

Price of Censorship: Censorship Resistance and Throughput Under Rational Concurrent Proposers

Aditya Saraf 

Cornell University, Ithaca, NY, USA

Ioannis Kaklamanis 

Yale University, IC3, New Haven, CT, USA

Sarisht Wadhwa 

Duke University, Durham, NC, USA

Fatima Elsheimy 

Yale University, New Haven, CT, USA

Abstract

Censorship resistance is the defining advantage of blockchains over their centralized counterparts. Yet block proposers censor transactions for many reasons, from legal consequences to economic incentives. We study *economically-incentivized* censorship, modeled by an adversary who bribes proposers to exclude a target transaction, and define the *economic censorship resistance* (eCR) of a transaction as the adversary's expected cost of successful censorship divided by the user's expected payment for inclusion. Single-proposer systems are structurally weak by this measure: under a first-price auction the adversary need only match the user's bid, and fee burning pushes eCR to a few percent of what the user pays.

We therefore turn to multiple concurrent proposers (MCP), where block capacity is divided among n proposers and the block is the union of their sub-blocks. While MCP can substantially increase the cost of censorship by requiring the adversary to bribe many proposers, it also introduces transaction duplication, reducing throughput. The resulting trade-off depends critically on the transaction fee mechanism (TFM), which determines how fees are shared among competing proposers.

We create a game theoretic model where validators construct blocks from a shared mempool, subject to an adversary's bribery attempt. We provide an algorithm that solves for the mixed equilibrium of a given mempool, which is characterized by the probability of including each transaction. This algorithm works for a wide class of TFMs, and allows us to calculate the expected throughput and censorship resistance for any bid distribution. We then use simulations to show how the eCR and throughput vary as the number of proposers increases. We compare three TFMs, finding that the *duplication-penalizing* TFM dominates the others across many settings. We also validate our findings with empirical Ethereum data.

We find that MCP systems enjoy much higher eCR than single proposer systems. While burnt user fees can significantly lower eCR, eCR scales linearly with the number of validators, and thus strong eCR can be achieved even in the presence of significant burn. Overall, our results show that MCP systems are, with the correct TFM, significantly more robust to censorship.

2012 ACM Subject Classification Security and privacy → Economics of security and privacy

Keywords and phrases blockchains, multiple concurrent proposers, censorship resistance

Digital Object Identifier 10.4230/LIPIcs.AFT.2026.26

Related Version *Full Version*: <https://arxiv.org/abs/2607.16995> [29]

Supplementary Material *Dataset (Source code and data)*: <https://zenodo.org/records/21831640>

Funding *Aditya Saraf*: AFOSR Grant FA9550-24-1-0267

Sarisht Wadhwa: NSF Grant CNS-2237814

Fatima Elsheimy: NSF Grant CNS-2247308, CNS-2153768

Acknowledgements Ioannis thanks IC3 for a research gift. Fatima thanks Center for Algorithms, Data, and Market Design at Yale (CADMY) and Protocol Labs.



© Aditya Saraf, Ioannis Kaklamanis, Sarisht Wadhwa, and Fatima Elsheimy;
licensed under Creative Commons License CC-BY 4.0

8th Conference on Advances in Financial Technologies (AFT 2026).

Editors: Aggelos Kiayias and Maria Kyropoulou; Article No. 26; pp. 26:1–26:24

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Blockchains are decentralized across many blocks, but within a single block they are anything but: the current proposer is a monopolist with complete control over the block’s contents. Decentralization is meant to give every transaction in the mempool a “fair chance” at inclusion, relative to its bid, unlike centralized systems, where the operator can simply refuse to process a payment or a post [15]. In practice, the single-slot monopoly undermines this promise. Wahrstätter et al. [38] measure that 46% of post-merge Ethereum blocks were produced by actors censoring OFAC-sanctioned transactions, and as of early 2026, live dashboards report that roughly 35% of Ethereum builders and half of all relays engage in some form of transaction censorship [11]. These measurements suggest that transaction censorship is a persistent feature of today’s blockchain ecosystem.

Some of this censorship is legally motivated, and its cost is mainly latency: a transaction censored by regulated proposers waits until a validator in a different jurisdiction proposes. Economically motivated censorship is more corrosive, because the censor can profit from every slot of delay. Consider a participant in an on-chain auction, a DeFi liquidation, or an arbitrage race: excluding a competitor’s transaction for even a few blocks can be worth orders of magnitude more than the fees that transaction pays. Since a bribe only needs to outbid the proposer’s *fee* from including the transaction, not the value the transaction protects, censorship is profitable precisely when value at stake far exceeds fees, which is the defining regime of maximal extractable value (MEV). An attacker can thus censor a competitor indefinitely by repeatedly offering each proposer a modest bribe.

In this paper, we focus on such *economically incentivized censorship*, which we model as an adversary who publicly offers every proposer a bribe for excluding a target transaction. We quantify a system’s resilience by its *economic censorship resistance* (eCR): the ratio between the adversary’s expected total expenditure for successful censorship and the fee the user expects to pay for inclusion. For example, if a user expects to pay 1 ETH, but an adversary must spend 5 ETH to censor the transaction, then the eCR is 5.

The eCR of existing designs is inherently limited. With a single proposer per block and a first-price auction transaction fee mechanism (TFM), the eCR is at most 1: the adversary simply matches the user’s bid. Fee burning makes matters strictly worse. On Ethereum, the proposer receives only the tip while the base fee is burned; in typical network conditions, tips are on the order of 0.5–2 gwei against base fees of 10–50 gwei, yielding $\text{eCR} \approx \frac{\text{tip}}{\text{base fee} + \text{tip}} \in [0.02, 0.1]$. An adversary can censor a transaction on today’s Ethereum by paying a few percent of what the user pays.

Multiple concurrent proposers. To achieve higher censorship resistance, we must break the proposer’s monopoly by having multiple concurrent proposers (MCP). MCP designs range from research protocols such as MirBFT [31], to DAG-based systems deployed in production by Sui and Aptos [24, 6, 4], to active proposals for Ethereum and Solana [21].

We propose a simple MCP model that captures many of these systems: the total block capacity is divided among n concurrent proposers, and the final block is the union of their sub-blocks. Some transactions may then be included by multiple proposers, which yields two complications. First, the TFM must dictate how a transaction’s fee is split among the proposers. [20] analyzes a TFM that pools all tips and divides them equally among proposers, regardless of who includes which transactions. [17] solicits a large and a small tip from each user: a transaction’s sole includer receives the large tip, while multiple includers each receive the small one. As the Ethereum calculation above suggests, this choice heavily affects

censorship resistance. We mostly analyze an extreme version of the TFM of [17], which we term the *duplication-penalizing* TFM: it solicits a single tip and pays it to the includer only if the transaction is included exactly once. Our equilibrium analysis, however, holds for a large family of TFMs including all these examples.

The second complication is that multiple inclusions cause a loss of *throughput*, which we define as the percentage of *unique* transactions in a block. There is a natural tension between censorship resistance and throughput: doubling the number of proposers increases the amount the adversary must spend, but also creates more opportunities for duplication. The TFM again matters, as different TFMs discourage duplication to different degrees.

MCP systems are adopted precisely for these two advertised benefits: higher throughput, from parallel block production, and censorship resistance, from the fact that no single proposer can unilaterally exclude a transaction. But a design cannot simply maximize both. The redundancy that protects transactions manifests, in equilibrium, as duplication that consumes the very capacity parallelism creates, and where a system lands on this frontier depends directly on its transaction fee mechanism. This dependence frames the central question of our work:

Which transaction fee mechanism best navigates the trade-off between censorship resistance and throughput, and how do both scale with the number of concurrent proposers?

Main contributions

To answer this question, we perform a game-theoretic analysis: due to the possibility of transaction duplication, a proposer’s utility from their block depends on the other proposers. Our threat model is deliberately simple and stated up front: the adversary announces a public, unconditional bribe paid to any proposer that excludes the target transaction; proposers are rational, do not collude (collusion is typically equivalent to a single proposer), and build their blocks simultaneously from a shared mempool; and user tips are exogenous. Richer adversaries (e.g., bribes conditional on successful censorship) and endogenous user bidding are natural extensions that we discuss but do not model.

We solve this game algorithmically for its unique symmetric mixed Nash equilibrium, which specifies the probability of including each transaction as a function of the bids. Our algorithm is a one-dimensional bisection, is efficient, and applies to a wide family of TFMs. A “cut-and-paste” argument then yields closed-form expressions for the bribe required to censor any target transaction with any target probability.

With this machinery, we compute the expected throughput, user payments, and cost of bribery for any bid distribution. On synthetic (Pareto) bids, we vary the number of proposers, the block congestion, the adversary’s target censorship probability, and the amount of burn, comparing three TFMs: duplication-penalizing, even-split, and collective. The duplication-penalizing TFM consistently dominates the other two on both eCR and throughput. We then validate these findings on the empirical tip distribution of 10,000 recent Ethereum blocks.

Quantitatively, we find that eCR scales linearly with the number of proposers, while throughput quickly plateaus near the optimistic bound achieved by uncoordinated uniform sampling; thus, beyond the one-time duplication cost of moving away from a single proposer, adding proposers buys censorship resistance at essentially no additional throughput cost. This scaling survives fee burning: at Ethereum’s early-2026 burn levels (base fee $\approx 8.7\times$ the average tip), a single proposer yields eCR ≈ 0.09 , whereas 49 concurrent proposers yield eCR ≈ 2.5 ; the adversary must outspend the user rather than pay a few percent of their fee. Overall, our results show that MCP, equipped with the right TFM, is a promising structural remedy to economic censorship.

