Maximizing Miner Revenue in Transaction Fee Mechanism Design

Ke Wu
CSD Department, Carnegie Mellon University, Pittsburgh, PA, USA

Elaine Shi
ECE and CSD Department, Carnegie Mellon University, Pittsburgh, PA, USA

Hao Chung
ECE Department, Carnegie Mellon University, Pittsburgh, PA, USA

Abstract

Transaction fee mechanism design is a new decentralized mechanism design problem where users bid for space on the blockchain. Several recent works showed that the transaction fee mechanism design fundamentally departs from classical mechanism design. They then systematically explored the mathematical landscape of this new decentralized mechanism design problem in two settings: in the plain setting where no cryptography is employed, and in a cryptography-assisted setting where the rules of the mechanism are enforced by a multi-party computation protocol. Unfortunately, in both settings, prior works showed that if we want the mechanism to incentivize honest behavior for both users as well as miners (possibly colluding with users), then the miner revenue has to be zero. Although adopting a relaxed, approximate notion of incentive compatibility gets around this zero miner-revenue limitation, the scaling of the miner revenue is nonetheless poor.

In this paper, we show that if we make a mild reasonable-world assumption that there are sufficiently many honest users, we can circumvent the known limitations on miner revenue, and design auctions that generate asymptotically optimal miner revenue. We also systematically explore the mathematical landscape of transaction fee mechanism design under the new reasonable-world assumptions, and demonstrate how such assumptions can alter the feasibility and infeasibility landscape.

2012 ACM Subject Classification Security and privacy → Cryptography

Keywords and phrases Blockchain, Mechanism Design, Transaction Fee

Digital Object Identifier 10.4230/LIPIcs.ITCS.2024.98

Related Version *Full Version*: `https://eprint.iacr.org/2023/283`

Funding Research supported by NSF awards 2212746, 2044679, 1704788, a Packard Fellowship, a generous gift from the late Nikolai Mushegian, a gift from Google, and an ACE center grant from Algorand Foundation.

1 Introduction

The transaction fee mechanism (TFM) [26, 35, 6, 7, 29, 30, 13, 8, 15, 37] is a new decentralized mechanism design problem that arises in a blockchain environment. Since the space on the blockchain is scarce, users must bid to get their transactions included and confirmed whenever a new block is minted. Earlier works observed that mechanism design in a decentralized environment departs fundamentally from classical mechanism design [26, 35, 6, 7, 29, 30, 13, 8, 15, 37]. The majority of classical mechanisms assume that the auctioneer is trusted and will honestly implement the prescribed mechanism. However, in a decentralized environment, the auction is implemented by a set of miners or consensus nodes[1] who are incentivized

[1] Throughout the paper, we call the consensus nodes "miners" regardless of whether the consensus protocol uses proof-of-work or proof-of-stake.

to take advantage of profitable deviations (if there are any) rather than implementing the mechanism honestly. As a simple example, while the Vickrey auction [32] (a.k.a., the second-price auction) is considered an awesome auction by classical standards, it is not a great fit for a decentralized environment as explained below. Suppose we confirm k bids and they all pay the $(k+1)$-th price to the miner. Then the miner could inject a fake bid that is slightly less than the k-th price, thus causing confirmed bids to pay essentially the k-th price. Alternatively, the same effect can also be achieved if the miner colludes with the $(k+1)$-th bidder and asks it to raise its bid to almost exactly the k-th price. The coalition can then split off the additional gains off the table, using *binding* side contracts that can be instantiated through the decentralized smart contracts that are available in blockchain environments.

This drives us to rethink what is a "dream" TFM in the absence of a fully trusted auctioneer. Recent works [7, 29, 30, 8] formulated the following desiderata:

- *User incentive compatibility (UIC)*: a user's best strategy is to bid truthfully, even when the user has observed others' bids.
- *Miner incentive compatibility (MIC)*: the miner's best strategy is to implement the mechanism honestly, even when the miner has observed all users' bids.
- *Side-contract-proofness (SCP)*: playing honestly maximizes the joint utility of a coalition consisting of the miner and at most c users, even after having observed all others' bids.[2]

Roughgarden showed that Ethereum's EIP-1559 mechanism can simultaneously achieve all three properties, as long as the block size is *infinite* [7, 29]. In practice, EIP-1559 tries to be in the "infinite block size" regime by estimating a reserve price based on the recent history. The reserve price is a minimum threshold used to filter the transactions to avoid congestion. For the case of finite block size (i.e., when congestions do occur), Chung and Shi [8] showed that it is impossible to satisfy all three properties at the same time without the use of cryptography. However, the subsequent work of Shi, Chung, Wu [31] shows that we can use the *MPC-assisted model* to circumvent this impossibility. In the MPC-assisted model, the TFM's rules are securely enforced by a multi-party computation protocol among the miners, thus taking away the miner's ability to unilaterally decide which transactions to include in the block. Variants of the MPC-assisted model are being developed by mainstream blockchain projects such as Ethereum. In particular, the community has been making an effort to build "encrypted mempools" [1], which can be viewed as a concrete instantiation of the MPC-assisted model[3]. Under the MPC-assisted model, [31] showed that one can indeed construct a TFM that simultaneously satisfies UIC, MIC, and SCP (for $c = 1$) under finite block size.

Limit on miner revenue

Unfortunately, no matter whether in the plain or in the MPC-assisted model, all these prior works [29, 8, 31] suffer from a "zero-miner-revenue limitation": all the payment from the users is burnt[4] and the miner obtains zero revenue. The works of Chung and Shi [8] and Shi, Chung, Wu [31] proved that this limitation is in fact inherent. Specifically, any TFM (either

[2] The formal definition of joint utility is given in Section 3.1.

[3] Exactly whether encrypted mempool realizes the MPC ideal functionality needed by Shi, Chung, Wu [31] requires a rigorous proof.

[4] We stress that the miners can still get a fixed block reward which incentivizes them to mine – in fact, Ethereum's EIP-1559 burns all base fees and pays the miner a fixed block reward. This fixed block reward does not affect the game-theoretic analysis and thus is typically ignored in the game-theoretic modeling [29, 8, 31].

in the plain model or MPC-assisted model) that simultaneously satisfies UIC, MIC, and SCP (even for $c = 1$) must have zero miner revenue. In both the plain and the MPC-assisted models, the zero miner-revenue limitation holds in a very strong sense: regardless of whether the block size is finite or infinite, and even when the miner colludes with at most $c = 1$ user.

Moreover, [31] additionally explored whether relaxing the *strict* incentive compatibility notion to *approximate* incentive compatibility can increase the miner revenue. They show a somewhat pessimistic, *unscalability of miner revenue* result. Specifically, with ϵ-incentive compatibility, the miner revenue cannot enjoy linear scaling w.r.t. the magnitude of the bids. Given the landscape, we ask the following natural question:

Can we circumvent the severe limitation on miner revenue, under some reasonable-world assumptions?

1.1 Our Results and Contributions

We are inspired by the philosophy adopted by a line of work at the intersection of cryptography and game theory [21, 24, 2, 27, 5, 4, 17, 16, 18, 23, 11, 20, 10, 33, 9, 28, 25, 14, 12]. In these works, the game theoretic properties hold as long as sufficiently many players are honest. Because the game theoretic guarantees ensure that honest behavior is an equilibrium, and that players are incentivized to behave honestly, this in turn reinforces the "sufficient honesty" assumption.

Therefore, we ask whether we can overcome the severe limitation on miner revenue also under some type of "sufficient honesty" assumption. Phrasing the precise "sufficient honesty" assumption, however, turns out to be technically subtle, partly because TFMs must work in an *open* setting where anyone can post a bid, and the mechanism is unaware of the number of bids a-priori. One naïve attempt is to assume that among the bids posted, half of them come from honest users. Unfortunately, this approach does not work. In Section 6 of the online version [34], we show that even under such an "honest majority bids" assumption, we would still suffer from an $O(1)$-miner revenue limitation.

Reasonable-world assumption: known lower bound on the number of honest users

Instead of the "honest majority bids" assumption, we make a subtly different assumption – we assume that there is an a-priori known lower bound h on the number of honest users. Note that this assumption also promises that at least h users will show up. We refer to this as the *known-h model*. In this model, we first observe that the zero miner-revenue limitation no longer holds. Instead, we can prove an $O(h)$-limit on the miner revenue as stated in the following theorem.

▶ **Theorem 1.1** (Informal: limit on miner revenue in the known-h model)**.** *In the known-h model, no MPC-assisted mechanism that simultaneously satisfies UIC, MIC, and SCP (even in the Bayesian setting) can achieve more than $h \cdot \mathbf{E}(\mathcal{D})$ expected miner revenue where $\mathbf{E}(\mathcal{D})$ denotes the expectation of the value distribution $\mathcal{D}$.*

More generally, in the known-h model, if the number of users is n, no MPC-assisted mechanism that simultaneously satisfies ϵ-UIC, ϵ-MIC, and ϵ-SCP (even in the Bayesian setting) can achieve more than $h \cdot \mathbf{E}(\mathcal{D}) + \frac{2(n-h)}{\rho}(\epsilon + C_{\mathcal{D}}\sqrt{\epsilon})$ expected miner revenue, where ρ is an upper bound on the fraction of miners controlled by the strategic coalition, and $C_{\mathcal{D}} = \mathbf{E}_{X\sim\mathcal{D}}[\sqrt{X}]$.

Furthermore, the above limitation holds no matter when the block size is finite or infinite, and even when the miners collude with at most $c = 1$ user.

