Designing Redistribution Mechanisms for Reducing Transaction Fees in Blockchains

Sankarshan Damle*
IIIT, Hyderabad
Hyderbad, India
sankarshan.damle@research.iiit.ac.in

Manisha Padala*
IISc, Bangalore
Bangalore, India
manishap@iisc.ac.in

Sujit Gujar
IIIT, Hyderabad
Hyderbad, India
sujit.gujar@iiit.ac.in

ABSTRACT

Blockchains deploy Transaction Fee Mechanisms (TFMs) to determine which user transactions to include in blocks and determine their payments (i.e., transaction fees). Increasing demand and scarce block resources have led to high user transaction fees. As these blockchains are a public resource, it may be preferable to reduce these transaction fees. To this end, we introduce Transaction Fee Redistribution Mechanisms (TFRMs) – redistributing VCG payments collected from such TFMs as rebates to minimize transaction fees. Classic redistribution mechanisms (RMs) achieve this while ensuring Allocative Efficiency (AE) and User Incentive Compatibility (UIC). Our first result shows the non-triviality of applying RM in TFMs. More concretely, we prove that it is impossible to reduce transaction fees when (i) transactions that are not confirmed do not receive rebates and (ii) the miner can strategically manipulate the mechanism. Driven by this, we propose *Robust* TFRM (R-TFRM): a mechanism that compromises on an honest miner's individual rationality to guarantee strictly positive rebates to the users. We then introduce *Robust* and *Rational* TFRM (R^2-TFRM) that uses trusted on-chain randomness that additionally guarantees miner's individual rationality (in expectation) and strictly positive rebates. Our results show that TFRMs provide a promising new direction for reducing transaction fees in public blockchains.

KEYWORDS

Transaction Fee Mechanism Design, Redistribution Mechanism

ACM Reference Format:
Sankarshan Damle, Manisha Padala, and Sujit Gujar. 2024. Designing Redistribution Mechanisms for Reducing Transaction Fees in Blockchains. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 9 pages.

1 INTRODUCTION

Public blockchains have achieved mainstream prominence with Bitcoin [30] and Ethereum [4] processing $> 1M$ transactions daily [39, 40]. Most commonly, public blockchains comprise a cryptographically linked series of blocks. Each block may consist of several individual transactions. *Miners*, tasked with block creation, add a subset of transactions from the set of outstanding transactions (referred to as *mempool*). To incentivize miners to add their transaction to the block, transaction creators (henceforth *users*) include a *fee* as a commission. The fee absorbs the users' *valuation* for their transaction being added to the block. Roughgarden [34] proposes *transaction fee mechanisms* (TFMs) to study the strategic interaction between the miner and the users.

Transaction Fee Mechanism (TFM). TFMs resemble a classic auction setting. Users place a bid to include their transactions in the block, and the miner mimics an auctioneer to select the subset, which maximizes its revenue. E.g., Bitcoin's TFM resembles a first-price auction, where the block's miner greedily adds the transactions with the highest bids. Unfortunately, an increasing demand, cryptocurrency's market volatility, and supply-demand economics have led to users' over-paying [2]. E.g., Messias et al. [27] show that 30% of Bitcoin fees are two orders of magnitude more than recommended.

Considering public blockchains as a shared resource, it's desirable not to impose charges for transaction confirmation. However, given the infeasibility of confirming every transaction due to resource constraints, one may prefer only to confirm transactions with higher value (pertaining to their importance to users). The absence of transaction fees could lead users to misrepresent the value of their transactions in order to secure confirmation. Therefore, this paper aims to design TFMs that minimize transaction fees while upholding other incentive-related properties.

Clearly, the minimization of transaction fees is at odds with the miner's objective of maximizing revenue. Thus, the task of designing TFMs to minimize fees is more intricate than in the classic auction setting, primarily because, in TFMs, miners have complete control over the transactions they include in their blocks [34].

Our Goal. We aim to design a TFM that satisfies certain game theoretic properties like (i) Allocative Efficiency (AE): confirmed transactions maximize the overall valuation, (ii) User Incentive Compatibility (UIC): users bid their true valuation, and Individual Rationality (IR): users receive non-negative payoff. At the same time, the TFM must actively reduce transaction fees for users, thereby enhancing the blockchain's appeal. Unfortunately, from the famous Green-Laffont Impossibility Theorem [26], we know that it is impossible to design a TFM that is both AE and UIC and which guarantees *zero net transaction fees* – in mechanism design commonly referred to as *strong budget balance*.

Given this, our objective is to design a TFM that is both AE and UIC while minimizing the transaction fees (or is *weakly budget balanced*). Motivated from Maskin et al. [26], the mechanism design literature proposes the use of *Groves' Redistribution Mechanism* (RM) for this purpose [7, 18, 26]. In Groves' RM, the VCG mechanism is

*Equal Contribution.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

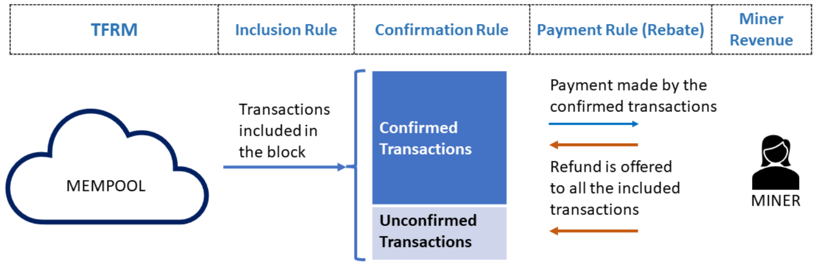

Figure 1: Overview of the framework for Transaction Fee Redistribution Mechanisms (TFRMs)

executed, and then the surplus money is redistributed among the users while preserving other game-theoretic properties.

Along similar lines, this paper introduces *Transaction Fee Redistribution Mechanisms* (TFRMs): a general class of TFMs based on RMs where the miner offers rebates from the transaction fees collected to the users while retaining AE, UIC, and IR. By offering users rebates, TFRMs, in effect, reduce the transaction fees paid by them. Figure 1 provides an overview.

TFRM: Challenges. Designing such TFRMs has the following primary challenges.

Miner IC (MIC): As miners possess complete control over the transactions included in their blocks [34], they may deviate from the intended TFM allocation rule (i.e., selecting a different subset from the mempool) and may introduce "fake" transactions (i.e., transactions created strategically to increase their revenue) into their blocks [34]. This is similar to *shill bidding* [33] in traditional auctions. Thus, it is imperative that a TFRM maintains AE, UIC, and low transaction fees even in the face of miner manipulation (or, alternately, in the presence of a *strategic miner*).

Roughgarden [34] introduces the notion of *miner IC* (MIC) to model the miner's strategic behavior. In auction theory, the Myerson-Satterthwaite impossibility theorem [29] states that it is impossible to design a mechanism that is AE, IR, weakly budget balanced, and IC for both sides of the market. Designing TFRMs has similarities with such a two-sided auction, and achieving both sides' IC (with other properties) is elusive.

