique is to compute the mean of all observations $\hat{o} = \Delta + \sum \psi_i / \ell$. Note that $\hat{o} = \Delta + \hat{\psi}$ where ψ has mean μ and variance σ^2/ℓ. In other words, using more observations reduces the variance of the "noise" and consequently using this aggregate in our budget mechanism will reduce its budget requirement. This works as expected when the observers are honest or weakly rational (i.e., report honestly under the assumption that they gain nothing from dishonest behavior).

Observers colluding with the leader. A problem with the previous strategy surfaces if the leader is colluding with observers—here even when the leader delays by $\Delta > 0$, colluding observers would report as if the leader did not delay. As standard, the idea now is to use the *median* observation. Under the assumption that the leader corrupts less than $\ell/2$ observers with the rest reporting honestly, it is easy to see that the median observation is guaranteed to be an honest report. Note that doing so effectively means that ℓ observers (with less than $\ell/2$ dishonest colluders) together simulate a single honest observer. In the worse case, there is no advantage in reducing variance, and our results from Section 4.4.1 directly apply.

5 General Model with Multiple Leaders

We now extend our general formalism to a setting with multiple leaders. and demonstrate the inherent advantage of such a protocol in incentivizing responsiveness.

5.1 Rewarding the Earliest Leader

One way to design multi-leader protocols is to allow multiple parties to propose their blocks, and then select the one confirmed by the underlying consensus protocol (e.g., the earliest one to receive enough attestations from other nodes) to be rewarded. We analyze such a protocol in the context of responsiveness in this section.

Basic description. Consider a protocol with K leaders who each propose a block. Each leader i chooses its action $\mathbf{a}_i = \Delta_i$ for its proposal and let $\vec{\Delta} = (\Delta_1, \ldots, \Delta_K)$ denote the action vector. We use $\vec{\mathbf{0}}$ to denote all leaders playing the "honest" strategy of $\Delta = 0$.

Let Ψ denote the distribution of latency required to aggregate enough votes on a proposal (for simplicity, we will just call this the network latency). In other words, when i chooses Δ_i, the timestamp associated with it can be thought of as $\Delta_i + \psi_i$ where $\psi_i \sim \Psi$. The latencies $\psi_1, \ldots, \psi_K$ for the K proposals are sampled i.i.d. from Ψ. The winning proposal is the one with the lowest timestamp $\Delta_j + \psi_j$; note that for real-valued timestamps, ties occur with probability 0 and may be arbitrary broken.

5.1.1 Analysis. Our goal now, is to analyze the equilibria in this game. We find that, similar to the two-choice model, all leaders playing $\Delta = 0$ will be a unique pure Nash equilibrium under suitable constraints on the growth of MEV. Further, this responsiveness can be incentivized by a mechanism with zero budget.

As in our previous results, $\mathcal{T}$ does not constrain us since any observed fees $t = \mathsf{T}(\mathbf{a})$ can always be burned. We therefore focus only on the unknown MEV utility with expected value $\widetilde{\mathsf{M}}(\Delta)$.

No strict dominant strategy. We first observe that, unlike the simple two-choice model, there is no strict dominant strategy. This is easy to see—for example, if other leaders chose a very large delay, the best response for leader i is to choose the largest possible Δ_i that will still allow it to win against the others. Doing so will give i the highest MEV. In other words, i's best response depends on others and there is no dominant strategy.

Honest equilibrium at zero delay. Nevertheless, we still find strong incentive for leaders to play honestly. We find that under appropriate MEV constraints, there exists an R that makes the honest strategy $\vec{\mathbf{0}}$ a unique pure Nash equilibrium. R will give $-t + \epsilon$ for arbitrarily small $\epsilon > 0$ to the winning proposal and 0 to others.

For lack of clutter, we state our result in Theorem 5.1 in terms of a single world but it naturally extends to Γ with multiple worlds.

Theorem 5.1. *Consider a setting with $K \geq 2$ leaders as described in Section 5.1. Define a function related to MEV-growth as follows,*

$$G(\Delta) = \frac{\widetilde{\mathsf{M}}(\Delta) - \widetilde{\mathsf{M}}(0)}{\widetilde{\mathsf{M}}(0)}$$

and suppose that

$$\sup_{\Delta>0} \frac{G(\Delta)}{(1-e^{-c_\Psi \Delta})} < K - 1.$$

where $c_\Psi \geq 0$ is a constant that depends on the network distribution Ψ (in particular, it is a lower-bound on Ψ's "hazard rate").

Then, there exists a responsive R *with zero budget such that $\vec{\Delta} = \vec{\mathbf{0}}$ (i.e., all leaders playing 0 delay) is the unique pure Nash equilibrium.*

The full proof is deferred to Appendix A.3. Note that as before, a simple corollary is that by increasing the number of leaders, we can always force the above inequality to be satisfied, and consequently incentivize the honest strategy of playing $\Delta = 0$.

Remark on heterogeneous networks. If leaders have heterogeneous latencies (e.g., one leader is better connected than others), a conservative approach is to use the same bound as Theorem 5.1, but the hazard rate corresponding to the slowest leader competing. Our example below shows that this should be sufficient in practice.

Practical example. We consider our concrete example from Section 4.4 in the context of multiple leaders. While MEV variance is large, conservatively, suppose that $\widetilde{\mathrm{M}}(0)$ is 0.041 ETH (c.f., [39]) and it grows linearly as before at 6.5×10^{-3} ETH/s [29]. Given hazard rate 0.5 (e.g., exponential network latency with mean 2 seconds), we only need $K = 2$ leaders to incentivize the honest equilibrium at $\vec{\mathbf{0}}$. Even using a conservative hazard rate of 0.1 (e.g., 10 second network latency for the slowest leader) requires only $K = 3$.