2 Related Work

Protocol-level censorship resistance. Classical analyses of blockchain protocols establish *liveness*: every valid transaction is eventually included [19, 43, 26]; see [2] for a survey of the censorship-resistance guarantees of widely-used consensus protocols. Eventual inclusion, however, is too weak for financial applications, motivating *short-term* censorship resistance in the Byzantine model. BigDipper [41] adds short-term censorship resistance to leader-based BFT, and Abraham, Efron, and Ren [1] prove inherent latency costs of censorship resistance. Closest to our setting, concurrent work by Elsheimy et al. [16] formalizes the censorship-resistance/throughput trade-off in multi-proposer BFT protocols with up to f Byzantine parties, and designs transaction-assignment rules that navigate this trade-off. This line of work bounds the *worst-case delay* that faulty parties can impose; in contrast, we assume all proposers are rational and instead *price* censorship under bribery. The two perspectives are complementary: Byzantine analyses cover arbitrary deviations by a bounded fraction of parties, while ours covers economically-motivated deviations by all parties.

Transaction fee mechanisms. Transaction fee mechanisms (TFMs) in blockchain systems have evolved from first-price auctions (as in Bitcoin) to hybrid models like Ethereum’s EIP-1559. Roughgarden’s [28] analysis of EIP-1559 showed that it steers the system toward market-clearing prices under demand, but also that no mechanism can be both fully strategy-proof and resistant to collusion or bribery. Chung and Shi [12] sharpened this into a formal impossibility: no non-trivial TFM can simultaneously be incentive-compatible for users and for the block producer while resisting producer–user collusion. These results define a benchmark for later designs: TFMs must trade off between efficiency, incentive-compatibility, and resistance to strategic censorship or bribery.

Recent research has begun to extend TFM design to multi-proposer settings. One line of work, such as FOCIL (Fork-Choice Enforced Inclusion Lists), uses committees of validators to construct transaction inclusion lists before block building [36]. However, how to fairly allocate fees among multiple proposers remained an open question. Stouka et al. [32] introduce a game-theoretic model for multi-proposer blockchains (motivated by Ethereum’s FOCIL design). They extend Roughgarden’s canonical model by adding stages where several proposers each choose which transactions to include and by introducing a bribing adversary that can pay proposers to censor a transaction. Using this framework, they design and evaluate concrete multi-proposer TFMs (e.g. “Double” and “Single” TFMs) that maintain the Bayesian incentive-compatibility guarantees of EIP-1559 while greatly raising the cost of censorship. In effect, Stouka et al. show that appropriate fee-splitting rules can preserve honest behavior and enforce inclusion even under bribery. Notably, their adversary can condition bribes on successful censorship, which is generally cheaper for the adversary than the unconditional bribes we analyze; extending our equilibrium computation to conditional bribes is a natural direction for future work (see Section 7).

Fox et al. [17] and others quantify the “cost of censorship” in economic terms and demonstrate that multi-proposer TFMs can raise this cost by distributing responsibility for inclusion. Some mechanisms reward proposers based on the number of independent inclusions of a transaction, effectively discouraging selective omission [17, 32]. At the application layer, Alpos et al. [3] design censorship-resistant sealed-bid auctions, complementing the auction-driven motivation of [17].

Bribery attacks and censorship in practice. Bribery attacks on consensus have a long history: Bonneau [8] observed that adversaries can *rent* rather than buy consensus power, McCorry et al. [25] constructed smart contracts that trustlessly bribe miners, Winzer et al. [40] analyzed temporary censorship attacks against rational miners, Sun et al. [33] give a practical quantification of bribery attacks, and Karakostas et al. [23] study the efficacy of counterincentives. Closest to our notion of economic censorship, Berger et al. [7] study bribery attacks against optimistic-rollup fraud proofs. Their setting considers repeated censorship over many sequential blocks, whereas we study the cost of censoring a transaction within a single block under rational proposer behavior. The two works are complementary: their analysis assumes a per-block censorship cost, while our eCR metric quantifies exactly that cost. Nor is censorship hypothetical: beyond the measurements of Wahrstätter et al. [38], empirical studies of Ethereum’s proposer–builder separation document widespread filtering by builders and relays [22], heavy concentration of the builder market [42], and the centralizing role of private order flow [39] – the latter directly motivating the private-mempool extension we discuss in Section 7.

In parallel, new multi-leader consensus protocols have been proposed to boost throughput and resilience. Building on DAG mempools and consensus [14, 30, 18], DAG-based or “leaderless” designs (such as those used in Avalanche [5], Sui’s Mysticeti [6], or Aptos [4]) allow every validator to propose blocks simultaneously, vastly improving throughput compared to single-leader chains. Garimidi, Neu, and Resnick [21] make the general case for multiple concurrent proposers and design a protocol achieving selective-censorship resistance together with transaction hiding. Garimidi et al. [20] formalize fee mechanisms for such leaderless protocols: they define a “strongly BPIC” equilibrium concept and propose the FPA-EQ mechanism, which splits first-price auction revenue equally among all contributing blocks. They prove FPA-EQ is strongly incentive-compatible and guarantees at least a 63.2% fraction of the optimal welfare at equilibrium. These works (and others on Avalanche and Narwhal) highlight that multiple concurrent proposals require fundamentally different TFM designs, but can achieve higher performance and improved censorship resistance if fees are shared appropriately.

Collectively, these works highlight a growing interest in mechanisms that jointly optimize censorship resistance, incentive compatibility, and throughput in decentralized settings. Our work contributes to this space by modeling bribery and duplication in MCP systems and analyzing the equilibrium effects of different TFMs.

Comparison to inclusion lists. Inclusion lists are a powerful tool for censorship resistance. They originate from a line of Ethereum research on censorship under proposer–builder separation [9], refined through forward inclusion lists [13], EIP-7547 [27], the “no free lunch” design [10], and committee-enforced inclusion sets [35], culminating in FOCIL [36], now drafted as EIP-7805 [34]. Their main strength lies in their strictness; proposers sacrifice all of their reward for censoring transactions on the inclusion list. This is also a weakness. Proposers in certain jurisdictions may be legally required to censor certain transactions. By preventing this, inclusion lists may result in more geographic centralization.

A related recent proposal is AUCIL (Auction-based Inclusion Lists) [37], which is designed for committee-based inclusion. In AUCIL, a committee of proposers each compiles an input list of transactions where the number of overlaps is based on the fee paid by the transaction, and then an on-chain auction algorithm merges these lists into a final committed inclusion set. The key idea is that by having each committee member assigned certain transactions (using a correlated equilibrium based algorithm), the protocol ensures that high-paying

transactions appear on at least one list. An adversary wishing to censor a transaction must therefore bribe all proposers who listed it, which is prohibitively expensive. In practice, AUCIL’s auction step selects as many of the committee input lists as possible to form the block, thereby “authentically” binding each proposer to its list and making censorship very costly. This mechanism complements the multi-proposer TFMs above by enforcing that committed transactions are actually included, further strengthening censorship resistance in multi-proposer or proposer-builder systems. As a point of comparison to our design, AUCIL employs the even-split TFM, one of the admissible TFMs analyzed in this work.

In contrast, MCP presents a weaker, but more flexible, tool for censorship resistance. Proposers can still censor transactions if legally necessary, but that makes such transactions even more lucrative for proposers who can include it, helping to avoid geographic centralization. Also, MCP is desired for other reasons; namely, it can reduce MEV by limiting the ability of proposers to control the exact ordering of the final block. In that sense, our work analyzes the “free” censorship resistance that comes along with MCP; beyond the consensus protocol and TFM, no additional mechanism is required.

3 Model and Definitions

We now explain our model. There are n proposers who collectively build a block of size B . Each proposer submits a sub-block of at most $\ell := B/n$ transactions (we assume for simplicity that n divides B). These transactions are chosen from a *shared mempool*. We denote this by the sorted list $\{t_1, t_2, \dots, t_m\}$, where t_i refers to the tip of transaction i , and $t_1 \geq t_2 \geq \dots \geq t_m$. We sometimes abuse notation to use t_i to refer to the tip or the transaction itself, when it is clear from the context.

The union of the proposers’ sub-blocks form the final list of executed transactions. The TFM then portions out the proposers’ payments as a function of the full list of sub-blocks. We consider the following three TFMs in this paper:

1. the *duplication-penalizing* TFM. If a transaction t_i is included by only one proposer, then the proposer who included it is paid the full t_i . If a transaction is included by more than one proposer, then the transaction is not charged, and no tip is given to the proposers.
2. the *even-split* TFM. If a transaction t_i is included by k different proposers, then it is charged t_i , and every proposer that includes t_i is paid t_i/k .
3. the *collective* TFM. All transactions are charged their tips, which are pooled and shared equally among *all* proposers.

Further, we assume that the TFM can specify a *base fee* β that is burnt. So, each included transaction pays a flat fee β as well as its payment from the auction.