User IC (UIC): Typically, RMs ensure UIC by offering rebates to everyone participating in the auction (irrespective of the allocation). In TFRMs, the transactions that are only part of the mempool (i.e., are not part of the block) are not available to the blockchain. Thus, unlike RMs, in TFRMs, we cannot offer rebates for each available transaction. As some transactions do not receive rebates, we can easily construct instances where the users of these transactions have an incentive to overbid to get included in the block and receive rebates. Thus, ensuring UIC in TFRM is non-trivial. As such, we propose *restricted UIC* (RUIC), which ensures that bidding truthfully is a weakly dominant strategy *only* for the users whose transactions are included in the block.

Our Contributions. Broadly, we (i) formally introduce TFRMs (refer to Figure 1 for an overview), (ii) analyze the challenges due to miner manipulation in vanilla-TFRMs, and (iii) introduce two novel TFRMs, namely R-TFRM and R^2-TFRM that are robust to miner manipulation. We discuss these in detail next.

(1) `Ideal-TFRM`. As we cannot offer rebates to all transactions in the mempool, we begin our analysis with an "`Ideal-TFRM`" that offers non-zero rebates only to confirmed transactions. Unfortunately, we show that it is impossible for `Ideal-TFRM` to satisfy UIC while offering non-zero rebates to confirmed transactions (Theorem 2).

(2) **TFRM: Effect of A Strategic Miner.** We shift our focus to TFRMs that provide rebates to all transactions included in the block. An RM's effectiveness is measured using the Redistribution Index (RI) [19], which is the fraction of the VCG surplus redistributed. To absorb the effect of strategic miners, we introduce *Resilient Redistribution Index* (RRI). RRI measures the fraction of redistributed funds under optimal miner manipulation. We prove that it is impossible to design a TFRM that satisfies AE, RUIC, and is IR for both users (IR$_u$) and miners (IR$_\mathbf{M}$), while guaranteeing strictly positive RRI (Theorem 3).

(3) **Robust TFRM (R-TFRM).** Given these impossibilities, we propose R-TFRM: a TFRM that guarantees strictly positive RRI *and* satisfies all user-specific properties. However, R-TFRM is not individually rational for the miner[1].
At its core, R-TFRM builds on an RM with VCG payments as transaction fees and a linear rebate function. The rebate function maximizes the worst-case rebate while satisfying RUIC, IR$_u$, and Approx-IR$_\mathbf{M}$. Designing such a rebate function is equivalent to solving the linear program in Figure 4 for its coefficients. We finally show that the payments are reduced by a fraction k/n, where k transactions are confirmed out of the n included in the block (Theorem 5). The fraction remains the same, even with miner manipulation, i.e., RRI is also k/n. In other words, each confirmed user sees a reduction by $(1 - k/n)$ in its transaction fee compared to the equivalent VCG-based TFM.

(4) **Robust and Rational TFRM (R^2-TFRM).** R-TFRM ensures positive RRI by compromising IR$_\mathbf{M}$. Another way of ensuring positive RRI is by randomly offering rebates to the users. Such an approach guarantees IR$_\mathbf{M}$, in expectation. R^2-TFRM uses this approach wherein each user receives the rebate given by R-TFRM with probability α and does not receive any rebate probability $1 - \alpha$. The randomization is carried out by the blockchain in a trusted manner [8]. Theorem 6 shows that for

[1]We remark that miners, or block proposers in general, often have alternate revenue streams (e.g., block rewards [30] or attestation rewards [36]). These rewards can primarily help absorb the reduction in revenue due to reduced transaction fees. They may also alleviate the lack of IR$_\mathbf{M}$ guarantee in R-TFRM.

$\alpha \in (0, \overline{\alpha})$ & $\overline{\alpha} < 1$, R^2-TFRM is AE, IR_u, and RUIC and $\text{IR}_{\mathbf{M}}$, in expectation. Further, it ensures an expected RRI of $\alpha \cdot k/n$.

2 RELATED WORK

The role of transaction fees in decentralized cryptocurrencies such as Bitcoin and Ethereum has been studied in relation to (i) transaction latency [23, 24, 31], (ii) fairness [2, 35] and most recently, as decentralized auction-based mechanisms [8, 15, 34, 41].

Transaction Fee Mechanism (TFM). Roughgarden [34] formulated the transaction creator and miner interaction as an auction setting. More concretely, the author expresses popular mechanisms in the TFM framework, including first-price, second-price, and EIP-1559. Roughgarden [34] introduces user incentive compatibility (UIC), miner incentive compatibility (MIC), and off-chain agreement (OCA) proof as desirable incentive properties for TFMs. E.g., EIP-1559 satisfies UIC, MIC, and OCA-proof – under specific constraints on the base fee. Ferreira et al. [15] present a dynamic posted-price TFM, which is UIC (if the network is not congested) and MIC. The authors also provide an equilibrium posted price based on the demand. Chung and Shi [8] show it is impossible to construct a TFM that simultaneously satisfies UIC, MIC, and OCA-proof (when only a single user and the miner collude). To address the impossibility, they introduce a *penalty* to the creator of a fake transaction, discounted by a parameter "γ." The authors present a randomized (based on trusted on-chain randomness) second-price auction that satisfies the three properties. Lastly, the authors in [41] relax UIC to Bayesian UIC (BUIC) to construct another second-price-based auction, satisfying BUIC, MIC, and OCA-proof.

Redistribution Mechanism (RM). Popular auction-based mechanisms like VCG and Groves [9, 16, 38] satisfy AE and UIC but do not satisfy SBB. Faltings [14] and Guo and Conitzer [20] achieve SBB by compromising on AE. Hartline and Roughgarden [22] propose a mechanism that maximizes the sum of the users' utility in expectation. de Clippel et al. [12] "destroy" some items to maximize the users' utilities, leading to approximate AE and SBB. Parkes et al. [32] propose an alternate approach by proposing an optimization problem, which is approximately AE, SBB.

Maskin et al. [26] first propose the idea of redistribution of the surplus as far as possible after preserving UIC and AE. Bailey [1], Cavallo [6], [28], and Guo and Conitzer [19] consider a setting of allocating k homogeneous objects among n competing users with unit demand. Guo and Conitzer [21] generalize their work in [19] to multi-unit demand to obtain worst-case optimal (WCO) RM.

In summary, the current TFM literature only focuses on the satisfiability of desirable incentive properties and *not* on reducing the user cost. Given that a decentralized cryptocurrency (e.g., Bitcoin or Ethereum) is a public resource, re-imaging a TFM as an RM will (i) continue to guarantee these properties but, crucially, (ii) minimize the cost paid by the user.

3 PRELIMINARIES

We now (i) formally introduce TFMs, (ii) relevant game-theoretic definitions, and (iii) summarize redistribution mechanisms.