In the above theorem, $\epsilon \geq 0$ is a parameter that measures the slack in the incentive compatibility notion. When $\epsilon = 0$, there is no slack, and we achieve strict incentive compatibility. One informal interpretation of the above theorem is the following: for ϵ incentive compatibility, Theorem 1.1 allows us to hope for a mechanism where roughly speaking, from each of h users, the miners can hope to get $\mathbf{E}(\mathcal{D})$ revenue which scales proportionally w.r.t. to the bid distribution $\mathcal{D}$. For each of the remaining users, the miners can potentially get some function that depends on ϵ and the bid distribution $\mathcal{D}$, but the term does not scale linearly w.r.t. the magnitude of the bid distribution for natural distributions.

The above Theorem 1.1 allows us to hope for a TFM in the known-h model that achieves revenue that scales with h as well as the magnitude of the bid distribution $\mathcal{D}$. So can we indeed design a mechanism with asymptotically optimal mine revenue matching Theorem 1.1?

Mechanisms for infinite block size

For the infinite block size regime, we propose two mechanisms in the MPC-assisted model:

- The first one, called *threshold-based mechanism*, is a simple and practical mechanism that satisfies almost-strict incentive compatibility except for a tiny slack ϵ that is exponentially small in h.
- The second one, called *LP-based mechanism* (since it uses linear programming), is a result of theoretical interest. It achieves *strict* incentive compatibility, but under one extra assumption (besides a-priori knowledge of h), that the number of fake bids injected by the strategic coalition is bounded.

Both mechanisms achieve asymptotically optimal miner revenue[5] w.r.t. Theorem 1.1. We assume that only honest users' true values are i.i.d. sampled from some distribution $\mathcal{D}$, whereas the strategic users' true values can be *arbitrary* non-negative real numbers. Next, we state the corresponding theorems for the two mechanisms below:

▶ **Theorem 1.2** (Informal: threshold-based mechanism)**.** *Suppose that honest users' values are sampled i.i.d. from some distribution $\mathcal{D}$. Then, there exists an MPC-assisted TFM in the known-h model that satisfies ex post UIC, Bayesian ϵ-MIC, and Bayesian ϵ-SCP (for any number of colluding users) for $\epsilon = O_{\mathcal{D}}(\exp(-\Omega(h)))$ where $O_{\mathcal{D}}(\cdot)$ hides terms related to the value distribution $\mathcal{D}$. Furthermore, the expected total miner revenue $\Theta(h) \cdot \mathsf{median}(\mathcal{D})$.*

Essentially, from each of h users, the miners can obtain revenue that scales linearly w.r.t. both the bid magnitude. By contrast, without the known-h assumption, for our choice of ϵ which is exponentially small in h, the miner revenue must be exponentially small in h as shown in prior work [31]. Observe also that the miner revenue is asymptotically optimal up to additive factors that are exponentially small in h due to Theorem 1.1.

▶ **Theorem 1.3** (Informal: LP-based mechanism)**.** *Suppose that honest users' values are sampled i.i.d. from some distribution $\mathcal{D}$. Then, there exists an MPC-assisted TFM that satisfies ex post UIC, Bayesian MIC, and Bayesian SCP where the MIC and SCP guarantees hold as long as the total number of bids d contributed by the strategic coalition satisfies $d \leq \frac{1}{8}\sqrt{\frac{h}{2\log h}}$. Further, the expected total miner revenue is $\Theta(h) \cdot \mathsf{median}(\mathcal{D})$.*

[5] We achieve asymptotic optimal miner revenue w.r.t. h assuming that the expectation and median of the distribution $\mathcal{D}$ is a constant independent of h.

In the above theorem, we need the extra assumption that the strategic coalition does not control too many bids. Effectively, this is assuming that the coalition cannot inject too many fake bids. Currently, we do not know whether this extra assumption (besides known-h) is needed to overcome the zero miner-revenue limitation. We leave this as an interesting open question.

Mechanisms for finite block size

For the finite block size case, we propose two mechanisms:

- We propose a simple mechanism called *diluted threshold-based mechanism* that achieves *approximate* incentive compatibility. Further, for sufficiently large h, the mechanism achieves asymptotically optimal miner revenue.
- For theoretical interest, we propose another mechanism called *LP-based mechanism with random selection* which achieves *strict* incentive compatibility and asymptotically optimal miner revenue – but under the additional assumptions that the coalition cannot inject too many fake bids, and moreover, the miners collude with at most $c = 1$ user. Jumping ahead, the $c = 1$ assumption will be later justified in Theorem 1.6.

We state the corresponding theorems for the two mechanisms below:

▶ **Theorem 1.4** (Informal: diluted threshold-based mechanism)**.** *Suppose the block size is k, and that honest users' values are sampled i.i.d. from some bounded distribution $\mathcal{D}$. Then, there exists an MPC-assisted TFM in the known-h model that satisfies ex post UIC, Bayesian ϵ-MIC, and Bayesian ϵ-SCP (for any number of colluding users) for $\epsilon = O_{\mathcal{D}}(\exp(-\Omega(h)))$. Furthermore, for sufficiently large h, the mechanism achieves expected total miner revenue $\Theta(k) \cdot \mathsf{median}(\mathcal{D})$.*

▶ **Theorem 1.5** (Informal: LP-based mechanism with random selection)**.** *Suppose the block size is k, and suppose that honest users' values are sampled i.i.d. from some distribution $\mathcal{D}$. Then, there exists an MPC-assisted TFM that satisfies ex post UIC, Bayesian MIC, and Bayesian SCP, where the MIC and SCP guarantees hold when 1) at most $c = 1$ user colludes when miners, and 2) the total number of bids d contributed by the strategic coalition satisfies $d \le \frac{1}{8}\sqrt{\frac{h}{2\log h}}$. Further, the expected total miner revenue is $\Theta(\min\{h, k\}) \cdot \mathsf{median}(\mathcal{D})$.*

We justify the $c = 1$ assumption in the LP-based mechanism with random selection by proving the following impossibility result: for finite block size, no "interesting" mechanism can simultaneously achieve UIC, MIC, and SCP for $c \ge 2$ even in the MPC-assisted model. Specifically,

▶ **Theorem 1.6** (Informal: finite block, $c \ge 2$)**.** *Even in the known-h model, any MPC-assisted TFM that simultaneously satisfies Bayesian UIC, Bayesian MIC, and Bayesian SCP for $c \ge 2$ must suffer from 0 expected social welfare for the users under a bid vector $\mathbf{b} \sim \mathcal{D}^\ell$ where $\ell > h$.*

Necessity of Bayesian equilibrium

All of our feasibility results, namely, Theorems 1.2–1.5, rely on a Bayesian notion of equilibrium (for the MIC and SCP guarantees). As argued by [31], the Bayesian notion of equilibrium is suitable for the MPC-assisted model since the users cannot observe others' bids before submitting their own.

We show that the reliance on Bayesian notions of equilibrium is necessary – had we insisted on an *ex post* notion of equilibrium in the MPC-assisted model, our additional reasonable-world assumptions would not help us overcome the previously known impossibility

results. More specifically, we show that any MPC-assisted mechanism that simultaneously achieves ex post UIC and SCP must suffer from zero miner revenue even in the known-h model. Similarly, for approximate but ex post notions of incentive compatibility, the same miner revenue limitation stated in [31] still applies even in the known-h model. Further, the above restrictions on miner revenue hold no matter whether the block size is finite or infinite.

1.2 Philosophical Discussions about Our Assumptions and Modeling

Known-h assumption

Our assumption about a known upper bound h on the number of honest users has the following justifications:

- First, as mentioned, the zero miner-revenue limitation in earlier works holds in a very strong sense, and even when we make a cryptographic style assumptions such as "a majority of the users or bids are honest" – see the online version [34] for more details.
- Second, TFMs must work in an *open setting* where anyone can post a bid and the mechanism or players do not know the number of bids in advance. This is also an important reason why TFMs depart from classical mechanism design. The precise assumption we need for circumventing the zero miner revenue limitation is an absolute lower bound h on the number of honest users. Simply assuming that the majority of the bids are honest is not sufficient (see Section 6 of the online version [34]).
- Finally, our definitional framework ensures that *honest behavior is an equilibrium*, and thus players are incentivized to behave honestly. This, in turn, reinforces the h-honest users assumption (as long as enough users show up). As mentioned earlier, the same philosophy has been adopted in a line of prior works at the intersection of game theory and cryptography [21, 24, 2, 27, 5, 4, 17, 16, 18, 23, 11, 20, 10, 33, 9, 28, 25, 14, 12].

Limited fake bids

Recall that besides the known-h assumption, our results that achieve strict incentive compatibility require an extra assumption that the number of fake bids injected by the coalition is limited. This assumption is motivated and justified by the following observations. To submit fake bids, the strategic player or coalition needs to have some coin or account with a non-zero balance. Given that the strategic player has a limited initial budget, it cannot control infinitely many accounts. Moreover, if one posts multiple conflicting transactions double-spending the same coin or units of currency, they can easily be detected and suppressed.

As mentioned, we currently do not know whether this limited fake bids assumption can be removed while still achieving strict incentive compatibility. We pose this as an open question.

Robustness w.r.t. parameter estimation

Among our proposed mechanisms, the ones that achieve approximate incentive compatibility, namely, threshold-based or diluted threshold-based mechanisms are simpler and more practical. Just like how Ethereum's EIP-1559 needs to estimate a suitable base fee, these mechanisms also need to estimate some parameters a-priori. In particular, our (diluted) threshold-based mechanism needs to know an estimate of h and the median of the value distribution $\mathcal{D}$ in advance. Just like Ethereum's EIP-1559, we can estimate these parameters from recent history of the blockchain and mempool. For example, one can estimate the total number of bids n from the degree of congestion observed in recent blocks; one can estimate the median of honest users' values from the pending transactions in the mempool. Now, if we are willing to assume that half of the n anticipated bids are honest (note that our mechanisms incentivize honest behavior), we can get an estimate of h.

