Bayesian-Nash-Incentive-Compatible Mechanism for Blockchain Transaction Fee Allocation

Zishuo Zhao
Department of ISE
University of Illinois Urbana-Champaign
Urbana, IL
zishuoz2@illinois.edu

Xi Chen
Stern School of Business
New York University
New York, NY
xc13@stern.nyu.edu

Yuan Zhou
Yau Mathematical Sciences Center
Tsinghua University
Beijing, China
yuan-zhou@tsinghua.edu.cn

Abstract

In blockchain systems, the design of transaction fee mechanisms is essential for stability and satisfactory for both miners and users. A recent work has proven the impossibility of collusion-proof mechanisms with non-zero miner revenue which is Dominate-Strategy-Incentive-Compatible (DSIC) for users. In our work, we relax the DSIC requirement for users to Bayesian-Nash-Incentive-Compatibility (BNIC), and design a so-called soft second-price mechanism to ensure a form of collusion-proofness with an asymptotic constant-factor approximation of optimal miner revenue. Our result breaks the zero-revenue barrier while preserving reasonable truthfulness and collusion-proof properties.

1 Introduction

Blockchains, like Bitcoin and Ethereum, are essentially distributed databases growing over time with history data saved in "blocks" as linked lists created by miners. In each block, users leverage blockchains to either store or verify information, such as money (cryptocurrency) transfers, texts, and modern data such as smart contracts. This on-chain information is usually refered to as "transactions'. Miners, in turn, try to "mine" a block by getting the access to write a block with certain efforts (computational power in PoW or cryptocurrency deposit in PoS) and then put information into the block.

To incentivize the miners to mine the block, they need to get paid in cryptocurrency, which consists of a mining reward and an additional reward extract from transaction fees paid by users. Generally, a block has limited size and takes social cost (e.g., PoW/PoS and storage space) to create, and users can benefit from transactions being confirmed on the blockchain. Thus, it is reasonable to charge transaction fees from confirmed transactions. To facilitate social efficiency of the system, we want to confirm transactions of high values. Therefore, many blockchains adopt certain bidding-confirmation transaction fee mechanism (TFM) such as auctions.

However, due to the online and anonymous properties of blockchains, transaction fee mechanism design faces a major concern of *credibility* Daskalakis et al. (2020). Compared to traditional auctions, the miner has a wider strategy space to conduct dishonest activities, including adding fake transactions, concealing users' bids, and colluding with users. Therefore, a desirable TFM should prevent these dishonest activities so that it can operate correctly and efficiently.

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

One advantage of blockchain is as follows. Since the blockchain is public, it is not possible for the miner to behave in a Byzantine way via commuting different bidding vectors to different users (see discussion from Daskalakis et al. (2020)). Besides, the paper by Daskalakis et al. (2020)) adopts a *secure commitment* scheme, which uses cryptographic protocols to guarantee that a bid cannot be modified after proposed, thus restricting the strategy space of the miner to merely adding fake transactions and conceal transactions, ruling out strategies for the miner to collude with users and make them change existing bids. However, the security of such cryptographic protocols is being constantly challenged by modern cryptanalysis techniques as quantum computation Anand et al. (2020) and subject to backdoor attacks Young and Yung (2005), rendering even currently secure protocols in risk of possible future attacks. Since blockchain systems are designed to operate into the future and there are heavy costs in updating their mechanisms, we are motivated to resolve the collusion issue via economic methods by preventing them from being profitable.

Basically, a desirable TFM should satisfy truthfulness. A strong version of truthfulness of the users can be specified as the User-Dominating-Strategy-Incentive-Compatibility (U-DSIC), which means that any single user will not benefit from deviation from truthful bidding even if she knows all bids of other users (as in Definition 3). In comparison, a weaker version of truthfulness of the users is the User-Bayesian-Nash-Incentive-Compatibility (U-BNIC), which means that when each user only knows the distribution of others' valuations, it is a Bayesian Nash equilibrium that all users truthfully bid their valuations (as in Definition 4). The truthfulness of the miner can be specified as the Miner-Incentive-Compatibility (MIC), which means that the miner will not benefit from untruthful behavior, e.g. injecting fake transactions or ignoring existing transactions. For the issue of collusion, the paper by Chung and Shi (2021) formulates collusion-proofness as c-Side-Contract-Proof (c-SCP): when the miner colludes with at most c users by asking them to change their bids, the coalition cannot gain increased total payoff by deviations from truthfully bidding their valuations (as in Definition 5).