Transaction Fee Mechanism (TFM) Model. We have a strategic but myopic[2] miner building a block B (with finite capacity) for the underlying blockchain. There are m transactions available to be confirmed in a *mempool* M. However, the block can hold only up to $n < m$ transactions. We assume that all transactions are of the same size. Among the n transactions included in the block, the miner confirms $k \leq n$ transactions[3].

Let each user i value the confirmation of its transaction at $\theta_i \in \mathbb{R}_{\geq 0}$. Each user i submits a bid $b_i \in \mathbb{R}_{\geq 0}$. We have $\Theta := [\theta_i]$ and $\mathbf{b} := [b_i]$. Given the bid profile and valuation profile, the transaction fee mechanism is characterized by the inclusion rule, confirmation rule, and payment rule, as defined below.

Definition 1 (Transaction Fee Mechanism (TFM) [8, 34].). *Given a bid profile* $\mathbf{b}$, *we define TFM as* $\mathcal{T} := (\mathbf{x}^I, \mathbf{x}^C, \mathbf{p})$ *where:*

- $\mathbf{x}^I$ *is a feasible block inclusion rule, i.e.,* $\sum_{i\in M} x_i^I(\mathbf{b}) \leq n$ *where* $x_i^I(\cdot) \in \{0,1\}$. *Let the set of included transactions be* $I = \{i | x_i^I = 1, i \in M\}$.
- $\mathbf{x}^C$ *is a feasible block confirmation rule, i.e.,* $\sum_{i\in M} x_i^C(\mathbf{b}) \leq k$ *where* $x_i^C(\cdot) \in \{0,1\}$. *Let the set of confirmed transaction be* $C = \{i : x_i^C = 1, i \in M\}$. *Trivially,* $C \subseteq I$.
- $\mathbf{p}$ *is the payment rule with the payment for each included transaction* i *be* p_i, *i.e.,* $\forall i \in \mathbf{x}^I$, *the payment is denoted by* $p_i(\mathbf{b}, \mathbf{x}^I, \mathbf{x}^C)$.

In TFMs, the included (but not confirmed) bids are often used as 'price-setting' bids [8]. We use the example of a second-price TFM to explain Definition 1 better.

Example 1 (Second-price TFM (SPA) [8, 34].). W.l.o.g., assume that $\mathbf{b} = (b_1, \ldots, b_m)$ are bids in decreasing order. Now, the inclusion rule is $x_i^I = 1$, $\forall i \in \{1, \ldots, n\}$ and zero otherwise, i.e., the top n transactions are included in the block. With $k = n - 1$, the confirmation rule is $x_i^C = 1$, $\forall i \in \{1, \ldots, k\}$ and zero otherwise. The top k (among n) transactions are confirmed, and the last included transaction is the price-setting transaction. Each confirmed user $i \in [k]$ pays $p_i = b_{k+1}$ to the miner and unconfirmed user ($i \in I \setminus C$) pays $p_i = 0$. The miner's net revenue is $k \cdot b_{k+1}$.

In order to define the desirable properties of a TFM, we first define user and miner utilities.

Utility Model. We first reiterate that we assume that the miners and transaction creators (or users) are myopic [8, 15, 34, 41]. For each user i, let its transaction i's valuation be $\theta_i \in \mathbb{R}_{\geq 0}$ with bid $b_i \in \mathbb{R}_{\geq 0}$. We have $\Theta := [\theta_i]$ and $\mathbf{b} := [b_i]$. Now, given $\mathcal{T} = (\mathbf{x}^I, \mathbf{x}^C, \mathbf{p})$, each user i's *quasi-linear* utility u_i is defined as:

$$u_i(\theta_i, \mathbf{b}) := \left(\mathbb{1}_{x_i^C=1} \cdot \theta_i\right) - p_i(\mathbf{b}, \mathbf{x}^I, \mathbf{x}^C) \tag{1}$$

We now define the miner's utility. Since the miner has complete control over the transactions, it adds to its block [34], it can add a set of "fake" transactions (say F) to deviate from the intended allocation rule $\mathbf{x} = (\mathbf{x}^C, \mathbf{x}^I)$. The miner's utility (say $u_{\mathbf{M}}$), given $\mathcal{T}$, and for the block B, is given by:

$$u_{\mathbf{M}}(F, \mathbf{b}) := \sum_{i \in B \cap M} p_i(\mathbf{b}, \mathbf{x}^I, \mathbf{x}^C) \tag{2}$$

That is, the miner's utility only depends on the set of transactions $B \cap M$ since for transactions in F, it is paying to itself.

[2]A miner is myopic if its utility is its net revenue from the current block [8, 15, 34].
[3]This is analogous to the 'homogeneous' (unit demand) setting in the RM literature [21].

3.1 TFMs: Desirable Properties

We now define the relevant incentive properties of a TFM (from [15, 34]). We begin by defining Individual Rationality.

Individual Rationality (IR). To incentivize participation, mechanism designers focus on IR.

Definition 2 ((Ex-post) Individual Rationality (IR)). *Given a TFM $\mathcal{T} = (\mathbf{x}^I, \mathbf{x}^C, \mathbf{p})$, we say that it satisfies IR for both the users and miners if their utility post participation in the mechanism is non-negative, i.e., $u_i(\cdot) \geq 0, \forall i \in M$ and $u_{\mathbf{M}}(\cdot) \geq 0$.*

Note. We denote a mechanism that is IR w.r.t. miner as $\text{IR}_{\mathbf{M}}$ and IR w.r.t. user as IR_u.

User Incentive Compatibility (UIC). To provide a good user experience, TFMs must satisfy UIC, i.e., they must incentivize users to report their true valuation as their bids (or transaction fees).

Definition 3 (User Incentive Compatibility (UIC) [34]). *Given a TFM $\mathcal{T} = (\mathbf{x}^I, \mathbf{x}^C, \mathbf{p})$, we say that each user i's strategy $b_i^\star = \theta_i$ satisfies UIC, if bidding $b_i^\star$ maximizes its utility u_i (Eq. 1), irrespective of the bids of others. More formally, $\forall i \in M$, we have*

$$u_i(\theta_i, b_i^\star = \theta_i, \mathbf{b}_{-i}) \geq u_i(\theta_i, b_i, \mathbf{b}_{-i}), \forall \theta_i, \forall \mathbf{b}_{-i},$$

where $\mathbf{b}_{-i}$ are the bids of all users excluding i.

As we show later, ensuring UIC while minimizing transaction fees in a TFM is challenging. Hence, we focus on the incentive compatibility of a 'restricted' set of users whose transactions are included in the block. We define restricted UIC (RUIC), which states that for all the included users, reporting truthfully is IC irrespective of what the remaining included users report.

Definition 4 (Restricted UIC (RUIC)). *Given a TFM $\mathcal{T} = (\mathbf{x}^I, \mathbf{x}^C, \mathbf{p})$, we say that RUIC is satisfied if, $\forall i$ included in the block i.e., $\forall i \in I$, we have,*

$$u_i(\theta_i, b_i^\star = \theta_i, \mathbf{b}_{I\setminus i}) \geq u_i(\theta_i, b_i, \mathbf{b}_{I\setminus i}), \forall \theta_i, \mathbf{b}_{I\setminus i}$$

where $\mathbf{b}_{I\setminus i}$ is the bids of users included in the block excluding user i.