One important observation is that our threshold-based mechanism and diluted threshold-based mechanisms are quite robust to errors in the estimates. As mentioned later in Remark 2.2, if we set the threshold to $\widehat{h}/4$ for some estimated $\widehat{h}$, and let $\widehat{m}$ be the estimated median, then the mechanisms will achieve approximate incentive compatibility as long as $h_{\text{real}} \cdot q_{\text{real}} \geq \widehat{h} \cdot \frac{(1+\delta)}{4}$ for some arbitrarily small constant $\delta > 0$, where h_{real} is the actual number of honest users, and q_{real} is the actual percentile of the estimate $\widehat{m}$. For example, if $h_{\text{real}} = 0.6h$, and $q_{\text{real}} = 40\%$, then our mechanisms still satisfy approximate incentive compatibility for an exponentially small ϵ.

Independent identically distributed true values assumption

All the mechanisms in our paper only assume that honest users' true values are i.i.d. sampled from the distribution $\mathcal{D}$. The strategic users can have arbitrary non-negative true values.

1.3 Additional Related Work

We now review some closely related recent works besides the prior works on transaction mechanism design [26, 35, 6, 7, 29, 30, 13, 8] already mentioned.

TFM in a Bayesian setting

The recent works of Zhao, Chen, and Zhou [37] and Gafni and Yaish [15] both consider TFM in a Bayesian setting. Although their works did not explicitly define the MPC-assisted model, from a practical standpoint, their results are in fact only relevant in an MPC-assisted (or a similar) model. As explained in the online version [34], plain-model TFMs that achieve *Bayesian* equilibrium also achieve *ex post* equilibrium, since in the plain-model game, the strategic player can decide its actions *after* having observed honest users' bids.

Gafni and Yaish [15] suggest a mechanism that satisfies Bayesian UIC, while also satisfying MIC and OCA-proof (short for offchain-agreement-proof) even if the miner knows everyone's bid. Further, their mechanism works in the finite-block setting while achieving asymptotical optimality in social welfare and revenue. We stress that their result does not contradict the zero miner-revenue limitation proven by [31] since their OCA-proofness notion (originally defined by Roughgarden [29, 30]) is of a different nature from our side-contract-proofness (SCP) notion (originally defined by Chung and Shi [8]). Roughly speaking, OCA-proofness requires that a strategic coalition cannot enter an off-chain contract that increases *everyone*'s utility (*not just those in the coalition*) relative to what's achievable on-chain. In comparison, SCP is the notion that directly captures the cryptocurrency community's outpouring concerns about Miner Extractable Value (MEV). In particular, middleman platforms such as Flashbot facilitate the collusion of miners and users, where the coalition plays strategically to profit themselves at the expense of other users. This is why we choose to use the SCP notion rather than OCA-proofness. Moreover, the reason why the cryptocurrency community is developing encrypted mempool techniques (which can be viewed as instantiations of the MPC-assisted model) is also because they care about SCP (i.e., resilience to MEV).

Zhao, Chen, and Zhou [37] suggest a mechanism that generates positive miner revenue while achieving Bayesian UIC and Bayesian 1-SCP even for the finite block setting. Their result does not contradict the 0-miner revenue limitation of [31], since Zhao, Chen, Zhou [37] consider only a restricted strategy space. In their work, a strategic user or a miner-user coalition can only deviate by bidding untruthfully; the coalition cannot inject fake bids, strategic users cannot drop out, and nor can strategic miners alter the inclusion rule. Due to their restricted strategy space, their results are only relevant under very stringent assumptions:

1) the TFM is implemented in the MPC-assisted (or similar) model; 2) the TFM is fully "permissioned" and allows only a set of pre-registered users to submit bids. In particular, the latter "permissioned" requirement is unrealistic for major decentralized cryptocurrencies today where any user can join and submit transactions.

Cryptography meets game theory

Prior to the advent of cryptocurrencies, a line of work [21, 24, 2, 27, 5, 4, 17, 16, 18, 23, 11, 20, 10, 33, 9, 28, 25, 14, 12] investigated how cryptography and game theory can help each other. For example, cryptography can help remove the trusted mediator assumption in correlated equilibria [11]. Adopting game-theoretic fairness can allow us to circumvent lower bounds pertaining to the more stringent cryptographic notions of fairness [21, 2, 22, 27, 10, 33]. Ferreira and Weinberg [14] and Essaidi, Ferreira and Weinberg [12] showed that cryptographic commitments can help us circumvent impossibilities pertaining to credible auctions. As Chung and Shi [8] explained in detail, credible auction is of a fundamentally different nature from transaction fee mechanism design.

2 Technical Roadmap

2.1 Transaction Fee Mechanism

Environment

For preciseness in our subsequent formal description, we define an (h, ρ, c, d)-environment, where h is the promised lower bound on the number of honest users, $\rho \in [0, 1]$ is the fraction of strategic miners, c is the maximum number of strategic users that collude with miners, and d is the maximum number of bids contributed by the strategic coalition, i.e. fake bids.

Universality

We can replace a subset of these variables with a wildcard $*$ if the mechanism achieves incentive compatibility no matter what the variable turns out to be. For example, a TFM that achieves incentive compatibility in an $(*, \rho, c, *)$-environment if it works when the maximum fraction of strategic miners is ρ, and the maximum number of colluding users is at most c, and regardless of of how many honest users there are and how many bids are contributed by strategic individuals or the coalition. In this case, we also say that the mechanism is universal in the parameters h and d. Using this notation, the mechanisms described by Shi, Chung, Wu [31] are universal in the parameters h and d. Similarly, the limitation on miner revenue they prove can also be interpreted as a limitation of mechanisms that are universal in h and d.

Transaction fee mechanism in the MPC-assisted model

As in previous works, we consider a single auction instance that decides which transactions can be confirmed in the next block. In this paper, "bids" and "transactions" are used interchangeably. A transaction fee mechanism (TFM) in the MPC-assisted model consists of the following randomized algorithms.

- *Confirmation rule* chooses a subset of at most k bids to confirm, where k denotes the block size.[6]
- *Payment rule* decides how much each confirmed bid pays.
- *Miner revenue rule* decides how much revenue the miners get.

We consider a single-parameter environment, i.e., each user has a transaction with the true value represented by a single, non-negative real number. A user's utility equals its true value minus its payment if its transaction gets confirmed in the block. An honest user submits one bid $b \in \mathbb{R}$ representing its true value, while an honest miner submits zero bids and follows the protocol honestly. Strategic users or miners, however, can choose to register one or more identities and submit arbitrary bids under these identities. Moreover, they can choose to drop out during the execution. When a strategic user or miner submits multiple bids, they can only obtain the value when the bid with the true value is confirmed, and other bids are considered fake bids. Fake bids have no intrinsic value to the users and the miners even when they are confirmed.

All mechanisms proposed in our paper work in the MPC-assisted model. We consider the same, generic MPC-assisted model as [31], where the miners jointly execute an MPC protocol that securely evaluates the rules of the TFM. Further, the miners equally divide up the revenue among themselves. A formal description of the model can be found in Section 3.

To illustrate the technical highlights, it is convenient to abstract out the multi-party computation protocol as an ideal functionality. Therefore, we can think of the game as follows:

- Each player (either user or miner) first submit zero to multiple bids to the ideal functionality.
- The ideal functionality then executes the rules of the TFM, outputs the set of confirmed bids, how much reach bid needs to pay, and how much the miners get based on the prescribed rules.

Notably, in the MPC-assisted model, players cannot see others' bids when posting their own. In particular, in an actual instantiation of the above functionality, the players would send bids either in a secret-shared or encrypted format. For this reason, Bayesian notions of equilibrium make sense in the MPC-assisted model. [31] suggested one way to securely realize the above ideal functionality. Meanwhile, efforts to build an "encrypted mempool" by the cryptocurrency community can also be viewed an alternative way to instantiate the ideal functionality – although whether the suggested approaches provably realize this ideal functionality is yet to be proven.

[31] argued that as a starting point, it makes sense to first explore such a generic functionality, since currently, we lack understanding how cryptography in general can help decentralized mechanism design from a game theoretic perspective. Once we understand this better, we can then focus on making customized optimizations for the actual computational tasks that need to be securely evaluated. Our work is motivated by the same philosophy, i.e., we focus on understanding the game theoretic aspects rather than concrete optimizations for realizing the MPC.

[6] In the MPC-assisted model, the mechanism is implemented by the ideal functionality, and the miners cannot decide which transactions are included and considered as the input of the mechanism. Since all mechanisms proposed in our paper work in the MPC-assisted model, we simplify the definition of TFM compared to [8], where the TFM in [8] also needs to specify the *inclusion rule* that tells the miner which transactions are included in the next block. However, a strategic miner can still inject some fake bids. The formal strategy space is defined in Section 3.

Universality in ρ in the idealized model

As mentioned, it is convenient to think of the MPC as an ideal functionality. In this idealized model, our mechanisms can achieve universality in ρ. However, when the ideal functionality is actually instantiated, e.g., with an honest-majority MPC protocol, the resulting protocol would require $\rho < 50\%$.

2.2 Infinite Block Setting

For the infinite block setting, we can achieve $\Theta(h)$ miner revenue in (h, ρ, c, d)-environments. To aid understanding, we first present a simple parity-based mechanism that works for $h = 1$, and then we present our main results.

Glimpse of hope

First, consider the special case where we are promised that there is at least $h = 1$ honest user. In this case, the following simple parity-based mechanism satisfies ex-post UIC, Bayesian MIC, and Bayesian SCP in $(1, *, *, *)$-environments.

MPC-assisted, parity-based mechanism

// Let m be the median of the distribution $\mathcal{D}$ such that $\Pr_{x\sim\mathcal{D}}[x \geq m] = 1/2$.