Even with super-linear MEV, we find that only a few leaders are sufficient. Concretely, consider the analysis from [39] which fit the polynomial $\widetilde{\mathrm{M}}(\Delta) = (-1.99\Delta^3 + 2.44\Delta^2 + 32.5\Delta + 40.7) \times 10^{-3}$ ETH to MEV bids submitted by builders in practice. For this, our result shows that 4 competing leaders are sufficient.

Note that the budget required by the mechanism is 0 in all cases.

5.2 An Extension: Joint Proposal Protocols

A different approach is to use the different proposals from leaders to *jointly* construct the final block. For example, the final block may consist of all transactions that are present in sufficiently many individual proposals. Such a design is commonly seen in so-called asynchronous "leaderless" protocols [8, 19, 28].

Transaction fee mechanism (TFM) design for such protocols, i.e., dividing fee revenue between different proposers, (c.f., [18, 37]) is an interesting recent direction; in our responsiveness context however, since we can burn all observable fee revenue, this is less relevant than the MEV received. Unfortunately, it is unclear within prior works how the unobservable MEV gets split. We propose one realistic formulation that we use as basis for our analysis.

Dividing MEV revenue in joint-proposal protocols. Consider a model as follows: as before, each leader i chooses a delay Δ_i, and gets assigned a timestamp $\Delta_i + (\psi_i \sim \Psi)$ that takes network latency into account. The protocol picks some $\kappa < K$ leaders, with the lowest timestamps, as "winners" to jointly build the final block.

The MEV is distributed (unobservable to the protocol) as follows: let $\iota_1, \ldots, \iota_\kappa$ denote indices of the κ earliest arriving leaders. Then, the earliest leader ι_1 gets MEV $\widetilde{\mathrm{M}}(\Delta_{\iota_1})$, leader ι_2 gets $\widetilde{\mathrm{M}}(\Delta_{\iota_2}) - \widetilde{\mathrm{M}}(\Delta_{\iota_1})$ and so on. In other words, each leader captures any "left-over" MEV up to the time of its proposal. If $\Delta_{\iota_2} < \Delta_{\iota_1}$ but leader ι_1 arrives earlier due to network delays, then then ι_2 gets no MEV.

Brief results. Compared to Section 5.1, the condition for a unique Nash equilibrium at $\vec{\mathbf{0}}$ becomes stricter for joint-proposals. An equivalent framing is that we may need a larger K to incentivize the honest behavior of not delaying. This highlights a tradeoff between *cooperation* and *competition*. We posit that choosing a larger κ can lead to better performance (since more leaders cooperate to build blocks), but worse competition, i.e., needs a larger K to incentivize responsiveness—this may force a bound on the ratio κ/K. We give brief intuition in Appendix A.3 and leave a full game-theoretic analysis to future work. We also provide a discussion in Section 5.3.

5.3 Discussion

Multi-leader consensus with *competition*. A key observation in our work is that competition is inherently useful for incentivizing responsiveness. This motivates the design of consensus protocols that incorporate competition between multiple block proposers.

While some existing protocols involve multiple parties proposing blocks (e.g., DAG-based protocols [6, 23, 36]), as mentioned earlier, their goals are different and accordingly the multiple proposers are generally not competing with each other.

For our purposes, a consensus protocol with competing leaders would ideally guarantee that earlier block proposals are more likely to be chosen for finalization than later ones. This property has the flavor of *order-fair* protocols [10, 24] that aim to order transactions in the way they were received by a majority of the network—a kind of "distributed first-come-first-serve." Here, an analogous property would be required for block proposals rather than user transactions.

We leave the exploration of such designs, as well as a broader consensus-level exploration of leader competition, to future work.

Competition vs cooperation. One concern with leader competition is wasted effort—if many leaders propose a block, but only one gets confirmed, any resources used by others are to some extent wasted in terms of the final output. Leader *cooperation*, on the other hand, can have better performance in the honest case, but it may lose its advantages under incentive attacks (such as timing games). Tradeoffs of this form exist for e.g., censorship-resistance [1, 13].

An interesting open direction is to design consensus protocols that balance both cooperation and competition between parties. This would allow for appropriately navigating tradeoffs between performance and incentive-compatibility.

Conclusion

We presented a framework for designing *responsive* mechanisms to incentivize responsiveness in consensus protocols and deter timing games. For single-leader protocols, we showed a general lower bound on the budget required for such responsive mechanisms, and explicit constructions for a number of different settings. We then investigated multi-leader designs to show their fundamental advantage in incentivizing responsiveness through leader competition.

References

[1] Ittai Abraham, Yuval Efron, and Ling Ren. 2025. The Latency Cost Of Censorship Resistance. Cryptology ePrint Archive, Paper 2025/2136. https://eprint.iacr.org/2025/2136

[2] Kaya Alpturer, Kushal Babel, and Aditya Saraf. 2025. Timing Games in Responsive Consensus Protocols. arXiv:2510.25144

[3] Kaya Alpturer and S. Matthew Weinberg. 2024. Optimal RANDAO Manipulation in Ethereum. In *AFT*. 10:1–10:21.

[4] Yackolley Amoussou-Guenou, Bruno Biais, Maria Potop-Butucaru, and Sara Tucci Piergiovanni. 2020. Rational vs Byzantine Players in Consensus-based Blockchains. In *AAMAS*. International Foundation for Autonomous Agents and Multiagent Systems, 43–51.

[5] Hagit Attiya, Cynthia Dwork, Nancy A. Lynch, and Larry J. Stockmeyer. 1994. Bounds on the Time to Reach Agreement in the Presence of Timing Uncertainty. *J. ACM* 41, 1 (1994), 122–152.