► **Remark 1 (User payments and revenue).** Our analysis takes tips as exogenous and studies proposer incentives; we do not model users’ bidding behavior or user–proposer collusion. Two features of the duplication-penalizing TFM deserve emphasis in this light. First, duplicated transactions are executed free of charge, so this TFM forgoes fee revenue relative to the even-split TFM, which charges every included transaction; our eCR metric reflects this through lower expected user payments. Second, a user able to coordinate with even a single proposer could deliberately trigger duplication to avoid paying. The two-tier tips of Fox et al. [17], of which our duplication-penalizing TFM is the extreme case, mitigate both effects by charging a small tip upon duplicated inclusion; such intermediate TFMs are admissible in the sense of Definition 8, so our equilibrium framework applies to them verbatim. We analyze the extreme case for clarity, and leave a user-side equilibrium analysis to future work.

We assume the presence of an adversary who tries to censor a single transaction, t_c . The adversary offers an unconditional bribe of b to all proposers who do not include t_c . This bribe is common knowledge; all proposers know that all other proposers have been offered the same bribe.

The utility of a proposer is the sum of their payments, both from transactions and possibly from the bribe. Each proposer's strategy space is all subsets of size B/n from the mempool. Our solution concept is a symmetric, mixed Nash equilibrium, which specifies the probability of sampling *each subset* of transactions. But for our purposes, the only relevant feature of the equilibrium is its *induced probabilities* over the transactions, rather than the full distribution over all subsets. We thus represent the equilibrium as $\mathbf{p} = (p_1, \dots, p_m)$, where p_i is the probability that the proposer includes transaction i in equilibrium.

For the following definitions, we fix the number of proposers, n , the total block size, B , the mempool, $\{tx_1, \dots, tx_m\}$, and the TFM, $\mathbf{T}$, in which each proposer that includes a transaction tx_i receives $\mathbf{T}(t_i, n_i)$, where t_i is the transaction's tip and n_i is the number of *other* proposers that also include it (so a sole includer receives $\mathbf{T}(t_i, 0)$). In certain cases, even without including the transaction, the proposer can gain a certain value. This induces a unique equilibrium profile, $\mathbf{p}$, as we will show later.

► **Definition 2** (fixed-mempool throughput). *The throughput of the current mempool is the expected number of unique transactions in the final block (given that all proposers play $\mathbf{p}$), divided by $\min(B, m)$.*

So, the maximum throughput occurs when all transactions are included, or, if that is not possible, when the block is full with no duplicate transactions.

Finally, we turn to economic censorship resistance, defining it first on a per-transaction basis, assuming a fixed mempool state.

► **Definition 3** (transaction cost). *The transaction cost of tx_i is its expected payment conditional on inclusion in the final block: the flat base fee β plus the expected auction payment of tx_i , in the equilibrium induced by the mempool state and the TFM, conditioned on the event that tx_i appears in the final block.*

This is the quantity that a user who ends up included actually pays; transactions that are excluded pay nothing, and normalizing by the conditional payment makes the eCR of transactions with different inclusion probabilities comparable. To define the cost of bribery, we first assume that the adversary's goal is for t_c to appear in the final block with probability at most $1 - s$. Thus, s is the adversary's desired rejection probability. For ease of exposition, we typically assume that $s = 0.9$, but we later show how our simulation results vary with s . We also rely on the fact that, for any bribe b , there is a unique equilibrium profile $\mathbf{p}^b$, which we show later.

► **Definition 4** (bribery cost). *Suppose an adversary wishes for t_c to be censored with probability s , and let b^* be the minimum per-proposer bribe such that, under the induced equilibrium profile $\mathbf{p}^{b^*}$, the probability that t_c is included in the final block is at most $1 - s$. The bribery cost of transaction c is the adversary's expected total expenditure, $n(1 - p_c^{b^*})b^*$: the bribe is offered to each of the n proposers, and each collects it (by excluding t_c) with probability $1 - p_c^{b^*}$.*

► **Definition 5** (per-tx eCR). *The economic censorship resistance of transaction t_i is its bribery cost divided by its transaction cost.*

The intuition behind our definition is that transactions which cost more are likely from users with higher valuations, and should thus be harder to censor. Put differently, a bribery cost of 10 and a transaction cost of 1 represents strong censorship resistance; the adversary must spend 10 times the user. If the transaction cost is itself 10, this is no longer as resistant to censorship.

Since all of these definitions depend on the state of the mempool, we add a distributional assumption. So suppose all t_i 's are drawn iid from distribution F .

► **Definition 6 (throughput).** *The expected throughput under F is the expectation, over mempools $\mathbf{t} \sim F^m$, of the throughput of $\mathbf{t}$.*

► **Definition 7 (eCR).** *The expected censorship resistance under F is the expectation, over mempools $\mathbf{t} \sim F^m$, of the average economic censorship resistance of a transaction $t_i \in \mathbf{t}$.*

4 A Generalized Equilibrium Algorithm

In this section, we do not consider a fixed TFM or bribe. Instead, we provide an algorithm that solves the symmetric mixed Nash equilibrium, given *payoff functions*, that describe how the utility of a proposer changes as they include a transaction, given the other proposers' probability of including that transaction. Our algorithm works for a range of admissible payoff functions, that include all three TFMs we consider, in the absence or presence of any bribe. Intuitively, an admissible payoff function is one where the utility of including transaction i decreases as the probability that others include that transaction increases. In the simplest case of the duplication-penalizing TFM, this holds as the utility of including a transaction is only nonzero if no one else includes it.

4.1 Symmetric Equilibrium via a Linear Program

We characterize symmetric mixed-strategy equilibria of the proposer game as optima of a linear program, and use linear-programming duality to derive a relation between the fee paid by transactions and the probability of its selection. This relation, that all transactions assigned interior probability share a common marginal payoff λ^* is then used to prove a symmetric (ϵ -approximate) equilibrium in Algorithm 1.

Setup. Consider a set of transactions $tx_1, \dots, tx_m$, each with a single tip value t_i . We seek a symmetric mixed strategy: a probability vector $\vec{p} = (p_1, \dots, p_m) \in [0, 1]^m$ used by every proposer, where p_i is the probability of including tx_i . If, in addition to the proposer under consideration, n_i other proposers include tx_i , the proposer's payment for it is $\mathbf{T}(t_i, n_i)$, following the convention of Section 3. We assume throughout that $m > \ell$; otherwise the block trivially fits the entire mempool.

For each transaction tx_i , let $\bar{\lambda}_i$ denote the maximum possible payoff to a proposer (usually when it is the sole includer of tx_i), and $\underline{\lambda}_i$ the minimum possible payoff (usually when all proposers include tx_i). Fix a profile of beliefs $\vec{p}^* \in [0, 1]^m$ representing what every other proposer will do.

From a proposer's perspective, the expected marginal payoff of including tx_i is a function of the probability that other proposers select this transaction, and t_i . Let u_i be the marginal utility for the proposer (utility for including the transaction minus utility for not including the transaction) when including the transaction tx_i . If every other proposer includes tx_i independently with probability p_i , then

$$u_i = \sum_{n_i=0}^{n-1} \binom{n-1}{n_i} p_i^{n_i} (1-p_i)^{n-1-n_i} \mathbf{T}(t_i, n_i) \quad (1)$$

where, with slight abuse of notation, $\mathbf{T}(t_i, n_i)$ in (1) denotes the *marginal* payoff – the difference in a proposer’s payment between including and not including tx_i – given that n_i others include it. For the duplication-penalizing and even-split TFMs this coincides with the payment itself; for the collective TFM it is the change in the proposer’s pool share (Example 9).

Best response as a linear program. Treating $\vec{p}^*$ (and hence each u_i) as fixed, a deviator’s best response maximizes total expected payoff subject to a block-size budget ℓ :

$$\begin{aligned} \max_{\vec{p}} \quad & \sum_{i=1}^m u_i p_i \\ \text{s.t.} \quad & \sum_{i=1}^m p_i = \ell, \\ & 0 \leq p_i \leq 1 \quad \forall i \in \{1, \dots, m\}. \end{aligned} \tag{2}$$

A symmetric equilibrium then corresponds to a fixed point at which $\vec{p} = \vec{p}^*$ is itself optimal for (2) with the marginal payoffs u_i induced by $\vec{p}^*$ via (1). (We impose the block-size budget with equality; since payments are non-negative, only a bribed transaction can have negative marginal utility, and as $m > \ell$ there are always at least ℓ transactions with non-negative marginal utility to fill the sub-block, so this is without loss of generality.)

► **Property 1** (Equilibrium structure). *There exists an optimal solution $\vec{p}^*$ to (2) and a scalar $\lambda^* \in \mathbb{R}$ such that, for every $i \in \{1, \dots, m\}$:*

1. *if $u_i > \lambda^*$, then $p_i^* = 1$;*
2. *if $u_i < \lambda^*$, then $p_i^* = 0$;*
3. *if $0 < p_i^* < 1$, then $u_i = \lambda^*$.*

Conversely, any feasible $\vec{p}$ for which some λ satisfies conditions 1–3 is optimal for (2).