Unfortunately, the paper by Chung and Shi (2021) proved a negative result that any TFM that is U-DSIC and 1-SCP has zero miner revenue, indicating that there does not exist a non-trivial "ideal" TFM which can incentivize everyone to behave honestly. Thus, they relax the payoff function (as "γ-strict utility") to make unconfirmed over-bidder still pay a γ fraction of the worst-case cost, i.e. if a bidder i has valuation v_i and her bid $b_i > v_i$, even if the transaction is not confirmed, she gets a payoff of $-\gamma(b_i - v_i)$, and a fake transaction is assumed to have zero valuation. Thus, if the confirmation probability is $a_i(b_i, \mathbf{b}_{-i})$ ($\mathbf{b}_{-i}$ denotes bids of all other users than i) and she should pay $p_i(b_i, \mathbf{b}_{-i})$ if the transaction gets confirmed, the utility she gets is:

$$u_i^{(\gamma)}(b_i, \mathbf{b}_{-i}; v_i) = a_i(b_i, \mathbf{b}_{-i})(v_i - p_i(b_i, \mathbf{b}_{-i})) - \gamma(1 - a_i(b_i, \mathbf{b}_{-i})) \max\{b_i - v_i, 0\}.$$

The above relaxed payoff function is justified by the authors of Chung and Shi (2021) by considering the bidding process of more than one blocks in a blockchain: even if an overbidding (including fake) transaction is not confirmed in the current block, the authors assume that the bid could still be collected and confirmed into future blocks, and in the worst case, the over-bidder would have to pay their full bid and get a utility of $-(b_i - v_i)$. In this setting, the authors have further developed a Burning Second Price TFM that satisfies U-DSIC, MIC and c-SCP in the notion of γ-strict utility.

Intuitively, the setting of γ-strict utility circumvents the impossibility result essentially by imposing more penalty to deviations in a way that coarsely resemble the nature of blockchains. However, such a relaxation might lead to further issues. For example, the choice of γ is essential in the mechanism, but finding an accurate γ can be difficult if not impossible: the probability that a currently unconfirmed transaction gets confirmed in future blocks is not a universal constant, as unconfirmed transactions with higher bids are more likely to be confirmed in the future than those with lower bids. In this sense, the authors have proposed an open question: whether there are other reasonable relaxations of the models and incentive compatibility specifications that can also circumvent the impossibility result.

In this work, we address this open question by getting back to the one-block setting where a transaction is only valid for the current block and expires if not confirmed immediately, a setting that is easy to implement via time stamps. In turn, we relax the U-DSIC requirement to U-BNIC, which assumes the users only have information of distributions of other users' valuations instead of all their bids. This relaxation is reasonable because in the distributed network of blockchain, it is impossible for a user to actually know all other users' bids, especially those who propose transactions after them but competing the same block. On this basis, we design a TFM that satisfies U-BNIC and 1-SCP for bounded *i.i.d.* valuation distributions with a constant-factor approximation of the optimal revenue, thus answering the open question of Chung and Shi (2021).

More specifically, we first prove an impossibility result that no deterministic U-BNIC + 1-SCP TFM (satisfying mild conditions) can achieve positive miner reward (Theorem 2), and then modify the second-price auction by introducing randomness in the allocation rule, develop our main mechanism for block size 1, and estimate the optimal parameters to achieve an asymptotic constant fraction approximation of the optimal miner revenue (as in Section 3). Finally, we extend our mechanism to general block size k with $n \geq \left(\frac{e}{e-1} + \Theta(1)\right) k$, and prove that the asymptotic constant fraction approximation also holds in this general case.

2 Basic Models

There are n users numbered by $1, 2, \dots, n$. For user i, w.l.o.g. we assume her valuation v_i is in $[0,1]$ and drawn from an *i.i.d.* distribution $V_i = V_0$ with pdf $\rho_i(\cdot) = \rho(\cdot)$. By the revelation principle Durlauf and Blume (2010), we only need to consider direct mechanisms in which users propose bids, the miner collects the bids and the system decides which transactions to confirm and processes the payments. Formally, we can model any Transaction Fee Mechanism as its allocation, payment and miner revenues, i.e.

Definition 1 (Transaction Fee Mechanism). *For a fixed number n of users, a Transaction Fee Mechanism is modeled as $M(\mathbf{a}, \mathbf{p}, r)$:*

- *The* allocation rule $\mathbf{a} : [0,1]^n \to [0,1]^n$ *maps the bid vector to the allocation vector indicating the probability each user's transaction to be confirmed.*
- *The* payment rule $\mathbf{p} : [0,1]^n \to \mathbb{R}^n$ *maps the bid vector to the payment vector indicating the payment of a user if her transaction is confirmed.*
- *The* miner revenue rule $r : [0,1]^n \to \mathbb{R}$ *maps the bid vector to the miner's revenue.*

In execution of the mechanism, the system just needs to let all users propose their bids $\{b_i\}$, draw the confirmed transactions according to probabilities from $\{a_i(b_i, \mathbf{b}_{-i})\}$, and then charge transaction fees from confirmed bidders according to $\{p_i(b_i, \mathbf{b}_{-i})\}$ and give the miner revenue $r(\mathbf{b})$ to the miner.