[6] Kushal Babel, Andrey Chursin, George Danezis, Anastasios Kichidis, Lefteris Kokoris-Kogias, Arun Koshy, Alberto Sonnino, and Mingwei Tian. 2025. Mysticeti: Reaching the Latency Limits with Uncertified DAGs. In *NDSS*.

[7] Kushal Babel, Philip Daian, Mahimna Kelkar, and Ari Juels. 2023. Clockwork Finance: Automated Analysis of Economic Security in Smart Contracts. In *IEEE S&P*. 2499–2516.

[8] Michael Ben-Or. 1985. Fast Asynchronous Byzantine Agreement (Extended Abstract). In *PODC*. ACM, 149–151.

[9] Jonah Brown-Cohen, Arvind Narayanan, Alexandros Psomas, and S. Matthew Weinberg. 2019. Formal Barriers to Longest-Chain Proof-of-Stake Protocols. In *EC*. ACM, 459–473.

[10] Christian Cachin, Jovana Micic, and Nathalie Steinhauer. 2022. Quick Order Fairness. In *FC*. 316–333.

[11] ChorusOne Team. 2025. *Timing Games on Solana: Validator Incentives, Network Impacts, and Agave's Hidden Inefficiencies.* https://x.com/ChorusOne/status/1955615856580301252

[12] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. 2020. Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable Value, and Consensus Instability. In *IEEE S&P*. 585–602.

[13] Fatima Elsheimy, Ioannis Kaklamanis, Sarisht Wadhwa, Charalampos Papamanthou, and Fan Zhang. 2026. Censorship Resistance vs Throughput in Multi-Proposer BFT Protocols. Cryptology ePrint Archive, Paper 2026/126. https://eprint.iacr.org/2026/126

[14] Ethereum Foundation. 2026. How The Merge impacted ETH supply. https://ethereum.org/roadmap/merge/issuance/.

[15] Ittay Eyal and Emin Gün Sirer. 2018. Majority is not enough: bitcoin mining is vulnerable. *Commun. ACM* 61, 7 (2018), 95–102.

[16] Flashbots and Robert Miller. 2022. *MEV and the Limits of Scaling.* https://writings.flashbots.net/mev-and-the-limits-of-scaling

[17] Robin Fritsch, Maria Inês Silva, Akaki Mamageishvili, Benjamin Livshits, and Edward W. Felten. 2025. MEV Capture Through Time-Advantaged Arbitrage. In *MARBLE (Lecture Notes in Operations Research)*. Springer, 1–22.

[18] Pranav Garimidi, Lioba Heimbach, and Tim Roughgarden. 2025. Transaction Fee Mechanism Design for Leaderless Blockchain Protocols. In *FC*. 20–35.

[19] Bingyong Guo, Zhenliang Lu, Qiang Tang, Jing Xu, and Zhenfeng Zhang. 2020. Dumbo: Faster Asynchronous BFT Protocols. In *CCS*. 803–818.

[20] Hanna Halaburda, Zhiguo He, and Jiasun Li. 2021. *An Economic Model of Consensus on Distributed Ledgers.* Technical Report.

[21] Mohammad Mussadiq Jalalzai, Kushal Babel, Jovan Komatovic, Tobias Klenze, Sourav Das, Fatima Elsheimy, Mike Setrin, John Bergschneider, and Babak Gilkalaye. 2026. MonadBFT: Fast, Responsive, Fork-Resistant Streamlined Consensus. arXiv:2502.20692 [cs.DC] https://arxiv.org/abs/2502.20692

[22] Dakai Kang, Suyash Gupta, Dahlia Malkhi, and Mohammad Sadoghi. 2025. HotStuff-1: Linear Consensus with One-Phase Speculation. *Proc. ACM Manag. Data* 3, 3 (2025), 171:1–171:29.

[23] Idit Keidar, Oded Naor, Ouri Poupko, and Ehud Shapiro. 2023. Cordial Miners: Fast and Efficient Consensus for Every Eventuality. In *DISC (LIPIcs)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 26:1–26:22.

[24] Mahimna Kelkar, Fan Zhang, Steven Goldfeder, and Ari Juels. 2020. Order-Fairness for Byzantine Consensus. In *CRYPTO*. 451–480.

[25] Quentin Kniep, Jakub Sliwinski, and Roger Wattenhofer. 2025. *Solana Alpenglow Consensus.* https://drive.google.com/file/d/1Rlr3PdHsBmPahOInP6-Pl0bMzdayltdV/view

[26] Jason Milionis, Ciamac C. Moallemi, and Tim Roughgarden. 2024. Automated Market Making and Arbitrage Profits in the Presence of Fees. In *FC (1) (Lecture Notes in Computer Science)*. Springer, 159–171.

[27] Andrew Miller, Ahmed E. Kosba, Jonathan Katz, and Elaine Shi. 2015. Nonoutsourceable Scratch-Off Puzzles to Discourage Bitcoin Mining Coalitions. In *CCS*. ACM, 680–691.

[28] Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. 2016. The Honey Badger of BFT Protocols. In *CCS*. 31–42.

[29] Burak Öz, Benjamin Kraner, Nicolò Vallarano, Bingle Stegmann Kruger, Florian Matthes, and Claudio Juan Tessone. 2023. Time Moves Faster When There is Nothing You Anticipate: The Role of Time in MEV Rewards. In *DeFi@CCS*. ACM, 1–8.

[30] Rafael Pass and Elaine Shi. 2017. Hybrid Consensus: Efficient Consensus in the Permissionless Model. In *DISC (LIPIcs)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 39:1–39:16.

[31] Rafael Pass and Elaine Shi. 2018. Thunderella: Blockchains with Optimistic Instant Confirmation. In *EUROCRYPT (2) (Lecture Notes in Computer Science)*. Springer, 3–33.