- All bids that are at least m get confirmed and pay m.
- If the number of confirmed bids is odd, then the total miner revenue is m; else the total miner revenue is 0.

In the above mechanism, as long as there is at least one honest bid, the expected miner revenue is always $m/2$ no matter how the coalition behaves. This is because the strategic coalition cannot predict whether the honest bid is bidding at least the median or not. With this key observation, it is not hard to see that the mechanism satisfies Bayesian MIC and Bayesian SCP (for an arbitrary c). Further, ex post UIC follows directly since the mechanism is a simple posted-price auction from a user's perspective.

Observe also the following subtlety: when the number of confirmed bids is odd, it implies that there is at least one confirmed bid. Therefore, the mechanism guarantees that the miner revenue does not exceed the total payment.

▶ Remark 2.1 (A note about the median assumption). In the above, we assumed that the bid distribution $\mathcal{D}$ has a median m such that $\Pr_{x\sim\mathcal{D}}[x \geq m] = 1/2$. In case the median m does not exactly equally divide the probability mass half and half, then it must be that $\Pr_{x\sim\mathcal{D}}[x > m] < 1/2$ and $\Pr_{x\sim\mathcal{D}}[x < m] < 1/2$. In this case, we can modify the above mechanism slightly as follows: if a user's bid is strictly greater than m, then it is confirmed; if a user's bid is exactly m, then we confirm it with some appropriate probability q; else the user's bid is not confirmed. We can always pick a q such that a bid randomly sampled from $\mathcal{D}$ is confirmed with probability exactly $1/2$. Finally, the miner revenue rule is still decided the same way as before.

Threshold-based mechanism

The parity-based mechanism overcomes the 0 miner-revenue limitation by assuming the existence of at least $h = 1$ honest user. However, the drawback is obvious: the total miner revenue is severely restricted and does not increase w.r.t. the number of bids. A natural question is whether we can achieve $O(h)$ expected miner revenue for general h.

We give an affirmative answer. We first present a simple, practical mechanism called the threshold-based mechanism that achieves almost-strict incentive compatibility except for a tiny slack ϵ that is exponentially small in h. Then, for theoretical interest, we present another mechanism that achieves strict incentive compatibility but requires an extra assumption on the number of bids contributed by strategic players.

MPC-assisted, threshold-based mechanism

// Let m be the median of the distribution $\mathcal{D}$ such that $\Pr_{x\sim\mathcal{D}}[x \geq m] = 1/2$.

- All bids that are at least m get confirmed and pay m.
- If the number of confirmed bids is at least $h/4$, then the miner revenue is $m \cdot h/4$; else the total miner revenue is 0.

Due to the standard Chernoff bound, except with $e^{-\Omega(h)}$ probability, the number of confirmed bids among the h (or more) honest bids is at least $h/4$. Therefore, the above mechanism achieves at least $m \cdot h/4 \cdot (1 - e^{-\Omega(h)})$ expected miner revenue. If the number of confirmed honest bids is $h/4$ or higher, then the coalition cannot increase the miner revenue no matter how it behaves. Only when the number of confirmed honest bids is less than $h/4$, is it possible for the coalition to influence the miner revenue by at most $m \cdot h/4$. Therefore, it is not hard to see that the mechanism satisfies ϵ-Bayesian MIC and ϵ-Bayesian SCP in $(h, *, *, *)$-environments, for $\epsilon = m \cdot h \cdot e^{-\Omega(h)}/4$.

Just like before, in case the median m does not exactly divide the probability mass half and half, we can use the same approach of Remark 2.1 to modify the mechanism and make it work.

▶ Remark 2.2 (On the robustness of parameter estimation). The threshold-based mechanism requires the mechanism to estimate h and the median m of the distribution $\mathcal{D}$. Just like how Ethereum EIP-1559 estimates its base price, we can estimate h and the median from recent history. In particular, from the congestion level in recent blocks, we can estimate a lower bound on the total number of bids. Now, assuming that at least half of them are honest (recall that our mechanism incentivizes honesty), we can get an estimate of h correspondingly. Similarly, we can estimate the median of the bid distribution from past history.

An advantage of the threshold-based mechanism is that it is quite tolerant of errors in the estimation. For example, if the estimated m is actually the 40-percentile of $\mathcal{D}$, and the actual number of honest users is only $0.7h$ where h is the estimate used by the mechanism, the expected number of users bidding at least m is at least $0.4 \cdot 0.7h = 0.28h$. In this case, we can still guarantee that except with exponentially small in h probability, at least $h/4$ users will bid at least m. Thus, the resulting mechanism would still be almost strictly incentive compatible except for a slack ϵ that is exponentially small in h.

LP-based mechanism

Although the ϵ slack in the threshold-based mechanism is exponentially small which is not a problem in practice, it is still theoretically interesting to ask whether we can get $\Theta(h)$ total miner revenue but with *strict* incentive compatibility. To achieve this, our idea is to devise a mechanism that is "close in distance" to the aforementioned threshold-based mechanism, but correcting the "error" such that we can achieve strict incentive compatibility.

Observe that the earlier threshold-based mechanism only needs an a-priori known lower bound on h, and it is universal in the parameters c and d. To achieve strict incentive compatibility, we additionally assume that the the number of bids contributed by the strategic coalition is upper bounded by some a-priori known parameter d.

Now, consider the following mechanism that relies on linear programming to correct the error in the earlier threshold-based mechanism. For simplicity, we assume that the honest bid distribution $\mathcal{D}$ has a median m such that $\Pr_{x\sim\mathcal{D}}[x \geq m] = 1/2$ – if not, we can again use the technique of Remark 2.1 to modify the mechanism and make it work.

MPC-assisted, LP-based mechanism

// Let m be the median of the distribution $\mathcal{D}$, i.e., $\Pr_{x\sim\mathcal{D}}[x \geq m] = 1/2$.

- All bids that are at least m get confirmed and pay m.
- Let n be the length of the bid vector, let $\mathbf{y} := (y_0, y_1, \ldots, y_n)$ be any feasible solution to the following linear program:

$$\forall i \in [n]: \ 0 \leq y_i \leq i \cdot m \tag{1}$$

$$\forall 0 \leq j \leq d: \ \sum_{i=0}^{n-d} q_i \cdot y_{i+j} = \frac{m \cdot h}{4} \tag{2}$$

  where q_i is the probability of observing i heads if we flip $n-d$ independent fair coins.
- The total miner revenue is y_s where s is the number of bids confirmed.

In the above, Equation (1) expresses a *budget feasibility* requirement, i.e., the total miner revenue cannot exceed the total user payment. Equation (2) expresses a *fixed-revenue requirement* stipulating that the miner revenue must be exactly $m \cdot h/4$ no matter how the strategic individual or coalition behaves (as long as it controls at most d bids). More specifically, Equation (2) contains one requirement for each $j \in [0, d]$: conditioned on the fact that among the (at most) d bids controlled by the strategic individual or coalition, exactly j of them are confirmed, the expected miner revenue must be exactly $m \cdot h/4$ where h is an a-priori known lower bound on the number of honest users.

▶ Remark 2.3. We know that the actual number of honest users that show up is at least $\max(n-d, h)$. So if $n - d > h$, it means that more honest users showed up than the anticipated number h. Observe that on the left-hand side of Equation (2), we are tossing coins for $n-d$ honest users' bids. However, it is important that the right-hand-side of Equation (2) use the a-priori known h rather than the observed $n-d$; otherwise, injecting extra (but up to $d-c$) fake 0-bids can increase the expected miner revenue, which violates MIC and SCP.

If the LP in the above mechanism indeed has a feasible solution, then we can prove that the resulting mechanism satisfies ex post UIC, Bayesian MIC, and Bayesian SCP in $(h, *, *, d)$-environments. The formal proofs are presented in Section 4.2.

The key technical challenge is to answer the question of why the LP has a feasible solution. Intuitively, the earlier threshold-based mechanism gives an "approximate" solution $\widehat{\mathbf{y}} := (\widehat{y}_0, \ldots, \widehat{y}_n)$ to the LP, where $\widehat{y}_i = 0$ for $i \leq (n-d)/4$ and $\widehat{y}_i = \frac{m \cdot h}{4}$ otherwise. With the approximate solution $\widehat{\mathbf{y}}$, the equality constraints in Equation (2) may be satisfied with some small error. We want to show that we can adjust the $\widehat{\mathbf{y}} := (y_0, y_1, \ldots, y_n)$ vector slightly such that we can correct the error, and yet without violating the budget feasibility constraints (Equation (1)).

To achieve this, we will take a constructive approach. We first guess that a feasible solution is of the form $\mathbf{y} = \widehat{\mathbf{y}} + \mathbf{e}$ where $\mathbf{e}$ is a correction vector that is zero everywhere except in the coordinates $\tau, \tau+1, \ldots, \tau+d$ for some appropriate choice of τ that is close to $(n-d)/2$. Henceforth, let $\boldsymbol{\delta} := \mathbf{e}[\tau : \tau + d] / \left(\frac{m \cdot h}{4}\right)$ be the non-zero coordinates of the correction vector $\mathbf{e}$ scaled by $\frac{m \cdot h}{4}$.