Proof sketch. The claim follows from linear-programming duality: introducing a dual variable $a_i \geq 0$ for each capacity constraint $p_i \leq 1$ and λ for the clearing constraint, complementary slackness forces $p_i^* = 1$ when $u_i > \lambda^*$, $p_i^* = 0$ when $u_i < \lambda^*$, and $u_i = \lambda^*$ on the interior; the converse direction verifies that any profile satisfying conditions 1–3 admits a dual-feasible certificate. The full proof appears in the full version [29]. ◀

Property 1 reduces the problem of finding a symmetric equilibrium to identifying the threshold λ^* and the corresponding inclusion probabilities. The next subsection presents an efficient binary-search algorithm for this task.

4.2 A Binary-Search Algorithm for Symmetric Equilibria

Building on Property 1, we now give an efficient algorithm to compute symmetric equilibria. The idea is to invert the marginal-payoff relationship (1): at a symmetric fixed point $\vec{p} = \vec{p}^*$, equation (1) couples each p_i to u_i , and Property 1 couples each interior p_i to a common threshold λ^* . Composing these, the equilibrium probability of each transaction is determined by λ^* alone, reducing the search to a one-dimensional bisection over λ .

Setup. For each transaction tx_i , let $f_i(\lambda)$ denote the symmetric inclusion probability that makes a proposer indifferent at marginal payoff λ : the value of p_i for which (1) yields $u_i = \lambda$ when every other proposer also plays p_i . Concretely, for the marginal-payoff form

of Example 9, f_i admits a closed form; in general it can be computed by inverting (1) numerically. By construction, f_i takes the value 1 at $\lambda = \underline{\lambda}_i$ (full competition) and 0 at $\lambda = \bar{\lambda}_i$ (monopoly inclusion), matching the bounds defined in Section 4.1.

► **Definition 8** (Admissible payoff function). *A function $f_i: \mathbb{R} \rightarrow \mathbb{R}$ is admissible with parameters $(\underline{\lambda}_i, \bar{\lambda}_i)$ if:*

- **(A1) Boundary values:** $f_i(\underline{\lambda}_i) = 1$ and $f_i(\bar{\lambda}_i) = 0$.
- **(A2) Continuity and strict monotonicity:** f_i is continuous and strictly decreasing on $[\underline{\lambda}_i, \bar{\lambda}_i]$.
- **(A3) Range:** $f_i(\lambda) > 1$ for $\lambda < \underline{\lambda}_i$ and $f_i(\lambda) < 0$ for $\lambda > \bar{\lambda}_i$.

► **Example 9** (Collective TFM). Suppose all fees from included transactions are pooled and divided equally among the n proposers, regardless of who included them. Proposer j 's total payoff is

$$U_j = \frac{1}{n} \sum_{i=1}^m T_i \cdot \mathbf{1}[tx_i \text{ included by at least one proposer}].$$

Because j receives the same share of T_i whether or not she personally includes tx_i (provided someone does), the *marginal* value of inclusion equals the expected pool contribution conditional on every other proposer omitting tx_i :

$$\Delta u_i(p_i) = \frac{T_i}{n} (1 - p_i)^{n-1}.$$

Setting $\Delta u_i(p_i) = \lambda$ and inverting yields

$$f_i(\lambda) = 1 - \left(\frac{n\lambda}{T_i} \right)^{\frac{1}{n-1}}, \quad \underline{\lambda}_i = 0, \quad \bar{\lambda}_i = \frac{T_i}{n}.$$

Verification of admissibility. **(A1)** $f_i(0) = 1$ and $f_i(T_i/n) = 0$. **(A2)** The map $\lambda \mapsto (n\lambda/T_i)^{1/(n-1)}$ is strictly increasing for $\lambda \geq 0$, so f_i is strictly decreasing on $[0, T_i/n]$. **(A3)** For $\lambda > T_i/n$ the inner term exceeds 1, so $f_i(\lambda) < 0$; the natural extension to $\lambda < 0$ gives $f_i(\lambda) > 1$.

► **Example 10** (Duplication-penalizing TFM). Under this TFM a transaction pays its tip T_i to its includer only when it is included exactly once, so the marginal value of including tx_i is nonzero only in the event that no other proposer includes it:

$$u_i(p_i) = (1 - p_i)^{n-1} T_i, \quad f_i(\lambda) = 1 - \left(\frac{\lambda}{T_i} \right)^{\frac{1}{n-1}}, \quad \underline{\lambda}_i = 0, \quad \bar{\lambda}_i = T_i.$$

This is the same functional form as Example 9 with T_i in place of T_i/n , and admissibility is verified identically.

► **Example 11** (Even-split TFM). Suppose a transaction included by k proposers is charged its full tip T_i , which is split evenly, so that each includer receives T_i/k . If every other proposer includes tx_i independently with probability p_i , the number of *other* includers is $K \sim \text{Bin}(n-1, p_i)$, and the marginal value of inclusion is

$$u_i(p_i) = T_i \cdot \mathbb{E} \left[\frac{1}{1+K} \right] = T_i \sum_{k=0}^{n-1} \binom{n-1}{k} \frac{p_i^k (1-p_i)^{n-1-k}}{k+1} = T_i \frac{1 - (1-p_i)^n}{n p_i},$$

where the last equality uses $\binom{n-1}{k} \frac{1}{k+1} = \frac{1}{n} \binom{n}{k+1}$ and the binomial theorem, and $u_i(0) := T_i$ by continuity.

u_i is strictly decreasing with $u_i(0) = T_i$ and $u_i(1) = T_i/n$, so its inverse f_i is admissible with parameters $(\underline{\lambda}_i, \bar{\lambda}_i) = (T_i/n, T_i)$; the verification of (A1)–(A3) appears in the full version [29].

Unlike the collective TFM, u_i here admits no closed-form inverse. This poses no obstacle: Algorithm 1 only requires *pointwise evaluation* of f_i , and since u_i is continuous and strictly decreasing, $f_i(\lambda)$ can be evaluated to any accuracy ε' by an inner bisection on $[0, 1]$, adding a factor $O(\log(1/\varepsilon'))$ to the per-evaluation cost in Theorem 14.

A tractable approximation. Replacing $\mathbb{E}[1/(1+K)]$ by $1/(1+\mathbb{E}[K])$ yields a closed-form admissible approximation of u_i that understates the value of inclusion (by Jensen’s inequality) and shares its boundary values; we use it only for intuition, and all reported numerical results are computed with the exact u_i . Details appear in the full version [29].

The clamped aggregate. For any candidate $\lambda \in \mathbb{R}$, define the *clamped total inclusion mass*

$$G(\lambda) = \sum_{i=1}^m \min(\max(f_i(\lambda), 0), 1).$$

By admissibility, the i th term equals 1 for $\lambda \leq \underline{\lambda}_i$, equals 0 for $\lambda \geq \bar{\lambda}_i$, and is strictly decreasing in between. Hence G is non-increasing on $\mathbb{R}$, with $G(\min_i \underline{\lambda}_i) = m$ and $G(\max_i \bar{\lambda}_i) = 0$. Combining Property 1 with the definition of f_i , finding a symmetric equilibrium reduces to locating λ^* with $G(\lambda^*) = \ell$, where ℓ is the target block size. Since G is monotone, a one-dimensional bisection suffices.

► **Remark 12 (Marginals suffice).** Two observations justify working with inclusion probabilities rather than full distributions over sub-blocks. First, since proposers randomize independently and a proposer’s payoff is additive over transactions, with the term for tx_i depending on the others only through the number of them that include tx_i (distributed $\text{Bin}(n-1, p_i)$), opponents’ strategies affect a proposer’s payoff only through their marginals; hence the best-response problem over distributions on ℓ -subsets reduces to the linear program (2) over marginals. Second, every feasible point of (2) is realizable: for any $\vec{p} \in [0, 1]^m$ with $\sum_i p_i = \ell \in \mathbb{N}$ there is a random ℓ -subset S with $\Pr[i \in S] = p_i$. (This follows from a standard systematic-sampling construction; see the full version [29].) Together, these imply that optimal marginal vectors correspond exactly to best responses in the original strategy space.

► **Theorem 13 (Existence and uniqueness).** *For any set of transactions $tx_1, \dots, tx_m$ with admissible payoff functions $f_1, \dots, f_m$ and any target block size $\ell \in (0, m)$, there exists λ^* with $G(\lambda^*) = \ell$, and $\mathbf{p}^* = (\hat{f}_1(\lambda^*), \dots, \hat{f}_m(\lambda^*))$ is a symmetric mixed Nash equilibrium. Moreover, the equilibrium profile is unique: any two thresholds with clamped inclusion mass ℓ induce the same probability vector $\mathbf{p}^*$.*

Proof sketch. *Existence:* G is continuous and non-increasing with $G(\min_i \underline{\lambda}_i) = m$ and $G(\max_i \bar{\lambda}_i) = 0$, so the intermediate value theorem yields λ^* with $G(\lambda^*) = \ell$, and the induced profile satisfies the conditions of Property 1 (realizable over ℓ -subsets by Remark 12). *Uniqueness:* any two thresholds with clamped mass ℓ induce identical probability vectors, since each $\hat{f}_i$ is non-increasing and the sums agree. The full proof appears in the full version [29]. ◀

► **Theorem 14 (Termination and runtime).** *Let ω_G be a modulus of continuity of G on $[\min_i \underline{\lambda}_i, \max_i \bar{\lambda}_i]$ (one exists, since G is continuous on this compact interval), and fix any $\delta_\varepsilon > 0$ with $\omega_G(\delta_\varepsilon) \leq \varepsilon$. Then Algorithm 1 terminates after at most $\lceil \log_2(\Delta/\delta_\varepsilon) \rceil$ iterations, where $\Delta = \max_i \bar{\lambda}_i - \min_i \underline{\lambda}_i$. Each iteration evaluates G once at cost of $O(m)$ evaluations of the f_i , giving total runtime $O(m \log(\Delta/\delta_\varepsilon))$.*

■ **Algorithm 1** Symmetric-equilibrium binary search.