[32] Kaihua Qin, Liyi Zhou, and Arthur Gervais. 2022. Quantifying Blockchain Extractable Value: How dark is the forest?. In *IEEE S&P*. 198–214.

[33] Ayelet Sapirshtein, Yonatan Sompolinsky, and Aviv Zohar. 2016. Optimal Selfish Mining Strategies in Bitcoin. In *Financial Cryptography (Lecture Notes in Computer Science)*. Springer, 515–532.

[34] Caspar Schwarz-Schilling. 2022. *Retroactive proposer rewards.* https://notes.ethereum.org/@casparschwa/S1vcyXZL9

[35] Caspar Schwarz-Schilling, Fahad Saleh, Thomas Thiery, Jennifer Pan, Nihar Shah, and Barnabé Monnot. 2023. Time Is Money: Strategic Timing Games in Proof-Of-Stake Protocols. In *AFT (LIPIcs)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 30:1–30:17.

[36] Alexander Spiegelman, Neil Giridharan, Alberto Sonnino, and Lefteris Kokoris-Kogias. 2022. Bullshark: DAG BFT Protocols Made Practical. In *CCS*. ACM, 2705–2718.

[37] Aikaterini-Panagiota Stouka, Julian Ma, and Thomas Thiery. 2025. Multiple Proposer Transaction Fee Mechanism Design: Robust Incentives Against Censorship and Bribery. arXiv:2505.13751

[38] Weijie Sun, Zihuan Xu, Wangze Ni, and Lei Chen. 2025. InTime: Towards Performance Predictability In Byzantine Fault Tolerant Proof-of-Stake Consensus. *Proc. ACM Manag. Data* 3, 1 (2025), 47:1–47:27.

[39] Anton Wahrstätter, Liyi Zhou, Kaihua Qin, Davor Svetinovic, and Arthur Gervais. 2023. Time to Bribe: Measuring Block Construction Market. *IACR Cryptol. ePrint Arch.* 2023 (2023), 760.

[40] Aviv Yaish, Gilad Stern, and Aviv Zohar. 2023. Uncle Maker: (Time)Stamping Out The Competition in Ethereum. In *CCS*. ACM, 135–149.

[41] Maofan Yin, Dahlia Malkhi, Michael K. Reiter, Guy Golan-Gueta, and Ittai Abraham. 2019. HotStuff: BFT Consensus with Linearity and Responsiveness. In *PODC*. ACM, 347–356.

[42] Liyi Zhou, Kaihua Qin, Christof Ferreira Torres, Duc V Le, and Arthur Gervais. 2021. High-Frequency Trading on Decentralized On-Chain Exchanges. In *IEEE S&P*. 428–445.

A Deferred Proofs

In this appendix, we detail proofs deferred from the main body.

A.1 Two-Choice Model (Section 3)

A.1.1 Detectable settings.

Proof of Theorem 3.4. For always responsiveness, recall that the first condition is $\mathsf{R}_{\mathsf{early}}(t_0) > -t_0 - m_0$ for all t_0, m_0. Since $m_0 \geq 0$, this can be rewritten as

$$\mathsf{R}_{\mathsf{early}}(t_0) > -t_0$$

Now, given $\mathsf{R}_{\mathsf{early}}$, the second responsiveness condition is equivalent to $\mathsf{R}_{\mathsf{late}}(t_1) < \mathsf{R}_{\mathsf{early}}(t_0) + (t_0 - t_1) - (m_1 - m_0)$ for all (t_0, t_1) and (m_0, m_1). Since $m_1 - m_0 \leq m_\Gamma^+$ is given, we can rewrite this as

$$\mathsf{R}_{\mathsf{late}}(t_1) < \mathsf{R}_{\mathsf{early}}(t_0) + (t_0 - t_1) - m_\Gamma^+$$

for all (t_0, t_1). Now, this is guaranteed exactly when $\mathsf{R}_{\mathsf{late}}(t_1)$ is strictly less than the infimum of the RHS over all t_0. This gives,

$$\mathsf{R}_{\mathsf{late}}(t_1) < \inf_{t_0 \leq t_1} \left(\mathsf{R}_{\mathsf{early}}(t_0) + t_0\right) - t_1 - m_\Gamma^+$$

for all t_1, which are exactly the conditions in the theorem statement for R to be responsive.

Now, consider the previously mentioned reward $\mathsf{R}_{\mathsf{early}}(t) = -t + \varepsilon$ and $\mathsf{R}_{\mathsf{late}}(t) = -t - m_\Gamma^+$ for any $\varepsilon > 0$. Observe that this has budget $m_\Gamma^+ + \varepsilon$, i.e., the infimum of any R of this form is m_Γ^+.

We now show that any reward with smaller budget than m_Γ^+ cannot be responsive. Consider any responsive R′. Then, observe that putting the two properties from the theorem statement together, we know $\mathsf{R}_{\mathsf{late}}(t_1) < -t_1 - m_\Gamma^+$ and consequently,

$$\mathbb{U}^{\mathsf{leader}}_{(\mathsf{T},\mathsf{M})}(\mathsf{late}) < m_1 - m_\Gamma^+$$

for all $m_1 \geq 0$. In turn, $\mathsf{budget}(\mathsf{R}) > m_\Gamma^+$ which is a contradiction. □

A.1.2 Undetectable setting with no MEV.

Theorem A.1. *Consider an undetectable setting Γ that only contains worlds with $\mathcal{M} = 0$ (i.e., no unknown exogenous utility / MEV). A reward function R is always responsive w.r.t. Γ if and only if it has the form $\mathsf{R}(t) = -t + g(t)$ where $g : \mathbb{R}^{\geq 0} \to \mathbb{R}^{\geq 0}$ is strictly decreasing.*

Proof. Consider a reward R that is always responsive. Since there is no MEV, the second condition for responsiveness is equivalent to $t_0 + \mathsf{R}(t_0) > t_1 + \mathsf{R}(t_1)$ for all (t_0, t_1).