By Equation (2), we know that the correction vector $\boldsymbol{\delta}$ must satisfy the following system of linear equations where $t := \frac{n-d}{4}$:

$$\begin{pmatrix} \binom{n-d}{\tau} & \binom{n-d}{\tau+1} & \cdots & \binom{n-d}{\tau+d} \\ \binom{n-d}{\tau-1} & \binom{n-d}{\tau} & \cdots & \binom{n-d}{\tau+d-1} \\ \vdots & \vdots & \ddots & \vdots \\ \binom{n-d}{\tau-d} & \binom{n-d}{\tau-d+1} & \cdots & \binom{n-d}{\tau} \end{pmatrix} \cdot \boldsymbol{\delta} = \begin{pmatrix} \sum_{i=0}^{t} \binom{n-d}{i} \\ \sum_{i=0}^{t-1} \binom{n-d}{i} \\ \vdots \\ \sum_{i=0}^{t-d} \binom{n-d}{i} \end{pmatrix}. \tag{3}$$

In Lemma 4.5, we prove that as long as $d \leq \frac{1}{8}\sqrt{\frac{h}{2\log h}}$, and that τ is an appropriate choice close to $n/2$, then the solution $\boldsymbol{\delta}$ to the linear system in Equation (3) has a small infinity norm – specifically, $\|\boldsymbol{\delta}\|_\infty \leq 1$ – such that the resulting $\mathbf{y}$ vector will respect the budget feasibility constraints, i.e., Equation (1). The actual proof of this bound is somewhat involved and thus deferred to Section 4.2. In particular, a key step is to bound the *smallest singular value* of the matrix in Equation (3) (henceforth denoted A) appropriately – to achieve this, we first bound A's determinant, and then use an inequality proven by [36] which relates the smallest singular value and the determinant.

2.3 Finite Block Setting

2.3.1 Strict Incentive Compatibility

Feasibility for $c = 1$

The LP-based mechanism confirms any bid that offers to pay at least m. Thus, total number of confirmed bids may be unbounded. Therefore, when the block size k is finite, we cannot directly run the LP-based mechanism. We suggest the following modification to the LP-based mechanism such that it works for the finite-block setting:

MPC-assisted, LP-based mechanism with random selection

// Let k be the block size, let m be the median $\mathcal{D}$ such that $\Pr_{x\sim\mathcal{D}}[x \geq m] = 1/2$.

- All bids offering at least m are candidates. If there are more than k candidates, randomly select k of them to confirm; else confirm all candidates. Every confirmed bid pays m.
- Let n be the length of the bid vector, let $\mathbf{y} = (y_0, y_1, \ldots, y_n)$ be any feasible solution to the following linear program:

$$\forall i \in [n]: \quad 0 \leq y_i \leq \min(i,k) \cdot m \tag{4}$$

$$\forall 0 \leq j \leq d: \quad \sum_{i=0}^{n-d} q_i \cdot y_{i+j} = \frac{m \cdot \min(h,k)}{4} \tag{5}$$

  where q_i is the probability of observing i heads if we flip $n-d$ independent fair coins.
- The total miner revenue is y_s where s is the number of *candidates*.

In comparison with the earlier LP-based mechanism, we modify the budget feasibility constraints (Equation (4)) to make sure that the total miner revenue is constrained by the actual number of confirmed bids which is now $\min(i,k)$ if the number of candidates is i. Further, we modify the expected miner revenue (Equation (5)) to be $\frac{m\cdot\min(h,k)}{4}$ which takes into account the block size k. In Section 5 of the online version [34], we prove that as long as $c \leq d \leq \frac{1}{8}\sqrt{\frac{h}{2\log h}}$, the above LP indeed has a feasible solution and the resulting mechanism satisfies ex post UIC, Bayesian MIC, and Bayesian SCP in $(h, *, 1, d)$-environments.

Infeasibility for $c \geq 2$

Unfortunately, the above approach fails for $c \geq 2$. In this case, two users Alice and Bob may be in the same coalition. Alice can now help Bob simply by dropping out and not posting a bid, thus effectively increasing Bob's chance of getting confirmed. In the event that Alice's true value is very small and Bob's true value is sufficiently large, this strategic action can increase the coalition's joint utility.

Interestingly, it turns out that this is no accident. In fact, we prove that for any $h \geq 1$, $\rho \in (0, 1)$, and $d \geq c \geq 2$, no "interesting" mechanism can simultaneously achieve Bayesian UIC, MIC and SCP in (h, ρ, c, d)-environments – any such mechanism must suffer from 0 total social welfare for the users if the actual number of bids received is greater than h (see the online version [34] for full details). We can regard Theorem 1.6 as a generalization of [31]'s Theorem 5.5: they show that any MPC-assisted mechanism that achieves Bayesian UIC, MIC, and SCP in $(*, \rho, c, *)$-environments for $c \geq 2$ must suffer from 0 social welfare for users.

Proof roadmap

We use the following blueprint to prove Theorem 1.6, and the full proof is given in the online version [34]. Below, consider any TFM that satisfies Bayesian UIC, MIC and Bayesian SCP in (h, ρ, c, d)-environments where $d \geq c \geq 2$.

1. First, using techniques inspired by Goldberg and Hartline [19] we prove the following: provided that there are at least h honest users (not including i and j) whose bids are sampled at random from $\mathcal{D}$, then a strategic user i changing its bid should not affect the utility of another user j, if user j's bid is also sampled at random from $\mathcal{D}$.
2. Next, we prove a strategic user i dropping out should not affect another user j's utility, assuming that at least h bids (excluding user i) sampled at random from $\mathcal{D}$.
3. Next, we show that in a world of at least h random bids (excluding user i) sampled from $\mathcal{D}$, user i's expected utility when its bid is sampled randomly from $\mathcal{D}$ depends only on i's identity, and does not depend on the identities of the other random bids. Therefore, henceforth we can use U_i to denote this expected utility.
4. Next, we show that for any two identities i, j, it must be that $U_i = U_j$, otherwise, it violates the assumption that the mechanism is weakly symmetric (see definition of weak symmetry below).
5. Next, we can show that $U_i = 0$: imagine a world with K bids sampled independently from $\mathcal{D}$ whose support is bounded. There must exist some user whose confirmation probability is upper bounded by k/K. This user's expected utility must be arbitrarily small when K is arbitrarily large. With a little more work, we can show that if the world consists of more than h bids sampled independently at random from $\mathcal{D}$, it must be that every user's expected utility is 0.

One technicality that arises in the full proof (see online version [34]) is the usage of the weak symmetry assumption. In particular, the proof would have been much easier if we could instead assume *strong symmetry* which, unfortunately, is too stringent. In strong symmetry, we assume that any two users who bid the same amount will receive the same treatment. While it is a good approach for gaining intuition about the proof, it is too stringent since there could well be more bids offering the same value than the block size k – in this case, a non-trivial mechanism would treat them differently, i.e., confirm some while rejecting others. Our actual proof of Theorem 1.6 needs only a *weak symmetry* assumption which is a standard assumption made in prior works [8, 31], that is, if two input bid vectors $\mathbf{b}, \mathbf{b}'$ of length n are

permutations of each other, then the joint distribution of the *set* of outcomes $\{(x_i, p_i)\}_{i\in[n]}$ must be identical. This implies that if the input bid vectors are permutations of each other, then the vector of expected utilities are permutations of each other too.

2.3.2 Approximate Incentive Compatibility

Because of the limitation shown in Theorem 1.6, we relax the notion to approximate incentive compatibility, and ask if we can achieve optimal miner revenue in the finite block setting. Consider the following TFM.

MPC-assisted, diluted threshold-based Mechanism

/ Let k be the block size, let m be the median of $\mathcal{D}$ such that $\Pr_{x\sim\mathcal{D}}[x \geq m] = 1/2$, let T be the maximum value of the distribution $\mathcal{D}$. */*

- Let $R := \max\left(2c\sqrt{\frac{kT}{\epsilon}}, k\right)$. All bids offering at least m are candidates. If the number of candidates $s \leq R$, randomly select $\frac{k}{R} \cdot s$ candidates to confirm; else, randomly select k candidates to confirm. Every confirmed bid pays mid.
- If $s \geq \frac{h}{4}$, then the total miner revenue is $\min(\frac{h}{4} \cdot \frac{k}{R}, k) \cdot \mathsf{mid}$. Otherwise, the miners get nothing.

Intuitively, here are modifying the earlier threshold-based mechanism to 1) make it compatible with finite block size, and 2) make sure that up to c users dropping out can only minimally increase their friend's probability of getting confirmed. In particular, resilience to drop-out is achieved by artificially diluting the probability that a user is confirmed when the number of eligible bids (i.e., offering at least m) is small. With the dilution, we guarantee that a coalition of c users cannot noticeably alter their own probability of getting confirmed, nor their friend's probability. This implies a strategic coalition has little influence over the expected utility of all users in the coalition. Moreover, we guarantee that a strategic coalition has very little influence on the miner revenue as well: similar to the threshold-based mechanism, except with $\exp(-\Omega(h))$ probability, the miner revenue is an a-priori fixed amount, that is, $\min(\frac{h}{4} \cdot \frac{k}{R}, k) \cdot m$. Summarizing the above, we can show that the mechanism satisfies ex post UIC, Bayesian ϵ-MIC, and Bayesian ϵ-SCP in $(h, *, c, *)$-environments, as long as $\epsilon \geq \mathsf{mid} \cdot \frac{h}{2} \cdot e^{-\frac{h}{16}}$.

Finally, for sufficiently large $h \geq \max(4k, 8c\sqrt{\frac{kT}{\epsilon}})$, the mechanism achieves $k \cdot m$ total miner revenue and $k \cdot C_{\mathcal{D}}$ user social welfare where $C_{\mathcal{D}}$ is defined in Theorem 1.1. For example, suppose we are willing to tolerate $\epsilon = 0.01T$, then we just need $h \geq \max(4k, 80c \cdot \sqrt{k})$ to achieve asymptotic optimality in miner revenue and social welfare. The full proof is deferred to the online version [34].