Require: Admissible functions $f_1, \dots, f_m$ with parameters $(\underline{\lambda}_i, \bar{\lambda}_i)$; block size ℓ ; tolerance $\varepsilon > 0$.

Ensure: Probabilities $(p_1^*, \dots, p_m^*)$ forming an ε -approximate symmetric equilibrium.

$G(\lambda)$ denotes the clamped total inclusion mass defined above.

```

1: function BINARYSEARCH( $f_1, \dots, f_m, \ell, \varepsilon$ )
2:   if  $m \leq \ell$  then return  $(1, 1, \dots, 1)$ 
3:   if  $\ell = 0$  then return  $(0, 0, \dots, 0)$ 
4:    $\lambda_{\min} \leftarrow \min_i \underline{\lambda}_i; \quad \lambda_{\max} \leftarrow \max_i \bar{\lambda}_i$            ▷  $G(\lambda_{\min}) = m \geq \ell, G(\lambda_{\max}) = 0 \leq \ell$ 
5:   loop
6:      $\lambda_{\text{mid}} \leftarrow (\lambda_{\min} + \lambda_{\max})/2$ 
7:     if  $|G(\lambda_{\text{mid}}) - \ell| \leq \varepsilon$  then
8:       break                                                     ▷ inclusion mass is within tolerance of  $\ell$ 
9:     else if  $G(\lambda_{\text{mid}}) > \ell$  then
10:       $\lambda_{\min} \leftarrow \lambda_{\text{mid}}$                                        ▷ too many inclusions; raise  $\lambda$ 
11:     else
12:       $\lambda_{\max} \leftarrow \lambda_{\text{mid}}$                                        ▷ too few inclusions; lower  $\lambda$ 
13:    $\lambda^* \leftarrow \lambda_{\text{mid}}$ 
14:   for  $i = 1, \dots, m$  do
15:      $p_i^* \leftarrow \min(\max(f_i(\lambda^*), 0), 1)$ 
16:   return  $(p_1^*, \dots, p_m^*)$ 

```

Proof. By Theorem 13 there exists λ^* with $G(\lambda^*) = \ell$, and the update rule maintains the invariant $\lambda_{\min} \leq \lambda^* \leq \lambda_{\max}$ (verified in the proof of Theorem 16). The bracket width halves each iteration: after k iterations it equals $\Delta/2^k$. Once $\Delta/2^k \leq \delta_\varepsilon$, the midpoint satisfies $|\lambda_{\text{mid}} - \lambda^*| \leq \delta_\varepsilon$, hence $|G(\lambda_{\text{mid}}) - \ell| = |G(\lambda_{\text{mid}}) - G(\lambda^*)| \leq \omega_G(\delta_\varepsilon) \leq \varepsilon$, and the termination condition is met. Each iteration computes G once, inspecting all m transactions; for TFMs whose f_i admits no closed form (e.g., the even-split TFM of Example 11), each evaluation additionally costs an inner bisection, multiplying the runtime by $O(\log(1/\varepsilon'))$ for inner accuracy ε' . ◀

With existence and uniqueness established, we can now prove that Algorithm 1 outputs an approximate equilibrium.

► **Definition 15** (ε -approximate equilibrium). *A probability vector $(p_1, \dots, p_m) \in [0, 1]^m$ is an ε -approximate symmetric mixed Nash equilibrium if there exists $\hat{\lambda} \in \mathbb{R}$ such that*

1. *the equilibrium conditions of Property 1 hold at $\hat{\lambda}$ with probabilities p_i , and*
2. $|\sum_{i=1}^m p_i - \ell| \leq \varepsilon$.

► **Theorem 16** (Correctness). *The output $(p_1^*, \dots, p_m^*)$ of Algorithm 1 is an ε -approximate symmetric mixed Nash equilibrium, where ε is the tolerance supplied to the algorithm.*

Proof sketch. By Theorem 13 there is a λ^* with $G(\lambda^*) = \ell$. The bisection invariant brackets λ^* , and after the stated number of iterations the interval is short enough that, by the modulus of continuity of G and monotonicity of each $\hat{f}_i$, the returned probabilities are ε -approximate in the stated sense. The full proof appears in the full version [29]. ◀

5 Deriving Economic Censorship Resistance

In this section, we describe a general “cut-and-paste” method to derive the bribe required to censor a target transaction with a given probability. Our method only requires computing the equilibrium probabilities for a *single* game. Using these probabilities, we can derive a closed-form expression for the bribe required, and by extension, for the economic censorship resistance (per Definition 5) of the target transaction. We then apply the cut-and-paste method to all three TFMs we study in this paper.

5.1 Cut-and-Paste Method

Setup. We recall some notation from earlier. Let $u_i(p_i)$ be the expected (increase in) utility of including tx_i assuming everyone else includes it with probability p_i . Given transaction tips $\mathbf{t} = (t_i)_{i \in [m]}$ and target probability mass ℓ , let $(\mathbf{p} = (p_i)_{i \in [m]}, \lambda) = \text{EQUILIBRIUM}(\mathbf{t}, \ell)$ be the equilibrium probabilities and common utility level, respectively. This means that $u_i(p_i) = \lambda$ for all $i \in [m]$. Also let $f_i(x)$ be the admissible function (per Definition 8) such that $p_i = f_i(\lambda)$.

Suppose the adversary $\mathcal{A}$ wishes to censor a target transaction tx_c (with tip t_c) for some $c \in [m]$. We are interested in determining the bribe amount b^* that $\mathcal{A}$ needs to offer to all parties, so that tx_c is censored with probability $s \in (0, 1]$. In order for $\mathcal{A}$ to achieve the target censorship probability s , each party must include tx_c with probability p_c^* such that $s = (1 - p_c^*)^n$, or equivalently, $p_c^* = 1 - s^{1/n}$. This holds regardless of the underlying TFM.

Cut-and-Paste Method. The method proceeds in two steps.

First, the *cut* step: solve for the equilibrium of the game with tx_c removed, reserving exactly p_c^* units of capacity to be filled in later. Let $(\mathbf{p}_{-c}, \lambda_{-c}) := \text{EQUILIBRIUM}(\mathbf{t}_{-c}, \ell - p_c^*)$ denote this equilibrium.

Second, the *paste* step: add tx_c back, with the adversary’s bribe b^* chosen so that, in the resulting equilibrium of the full game $G(\mathbf{t}, \ell)$, (1) tx_c is included with the target probability p_c^* ; and (2) every $tx_j \neq tx_c$ retains the probability p_j it had in $\mathbf{p}_{-c}$. We carry out this analysis in Section 5.2 and obtain the closed-form bribe $b^* = u_c(p_c^*) - \lambda_{-c}$.

5.2 Correctness of the Cut-and-Paste Method

We prove that the cut-and-paste method produces an equilibrium of $G(\mathbf{t}, \ell)$ under bribe b^* in which tx_c attains probability p_c^* and every $tx_j \neq tx_c$ retains its probability from the post-cut equilibrium $\mathbf{p}_{-c}$.

The argument relies on a structural property of equilibria captured by the following lemma: a new transaction can be added to the game without disturbing the existing inclusion probabilities or utility level, provided capacity is expanded by exactly the amount the new transaction consumes at level λ . The paste step will exploit this to add tx_c back into the post-cut equilibrium without re-solving the full system.

► **Lemma 17** (Marginal extension of equilibrium). *Let $\mathbf{p} = \{p_1, \dots, p_m\}$ be a symmetric Nash equilibrium of the game $G(\mathbf{t}, \ell)$ at utility level λ , where $\mathbf{t} = \{tx_i\}_{i \in [m]}$. Let tx_{m+1} be a new transaction with admissible function f_{m+1} of parameters $\underline{\lambda}_{m+1}, \bar{\lambda}_{m+1}$, and let $\hat{f}_{m+1}$ be its associated clamped function. Define $\mathbf{t}' := \mathbf{t} \cup \{tx_{m+1}\}$ and $\ell' := \ell + \hat{f}_{m+1}(\lambda)$. Then $\mathbf{p}' := \mathbf{p} \cup \{\hat{f}_{m+1}(\lambda)\}$ is the unique symmetric mixed Nash equilibrium of $G(\mathbf{t}', \ell')$, supported at the same utility level λ . In particular, p_i is unchanged for every $i \in [m]$, and tx_{m+1} attains $p_{m+1} = \hat{f}_{m+1}(\lambda)$.*

Proof sketch. One verifies that appending $\hat{f}_{m+1}(\lambda)$ to the original profile satisfies all equilibrium conditions of $G(\mathbf{t}', \ell')$ at the same level λ : the per-transaction identities are local and unaffected by enlarging the transaction set, and the clearing condition holds by construction of the enlarged budget. Details appear in the full version [29]. ◀