Now, define $g(t) = \mathsf{R}(t) + t$ and rewrite the previous condition as $g(t_0) > g(t_1)$ where $t_0 < t_1$. In other words, g is strictly decreasing.

Note further that the first condition for responsiveness is equivalent to $g(t) > 0$. In effect, R is responsive if and only if it has the form $\mathsf{R}(t) = -t + g(t)$ where g is strictly decreasing. □

A.2 General Model for a Single Leader (Section 4)

This appendix expands on results from Section 4.

A.2.1 Unbounded MEV leads to infinite budget.

Lemma A.2. *No reward R can be responsive in settings where the expected MEV m^*_Γ (see Section 4.1) is unbounded.*

Proof. It is sufficient to prove this impossibility in the setting easiest for the mechanism—that is, when Δ is perfectly detectable ($\chi_\Delta = \Delta$). First, observe that the only way m^*_Γ is unbounded is if $\lim_{\Delta\to\infty} m^*_\Gamma(\Delta) = \infty$.

Let R be any reward with a finite budget B. Now, since m^*_Γ is unbounded, there exists a sufficiently large Δ' such that $m^*_\Gamma(\Delta') - m^*_\Gamma(0) > B$, and as a result, some world in Γ where $\widetilde{\mathsf{M}}(\Delta') - \widetilde{\mathsf{M}}(0) > B$. Therefore, it is not possible to prevent the leader from choosing Δ' over a delay of 0 with a budget of only B. □

A.2.2 Fully detectable setting.

Proof of Theorem 4.2. Consider the following R:

$$\mathsf{R}(t, d) = \begin{cases} -t + \epsilon & \text{if } d \leq \delta \\ -t - b_\delta & \text{if } d > \delta \end{cases}$$

where $\epsilon > 0$ is arbitrarily small. To see that this is δ-responsive, observe that by construction, for any $\Delta_1 \leq \delta$ and $\Delta_2 > \delta$, we have

$$\mathbb{EU}^{\text{leader}}(\Delta_2) < 0 < \mathbb{EU}^{\text{leader}}(\Delta_1)$$

which directly gives δ-responsiveness. Further, it is easy to see that the budget is at most b_δ since this is the maximum penalty imposed on the leader. □

A.2.3 Setting with threshold detection.

Theorem A.3. *Consider the threshold detection setting (see Section 4.2). There exists a δ-responsive reward R with budget B where*

$$B = \sup_{w\in\Gamma} \sup_{\Delta>\delta} \frac{m^*_w(\Delta) - m^*_w(\delta)}{\Pr_{d\sim\chi_\Delta}[d = L]}$$

Proof. Concretely, define R as:

$$\mathsf{R}(t, d) = \begin{cases} -t + \epsilon & \text{if } d = 0 \text{ (early)} \\ -t - B & \text{if } d = D \text{ (late)} \end{cases}$$

where $\epsilon > 0$ is arbitrarily small.

First, observe that the leader receives positive utility for $\Delta \leq \delta$ (since it is always detected as early).

Now, for any world w and $\Delta > \delta$, the leader is detected as late with probability $p_{\text{late}} = \Pr_{d\sim\chi_\Delta}[d = L]$ and penalized B. Therefore, it's expected utility is:

$$\begin{aligned} \mathbb{EU}^{\text{leader}}(\Delta) &\leq m^*_w(\Delta) - p_{\text{late}} \cdot B \\ &< m^*_w(\Delta) - [m^*_w(\Delta) - m^*_w(\delta)] \\ &= m^*_w(\delta) \end{aligned}$$

Then, the leader prefers to play δ over any $\Delta > \delta$. In other words, R is δ-responsive. □

A.2.4 No finite budget with only network mean known.

Proof of Theorem 4.4. Consider a particular world w.

Now, fix any $\Delta^\circ > \delta$ with $m^*_w(\Delta^\circ) - m^*_w(0) > 0$ and any fixed z. The idea now is to design χ_{Δ° and χ_0 to have almost all probability mass at z while still keeping the mean tied to the given τ.

In particular, for any $C > |\tau(\Delta^\circ) - \tau(0)|$, define χ^C_0 and $\chi^C_{\Delta^\circ}$ as follows:

- χ^C_0 outputs z with probability $1 - 1/C$ and y_0 with probability $1/C$ where y_0 is set to make the mean $\tau(0)$, i.e.,
$$y_0 = C\tau(0) - Cz + z.$$
- $\chi^C_{\Delta^\circ}$ similarly outputs z with probability $1 - 1/C$ and y_1 otherwise where y_1 is set to make the mean $\tau(\Delta^\circ)$.

Note that since there are no restrictions other than on the mean, for all C, Γ contains a world identical to w but with the family $\{\chi^C\}$.

Now, it is easy to see that $\mathsf{TVD}(\chi^C_0, \chi^C_{\Delta^\circ}) \leq 1/C$ and therefore the infimum over all C is 0. As a consequence, by using the lower bound from Theorem 4.3, we conclude that infinite budget is required for any responsive R. □

A.2.5 Budget with known mean and bounded variance.

Proof of Theorem 4.5. Recall that our setting is such that $\chi_\Delta = \Delta + \varepsilon$ with $\mathbb{E}[\varepsilon] = 0$ and $\mathrm{Var}(\varepsilon) \leq \sigma^2$. Let γ be as defined in the theorem statement and define $c(d)$ as $d + 4\sigma$ "clipped" to the range $[0, 8\sigma]$, i.e.,

$$c(d) = \max(0, \min(d + 4\sigma, 8\sigma)).$$

Recall that M^*_Γ is the supremum of the function $m^*_\Gamma(\Delta)$. Consider arbitrarily small $\nu > 0$ (this is just for making the utility for $\Delta = 0$ strictly positive) and define the reward mechanism as

$$\mathsf{R}(t, d) = -t - \alpha \cdot c(d) + \nu$$

where $\alpha = \max(2\gamma, M^*_\Gamma/\sigma)$. We show that R achieves 0-responsiveness with budget given by $B = \max(16\sigma\gamma, 8M^*_\Gamma)$.