Additional definitions are deferred to Appendix B.

3 The Proposed Mechanism for Block Size 1

We first consider the case with block size $k = 1$, where exactly one transaction is confirmed, to give a simple and intuitive understanding of our mechanism. In Appendix C, we further extend our setting to general block size k for the main mechanism.

3.1 Warm Up: Soft Second-Price Mechanism

The second-price auction mechanism has been widely used in traditional auctions , in which the highest bidder gets confirmed but pays the second highest bid. However, as we prove that any deterministic TFM which is U-BNIC and 1-SCP satisfying mild assumptions has non-positive miner revenue (Appendix A), we try to introduce randomness into the allocation rule.

As a basis of our main mechanism, we firstly develop a U-DSIC and 1-SCP mechanism named soft second-price mechanism. It adopts the logit choice model as the allocation rule. After fixing the allocation rule, the payment rule can be correspondingly fixed by the Myerson's Lemma (Lemma 3) Myerson (1981), and from an impossibility result (1-SCP + U-DSIC ⇒ Zero Miner Revenue) Chung and Shi (2021) we set the miner revenue to zero. The resulting mechanism is shown as:

$$a_i(b_i, \mathbf{b}_{-i}) = \frac{e^{mb_i}}{\sum_{j=1}^{n} e^{mb_j}} \tag{1}$$

$$p_i(b_i, \mathbf{b}_{-i}) = b_i - \frac{\sum_{j=1}^{n} e^{mb_j}}{me^{mb_i}} \cdot \ln \frac{\sum_{j=1}^{n} e^{mb_j}}{1 + \sum_{j \neq i} e^{mb_j}} \tag{2}$$

$$r(\mathbf{b}) = 0. \tag{3}$$

The proof that this mechanism is U-DSIC and 1-SCP is deferred in Appendix E.1. Although the soft second-price mechanism has zero miner revenue, we can modify it in a way that preserves U-BNIC and 1-SCP properties and yields positive expected miner revenue, as in Section 3.2.

3.2 Our Proposed Mechanism for Block Size 1

We notice that the Soft Second-Price Mechanism is U-DSIC, and thus it is U-BNIC too. Here, our main idea is that when we perturb the payment function in such a way that the expected payment for any user is preserved, the U-BNIC property is also preserved.

Therefore, if $M = (\mathbf{a}, \mathbf{p}, 0)$ is a U-DSIC and 1-SCP mechanism, we can construct a mechanism $\tilde{M} = (\mathbf{a}, \tilde{\mathbf{p}}, \tilde{r})$, in which we denote $\tilde{p}_i(b_i, \mathbf{b}_{-i}) = p_i(b_i, \mathbf{b}_{-i}) + \frac{\theta_i(b_i, \mathbf{b}_{-i})}{a_i(b_i, \mathbf{b}_{-i})}$. Then we observe that: (proof deferred in Appendix E.2)

Observation 1. *$\tilde{M}$ is U-BNIC if*

$$\mathbb{E}_{\mathbf{b}_{-i} \sim V_{-i}}[\theta_i(b_i, \mathbf{b}_{-i})] = 0. \tag{4}$$

As we have characterized a sufficient condition for U-BNIC, now we consider the condition for 1-SCP. From Lemma 4 (in Appendix) we know that for an 1-SCP mechanism $(\mathbf{a}, \tilde{\mathbf{p}}, \tilde{\mathbf{r}})$, if we fix $\mathbf{b}_{-i}$, the difference of $a_i(\cdot, \mathbf{b}_{-i})\tilde{p}_i(\cdot, \mathbf{b}_{-i})$ and $\tilde{r}_i(\cdot, \mathbf{b}_{-i})$ is a constant. Furthermore, since $a(\cdot, \mathbf{b}_{-i})$ is monotonic increasing, if we want $\tilde{M}$ to be 1-SCP, from Lemma 4 we only need:

$$\theta_i(b_i, \mathbf{b}_{-i}) - \theta_i(0, \mathbf{b}_{-i}) = \tilde{r}(b_i, \mathbf{b}_{-i}) - \tilde{r}(0, \mathbf{b}_{-i}). \tag{5}$$