Other Properties. Outside of these common TFM properties, we also define additional properties, namely allocative efficiency (AE) and weakly/strongly budget balance (WBB/SBB).

Definition 5 (Allocative Efficiency (AE)). *We say that a TFM $\mathcal{T} = (\mathbf{x}^I, \mathbf{x}^C, \mathbf{p})$ satisfies AE if given Θ, the mechanism confirms the transactions which maximize the overall valuation. That is, for any given Θ and every feasible allocation $\mathbf{x}^C$, we have: $\mathbf{x}^\star := \arg\max_{\mathbf{x}^C} \sum_{i \in M} x_i^C \cdot \theta_i$*

Definition 6 (Weakly/Strongly Budget Balance (WBB/SBB)). *We say that a TFM $\mathcal{T} = (\mathbf{x}^I, \mathbf{x}^C, \mathbf{p})$ satisfies WBB if the total payment to the miner is non-negative, i.e., $\sum_{i \in I} p_i \geq 0$. When the equality holds, a TFM is strongly budget balanced (SBB), i.e., $\sum_{i \in I} p_i = 0$.*

3.2 Groves' Redistribution Mechanism (RM)

Towards minimizing the user cost in a TFM, we employ *Redistribution Mechanisms* (RMs) [26]. In RM, the users are charged VCG payments, and the money is redistributed back to users while ensuring UIC. The redistribution is decided by constructing an appropriate rebate function, $g : \mathbf{b} \to \mathbb{R}$. We desire rebate functions that ensure maximum rebate (or, equivalently, minimize the transaction fees in TFM). To utilize an RM as a TFM, we require that it satisfies UIC/RUIC and (ex-post) IR for the users and the miner. An RM is IR for users when each user's overall payment, including the rebate, provides a non-negative utility. Likewise, we say that an RM is IR for the miner when the total rebate is less than the payment (transaction fees) received. We also want the RM to be *anonymous* [21].

Definition 7 (Anonymity). *An RM satisfies anonymity if the rebate function is the same for all the users, i.e., $\forall i, j \in [n]$ and $i \neq j$, $g_i(\cdot) = g_j(\cdot) = g(\cdot)$.*

Note that an anonymous rebate function may still result in different redistribution payments to different users as the input to the function may be arbitrarily different.

Rebate Function. We aim to design an appropriate rebate function for an anonymous RM such that incentive properties from Section 3.1 hold. The rebate function must also redistribute most of the payments (VCG payments) as much as possible to minimize the user cost. We begin by providing the following characterization for designing UIC rebate functions.

Theorem 1 ([17]). *In an RM, any deterministic, anonymous rebate function $g(\cdot)$ is UIC iff the rebate for user i is defined as $r_i := g(b_1, b_2, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n)$, $\forall i \in [n]$, where $b_1 \geq b_2 \geq \ldots \geq b_n$.*

The rebate function is UIC if the rebate for a user i is independent of its own bid. In general, it could take any form. E.g., the linear rebate function is defined as,

Definition 8 (Linear Rebate Function [19]). *The rebates to a user i follow a linear rebate function if the rebate is a linear combination of the bid vectors of all the remaining users. That is, $r_i = c_0 + c_1 b_1 + \ldots + c_{i-1} b_{i-1} + c_i b_{i+1} + \ldots + c_{n-1} b_{n-1}$ where $c_j \in \mathbb{R}$, $\forall j$.*

Given a rebate function, the fraction of VCG payment redistributed depends on the input bids. Thus, we study the worst-case and average-case performance of the rebate functions.

Definition 9 (Worst-case/Average Redistribution Index (RI) [19]). *The Worst-case or Average RI of an RM is defined as the worst-case or average-case fraction of VCG surplus that gets redistributed among the users, respectively. That is, given that $p(\mathbf{b}) = \sum_i p_i(\mathbf{b})$, the total VCG payment collected:*

$$e_{wc} = \inf_{\mathbf{b}: p(\mathbf{b}) \neq 0} \frac{\sum_i r_i}{p(\mathbf{b})} \text{ and } e_{avg} = \mathbb{E}_{\mathbf{b}: p(\mathbf{b}) \neq 0} \left[\frac{\sum_i r_i}{p(\mathbf{b})} \right] \quad (3)$$

Guo and Conitzer [21] propose *Worst-case Optimal (WCO)*, an RM that uniquely maximizes the worst-case RI (among all RMs that are deterministic, anonymous, and satisfy UIC, AE, and IR) when the items are homogeneous with unit demand [21, Theorem 1].

TFM <> RM. We assume that each block has k slots available for confirmation, and any user equally values transaction confirmation at every slot. Each transaction only requires one slot for confirmation. Thus, this is a homogeneous setting with unit demand.

4 IDEAL-TFRM: IMPOSSIBILITY OF ACHIEVING STRICTLY POSITIVE RI

We now present a first attempt at implementing an RM for minimizing transaction fees using the second-price TFM (refer to Example 1).

In a second-price TFM, the transactions/bids are sorted in decreasing order. The top k bids are confirmed with the $n = (k+1)^{th}$ transaction as the price-setting one. Each confirmed transaction pays $p = b_{k+1}$; with the net miner revenue as $k \cdot b_{k+1}$. To minimize the user fees, we must "redistribute" the collected surplus.

It may be preferable to only provide rebates to users whose transactions are confirmed (i.e., are among the top k bids) since each such user pays b_{k+1}. If we also provide rebates to the remaining $n-k$ users, the remaining transactions in the mempool may prefer to overbid just enough also to get included in the block. By doing so, they can grab rebates for free[4]. Thus, to achieve UIC, we must not provide rebates to unconfirmed transactions. With this motivation, we propose the following.

Ideal-TFRM. The goal is to maximize the fraction of VCG payments redistributed to the users, denoted by f, while ensuring non-zero rebates only to confirmed transactions. Further, in `Ideal-TFRM`, we would like the rebate offered to each user to be less than the payment it makes. Eq. 4 captures this optimization.

$$\left.\begin{aligned} \max_{r_i, i\in \mathbf{x}} f \ \textbf{s.t.} \ \sum_{i\in C} r_i \geq f \cdot \sum_{i \in C} p_i \\ \textbf{and } p_i \geq r_i, \ \forall i \in C \textbf{ and } r_i = 0, \ \forall i \in I \setminus C \end{aligned}\right\} \quad (4)$$

Here, $r_i = g(b_1, b_2, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n)$ is the rebate (Definition 8) for each user $i \in M$ with p_i as the VCG payment. The goal is to find an optimal $g(\cdot)$ such that e_{wc} is maximized.