First, define $P(d)$ as the term that depends on $c(d)$ such that $\mathsf{R}(t, d) = -t - P(d) + a$.

Case 1: $\Delta \in (0, 2\sigma)$**.** Define the event E such that $\varepsilon \in [-4\sigma, 4\sigma - \Delta]$. In E, both ε and $\Delta + \varepsilon$ lie in $[-4\sigma, 4\sigma]$ where $P(d)$ is linear with slope α. Therefore, in this region, $P(\Delta + \varepsilon) - P(\varepsilon) = \alpha\Delta$

Further, even on the complement region denoted by E°, P is non-decreasing. We now use Chebyshev's inequality to get,

$$\Pr[\varepsilon < -4\sigma] \le \Pr[|\varepsilon| > 4\sigma] \le \frac{\sigma^2}{(4\sigma)^2} = \frac{1}{16}.$$

For $\Delta \le 2\sigma$, i.e., $4\sigma - \Delta \ge 2\sigma$, we have,

$$\Pr[\varepsilon > 4\sigma - \Delta] \le \Pr[\varepsilon > 2\sigma] \le \Pr[|\varepsilon| > 2\sigma] \le \frac{\sigma^2}{(2\sigma)^2} = \frac{1}{4}.$$

Taking the union bound gives us the probability of the complement E° to be at most $5/16$. In other words, $\Pr[\mathsf{E}] \ge 1 - (5/16) > 1/2$.

Now,

$$\mathbb{E}[P(\Delta + \varepsilon) - P(\varepsilon)] > \alpha\Delta \cdot \Pr[\mathsf{E}] + 0 \cdot \Pr[\mathsf{E}^\circ] \ge \gamma\Delta$$

Note that $\gamma\Delta \ge m^*_w(\Delta) - m^*_w(0)$ from the definition of γ. Moreover, γ is finite since we are given that the derivative of $m^*_w(\cdot)$ at $\Delta = 0$ is finite. As a consequence,

$$\begin{aligned}\mathbb{EU}^{\text{leader}}(0) &= m^*_w(0) - \mathbb{E}_\varepsilon[P(\epsilon)] + \nu \\ &> m^*_w(\Delta) - \mathbb{E}_\varepsilon[P(\Delta + \varepsilon)] + \nu \\ &= \mathbb{EU}^{\text{leader}}(\Delta)\end{aligned}$$

i.e., R is 0-responsive for the case of $\Delta \in (0, 2\sigma)$.

Case 2: $\Delta > 2\sigma$**.** For any $\Delta > 2\sigma$, again since $P(d)$ is non-decreasing, we get

$$\mathbb{E}[P(\Delta + \varepsilon)] - P(\varepsilon)] \ge \mathbb{E}[P(2\sigma + \varepsilon) - P(\varepsilon)]$$

which is $\ge \frac{\alpha \cdot 2\sigma}{2} = \alpha\sigma$ from the analysis in case (1).

For 0-responsiveness in the region $\Delta > 2\sigma$, observe that we need to have $\alpha\sigma$ be larger than $m^*_w(\Delta) - m^*_w(0)$. Since $m^*_w(\Delta) - m^*_w(0) \le M^*_\Gamma$, it is sufficient to have $\alpha\sigma \ge M^*_\Gamma$, which holds by construction.

Putting the two cases together completes the proof. Notice that the budget is exactly $8\sigma\alpha = \max\{16\sigma\gamma, 8M^*_\Gamma\}$. □

A.2.6 Matching upper bound with known network distribution.

Theorem A.4. *Consider Γ with a single world where $\chi_\Delta = \Delta + (\psi \sim \Psi)$ and the pdf f_Ψ of Ψ is known. Define B^* that is any arbitrarily small value above the lower-bound from Theorem 4.3 and suppose that it is achieved uniquely at Δ^*. Let $\mathbf{I}_d$ denote the indicator variable that $f_\Psi(d - \Delta^*) > f_\Psi(d)$. Define*

$$R(t, d) = -t - \mathbf{I}_d \cdot B^*$$

Then, $R(t, d)$ is 0-responsive.

Proof. Fix any $\Delta^\circ > 0$ and define the region

$$H_{\Delta^\circ} = \{d : f_\Psi(d - \Delta^\circ) > f_\Psi(d)\}.$$

Now,

$$\begin{aligned}&\mathbb{E}_{d \sim \chi^\circ_\Delta}[P(d)] - \mathbb{E}_{d \sim \chi_0}[P(d)] \\ &\quad = B^* \cdot (\Pr[\psi + \Delta^\circ \in H_{\Delta^\circ}] - \Pr[\psi \in H_{\Delta^\circ}]) \\ &\quad = B^* \cdot \mathsf{TVD}(\chi_0, \chi_{\Delta^\circ}) \\ &\quad > m^*_w(\Delta^\circ) - m^*(0).\end{aligned}$$

where the second step is because H_{Δ° is the maximal positive region exactly the same way as the proof of Theorem 4.3. This implies that $\mathbb{EU}(0) > \mathbb{EU}(\Delta^\circ)$ and since this holds for any $\Delta^\circ > 0$, we conclude that R is 0-responsive. □

A.3 General Multi-Leader Model (Section 5)

A.3.1 Honest Equilibrium at $\vec{\mathbf{0}}$ for earliest proposal designs.