► **Theorem 18** (Correctness of cut-and-paste). *Fix a target probability $p_c^* \in (0, 1)$ and let $(\mathbf{p}_{-c}, \lambda_{-c}) = \text{EQUILIBRIUM}(\mathbf{t}_{-c}, \ell - p_c^*)$ be the post-cut equilibrium. Under the per-proposer bribe*

$$b^* = u_c(p_c^*) - \lambda_{-c}$$

the game $G(\mathbf{t}, \ell)$ has a symmetric equilibrium in which tx_c is included with probability p_c^ and every $tx_j \neq tx_c$ retains its probability from $\mathbf{p}_{-c}$. The boundary case $p_c^* = 0$ (full censorship, $s = 1$) is handled by continuity, giving $b^* = \bar{\lambda}_c - \lambda_{-c}$.*

Proof sketch. A bribe b translates the demand of tx_c : the shifted inverse satisfies the *bribe-shift identity* $f'_c(\lambda, b) = f_c(\lambda + b)$, and horizontal translation preserves admissibility. Applying Lemma 17 to the post-cut equilibrium with tx_c (equipped with the shifted demand) as the new transaction, the augmented equilibrium coincides with that of $G(\mathbf{t}, \ell)$ at target probability p_c^* precisely when $\hat{f}'_c(\lambda_{-c}, b) = p_c^*$; by the bribe-shift identity this reads $\hat{f}'_c(\lambda_{-c} + b) = p_c^*$, whose solution is $b^* = u_c(p_c^*) - \lambda_{-c}$. The full derivation appears in the full version [29]. ◀

Finally, recall from Definition 4 that the bribery cost is the adversary's *expected total expenditure*: the per-proposer bribe b^* is collected by each proposer with probability $1 - p_c^* = s^{1/n}$, so the bribery cost equals $n s^{1/n} b^*$. All eCR values we report (Section 6) use this total.

5.3 Applying Cut-and-Paste to Concrete TFMs

We now apply the cut-and-paste recipe to all three TFMs.

Duplication-Penalizing TFM. For this TFM, recall from Example 10 that

$$u_i(p_i) = (1 - p_i)^{n-1} \cdot t_i \quad f_i(\lambda) = 1 - \left(\frac{\lambda}{t_i} \right)^{\frac{1}{n-1}}.$$

Since f_i is admissible, substituting into $b = u_i(p_i) - \lambda$ gives $b = s^{\frac{n-1}{n}} \cdot t_i - \lambda$.

Even-Split TFM. For this TFM, recall from Example 11 that

$$u_i(p_i) = t_i \cdot \frac{1 - (1 - p_i)^n}{n p_i},$$

whose inverse f_i is admissible with parameters $(\lambda_i, \bar{\lambda}_i) = (t_i/n, t_i)$, though it has no closed form. The bribe, however, does: substituting $p_c^* = 1 - s^{1/n}$ into $b = u_c(p_c^*) - \lambda_{-c}$ and using $(1 - p_c^*)^n = s$ gives $b = \frac{t_c(1-s)}{n(1-s^{1/n})} - \lambda_{-c}$. As a sanity check, as $n \rightarrow \infty$ we have $n(1 - s^{1/n}) \rightarrow \ln(1/s)$, so the per-proposer bribe approaches $t_c(1 - s)/\ln(1/s) - \lambda_{-c}$: even with many proposers, each proposer's marginal stake in tx_c remains a constant fraction of t_c , unlike the collective TFM below where it vanishes as $1/n$.

Collective TFM. For this TFM, recall from Example 9 that

$$u_i(p_i) = (1 - p_i)^{n-1} \cdot \frac{t_i}{n} \quad \text{and} \quad f_i(\lambda) = 1 - \left(\frac{\lambda}{\frac{t_i}{n}} \right)^{\frac{1}{n-1}}.$$

Since f_i is admissible, substituting into $b = u_i(p_i) - \lambda$ gives $b = s^{\frac{n-1}{n}} \cdot \frac{t_i}{n} - \lambda$.

Note that this reduces to the duplication-penalizing TFM, with just a $1/n$ factor applied to all tips. This reflects the fact that the mechanisms are game-theoretically quite similar. For both mechanisms, the only factor relevant to the validator’s utility is whether anyone else has included the transaction. If so, including it in their sub-block doesn’t increase their utility; if not, their utility increases by t_i or t_i/n .

6 Comparing TFMs on Idealized and Real-World Data


In this section, we run Monte-Carlo simulations using Algorithm 1 and the cut-and-paste recipe to plot the approximate throughput and eCR as the number of proposers increases. We show how these values vary with congestion, burn, and the adversary’s desired censorship probability. We compare across the three TFMs, finding that the duplication-penalizing TFM performs best across the board.

6.1 Comparing TFMs

We first describe our baseline parameters and simulation method. We simulate a mempool of 250 transactions for a block that can hold any 200 of these. The bids of the transactions are drawn iid from $\text{PARETO}(1, 2.5)$, to simulate the standard environment of many low bidders and a few high bidders. We perform 1000 trials, each with different mempools, and for each trial, we use our “cut-and-paste” method to compute its economic censorship resistance and throughput, assuming that an adversary wants to achieve a 90% censorship probability. We graph the average of the trials as the number of validators, n , varies.

Figure 1 compares the three TFMs using the baseline parameters, showing that the duplication-penalizing TFM achieves the highest eCR, and ties the collective TFM for the highest throughput. Recall from Section 5 that the *collective* TFM reduces to the duplication-penalizing TFM with all tips scaled by $1/n$; since a uniform rescaling of the tips merely rescales λ^* , the two TFMs induce identical equilibrium probabilities, and thus identical throughput. They differ starkly in eCR, however: under the collective TFM each proposer’s marginal stake in any single transaction is only t_i/n , so the adversary’s bribes shrink as n grows, while users continue to pay their full tips. This is why the collective TFM’s eCR remains near zero regardless of n .

The even-split TFM has lower throughput than the other two: including an already-included transaction still pays t_i/k , so duplication is penalized only weakly, and the equilibrium concentrates more probability mass on high-tip transactions. Its eCR sits between the other two TFMs. Compared to the duplication-penalizing TFM, the gap is driven by the price to users: the even-split TFM always charges users their tip, while the duplication-penalizing TFM only does so if their transaction is singly-included. As a result, the duplication-penalizing TFM has significantly higher eCR, as the bribes required are much larger relative to the expected user payment.




■ **Figure 1** Comparison of three TFMs, with the baseline parameters of $m = 250$, $B = 200$, $s = 0.9$, and bids drawn iid from $\text{PARETO}(1, 2.5)$.

Throughput. Note that the throughput quickly converges in both cases. For context, a single proposer trivially achieves throughput 1 by deduplicating its own block; moving to MCP thus incurs a one-time duplication cost (roughly 0.65 at our baseline under the duplication-penalizing TFM). The key observation is that this cost does not compound: it stabilizes almost immediately, so increasing n further buys censorship resistance at essentially no additional throughput cost. To contextualize the throughput, we compare these algorithms to an optimistic baseline. Suppose that the validators had the explicit goal of maximizing throughput, but were not able to coordinate. It's easy to see that the optimal strategy would be for all validators to choose each transaction uniformly at random. Now, assuming that $m \geq B$, i.e., there are more transactions than block space (so that $\min(B, m) = B$ in Definition 2), the throughput can be written as:

$$\frac{1}{B} \sum_{i \leq m} 1 - (1 - p_i)^n \approx \frac{1}{B} \sum_{i \leq m} 1 - e^{-np_i}$$

The approximation holds well as long as p_i is small, which happens when n is moderate, as that decreases the sub-block size. Let $y_i = np_i$, the expected number of copies of transaction i . Under the uniform distribution, each transaction is chosen with $p_i = \ell/m$, where $\ell = B/n$ is the sub-block size. Thus, $y_i = np_i = B/m$, and the throughput of the uniform distribution is approximately $\frac{m}{B}(1 - e^{-B/m})$. This represents an *upper bound* on throughput – profit incentives will decrease the actual throughput, as validators are more likely to include more valuable transactions.




■ **Figure 2** Comparing throughput to the optimistic upper bound of a uniform distribution. We use the baseline parameters of $m = 250$, $B = 200$, $s = 0.9$, and bids drawn iid from $\text{PARETO}(1, 2.5)$.

In Figure 2, we plot the throughput of the duplication-penalizing and even-split TFMs, as well as their “gap from uniform,” in terms of their $\sum_{i \leq m} |y_i - B/m|$. This represents their aggregate excess duplication, where “excess” is defined relative to the uniform distribution. The main takeaway is that the duplication-penalizing TFM maintains a throughput very close to the optimistic upper-bound of a uniform distribution.


Burn. So far, we have been assuming no burn; users pay only the payment from the TFM. Burn hurts eCR significantly, since higher burn means that adversaries can bribe proposers for much less than the users pay. To determine a reasonable amount of burn, we use data from Ethereum in January 2026, and find that the base fee was 8.7x the average tip. Figure 3 plots the eCR of the three TFMs assuming a burn multiple of 4 and 8.7. All users are charged the same amount of burn.

Though burn significantly hampers the eCR, we see that eCR still increases roughly linearly with n , and so increasing the number of validators can still result in high eCR. For example, with a burn multiple of 8.7 and 49 validators, the eCR is approximately 2.5, so an adversary would need to pay 2.5x more than the user to censor the average transaction. The figure also plots the single-proposer baseline, which suffers greatly as burn increases. At a burn multiple of 8.7, the single-proposer eCR is approximately 0.09, so an adversary can censor the average transaction by paying only 9% of what the user pays.