Unfortunately, we now show that for both the worst and average-case, `Ideal-TFRM` admits zero rebates for the users with confirmed transactions, i.e., $r_i = 0, \ \forall i \in \mathbf{x}^C$ while guaranteeing UIC.

Worst-case Rebate. Theorem 2 formally[5] shows the impossibility of simultaneously guaranteeing UIC and minimizing transaction fees in `Ideal-TFRM`.

THEOREM 2 (`Ideal-TFRM` IMPOSSIBILITY). *If $r^\star$ is an anonymous rebate function that satisfies Theorem 1, no* `Ideal-TFRM` *can guarantee a non-zero redistribution index (RI) in the worst case, i.e., $e_{wc} = 0$.*

PROOF SKETCH. Informally, if the rebate function is anonymous (Definition 7) and the user with the highest valued transaction is confirmed and receives a positive rebate. Then, it is easy to show that there exists a different bid profile for which the last unconfirmed transaction must receive the same positive rebate. Therefore, the only way to ensure zero rebates for unconfirmed transactions is also to ensure zero rebates for every confirmed transaction. Thus, the worst-case RI is zero. □

Average-case Rebate. Theorem 2 shows that $e_{\text{wc}} = 0$ in `Ideal-TFRM`, for a linear rebate function. We now aim to find a *non-linear* rebate function that maximizes e_{avg} in `Ideal-TFRM`. However, it is analytically intractable to characterize similar results to show the outcome of a rebate function that maximizes e_{avg}. As such, we simulate the optimization in Eq. 4 as a Neural Network (NN), similar to [13, 25, 37].

Architecture & Setup. We consider a typical 3-layer feed-forward NN with bias, ReLU activation, and with AdamW optimizer. The input to our NN is the n-dimensional bid vector $\mathbf{b}_I$ sampled from a specific distribution. Each hidden layer comprises $2n$ neurons, with n as the output layer's dimension. Given $\mathbf{b}_I$, the NN computes the payments and rebates to the confirmed and included transactions.

Loss Function. For optimization, our loss function is a weighted sum of the following three quantities: (i) average rebate to the n bidders (denote as r_{avg}), (ii) feasibility, i.e., $\sum_i r_i \leq \sum_i p_i$ (denote as r_{feas}) and (iii) zero-rebate, i.e., $r_i = 0, \forall i \in I \setminus C$ (denote as r_{zero}). More concretely, for weights $\beta_1, \beta_2 \in (0,1)$ the loss function takes the form: Loss $= r_{\text{avg}} + \beta_1 \cdot r_{\text{feas}} + \beta_2 \cdot r_{\text{zero}}$.

Training Details. We keep $n = 10$ and $k = 7$. For the optimizer, we choose a fixed learning rate $\eta = 5e-4$. The batch size is 1000, and we train for 50,000 epochs.

Results. We observe that $e_{\text{avg}} \approx 0$ when transactions are sampled from $U[0,1]$ and $\mathcal{N}(0,1)$. That is, the average case rebate to confirmed transactions is zero, even with non-linear rebate functions.

We conclude that it is impossible to design a TFRM with a linear rebate function that is UIC (Theorem 1) and offers a non-zero rebate to any user. Our experiments also highlight that this may also be unlikely for non-linear rebate functions. Therefore, the next section introduces the general TFRM framework, where we focus on restricted UIC.

(1) *Inclusion Rule* $(\mathbf{x}^I)$. Select highest n transactions from the mempool, M. W.l.o.g., assume that these n transactions are ordered as $b_1 \geq b_2 \geq \ldots \geq b_n \implies x_i^I = 1, \ \forall i \in \{1, \ldots, n\}$.
(2) *Confirmation Rule* $(\mathbf{x}^C)$. Select highest k bids from the n included, $x_i^C = 1, \forall i \in [k]$, where $k \leq n-2$.
(3) *Payment Rule* $(\mathbf{p})$. Each confirmed user i (i.e., $i \in C$) pays $p_i = b_{k+1} - r_i$. Each included but not confirmed user j ($j \in I \setminus C$) pays $p_j = -r_j$.
(4) *Miner Revenue Rule.* The miner receives the net revenue of $\sum_{i=1}^n p_i$.

Figure 2: TFRM: General Framework

5 TRANSACTION FEE REDISTRIBUTION MECHANISM (TFRM)

As both $e_{\text{avg}} \approx e_{\text{wc}} = 0$ for `Ideal-TFRM`, we must also provide rebates to users whose transactions are included but not confirmed. With this, we present the general TFRM framework in Figure 2.

In a TFRM, out of the m outstanding transactions in the mempool, we include the n highest bids in the block (denoted by the set I). Among the bids in the block, we confirm the k highest bids (denoted by the set C) where $n \geq k+2$. The remaining bids (denoted by the set P) are included but not confirmed; we refer to them as price-setting transactions. That is, $|I| = n, |C| = k$ and $|P| = n-k$. W.l.o.g., we assume that $b_1 \geq b_2 \geq \ldots \geq b_n$. Hence, $C = \{b_1, \ldots b_k\}$ and $P = \{b_{k+1}, \ldots, b_n\}$. Each user i's payment is computed based on the VCG payments and a rebate function, with r_i as the rebate to user i. Since k bids are confirmed, the VCG payment for the confirmed transactions is the $(k+1)^{th}$ highest bid, i.e., b_{k+1}.

TFRM: RUIC. While the rebate function satisfies Theorem 1, note that TFRM is not UIC. E.g., users not part of the block may report $b > \theta$ to grab the additional rebate, as only confirmed bids pay the

[4]Assigning future costs to transactions not confirmed, e.g., as in [8] (Section 2), may help overcome such manipulation. We leave the analysis for future work.
[5]We refer the reader to [10] for the formal proofs of the results presented in this work.

transaction fee. However, TFRM satisfies RUIC, i.e., it is UIC for users included in the block to bid their true valuation. All included users, confirmed or not, are offered rebates, implying that k slots offered to n users are equivalent to the allocation of k resources among n users. Thus, by Theorem 1, TFRM is RUIC.

We now show that in the presence of a strategic miner, the TFRM in Figure 2 results in net zero rebates to confirmed users.

TFRM: Effect of Strategic Miners. In general, strategic miners may introduce fake transactions to increase their revenue. In the second-price TFM itself (Example 1), the miner may introduce a fake bid $\hat{b}_{k+1} = b_k$ to increase its revenue to $k \cdot b_k$ from the intended $k \cdot b_{k+1}$. Similarly, fake transactions can also affect the rebate offered by a TFRM. The miner's deviation may result in the following: (i) Fake bids affect the rebate of all users, potentially reducing the rebate, and (ii) As a fake bidder, the miner pays the rebate to itself.