2.4 Additional Results

Limit on miner revenue

The proof Theorem 1.1 is given in the online version [34]. Specifically, we prove that $\Theta(h)$ revenue is optimal in (h, ρ, c, d)-environments for strict incentive compatibility; and further, we generalize the bound to approximate incentive compatibility as well. The proof is a generalization of the techniques proposed by Shi, Chung, Wu [31]. Specifically, they proved that any mechanism that satisfies Bayesian UIC, MIC, and SCP in $(*, \rho, 1, *)$-environments must suffer from 0 miner revenue. In their proof, they argue that if we remove one bid, the miner revenue must be unaffected. In our case, because the mechanism is promised a lower

bound h on the number of honest uesrs, we can repeat this argument till there are h honest bids left, and no more. This gives rise to an $O(h)$ limit on miner revenue. The proof for the approximate incentive compatibility is also a generalization of [31]'s techniques, but more technically involved since even Myerson's lemma does not hold for approximate incentive compatibility.

Necessity of Bayesian equilibrium

As mentioned, our reasonable-world assumptions (formalized through the definition of an (h, ρ, c, d)-environment) would not have helped had we insisted on ex post notions of equilibrium (for all of UIC, MIC, and SCP). In Section 6 of the online version [34], we explain why for ex post notions of incentive compatibility, even mechanisms in the (h, ρ, c, d)-environment are subject to the same miner-revenue limitations of universal mechanisms.

3 Model and Definitions

Imagine that there are n_0 users, and each user has a transaction that wants to be confirmed. For $i \in [n_0]$, let v_i be user i's true valuation of getting its transaction confirmed. We want to design a transaction fee mechanism (TFM) such that no individual user or miner or a coalition thereof have any incentive to deviate from honest behavior. Throughout the paper, we consider a single-parameter environment, i.e., each user's bid is represented by a single, non-negative real number.

Chung and Shi [8]'s results ruled out the existence of interesting TFMs in the plain model without cryptography. The subsequent work of Shi, Chung, and Wu [31] considered how cryptography can help circumvent these strong impossibilities. Below, we review the MPC-assisted model proposed by [31].

3.1 MPC-Assisted Model

The definition of TFM has been given in Section 2.1. Here, we formalize the MPC-assisted model and the strategy spaces for users and miners.

Ideal-world game

Recently, blockchain projects such as Ethereum are developing "encrypted mempool" techniques which can be viewed as concrete protocols that realize an MPC-assisted model for TFM. However, for understanding the game theoretic landscape, it helps to abstract out the cryptography and think of it as a trusted *ideal functionality* (henceforth denoted $\mathcal{F}_{\text{mpc}}$) that always honestly implements the rules of the TFM.

With the ideal functionality $\mathcal{F}_{\text{mpc}}$, we can imagine the following game that captures an instance of the TFM:

1. Each user registers zero, one, or more identities with $\mathcal{F}_{\text{mpc}}$, and submits exactly one bid on behalf of each identity.
2. Using the vector of input bids as input, $\mathcal{F}_{\text{mpc}}$ executes the rules of the TFM. $\mathcal{F}_{\text{mpc}}$ now sends to all miners and users the output of the mechanism, including the set of bids that are confirmed, how much each confirmed bid pays, and how much revenue the miner gets.

We make a couple of standard assumptions:

- *Individual rationality*: each confirmed bid should pay no more than the bid itself;
- *Budget feasibility*: the miner revenue should not exceed the total payment from all confirmed bids.

Using standard techniques in cryptography, we can instantiate the ideal functionality $\mathcal{F}_{\text{mpc}}$ using an actual cryptographic protocol among the miners and users (see Appendix D of [31]). A detailed discussion about the instantiation is given in the online version [34].

Strategy space and utility

An honest user will always register a single identity and submit only one bid reflecting its true value. A strategic user or miner (possibly colluding with others) can adopt the following strategies or a combination thereof:

- *Bid untruthfully*: a strategic user can misreport its value;
- *Inject fake bids*: a strategic user or miner can *inject fake bids* by registering fake identities;
- *Drop out*: a strategic user can also *drop out* by not registering its real identity.

In the real-world cryptographic instantiation, strategic miners can also deviate from the honest MPC protocol. However, as mentioned, the MPC protocol retains security (i.e., can be simulated by the ideal-world game) as long as at least one miner is honest. Therefore, we need not explicitly capture this deviation in the ideal-world game. Finally, strategic miners can cause the MPC protocol to abort without producing output, and as mentioned, this deviation never makes sense since it results in a utility of 0; thus we also need not explicitly capture it in the ideal-world game.

Let v_i denote user i's true value. If user i's transaction is confirmed and its payment is p_i, then its utility is defined as $v_i - p_i$. The miner's utility is simply its revenue. The joint utility of a coalition $\mathcal{C}$ is defined as the sum of the utilities of all members in $\mathcal{C}$.

3.2 Defining Incentive Compatibility

In the plain model without cryptography, users submit their bids in the clear over a broadcast channel, and a strategic coalition can decide its strategy after observing the remaining honest users' bids. By contrast, in the MPC-assisted model, bids are submitted to the ideal functionality $\mathcal{F}_{\text{mpc}}$ (in the actual cryptographic instantiation, the users verifiably secret-share their bids among the miners). This means that the strategic coalition must now submit its bids without having observed other users' bids. Therefore, in the MPC model, it makes sense to consider a *Bayesian* notion of equilibrium rather than an *ex post* notion. In an *ex post* setting, we require that a strategic individual or coalition's best response is to act honestly even after having observed others' actions. In a *Bayesian* setting, we assume that every honest user's bid is sampled independently from some distribution $\mathcal{D}$, and we require that acting honestly maximizes the strategic individual or coalition's expected utility where the expectation is taken over not just the random coins of the TFM itself, but also over the randomness in sampling the honest users' bids.

Notations

Henceforth, we use the notation $\mathbf{b}$ to denote a bid vector. Since we allow strategic players to inject fake bids or drop out, the length of $\mathbf{b}$ need not be the same as the number of users n_0. We use the notation $\mathcal{C}$ to denote a coalition, and we use $\mathbf{b}_{-\mathcal{C}}$ to denote the bid vector belonging to honest users outside $\mathcal{C}$. We use the notation $\mathcal{D}_{-\mathcal{C}}$ to denote the joint distribution of $\mathbf{b}_{-\mathcal{C}}$, that is, $\mathcal{D}_{-\mathcal{C}} = \mathcal{D}^{h_0}$ where h_0 denotes the number of honest users outside $\mathcal{C}$. Similarly, if i is an individual strategic user, then the notation $\mathbf{b}_{-i}$ denotes the bid vector belonging the remaining honest users in $[n_0]\backslash\{i\}$. We use the notation $\mathcal{D}_{-i}$ to denote the joint distribution of the honest bid vector $\mathbf{b}_{-i}$.

For generality, we define *approximate* incentive compatibility parameterized by an additive slack ϵ. The case of *strict* incentive compatibility can be viewed as the special case when $\epsilon = 0$.

Bayesian incentive compatibility

▶ **Definition 3.1** (Bayesian incentive compatibility)**.** *We say that an MPC-assisted TFM satisfies Bayesian ϵ-incentive compatibility for a coalition or individual $\mathcal{C}$, iff for any $\mathbf{v}_\mathcal{C}$ denoting the true values of users in $\mathcal{C}$, sample $\mathbf{b}_{-\mathcal{C}} \sim \mathcal{D}_{-\mathcal{C}}$, then, no strategy can increase $\mathcal{C}$'s expected utility by more than ϵ in comparison with honest behavior, where the expectation is taken over randomness of the honest users' bids $\mathbf{b}_{-\mathcal{C}}$, as well as random coins consumed by the TFM. Specifically, we define the following notions depending on who is the strategic individual or coalition:*

- *User incentive compatibility (UIC). We say that an MPC-assisted TFM satisfies Bayesian ϵ-UIC in some environment $\mathcal{E}$, iff for any n, for any user $i \in [n]$, for any true value $v_i \in \mathbb{R}^{\geq 0}$ of user i, for any strategic bid vector $\mathbf{b}_i$ from user i which could be empty or consist of multiple bids, the following holds as long as the conditions required by the environment $\mathcal{E}$ are respected:*

$$\mathop{\mathbf{E}}_{\mathbf{b}_{-i} \sim \mathcal{D}_{-i}} \left[\mathsf{util}^i(\mathbf{b}_{-i}, v_i)\right] \geq \mathop{\mathbf{E}}_{\mathbf{b}_{-i} \sim \mathcal{D}_{-i}} \left[\mathsf{util}^i(\mathbf{b}_{-i}, \mathbf{b}_i)\right] - \epsilon$$

  where $\mathsf{util}^i(\mathbf{b})$ denotes the expected utility (taken over the random coins of the TFM) of user i when the bid vector is $\mathbf{b}$.
- *Miner incentive compatibility (MIC). We say that an MPC-assisted TFM satisfies Bayesian ϵ-MIC in some environment $\mathcal{E}$, iff for any miner coalition $\mathcal{C}$, for any strategic bid vector $\mathbf{b}'$ injected by the miner, the following holds as long as the conditions required by the environment $\mathcal{E}$ are respected:*

$$\mathop{\mathbf{E}}_{\mathbf{b}_{-\mathcal{C}} \sim \mathcal{D}_{-\mathcal{C}}} \left[\mathsf{util}^\mathcal{C}(\mathbf{b}_{-\mathcal{C}})\right] \geq \mathop{\mathbf{E}}_{\mathbf{b}_{-\mathcal{C}} \sim \mathcal{D}_{-\mathcal{C}}} \left[\mathsf{util}^\mathcal{C}(\mathbf{b}_{-\mathcal{C}}, \mathbf{b}')\right] - \epsilon$$

  where $\mathsf{util}^\mathcal{C}(\mathbf{b})$ denotes the expected utility (taken over the random coins of the TFM) of the coalition $\mathcal{C}$ when the input bid vector is $\mathbf{b}$.
- *Side-contract-proofness (SCP). We say that an MPC-assisted TFM satisfies Bayesian ϵ-SCP in some environment $\mathcal{E}$, iff for any miner-user coalition, for any true value vector $\mathbf{v}_\mathcal{C}$ of users in $\mathcal{C}$, for any strategic bid vector $\mathbf{b}_\mathcal{C}$ of the coalition (whose length may not be equal to the number of users in $\mathcal{C}$), the following holds as long as the requirements of the environment $\mathcal{E}$ are respected:*

$$\mathop{\mathbf{E}}_{\mathbf{b}_{-\mathcal{C}} \sim \mathcal{D}_{-\mathcal{C}}} \left[\mathsf{util}^\mathcal{C}(\mathbf{b}_{-\mathcal{C}}, \mathbf{v}_\mathcal{C})\right] \geq \mathop{\mathbf{E}}_{\mathbf{b}_{-\mathcal{C}} \sim \mathcal{D}^{-\mathcal{C}}} \left[\mathsf{util}^\mathcal{C}(\mathbf{b}_{-\mathcal{C}}, \mathbf{b}_\mathcal{C})\right] - \epsilon$$

Henceforth, if a mechanism satisfies Bayesian ϵ-UIC for $\epsilon = 0$ (i.e., the *strict* incentive compatibility case), we often omit writing the ϵ, and simply say that the mechanism satisfies Bayesian UIC. The terms "Bayesian MIC", and "Bayesian SCP" are similarly defined.