Further, note that increasing the burn brings the duplication-penalizing and even-split TFMs closer together. With no burn and 49 validators, the eCR was approximately 35 and 10, for the duplication-penalizing and even-split TFMs respectively. At a burn multiple of 8.7, the eCR was approximately 2.5 and 1.8. This is because the burn is a flat fee charged to all transactions, and so the differences in user payments under the two TFMs are diluted.



■ **Figure 3** Comparing eCR assuming that the burn is either 4x or 8.7x the average tip. We use the baseline parameters of $m = 250$, $B = 200$, $s = 0.9$, and the bids are drawn iid from $\text{PARETO}(1, 2.5)$.



■ **Figure 4** Comparing eCR as congestion varies, using the baseline parameters of $B = 200$, $s = 0.9$, and bids drawn iid from $\text{PARETO}(1, 2.5)$. The number of transactions m varies in the different curves.

6.2 Congestion and Adversary Incentives

Here, we evaluate the impact of congestion and the adversary's goal on the censorship resistance. We only consider the duplication-penalizing TFM, as it beats the other two in both eCR and throughput.

For congestion, with the baseline parameters, we have 250 transactions with a block size of 200. In Figure 4, we plot the eCR for both more congested and less congested cases. We fix the block size at 200, and vary the number of transactions. As congestion increases, the eCR goes down, because the opportunity cost of censoring a transaction decreases (it's more likely that there is another good transaction to take its place).

In Figure 5, we plot the eCR for a range of adversary censorship probabilities. Our baseline assumes that the adversary wants to censor a transaction with 90% probability; we plot for both more lenient and more strict adversaries. As the adversary's desired censorship probability decreases, the eCR also decreases, as smaller bribes are required.




Figure 5 Comparing eCR as the adversary’s censorship probability varies, using the baseline parameters of $B = 200, m = 250$, and bids drawn iid from $\text{PARETO}(1, 2.5)$. The adversary’s censorship probability varies in the different curves.

6.3 Analysis on Empirical Data

We next perform the comparison on real Ethereum transaction data. We use a window of 10,000 recent Ethereum blocks, covering block numbers 24,341,338 through 24,351,341. For each block $\mathcal{B}$, we compute each transaction’s tip as $\text{gasPrice} - \text{baseFeePerGas}$ and keep only transactions with strictly positive tip, since these are the transactions that are afforded any eCR. This yields a block-specific set of valid transactions, whose size we denote by $m_{\mathcal{B}}$.


For each block $\mathcal{B}$, we fix the target total block size at $B_{\mathcal{B}} = 0.8 m_{\mathcal{B}}$, so that the protocol aims to include roughly 80% of the positive-tip transactions in that block. We then vary the number of validators n from 2 to 50, and for each value of n set the per-validator probability-mass budget to $\ell_{\mathcal{B}} = B_{\mathcal{B}}/n$, so that total capacity $n\ell_{\mathcal{B}}$ remains fixed as n changes. We run this procedure for the three TFMs in our study: duplication-penalizing, collective, and even-split.

Given the positive-tip transactions of a block $\mathcal{B}$ and a choice of n , we first solve for the baseline single-proposer equilibrium probabilities $(p_i)_i$ under the relevant TFM. From these probabilities we compute throughput by taking the expected number of distinct transactions included by at least one of the n validators and normalizing by the target block size. To measure economic censorship resistance, we take a sample of 20 target transactions per block. For each sampled target transaction tx_i , we use the cut-and-paste method to compute the per-validator bribe needed to achieve a target *total* censorship probability of 90%. This gives a post-bribe target inclusion probability p'_i and a per-censor bribe b_i , from which we compute

$$\text{eCR}_i = \frac{n(1 - p'_i) b_i}{P_i},$$

where P_i is the baseline expected user payment conditional on inclusion, in accordance with Definitions 4 and 5.

Finally, for each block $\mathcal{B}$ and each n , we average eCR_i over the sampled target transactions in that block, and separately record the corresponding throughput. We then average these block-level quantities over the full 10,000-block dataset to obtain the curves shown in Figure 6.



■ **Figure 6** Comparison of three TFMs on recent Ethereum block data, with $B = 0.8m$, $s = 0.9$, and bids given by observed positive transaction tips.

One caveat is in order: this dataset conditions on inclusion. Transactions that were successfully censored never appear on-chain, so the empirical tip distribution may understate the mass of low-tip or contested transactions. We therefore view these experiments as evaluating the TFMs on realistic bid shapes, rather than as a measurement of censorship itself.

Takeaways from empirical analysis. Similar to the analysis on the synthetic data, the analysis on the empirical data shows that the duplication-penalizing TFM dominates the other two in both throughput and economic censorship resistance. The main difference between the empirical data and our synthetic data is that the empirical tip distribution is more skewed; a larger percentage of the tips are closer to 0 in reality. This reduces the throughput, because the probability mass concentrates on the few high value tips. Lower throughput means more duplications, and so the duplication-penalizing TFM has much higher censorship resistance, since the average user payment is low. The link to reproduce any empirical result can be found at zenodo record number 21831640.

7 Conclusion

We proposed a novel definition of economic censorship resistance (eCR), and modeled the incentives of rational proposers in the presence of a censoring adversary. We found an efficient algorithm to determine the equilibrium of our game, and our algorithm works for a variety of TFMs. We used this algorithm to compare the eCR and throughput of three TFMs on real and synthetic data. Overall, we found that:

1. MCP systems enjoy much higher eCR than single proposer systems
2. Burnt user fees can significantly lower eCR
3. eCR scales linearly with the number of validators, and thus strong eCR can be achieved even in the presence of significant burn
4. Of the TFMs we study, the duplication-penalizing TFM dominates in both throughput and eCR, across a wide range of parameters

We suggest three avenues for future work. The first is to analyze a wider range of TFMs. In particular, one may be able to prove that the duplication-penalizing TFM is optimal among a family of TFMs, or to prove some theoretical upper bounds on censorship, as we informally did for throughput. A user-side equilibrium analysis – endogenous tips, and robustness to user–proposer collusion under the duplication-penalizing TFM (Remark 1) – is another natural target. This would complement our empirical analysis, and provide guidance on how to best implement an MCP system.

Second, one could analyze different bribery attacks. We consider a simple setting, where the censoring-adversary must declare a transaction to be censored, along with a bribe for any proposer who doesn't include that transaction. But censoring behavior may take other forms. For example, the bribe could be paid out only if *no* proposer includes that transaction.

Finally, we assume that all proposers have access to a shared mempool. But some proposers may have private access to some transactions. The simplest example of this is private order flow, but this may also occur for geographical reasons. For example, proposers located far from the rest of the network may receive transactions from their region before the rest of the network, and thus have semi-private access in some timing windows. Our model, and the duplication-penalizing TFM in particular, heavily rewards proposers with private transactions. Thus, MCP systems may serve as a counterbalance to the latency pressures that encourage geographical co-location.

References

- 1 Ittai Abraham, Yuval Efron, and Ling Ren. The latency cost of censorship resistance. Cryptology ePrint Archive, Paper 2025/2136, 2025. URL: <https://eprint.iacr.org/2025/2136>.
- 2 Orestis Alpos, Bernardo David, Nikolas Kamarinakis, and Dionysis Zindros. An overview of efficiency and censorship resistance in consensus protocols. *CoRR*, abs/2504.03588, 2025. [arXiv:2504.03588](https://arxiv.org/abs/2504.03588).
- 3 Orestis Alpos, Lioba Heimbach, Kartik Nayak, and Sarisht Wadhwa. Censorship-resistant sealed-bid auctions on blockchains. Cryptology ePrint Archive, Paper 2025/2127, 2025. URL: <https://eprint.iacr.org/2025/2127>.
- 4 Aptos Foundation. Aptos White Paper. <https://aptos.dev/network/blockchain/aptos-white-paper>, 2025. Accessed 2025-11.
- 5 Ava Labs. Avalanche consensus protocol. <https://docs.avax.network/overview/avalanche-consensus>, 2020.
- 6 Kushal Babel, Andrey Chursin, George Danezis, Lefteris Kokoris-Kogias, and Alberto Sonnino. Mysticeti: Low-Latency DAG Consensus with Fast Commit Path, February 2024. [arXiv:2310.14821](https://arxiv.org/abs/2310.14821) [cs]. doi:10.48550/arXiv.2310.14821.
- 7 Ben Berger, Edward W. Felten, Akaki Mamagishvili, and Benny Sudakov. Economic censorship games in fraud proofs. In *Proceedings of the 26th ACM Conference on Economics and Computation*, EC '25, pages 670–687, New York, NY, USA, 2025. Association for Computing Machinery. doi:10.1145/3736252.3742611.