Thus, designing TFRMs to minimize transaction fees may only work if made resilient to such strategic manipulations. Unlike RMs, TFRMs must quantify the rebate redistributed to the genuine users. Towards this, we define the following metric:

Definition 10 (Resilient Redistribution Index (RRI)). *Given that the miner manipulates the bids* $\mathbf{b}$ *to* $\hat{\mathbf{b}}$*, RRI is the fraction of the received payments that are redistributed in the worst case to the actual users. Given C confirmed and P unconfirmed users, let* $S \subseteq I$ *be the subset of users that are not impersonated by the miner. Then:* $\hat{e}_{wc} = \inf_{\hat{\mathbf{b}}, p(\cdot;\hat{\mathbf{b}}) \neq 0} \frac{\sum_{i \in S} r_i}{p(\cdot;\hat{\mathbf{b}})}$.

TFRM: Impossibility of Strictly Positive RRI. It is desirable to have a TFRM that is AE, RUIC, $\mathrm{IR}_{\mathbf{M}}$ and IR_u while ensuring RRI > 0 in the worst-case, i.e., $\hat{e}_{\mathrm{wc}} > 0$. Unfortunately, Theorem 3 proves that it is impossible to design such a TFRM with strategic miners.

Theorem 3 (TFRM: RRI Impossibility). *Given a strategic miner, it is impossible to design a TFRM with a linear rebate function that is RUIC, AE, both* IR_u *and* $\mathrm{IR}_{\mathbf{M}}$*, and guarantees a strictly positive RRI, i.e.,* $\hat{e}_{wc} > 0$.

Proof Sketch. A linear rebate that is RUIC and IR_u and $\mathrm{IR}_{\mathbf{M}}$ must depend only on bids $b_{k+2}, \ldots, b_n$. Changing these bids does not affect the payment received by the miner. Hence, the miner can replace these bids with fake bids such that the rebate is zero without any change in the payments. □

These results establish that preventing user manipulation entirely is not possible in the TFRM framework. Therefore, we focus on ensuring Restricted UIC (RUIC), which ensures that users of transactions included in the block will not misreport their values. Further, we know from Theorem 3 that even with RUIC, the miner can easily manipulate any known RM that satisfies all the desirable properties. The theorem also shows that the manipulation will lead to strictly zero rebate. Thus, in the next section, we propose Robust TFRM (R-TFRM), which relaxes $\mathrm{IR}_{\mathbf{M}}$, to ensure a positive rebate even with miner manipulation.

6 R-TFRM: A TFRM ROBUST TO MINER MANIPULATION

To ensure strictly positive RRI, we compromise on $\mathrm{IR}_{\mathbf{M}}$, i.e., the utility of an honest miner may be negative. However, we ensure

Maximize: f ▷RRI
Subject to: For every $\mathbf{b} \geq 0$,

$$r_n \geq 0 \quad \triangleright \mathrm{IR}_u$$

$$\sum_{i=1}^{n} r_i \leq k \cdot b_k \quad \triangleright \text{Approx-IR}_m$$

$$\sum_{i=1}^{k} r_i \geq f \cdot k \cdot b_{k+1} \quad \triangleright \text{Worst-case Fraction}$$

where $r_i = c_0 + c_1 b_1 + \ldots + c_{i-1} b_{i-1} + c_i b_{i+1} + \ldots + c_{n-1} b_n$

Figure 3: R-TFRM: Linear Program for Rebate Function with Approx-IR$_m$

Maximize: f
Subject To: $\sum_{j=k}^{i} c_j \geq f, \ \forall i \in \{k, \ldots, n-1\}$

$$(n-k) \cdot c_k \leq k \text{ and } n \sum_{j=k}^{n-1} c_j \leq k$$

$$n \sum_{j=k}^{k+i-1} c_j + (n-k-i) \cdot c_{k+i} \leq k, \ \forall i \in [n-k-1]$$

Figure 4: R-TFRM: Linear Program Independent of the Bid Vector b

that when the miner is strategic, it can always guarantee itself a non-zero utility. We denote such a TFRM that is resilient to miner manipulation by Robust TFRM (R-TFRM). Designing R-TFRM involves constructing an appropriate rebate function. We focus on a linear rebate function that maximizes the worst-case redistribution index RRI while ensuring IR_u. We still want RUIC; hence, we use the rebate function as given in Theorem 1.

IR_u Constraints. Each included user must have a non-negative utility, i.e., $u_i \geq 0, \forall i \in I$. W.l.o.g., we assume that $b_1 \geq b_2 \geq \ldots b_n$ and IR for user n is ensured when $r_n \geq 0$ as $u_n = r_n$. In Claim 1, we show that $r_i \geq 0, \forall i$ if $r_n \geq 0$; hence R-TFRM will be IR_u.

Claim 1. *R-TFRM with n included transactions and rebates* $(r_1, \ldots, r_n)$ *is user IR* (IR_u) *if* $r_n \geq 0$.

Approx-$\mathrm{IR}_{\mathbf{M}}$ Constraints. In any classic RM, to ensure $\mathrm{IR}_{\mathbf{M}}$, there is an additional constraint to ensure that the total rebate is less than the VCG payments, i.e., $\sum_i r_i \leq k \cdot b_{k+1}$. We modify this constraint to ensure $\sum_i r_i \leq k \cdot b_k$. This change is because a strategic miner can manipulate the VCG auction to insert a fake bid $\hat{b}_{k+1} = b_k$. Figure 3 describes the linear program to solve for such a rebate while maximizing the RRI fraction f.

We now aim to write the linear program in Figure 3 such that it is only dependent on $n, k, c_i's$ and independent of the bid vector, $\mathbf{b}$. For this purpose, we now state the following claims.

Claim 2. *If* $c_0, \ldots, c_{n-1}$ *satisfy* IR_u *and Approx-*$\mathrm{IR}_{\mathbf{M}}$*, then* $c_i = 0$ *for* $i = 0, \ldots, k-1$.

Claim 3. *The IR_u constraint $r_n \geq 0$ and the worst-case fraction constraint (refer to Figure 3) is equivalent to having $\sum_{j=k}^{i} c_j \geq f$, $\forall i \in \{k, \ldots, n-1\}$.*

Claim 4. *The Approx-$IR_{\mathbf{M}}$ constraint can be replaced by:*

$$(n-k)c_k \leq k \;\&\; n\sum_{j=k}^{n-1} c_j \leq k \textbf{ and}$$

$$n\sum_{j=k}^{k+i-1} c_j + (n-k-i)c_{k+i} \leq k,\; i \in \{1, \ldots, n-k-1\}$$

Using Claims 3 and 4, we reformulate the linear program in Figure 3 so that it is independent of the bid vectors. Figure 4 presents this reformulated LP.

Optimal worst-case Redistribution Fraction. We next provide the analytical solution to the linear program in Figure 4 and thereby also state the optimal worst-case fraction redistributed.