Proof of Theorem 5.1. Let $f = f_\Psi$ and $F = F_\Psi$ denote the (potentially unknown) pdf and CDF of Ψ and $\overline{F} = 1 - F_\Psi$ denote the "survival function." The hazard rate is given by $h(x) = f/\overline{F}$ with $c_\Psi \le h(x)$ for all x.

We split the proof into four parts as follows: in the first two parts, we provide some helpful setup. Next, we establish existence of the Nash equilibrium at $\vec{\mathbf{0}}$. Finally, we prove uniqueness.

(1) Setup

We start by defining two useful functions, $P(\Delta)$ and $Q(\Delta)$ as follows:

- Let $P(\Delta)$ denote the probability that a player (WLG, player 1) with delay Δ beats $K-1$ others who each play 0. In other words, this is the event that $\psi_i > \Delta + \psi_1$ for all $i \ne 1$. It is easy to see that
$$P(\Delta) = \int_0^\infty f(x) \cdot \overline{F}(x + \Delta)^{K-1}\, dx$$
and $P(0) = 1/K$.
- Let $Q(\Delta)$ denote the probability that player 1 with delay 0 beats $K-1$ others who play Δ. In other words, this is the event that $\psi_1 < \Delta + \psi_i$ for all $i \ne 1$. It is easy to see that
$$Q(\Delta) = F(\Delta) + \int_0^\infty f(x + \Delta) \cdot \overline{F}(x)^{K-1}\, dx$$
and $Q(0) = 1/K$.

Next, our goal is to bound $P(\Delta)$ and $Q(\Delta)$ using the hazard rate. First, observe that for all $x, \Delta \ge 0$:

$$\begin{aligned}\overline{F}(x + \Delta) &= \overline{F}(x) \cdot \exp\left(-\int_x^{x+\Delta} h(y)\, dy\right) \\ &\le \overline{F}(x) \cdot e^{-c_\Psi \Delta}\end{aligned}$$

since $h(y) \ge c_\Psi$ for all y.

Now, we can bound $P(\Delta)$ using this as

$$\begin{aligned}P(\Delta) &\le e^{-(K-1)c_\Psi\Delta} \int_0^\infty f(x) \cdot \overline{F}(x)^{K-1}\, dx \\ &= e^{-(K-1)c_\Psi\Delta} \cdot P(0) = \frac{e^{-(K-1)c_\Psi\Delta}}{K}. \qquad (2)\end{aligned}$$

For $Q(\Delta)$, we proceed as follows: let Z be the minimum of $\psi_{j\ne1}$. Then, using the bound on $\overline{F}$, we get

$$\begin{aligned}Q(\Delta) &= \Pr[\psi_1 < \Delta + Z] = 1 - \mathbb{E}[\overline{F}(\Delta + Z)] \\ &\ge 1 - e^{-c_\Psi\Delta} \cdot \mathbb{E}[\overline{F}(Z)]\end{aligned}$$

Observe that each $\overline{F}(\psi_j)$ is a probability integral transform and therefore uniform on $[0, 1]$. Further, since $\overline{F}$ is non-increasing, we have $\overline{F}(Z) = \overline{F}(\min_{j\ne1} \psi_j) = \max_{j\ne1} \overline{F}(\psi_j)$ (i.e., the maximum of $K-1$ i.i.d uniform random variables). In other words, we have $\mathbb{E}[\overline{F}(Z)] = (K-1)/K$ allowing us to bound

$$Q(\Delta) \ge 1 - \frac{K-1}{K} e^{-c_\Psi\Delta} \qquad (3)$$

With this setup, we are now ready to proceed with showing responsiveness under specific constraints on $P(\Delta)$ and $Q(\Delta)$. We will use the reward function $\mathsf{R}(t)$ such that it gives $-t + \varepsilon$ to the winning (i.e., earliest arriving) leader, and 0 to everyone else.

We now show the incentives that will make everyone play honestly with no delay, i.e., $\vec{\mathbf{a}} = \vec{\mathbf{0}}$. We first establish this as a Nash equilibrium under suitable conditions on $P(\Delta)$ and then prove its uniqueness under suitable conditions on $Q(\Delta)$. Then, we put the two conditions together to establish our results.

(2) Rewriting given MEV condition

Here, we show that the given MEV condition satisfies two equations for all $\Delta > 0$:

$$\frac{\widetilde{\mathsf{M}}(\Delta)}{\widetilde{\mathsf{M}}(0)} \leq e^{(K-1)c_\Psi \Delta} \tag{4}$$

which we use for existence of the Nash equilibrium and

$$\frac{\widetilde{\mathsf{M}}(\Delta)}{\widetilde{\mathsf{M}}(0)} \leq K - (K-1)e^{-c_\Psi \Delta} \tag{5}$$

which we use for uniqueness.

Claim proof: We will prove the first inequality follows from the second by showing

$$K - (K-1)e^{-c_\Psi \Delta} \leq e^{(K-1)c_\Psi \Delta}$$

for all $\Delta \geq 0$, $K \geq 2$. Start by letting $u = c_\Psi \Delta \geq 0$ and rearranging to get

$$\phi(u) = e^{(K-1)u} - K + (K-1)e^{-u}$$

We need to show $\phi(u) \geq 0$. Note the boundary condition $\phi(0) = 0$. Next, we can compute the derivative as

$$\phi'(u) = (K-1)e^{(K-1)u} - (K-1)e^{-u} = (K-1)[e^{(K-1)u} - e^{-u}].$$

Now, since $(K-1)u$ is positive and $-u$ is negative, $e^{(K-1)u} - e^{-u} \geq 0$ and therefore $\phi'(u) \geq 0$. In turn, since ϕ is non-decreasing with boundary $\phi(0) = 0$, we get $\phi(u) \geq 0$ for all u.