Notice that we only require honest users' true values are i.i.d. sampled, while the strategic players' true values can be arbitrary.

Ex post incentive compatibility

In the interest of space, we define ex post incentive compatibility in the online version [34].

4 Feasibility for Infinite Block Size

4.1 MPC-Assisted, Threshold-Based Mechanism

We assume that honest users' bids are drawn i.i.d. from some distribution $\mathcal{D}$ with the median mid such that $\Pr_{x\sim\mathcal{D}}[x \geq \mathsf{mid}] = 1/2$ (see Remark 2.1). For convenience, we repeat the MPC-assisted, threshold-based mechanism, which has been introduced in Section 2.2.

MPC-assisted, threshold-based mechanism

Parameters: lower bound h on the number of honest users, the distribution median mid.

Mechanism:

- *Confirmation rule.* Given a bid vector $\mathbf{b} = (b_1, \ldots, b_\ell)$, for each bid b_i, confirm b_i if $b_i \geq \mathsf{mid}$.
- *Payment rule.* Each confirmed bid pays mid.
- *Miner revenue rule.* Let s be the number of confirmed bids. If $s \geq \frac{h}{4}$, miner gets $\frac{h}{4} \cdot \mathsf{mid}$. Otherwise, the miner gets nothing.

▶ **Theorem 4.1.** *Fix any $h \geq 1$. The MPC-assisted, threshold-based mechanism satisfies ex post UIC, Bayesian ϵ-MIC and Bayesian ϵ-SCP in an $(h, *, *, *)$-environment, where $\epsilon = \frac{h}{4} \cdot \mathsf{mid} \cdot e^{-\frac{h}{16}}$.*

Proof. We prove three properties individually.

UIC

Since the confirmation and the payment of each bid are independent of other bids, injecting fake bids or dropping out does not help to increase each user's utility. For a fixed user i, suppose its true value is v. When it bids b, its expected utility is $v - \mathsf{mid}$ if $b \geq m$. By direct calculation, the expected utility is maximized when $b = v$ no matter what other bids are. Thus, the mechanism satisfies UIC.

ϵ-MIC

Recall that the only strategy that a strategic miner can apply is injecting some fake bids. Because injecting fake bids smaller than mid does not influence the colluding miners' utility, we only consider injecting bids at least mid. Let X denote the random variable representing the number of honest bids at least mid. The only situation where the colluding miners can increase their expected gain by injecting fake bids is when $X < \frac{h}{4}$. By the following Chernoff Bound,

▶ **Lemma 4.2** (Chernoff bound, Corollary A.1.14 [3]). *Let $X_1, \ldots, X_n$ be independent Bernoulli random variables. Let $\mu = \mathbf{E}\left[\sum_{i=1}^n X_i\right]$. Then, for any $\epsilon \in (0, 1)$, it holds that $\Pr\left[\sum_{i=1}^n X_i \leq (1-\epsilon)\mu\right] \leq e^{-\epsilon^2\mu/2}$.*

We have $\Pr\left[X < \frac{h}{4}\right] \leq e^{-\frac{h}{16}}$. Therefore, the colluding miners can gain at most $\frac{h}{4} \cdot \mathsf{mid} \cdot e^{-\frac{h}{16}}$ more expected revenue by injecting fake bids.

ϵ-SCP

Since the confirmation and the payment of each bid are independent of other bids, and the mechanism is strict UIC, the coalition cannot increase colluding users' utilities. Therefore, by deviating from the mechanism, the coalition can only try to increase the expected total

miner revenue. By a similar argument as MIC, the coalition can only increase the expected total miner revenue when $X < \frac{h}{4}$, which happens with a probability no more than $e^{-\frac{h}{16}}$. It follows that no matter how the coalition deviates, the expected miner's revenue can increase by at most $\rho \cdot \frac{h}{4} \cdot \mathsf{mid} \cdot e^{-\frac{h}{16}}$. ◀

4.2 Analysis of the LP-Based Mechanism

We now prove that the MPC-assisted LP-based mechanism satisfies strict incentive compatibility in an $(h, *, c, d)$-environment. Suppose that honest users' values are sampled i.i.d. from some distribution $\mathcal{D}$. Recall that mid denotes the median of the bid distribution $\mathcal{D}$, and $C_{\mathcal{D}} = \mathbf{E}_{x\sim\mathcal{D}}[x - m | x \geq m]$ is another constant related to the distribution $\mathcal{D}$. Without loss of generality, we assume $\Pr[x \geq m] = \frac{1}{2}$ (see Remark 2.1).

▶ **Theorem 4.3** (Theorem 1.3 restated)**.** *Suppose that the block size is infinite. Fix any*[7] $h \geq 2$*, and any* $d \leq \frac{1}{8}\sqrt{\frac{h}{2\log h}}$*, the MPC-assisted, LP-based mechanism guarantees ex post UIC, Bayesian MIC, and Bayesian SCP in an* $(h, *, *, d)$*-environment, and meanwhile, the mechanism achieves* $\Theta(h \cdot \mathsf{mid})$ *expected miner revenue, and at least* $\Theta(\widetilde{h} \cdot C_{\mathcal{D}})$ *expected social welfare for the users where* $\widetilde{h} \geq h$ *is the the actual number of honest users that show up.*

We prove Theorem 4.3 in two steps. First, we show that if the linear program defined in Equations (1) and (2) has a feasible solution, then the resulting mechanism satisfies incentive compatibility, as formally stated below:

▶ **Lemma 4.4.** *When the linear program defined in Equations* (1) *and* (2) *has a feasible solution, the LP-based mechanism satisfies ex post UIC, Bayesian MIC, and Bayesian SCP in an* $(h, *, *, d)$*-environment. Moreover, the expected miner revenue is* $\frac{h \cdot m}{4}$*, and the user social welfare is* $\Theta(\widetilde{h} \cdot C_{\mathcal{D}})$*.*

The proof of Lemma 4.4 requires an exhaustive case-by-case analysis, and a key observation is that Equation (2) ensures that the expected miner revenue is always $\frac{h \cdot m}{4}$ as long as no more than d bids are contributed by the strategic coalition. We defer the formal proof of Lemma 4.4 to the online version [34]. In the main body, we focus on proving the more challenging step, that is, as long as $d \leq \frac{1}{8}\sqrt{\frac{h}{2\log h}}$, the linear program indeed has a feasible solution, formally stated below.

▶ **Lemma 4.5.** *For* $h \geq 2$ *and* $d \leq \frac{1}{8}\sqrt{\frac{h}{2\log h}}$*, the linear program specified by Equations* (1) *and* (2) *is guaranteed to have a feasible solution.*

Proof. We will give a constructive solution to the linear program Equations (1) and (2). Let $\gamma := n - d$ denote the number of bids that are sampled randomly from $\mathcal{D}$. Let $t = \lfloor \frac{\gamma}{4} \rfloor$, and $\overline{\mu}$ be our target expected miner revenue $\frac{m \cdot h}{4}$. We start from an "approximate" solution $\widehat{\mathbf{y}} = (\widehat{y}_0, \ldots, \widehat{y}_n) \in \mathbb{R}^{n+1}$ such that $\widehat{y}_i = 0$ for any $i \leq t$, and $\widehat{y}_i = \overline{\mu}$ for any $i > t$. Our goal is to find a correction $\mathbf{e} = (e_0, \ldots, e_n) \in \mathbb{R}^{n+1}$ that is zero everywhere except for the indices $i \in [z + d, z + 2d]$ for some $z \geq \frac{\gamma}{2}$ such that $\widehat{\mathbf{y}} + \mathbf{e}$ is a feasible solution to the linear program Equations (1) and (2). Henceforth, let $\boldsymbol{\delta} := \mathbf{e}[z + d, z + 2d]/\overline{\mu}$ be the non-zero coordinates of the correction, scaled by $\overline{\mu}$. Then $\boldsymbol{\delta}$ must satisfy the linear system $A(z)\boldsymbol{\delta} = \boldsymbol{\Delta}$, where $A(z)$ and $\boldsymbol{\Delta}$ are defined as follows:

[7] For the special case $h = 1$, we can just use the parity-based mechanism of Section 2.2.

$$A(z) = \begin{pmatrix} \binom{\gamma}{z+d} & \binom{\gamma}{z+d+1} & \cdots & \binom{\gamma}{z+2d} \\ \binom{\gamma}{z+d-1} & \binom{\gamma}{z+d} & \cdots & \binom{\gamma}{z+2d-1} \\ \vdots & \vdots & \ddots & \vdots \\ \binom{\gamma}{z} & \binom{\gamma}{z+1} & \cdots & \binom{\gamma}{z+d} \end{pmatrix}, \qquad \boldsymbol{\Delta} = \begin{pmatrix} \sum_{i=0}^{t} \binom{\gamma}{i} \\ \sum_{i=0}^{t-1} \binom{\gamma}{i} \\ \vdots \\ \sum_{i=0}^{t-d} \binom{\gamma}{i} \end{pmatrix}.$$

If there exists a $z^* \in [\lceil \frac{\gamma}{2} \rceil, \lceil \frac{\gamma}{2} \rceil + 2d^2]$ such that this linear system $A(z^*)\boldsymbol{\delta} = \boldsymbol{\Delta}$ has a solution $\boldsymbol{\delta}$, then choosing $\mathbf{e}$ such that $\mathbf{e}[z^* + d : z^* + 2d] = \overline{\mu} \cdot \boldsymbol{\delta}$ gives a solution $\widehat{\mathbf{y}} + \mathbf{e}$ that satisfies Equation (2).