Theorem 4. *For any n and k such that $n \geq k+2$, the R-TFRM mechanism is unique. The fraction redistributed to the top-k users in the worst case is given by: $f^* = \frac{k}{n}$. In R-TFRM, the rebate function is characterized by the following: $c_k^* = \frac{k}{n}$ and $c_i^* = 0$, $\forall i \neq k$.*

Observe that the total redistribution to the users, when the miner is honest for R-TFRM, is given by

$$\sum_{i=1}^{n} r_i = \frac{k}{n}\left[(k \cdot b_{k+1}) + (n-k)b_k\right] \tag{5}$$

This value may exceed $k \cdot b_{k+1}$, thus violating $IR_{\mathbf{M}}$. However, it satisfies Approx-$IR_{\mathbf{M}}$ as $\sum_{i=1}^{n} r_i \leq k \cdot b_k$ (refer Figure 3). R-TFRM is similar to the Bailey-Cavallo mechanism [1]. The primary difference is due to the Approx-$IR_{\mathbf{M}}$ constraint, which makes $c_k = k/n$ instead of c_{k+1} as in Bailey-Cavallo.

6.1 R-TFRM: Analyzing Impact of Miner Manipulation on Rebate and Miner Revenue

With an honest miner, R-TFRM maximizes the worst-case redistribution index such that it is AE, RUIC, IR_u, and Approx-$IR_{\mathbf{M}}$. We now analyze the effect of miner manipulation on R-TFRM. Previously, we saw that it is impossible to ensure non-zero RRI (Theorem 3), but with R-TFRM we show that RRI is strictly positive even with miner manipulation.

Reduction in Transaction Fees. The rebate function for R-TFRM is characterized by the constants given in Theorem 4. With these, we now calculate RRI (Definition 10), i.e, $\hat{e}_{\text{wc}}$, for R-TFRM. Theorem 5 shows that irrespective of miner manipulation, $c_k^* = \frac{k}{n}$ fraction of payments will be returned.

Theorem 5. *Consider n included transactions with the set C as confirmed transactions such that $|C| = k$ with the remaining $n-k$ as price-setting transactions. Irrespective of any miner manipulation, R-TFRM ensures strictly positive RRI or $\hat{e}_{wc} = c_k^* = \frac{k}{n}$.*

From Theorem 4 and Theorem 5, we see that in R-TFRM, the fraction of payments redistributed to the top-k users, i.e., k/n, is the same for honest and strategic miner. This implies that R-TFRM is resilient to miner manipulation while being worst-case optimal.

Utility of Strategic Miner. From Theorem 5, if a miner is strategic and impersonates the price-setting transactions, the miner will receive positive utility. The miner will preferably set the fake bid $\hat{b}_{k+1}$ close to b_k. Hence, the maximum utility to a miner that deviates by impersonating the price-setting bids is: $u_{\mathbf{M}} = (1 - k/n) \cdot k \cdot b_k$. As we assume $n \geq k+2$, the miner's maximum utility is minimized for $k = n-2$.

The fraction redistributed to the genuine users is still $\frac{k}{n}$ of the payments received even when the miner impersonates the confirmed transactions. We illustrate this with an example next.

Example 2. It is possible for the miner to insert fake transactions with high enough bids such that they are confirmed. Consider $n = 5$ and $k = 3$ where $b_1 = b_2 = 100$, $b_3 = 10$ and $b_4 = b_5 = 4$. If the miner is only impersonating the price setting transactions, then it puts $\hat{b}_4 = b_3$ and arbitrary $\hat{b}_5 < b_3$, then its overall utility is $\left(1 - \frac{k}{n}\right) kb_k = 12$. Whereas if the miner is given more flexibility to insert a fake transaction within the confirmed and unconfirmed bids, it receives more payments. For e.g., let $\hat{b}_1 = 200$ and $\hat{b}_3 = 100$ hence the ordered bids are $\hat{b}_1 \geq b_1 \geq b_2 \geq \hat{b}_3$. Therefore, effectively, the first two transactions are confirmed, and they pay 100 each. Further, due to R-TFRM, it returns a rebate of $\frac{k}{n}100$ to each of the two users, thus obtaining an overall utility of $(1 - 3/5)\, 200 = 80$.

7 R^2-TFRM: ROBUST AND RATIONAL TFRM

R-TFRM compromises Miner IR to ensure positive RRI. We now introduce randomness in R-TFRM to obtain a mechanism that ensures positive utility to an honest miner, i.e., satisfies IR_m. Towards this, we propose Robust and Rational TFRM, R^2-TFRM (Figure 5). In R^2-TFRM, the rebate is not guaranteed for every included transaction. Instead, an included transaction gets a rebate with probability α, $\alpha \in [0, 1]$ where the rebate value is calculated using R-TFRM. Hence R^2-TFRM reduces to R-TFRM when $\alpha = 1$. On the other extreme, when $\alpha = 0$, R^2-TFRM reduces to a second price auction.

R^2-TFRM: On-chain Randomness. As stated, each transaction receives a rebate with a probability α. Similar to other TFMs [8], we employ trusted on-chain randomness for this randomization. Researchers have proposed such trusted randomized protocols using various cryptographic primitives [3, 11]. Significantly, the miner of the block cannot exert any influence on this randomization.

R^2-TFRM: Incentive and RRI Guarantees. Theorem 6 proves that R^2-TFRM mimics the incentive guarantees of R-TFRM. Moreover, for an appropriate α, R^2-TFRM is also $IR_{\mathbf{M}}$ in expectation.

Theorem 6. *For any n and k such that $n \geq k+2$ and any bid profile $\mathbf{b} = (b_1, \ldots, b_n)$, and probability $\alpha \in (0, 1)$ R^2-TFRM has an expected redistribution fraction (expectation over α) $f^* = \alpha \cdot \frac{k}{n}$. Further it satisfies AE, RUIC, IR_u, and is IR_m when $\alpha \leq \overline{\alpha} = \frac{n}{k+(n-k)b_k/b_{k+1}}$.*

Proof Sketch. We divide the proof into two parts.

- Firstly, notice that R^2-TFRM is AE since the top k bids are confirmed and IR_u, since payments are governed by VCG and all the rebates, are positive. Although in R^2-TFRM every user is not deterministically given rebates, so it is RUIC in expectation, where the expectation is over α. This is because a user's rebate is independent of its own bid.

(1) *Inclusion Rule* ($\mathbf{x}^I$). Select highest n transactions from the mempool, M. W.l.o.g., assume that these n transactions are ordered as $b_1 \geq b_2 \geq \ldots \geq b_n \implies x_i^I = 1,\ \forall i \in \{1, \ldots, n\}$.
(2) *Confirmation Rule* ($\mathbf{x}^C$). Select highest k bids from the n included, $x_i^C = 1, \forall i \in [k]$, where $k \leq n-2$.
(3) *Payment Rule* ($\mathbf{p}$). Each confirmed user i (i.e., $i \in C$) pays

$$p_i = \begin{cases} b_{k+1} - r_i & \text{w.p. } \alpha \\ b_{k+1} & \text{w.p. } (1-\alpha) \end{cases}$$

Each included but not confirmed user j (i.e., $j \in I \setminus C$) pays

$$p_j = \begin{cases} -r_j & \text{w.p. } \alpha \\ 0 & \text{w.p. } (1-\alpha) \end{cases}$$

The rebate r_i is given by Theorem 4.
(4) *Miner Revenue Rule.* The miner receives the net revenue of $\sum_{i=1}^{n} p_i$.