Now, we show that the second inequality is just a simple rearranging of the given MEV condition. Observe that we can rewrite the given condition as

$$\frac{\widetilde{\mathsf{M}}(\Delta) - \widetilde{\mathsf{M}}(0)}{\widetilde{\mathsf{M}}(0)} < K - Ke^{-c_\Psi} + e^{-c_\Psi \Delta} - 1$$

Rearranging, we get

$$\frac{\widetilde{\mathsf{M}}(\Delta)}{\widetilde{\mathsf{M}}(0)} < K - Ke^{-c_\Psi} + e^{-c_\Psi \Delta} = K - (K-1)e^{-c_\Psi \Delta}$$

(3) Existence:

We are now ready to proceed with our game-theory analysis. For existence, consider the profile $\vec{\mathbf{a}} = \vec{\mathbf{0}}$ and WLG, suppose that player 1 wishes to defect, i.e., $\mathbf{a}_1 = \Delta_1$ and others play 0. Then, it's expected utility is

$$\begin{aligned} U(\Delta_1; \vec{\mathbf{0}}) &= P(\Delta_1) \cdot [\mathsf{T}(\Delta_1) + \widetilde{\mathsf{M}}(\Delta_1) + \mathsf{R}(\mathsf{T}(\Delta_1))] + (1 - P(\Delta_1)) \cdot 0 \\ &= P(\Delta_1) \cdot \widetilde{\mathsf{M}}(\Delta_1) \end{aligned}$$

Observe that $\vec{\mathbf{0}}$ will be a Nash equilibrium if $U(\Delta_1; \vec{\mathbf{0}}) > U(\vec{\mathbf{0}})$ for all $\Delta_1 > 0$. That is, if

$$P(\Delta_1) \cdot \widetilde{\mathsf{M}}(\Delta_1) < \frac{\widetilde{\mathsf{M}}(0)}{K}$$

holds for all $\Delta_1 > 0$. Note that $P(\Delta_1)$ also depends on K. To show this is true, we combine Eqs. (2) and (4) directly to get:

$$P(\Delta_1) \cdot \frac{\widetilde{\mathsf{M}}(\Delta_1)}{\widetilde{\mathsf{M}}(0)} \cdot K < \frac{e^{-(K-1)c_\Psi \Delta_1}}{K} \cdot e^{c_\Psi (K-1)\Delta_1} \cdot K = 1$$

for all Δ_1 which is exactly the sufficient condition for $\vec{\mathbf{0}}$ to be a Nash equilibrium.

(4) Uniqueness

Our next goal is to show that this honest strategy $\vec{\mathbf{0}}$ is a unique pure Nash equilibrium. For this, consider any symmetric profile $\vec{\mathbf{a}} = \vec{\Delta}^* = (\Delta^*, \cdots, \Delta^*)$ with $\Delta^* \neq 0$. By symmetry, each player wins with probability $1/K$, so the equilibrium utility is $\widetilde{\mathsf{M}}(\Delta^*)/K$. While this will be larger than $\widetilde{\mathsf{M}}(0)/K$ (i.e., the utility in $\vec{\mathbf{0}}$), we show that players will want to deviate from $\vec{\Delta}^*$, making it not an equilibrium.

For this, consider player 1 deviating to $\Delta_1 = 0$. Then, it's utility is

$$\begin{aligned} U_1(0; \vec{\Delta}^*) &= Q(\Delta^*) \cdot [\mathsf{T}(0) + \widetilde{\mathsf{M}}(0) + \mathsf{R}(\mathsf{T}(0))] + (1 - Q(\Delta^*)) \cdot 0 \\ &= \widetilde{\mathsf{M}}(0) \cdot Q(\Delta^*). \end{aligned}$$

Now, for $\vec{\Delta}^*$ to *not* be a Nash equilibrium, we need,

$$\widetilde{\mathsf{M}}(0) \cdot Q(\Delta^*) > \frac{\widetilde{\mathsf{M}}(\Delta^*)}{K}$$

To show this is true, we combine Eqs. (3) and (5) directly to get:

$$\begin{aligned} &\frac{\widetilde{\mathsf{M}}(0)}{\widetilde{\mathsf{M}}(\Delta^*)} \cdot Q(\Delta^*) \cdot K \\ &\qquad > \frac{1}{K - (K-1)e^{-c_\Psi \Delta}} \cdot \frac{K - (K-1)e^{-c_\Psi \Delta}}{K} \cdot K = 1 \end{aligned}$$

Since this holds for all $\Delta^* > 0$, no pure strategy profile $\vec{\Delta}^*$ can be a Nash equilibrium. This establishes uniqueness of the honest equilibrium at $\vec{\mathbf{0}}$ and completes the proof. □

A.3.2 Joint Proposal Designs. We only provide brief intuition here.

First, consider the uniqueness analysis from the proof of Theorem 5.1. Under the profile $\vec{\Delta}^*$, the conditions under which a player wants to deviate are exactly as before—it will receive the same reward for deviating to 0 and does not care about leftover MEV reward captured by other players in a joint proposal. This means that the same MEV condition from Theorem 5.1 will also establish here, the non-existence of equilibrium at $\vec{\Delta}^* \neq \vec{\mathbf{0}}$.

However, the condition for $\vec{\mathbf{0}}$ itself to be an equilibrium is stricter. To see why, consider the profile $\vec{\mathbf{0}}$ and suppose that player 1 wishes to deviate to playing Δ_1. To receive non-zero utility, its timestamp $\Delta_1 + \psi_1$ only needs to be lower than smaller than $K - \kappa$ others (instead of all K in the case with just one winner). This means that it is incentivized to deviate under a larger set of parameters, making the Nash equilibrium condition stricter.