When the distributions of all users' valuations are i.i.d, i.e. $V = V_1 \times V_2 \times \cdots \times V_n$ and $\forall V_i$ has identical pdf $\rho : [0, 1] \to [0, +\infty)$, we denote

$$c_\rho = \int_0^1 \rho^2(t)dt, \tag{6}$$

then we can construct $(\theta, \tilde{r})$ as following:

$$\theta_i(b_i, \mathbf{b}_{-i}) = -\frac{1}{2}hb_i^2 \left(\frac{\sum_{j \neq i} b_j^2}{c_\rho(n-1)} - 1 \right) \tag{7}$$

$$\tilde{r}(\mathbf{b}) = \frac{1}{2}h \left(\sum_{i=1}^n b_i^2 - \frac{\sum_{1 \leq i < j \leq n} b_i^2 b_j^2}{c_\rho(n-1)} \right) \tag{8}$$

The corresponding mechanism $\tilde{M} = (a, \tilde{p}, \tilde{r})$ can be represented as:

$$\boxed{\begin{aligned}
a_i(b_i, \mathbf{b}_{-i}) &= \frac{e^{mb_i}}{\sum_{j=1}^n e^{mb_j}} \\
\tilde{p}_i(b_i, \mathbf{b}_{-i}) &= b_i - \frac{\sum_{j=1}^n e^{mb_j}}{me^{mb_i}} \left(\ln \frac{\sum_{j=1}^n e^{mb_j}}{1 + \sum_{j \neq i} e^{mb_j}} + \frac{1}{2}hmb_i^2 \left(\frac{\sum_{j \neq i} b_j^2}{c_\rho(n-1)} - 1 \right) \right) \\
\tilde{r}(\mathbf{b}) &= \frac{1}{2}h \left(\sum_{i=1}^n b_i^2 - \frac{\sum_{1 \leq i < j \leq n} b_i^2 b_j^2}{c_\rho(n-1)} \right)
\end{aligned}}$$

and it is U-BNIC and 1-SCP for $h \in [0, +\infty)$. We can compute that

$$E_{\mathbf{b} \sim V}[r(\mathbf{b})] = \frac{1}{4}hnc_\rho > 0. \tag{9}$$

However, we have to be careful to the value of h. Intuitively, the value of h describes the extent of perturbation from the original U-DSIC mechanism, and when the perturbation is too large, the *individual rationality* ($p_i(b_i, \mathbf{b}_{-i}) \leq b_i$) and *budget feasibility* properties may not hold. For best miner revenue, we want to make h as large as possible while keeping the mechanism feasible. Fortunately, for fixed c_ρ, we have an estimation of optimally feasible h that enables constant approximation ratio of the optimal revenue while preserving IR and BF constraints, as (proof deferred to Appendix E.3):

Theorem 1. *For $n \geq 1$ and $M = (\mathbf{a}, \mathbf{p}, 0)$ given by (1-3), let $m = 1$, we have $h_* = h_*(n, c_\rho) > 0$ such that $\forall h \in [0, h_*]$, the corresponding mechanism $\tilde{M} = (\mathbf{a}, \tilde{\mathbf{p}}, \tilde{r})$ is UIR and BF, and for $n \to +\infty$,*

$$h_*(n, c_\rho) = \Omega(c_\rho / n). \tag{10}$$

Acknowledgments

Zishuo Zhao would appreciate the valuable and delightful discussions in Yao Class Seminar #30 for partially inspiring his idea to design Bayesian-Nash Incentive Compatible mechanisms.

References

Anand, R., Maitra, A., and Mukhopadhyay, S. (2020). Evaluation of quantum cryptanalysis on speck. In *International Conference on Cryptology in India*, pages 395–413. Springer.

Chung, H. and Shi, E. (2021). Foundations of transaction fee mechanism design.

Daskalakis, C., Fishelson, M., Lucier, B., Syrgkanis, V., and Velusamy, S. (2020). Simple, credible, and approximately-optimal auctions. In *Proceedings of the 21st ACM Conference on Economics and Computation*, EC '20. Association for Computing Machinery.

Durlauf, S. N. and Blume, L. E. (2010). Revelation principle. In *Game Theory*, pages 312–318. Springer.

Leonardos, S., Monnot, B., Reijsbergen, D., Skoulakis, E., and Piliouras, G. (2021). Dynamical analysis of the eip-1559 ethereum fee market. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, pages 114–126.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of operations research*, 6(1):58–73.

Young, A. and Yung, M. (2005). A space efficient backdoor in rsa and its applications. In *International Workshop on Selected Areas in Cryptography*, pages 128–143. Springer.