▷ Claim 4.6. There exists a $z^* \in [\lceil \frac{\gamma}{2} \rceil, \lceil \frac{\gamma}{2} \rceil + 2d^2]$ such that the matrix $A(z^*)$ is non-singular, and

$$\|A(z^*)^{-1}\|_\infty \leq \frac{(z^* + 2d)^{2d(d+1)}}{\binom{\gamma}{z^*}} \cdot \left(\frac{d+1}{\sqrt{d}}\right)^d. \tag{6}$$

When choosing this z^*, we have a unique solution $\boldsymbol{\delta} = A(z^*)^{-1}\boldsymbol{\Delta}$. Moreover, under the given parameter range, the solution $\boldsymbol{\delta}$ has bounded infinity norm:

▷ Claim 4.7. For $h \geq 2$ and $d \leq \frac{1}{8}\sqrt{\frac{h}{2\log h}}$, we have $\|\boldsymbol{\delta}\|_\infty \leq 1$.

For now, we assume that Claim 4.6 and Claim 4.7 are true, and we show how they lead to Lemma 4.5. The proofs of the two claims appear are given in the online version [34]. To prove Lemma 4.5, it suffices to show that $\widehat{\mathbf{y}} + \mathbf{e}$ indeed satisfies the budget feasibility specified by Equation (1). Since for all $i \notin [z^* + d, z^* + 2d]$, we have $\widehat{y}_i + e_i = \widehat{y}_i \leq i \cdot \mathsf{mid}$, so we only need to show that for the correction position $z^* + d, \ldots, z^* + 2d$, the budget feasibility is satisfied. Substituting $\|\boldsymbol{\delta}\|_\infty \leq 1$, for each $i \in [z^* + d, z^* + 2d]$, we have $|e_i| \leq \overline{\mu}$. This implies that $\widehat{y}_i + e_i \geq \overline{\mu} - \overline{\mu} = 0$. Moreover, $\widehat{y}_i + e_i \leq 2\overline{\mu} \leq \frac{\gamma}{2} \cdot m \leq i \cdot m$. Lemma 4.5 thus follows. ◀

Deferred Materials

In the full online version [34], we further give the complete characterization for finite block size and the upper bounds on the miner revenue under different settings.

References

1. Encrypted mempools. talk by Justin Drake, Ethereum Foundation, `https://www.youtube.com/watch?v=XRMOCpGY3sw`.
2. Ittai Abraham, Danny Dolev, Rica Gonen, and Joseph Halpern. Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation. In *PODC*, 2006.
3. Noga Alon and Joel H Spencer. *The probabilistic method*. John Wiley & Sons, 2016.
4. Gilad Asharov, Ran Canetti, and Carmit Hazay. Towards a game theoretic view of secure computation. In *Eurocrypt*, 2011.
5. Gilad Asharov and Yehuda Lindell. Utility dependence in correct and fair rational secret sharing. *Journal of Cryptology*, 24(1), 2011.
6. Soumya Basu, David A. Easley, Maureen O'Hara, and Emin Gün Sirer. Towards a functional fee market for cryptocurrencies. *CoRR*, abs/1901.06830, 2019. `arXiv:1901.06830`.

7 Vitalik Buterin, Eric Conner, Rick Dudley, Matthew Slipper, and Ian Norden. Ethereum improvement proposal 1559: Fee market change for eth 1.0 chain. https://github.com/ethereum/EIPs/blob/master/EIPS/eip-1559.md.

8 Hao Chung and Elaine Shi. Foundations of transaction fee mechanism design. In *SODA*, 2023.

9 Kai-Min Chung, T-H. Hubert Chan, Ting Wen, and Elaine Shi. Game-theoretic fairness meets multi-party protocols: The case of leader election. In *CRYPTO*. Springer-Verlag, 2021.

10 Kai-Min Chung, Yue Guo, Wei-Kai Lin, Rafael Pass, and Elaine Shi. Game theoretic notions of fairness in multi-party coin toss. In *TCC*, volume 11239, pages 563–596, 2018.

11 Yevgeniy Dodis and Tal Rabin. Cryptography and game theory. In *AGT*, 2007.

12 Meryem Essaidi, Matheus V. X. Ferreira, and S. Matthew Weinberg. Credible, strategyproof, optimal, and bounded expected-round single-item auctions for all distributions. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31 - February 3, 2022, Berkeley, CA, USA*, volume 215 of *LIPIcs*, pages 66:1–66:19, 2022.

13 Matheus V. X. Ferreira, Daniel J. Moroz, David C. Parkes, and Mitchell Stern. Dynamic posted-price mechanisms for the blockchain transaction-fee market. *CoRR*, abs/2103.14144, 2021. arXiv:2103.14144.

14 Matheus V. X. Ferreira and S. Matthew Weinberg. Credible, truthful, and two-round (optimal) auctions via cryptographic commitments. In Péter Biró, Jason D. Hartline, Michael Ostrovsky, and Ariel D. Procaccia, editors, *EC '20: The 21st ACM Conference on Economics and Computation, Virtual Event, Hungary, July 13-17, 2020*, pages 683–712. ACM, 2020.

15 Yotam Gafni and Aviv Yaish. Greedy transaction fee mechanisms for (non-)myopic miners, 2022. doi:10.48550/arXiv.2210.07793.

16 Juan Garay, Jonathan Katz, Björn Tackmann, and Vassilis Zikas. How fair is your protocol? a utility-based approach to protocol optimality. In *PODC*, 2015.

17 Juan A. Garay, Jonathan Katz, Ueli Maurer, Björn Tackmann, and Vassilis Zikas. Rational protocol design: Cryptography against incentive-driven adversaries. In *FOCS*, 2013.

18 Juan A. Garay, Björn Tackmann, and Vassilis Zikas. Fair distributed computation of reactive functions. In *DISC*, volume 9363, pages 497–512, 2015.

19 Andrew V. Goldberg and Jason D. Hartline. Collusion-resistant mechanisms for single-parameter agents. In *SODA 2005*, pages 620–629, 2005.

20 Ronen Gradwohl, Noam Livne, and Alon Rosen. Sequential rationality in cryptographic protocols. In *FOCS*, 2010.

21 Joseph Halpern and Vanessa Teague. Rational secret sharing and multiparty computation. In *STOC*, 2004.

22 Sergei Izmalkov, Silvio Micali, and Matt Lepinski. Rational secure computation and ideal mechanism design. In *FOCS*, 2005.

23 Jonathan Katz. Bridging game theory and cryptography: Recent results and future directions. In *TCC*, 2008.

24 Gillat Kol and Moni Naor. Cryptography and game theory: Designing protocols for exchanging information. In *TCC*, 2008.

25 Ilan Komargodski, Shin'ichiro Matsuo, Elaine Shi, and Ke Wu. log*-round game-theoretically-fair leader election. In *CRYPTO*, 2022.

26 Ron Lavi, Or Sattath, and Aviv Zohar. Redesigning bitcoin's fee market. In *The World Wide Web Conference, WWW 2019*, pages 2950–2956, 2019.

27 Shien Jin Ong, David C. Parkes, Alon Rosen, and Salil P. Vadhan. Fairness with an honest minority and a rational majority. In *TCC*, 2009.

28 Rafael Pass and Elaine Shi. Fruitchains: A fair blockchain. In *PODC*, 2017.

29 Tim Roughgarden. Transaction fee mechanism design for the Ethereum blockchain: An economic analysis of EIP-1559. Manuscript, https://timroughgarden.org/papers/eip1559.pdf, 2020.

30 Tim Roughgarden. Transaction fee mechanism design. In *EC*, 2021.

31 Elaine Shi, Hao Chung, and Ke Wu. What can cryptography do for decentralized mechanism design. In *ITCS*, 2023.

32 William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

33 Ke Wu, Gilad Asharov, and Elaine Shi. A complete characterization of game-theoretically fair, multi-party coin toss. In *Eurocrypt*, 2022.

34 Ke Wu, Elaine Shi, and Hao Chung. Maximizing miner revenue in transaction fee mechanism design. Cryptology ePrint Archive, Paper 2023/283, 2023. URL: `https://eprint.iacr.org/2023/283`.

35 Andrew Chi-Chih Yao. An Incentive Analysis of Some Bitcoin Fee Designs (Invited Talk). In *ICALP 2020*, 2020. `doi:10.4230/LIPIcs.ICALP.2020.1`.

36 Yi-Sheng Yu and Dun-He Gu. A note on a lower bound for the smallest singular value. *Linear algebra and its Applications*, 253(1-3):25–38, 1997.

37 Zishuo Zhao, Xi Chen, and Yuan Zhou. Bayesian-nash-incentive-compatible mechanism for blockchain transaction fee allocation, 2022. `arXiv:2209.13099`.