- 8 Joseph Bonneau. Why buy when you can rent? Bribery attacks on Bitcoin-style consensus. In *Financial Cryptography and Data Security – FC 2016 International Workshops (BITCOIN)*, volume 9604 of *Lecture Notes in Computer Science*, pages 19–26. Springer, 2016. doi:10.1007/978-3-662-53357-4_2.
- 9 Vitalik Buterin. State of research: Increasing censorship resistance of transactions under proposer/builder separation (PBS). <https://notes.ethereum.org/s3JToeApTx6CKLJt8AbhFQ>, 2021. Ethereum research notes.
- 10 Vitalik Buterin and Mike Neuder. No free lunch – a new inclusion list design. <https://ethresear.ch/t/no-free-lunch-a-new-inclusion-list-design/16389>, 2023. Ethereum Research post.
- 11 Censorship.pics. <https://censorship.pics>, 2025. Accessed 2025.
- 12 Hao Chung and Elaine Shi. Foundations of transaction fee mechanism design. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3856–3899. SIAM, 2023. doi:10.1137/1.9781611977554.ch150.
- 13 Francesco D’Amato. Forward inclusion lists. <https://notes.ethereum.org/@fradamt/forward-inclusion-lists>, 2022. Ethereum research notes.
- 14 George Danezis, Lefteris Kokoris-Kogias, Alberto Sonnino, and Alexander Spiegelman. Narwhal and Tusk: A DAG-based mempool and efficient BFT consensus. In *Proceedings of the Seventeenth European Conference on Computer Systems (EuroSys ’22)*, pages 34–50. ACM, 2022. doi:10.1145/3492321.3519594.
- 15 Electronic Frontier Foundation. Financial censorship. <https://www.eff.org/issues/financial-censorship>, 2024.
- 16 Fatima Elsheimy, Ioannis Kaklamanis, Sarisht Wadhwa, Charalampos Papamanthou, and Fan Zhang. Censorship resistance vs throughput in multi-proposer BFT protocols. Cryptology ePrint Archive, Paper 2026/126, 2026. URL: <https://eprint.iacr.org/2026/126>.
- 17 Elijah Fox, Mallesh M. Pai, and Max Resnick. Censorship Resistance in On-Chain Auctions. In Joseph Bonneau and S. Matthew Weinberg, editors, *5th Conference on Advances in Financial Technologies (AFT 2023)*, volume 282 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 19:1–19:20, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.AFT.2023.19.
- 18 Fangyu Gai, Jianyu Niu, Ivan Beschastnikh, Chen Feng, and Sheng Wang. Scaling blockchain consensus via a robust shared mempool (stratus). In *Proc. of IEEE ICDE 2023*, pages 530–543. IEEE, 2023. doi:10.1109/ICDE55515.2023.00047.
- 19 Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The Bitcoin backbone protocol: Analysis and applications. In *Advances in Cryptology – EUROCRYPT 2015*, volume 9057 of *Lecture Notes in Computer Science*, pages 281–310. Springer, 2015. doi:10.1007/978-3-662-46803-6_10.
- 20 Pranav Garimidi, Lioba Heimbach, and Tim Roughgarden. Transaction fee mechanism design for leaderless blockchain protocols, 2025. doi:10.48550/arXiv.2505.17885.
- 21 Pranav Garimidi, Joachim Neu, and Max Resnick. Multiple concurrent proposers: Why and how. Cryptology ePrint Archive, Paper 2025/1772, 2025. URL: <https://eprint.iacr.org/2025/1772>.
- 22 Lioba Heimbach, Lucianna Kiffer, Christof Ferreira Torres, and Roger Wattenhofer. Ethereum’s proposer-builder separation: Promises and realities. In *Proceedings of the 2023 ACM Internet Measurement Conference (IMC ’23)*, pages 406–420. Association for Computing Machinery, 2023. doi:10.1145/3618257.3624824.
- 23 Dimitris Karakostas, Aggelos Kiayias, and Thomas Zacharias. Blockchain bribing attacks and the efficacy of counterincentives. In *Proc. of CCS 2024*, pages 1031–1045, 2024. doi:10.1145/3658644.3670330.
- 24 Mysten Labs. Sui: A platform for high-performance smart contracts. <https://github.com/MystenLabs/sui/blob/main/doc/paper/sui.pdf>, 2022. Accessed 2025-11.

- 25 Patrick McCorry, Alex Hicks, and Sarah Meiklejohn. Smart contracts for bribing miners. In Aviv Zohar, Yael Kalai, and David McGrew, editors, *Financial Cryptography 2018 Workshops*, volume 10958 of *Lecture Notes in Computer Science*, pages 3–18. Springer, 2019. doi:10.1007/978-3-662-58820-8_1.
- 26 Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. The honey badger of bft protocols. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, pages 31–42, New York, NY, USA, 2016. Association for Computing Machinery. doi:10.1145/2976749.2978399.
- 27 Mike Neuder, Vitalik Buterin, Francesco D'Amato, Terence Tsao, Potuz, and Manav. EIP-7547: Inclusion lists. <https://eips.ethereum.org/EIPS/eip-7547>, 2023. Ethereum Improvement Proposal (draft).
- 28 Tim Roughgarden. Transaction fee mechanism design. *CoRR*, abs/2106.01340, 2021. arXiv:2106.01340.
- 29 Aditya Saraf, Ioannis Kaklamanis, Sarisht Wadhwa, and Fatima Elsheimy. Censorship resistance and throughput with multiple concurrent proposers, 2026. doi:10.48550/arXiv.2607.16995.
- 30 Alexander Spiegelman, Neil Giridharan, Alberto Sonnino, and Lefteris Kokoris-Kogias. Bullshark: DAG BFT Protocols Made Practical. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, CCS '22, pages 2705–2718, New York, NY, USA, November 2022. Association for Computing Machinery. doi:10.1145/3548606.3559361.
- 31 Chrysoula Stathakopoulou, Tudor David, Matej Pavlovic, and Marko Vukolić. Mir-bft: High-throughput robust bft for decentralized networks, 2021. arXiv:1906.05552.
- 32 Aikaterini-Panagiota Stouka, Julian Ma, and Thomas Thiery. Multiple proposer transaction fee mechanism design: Robust incentives against censorship and bribery, 2025. arXiv:2505.13751v2 [cs.GT]. doi:10.48550/arXiv.2505.13751.
- 33 Hanyi Sun, Na Ruan, and Chunhua Su. How to model the bribery attack: A practical quantification method in blockchain. In Liqun Chen, Ninghui Li, Kaitai Liang, and Steve A. Schneider, editors, *Computer Security – ESORICS 2020, Proceedings, Part II*, volume 12309 of *Lecture Notes in Computer Science*, pages 569–589. Springer, 2020. doi:10.1007/978-3-030-59013-0_28.
- 34 Thomas Thiery, Francesco D'Amato, et al. EIP-7805: Fork-choice enforced inclusion lists (FOCIL). <https://eips.ethereum.org/EIPS/eip-7805>, 2024. Ethereum Improvement Proposal (draft).
- 35 Thomas Thiery, Francesco D'Amato, and Barnabé Monnot. The more, the less censored: Committee-enforced inclusion sets (COMIS) on Ethereum. <https://ethresear.ch/t/18835>, 2024. Ethereum Research post.
- 36 Thomas Thiery, Barnabe Monnot, Luca Zanolini, and Julian Ma. Fork-choice enforced inclusion lists (focil). <https://ethresear.ch/t/19870>, 2024. Ethereum Research post.
- 37 Sarisht Wadhwa, Julian Ma, Thomas Thiery, Barnabe Monnot, Luca Zanolini, Fan Zhang, and Kartik Nayak. Aucil: An inclusion list design for rational parties. *Cryptology ePrint Archive*, 2025.
- 38 Anton Wahrstätter, Jens Ernstberger, Aviv Yaish, Liyi Zhou, Kaihua Qin, Taro Tsuchiya, Sebastian Steinhorst, Davor Svetinovic, Nicolas Christin, Mikolaj Barczentewicz, and Arthur Gervais. Blockchain censorship. In *Proceedings of the 33rd ACM Web Conference (WWW '24)*, pages 1632–1643, Singapore, 2024. Association for Computing Machinery. doi:10.1145/3589334.3645431.
- 39 Shuzheng Wang, Yue Huang, Wenqin Zhang, Yuming Huang, Xuechao Wang, and Jing Tang. Private order flows and builder bidding dynamics: The road to monopoly in Ethereum's block building market. In *Proceedings of the ACM Web Conference 2025 (WWW '25)*. Association for Computing Machinery, 2025. doi:10.1145/3696410.3714754.
- 40 Fredrik Winzer, Benjamin Herd, and Sebastian Faust. Temporary censorship attacks in the presence of rational miners. In *IEEE Europe. Symposium on Security and Privacy Workshops (EuroSP Workshops) 2019*, pages 357–366. IEEE, 2019. doi:10.1109/EuroSPW.2019.00046.

- 41 Bowen Xue, Soubhik Deb, and Sreeram Kannan. Bigdipper: A hyperscale bft system with short term censorship resistance, 2023. doi:10.48550/arXiv.2307.10185.
- 42 Sen Yang, Kartik Nayak, and Fan Zhang. Decentralization of Ethereum’s builder market. In *Proceedings of the 46th IEEE Symposium on Security and Privacy (S&P)*. IEEE, 2025. doi:10.48550/arXiv.2405.01329.
- 43 Maofan Yin, Dahlia Malkhi, Michael K. Reiter, Guy Golan Gueta, and Ittai Abraham. Hotstuff: Bft consensus with linearity and responsiveness. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing (PODC ’19)*, PODC ’19, pages 347–356, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3293611.3331591.