Figure 5: $\mathbf{R^2}$-TFRM: Robust and Rational TFRM

- Next, we now formally show that R^2-TFRM is $\mathrm{IR}_{\mathbf{M}}$ in expectation. The payment obtained by an honest miner is given by kb_{k+1}. The expected rebate paid by the honest miner is given by $\alpha\left[k\frac{k}{n}b_{k+1} + (n-k)\frac{k}{n}b_k\right]$. If $\alpha = 1$, the refund is equal to the refund in R-TFRM given by Eq. 5. In order to ensure IR_{M} for the honest miner, the following must be true,

$$kb_{k+1} \geq \alpha\left[k\frac{k}{n}b_{k+1} + (n-k)\frac{k}{n}b_k\right] \implies \alpha \quad \leq \frac{nb_{k+1}}{kb_{k+1} + (n-k)b_k}$$

Therefore, R^2-TFRM satisfies $\mathrm{IR}_{\mathbf{M}}$ when $\alpha = \frac{nb_{k+1}}{kb_{k+1}+(n-k)b_k}$ and the total rebate is $k \cdot b_{k+1}$.

This completes the proof of the theorem. □

$b_{k+1} \to b_k \implies$ **R-TFRM and $\mathbf{R^2}$-TFRM Become Equivalent.** From Theorem 6, an honest miner obtains non-negative utility when,

$$\alpha \leq \frac{n}{k + (n-k)b_f} \tag{6}$$

where, $b_f = \frac{b_k}{b_{k+1}}$ is the bid ratio. W.l.o.g as the bids are ordered (Figure 5) $b_f \geq 1$. When $b_f = 1$, i.e., $b_k = b_{k+1}$ we have $\alpha = 1$ and R-TFRM $\iff$ R^2-TFRM. This implies that every user receives a rebate, and the miner IR is not violated.

This may seem to contradict R-TFRM not being $\mathrm{IR}_{\mathbf{M}}$; however, we can think of $b_{k+1} \to b_k$ as one of the deviations of the strategic miner. And R-TFRM is $\mathrm{IR}_{\mathbf{M}}$ when the miner is strategic. Furthermore, as b_f increases, the upper bound on α becomes smaller. To still guarantee $\mathrm{IR}_{\mathbf{M}}$, the overall rebate (i.e., $\alpha \cdot \frac{k}{n}$) decreases.

Observe that, in the worst case, trivially, $\alpha = 0$. Generally, the bound on α will depend on the bid distribution. Sampling the bids uniformly, i.e., $\mathbf{b} \sim U[0,1]$, the expected value of α (using order statistics) is, $\mathbb{E}_{\mathbf{b}}[\alpha] = n/(n+1)$ (i.e., $\alpha \to 1$ for large values of n).

$\mathbf{R^2}$-TFRM: Analyzing Miner Manipulation. Like R-TFRM, R^2-TFRM also ensures strictly positive RRI even with miner manipulation as formally stated in Theorem 7.

Theorem 7. *Consider n included transactions with the set C as confirmed transactions such that $|C| = k$ with the remaining $n-k$ as price-setting transactions. Irrespective of any miner manipulation R^2-TFRM ensures: expected RRI or $\mathrm{E}_\alpha[\hat{e}_{wc}] = \alpha \cdot \frac{k}{n}$.*

Proof Sketch. The result follows from the similar result in Theorem 5 for R-TFRM and the fact that the miner has no control over the randomization in R^2-TFRM. □

We observe that irrespective of the miner manipulation, the users receive back the fraction $\alpha \cdot \frac{k}{n}$ of the payment made on expectation. Based on what fake transactions the miner inserts, the payments change, but the refund fraction remains the same as for the case when the miner is honest.

Discussion on TFRMs

Given that the current TFM literature assumes that users and miners are myopic, we believe that redistributing the surplus is an effective method to reduce the net payments paid by the users. Another desirable property in TFMs is that of *predictable* transaction fees, i.e., reducing the volatility of the fees paid by the users. For instance, EIP-1559 [5] uses a deterministic function based on the previous block consumption to calculate a network-determined minimum threshold fee (called *base* fee) paid by each user. The users can also pay a *priority* fee over the base fee to incentivize the miners to include their transactions. The base fee aims to reduce the fee volatility and aims to arrive at a market clearing price. Crucially, this minimum threshold fee is *burned* – transferred to an unspendable address – implying that the miner does not receive this fee as revenue. We see that burning some fraction of the fee is necessary to guarantee such properties. In such mechanisms, the priority fee calculated post-burning can be further reduced by employing TFRMs.

8 CONCLUSION

In this paper, we argued the importance of minimizing user costs in a TFM. Our key idea is to employ a redistribution mechanism-based approach for determining the transaction fees in TFM, namely, TFRM. Due to strategic miner manipulation, we first show that guaranteeing a strictly positive rebate in a TFRM and other desirable properties is impossible. Hence, we propose R-TFRM, which ensures strictly positive rebates even in the worst case but compromises on miner's IR. However, we show that in R-TFRM, a strategic miner will never incur negative utility while still guaranteeing strictly positive rebates to the users. We also propose R^2-TFRM which uses blockchain's inherent randomness to guarantee a strictly positive rebate to the users while also respecting the miner's IR.

Future Work. Future directions can explore TFRMs with randomized rebate functions, which may likely satisfy stronger notions of IC and IR. Another approach may be to explore non-linear rebate functions, which may provide a better redistribution index on average. In addition, unlike this work, future work can also explore transactions with varying sizes.

ACKNOWLEDGMENTS

This work is partially funded by MeitY and the Ripple-IIITH CoE. The authors also thank the anonymous reviewers for their thoughtful feedback.

REFERENCES

[1] Martin J Bailey. 1997. The Demand Revealing Process: To Distribute the Surplus. *Public Choice* 91, 2 (1997), 107–26. https://EconPapers.repec.org/RePEc:kap:pubcho:v:91:y:1997:i:2:p:107-26

[2] Soumya Basu, David Easley, Maureen O'Hara, and Emin Sirer. 2019. StableFees: A Predictable Fee Market for Cryptocurrencie. *Available at SSRN 3318327* (2019).

[3] Adithya Bhat, Nibesh Shrestha, Zhongtang Luo, Aniket Kate, and Kartik Nayak. 2021. Randpiper–reconfiguration-friendly random beacons with quadratic communication. In *ACM CCS*. 3502–3524.

[4] Vitalik Buterin et al. 2014. A next-generation smart contract and decentralized application platform. *White Paper* 3, 37 (2014), 2–1.

[5] Vitalik Buterin, Eric Conner, Rick Dudley, Matthew Slipper, Ian Norden, and Abdelhamid Bakhta. 2019. EIP-1559: Fee market change for ETH 1.0 chain. eips.ethereum.org/EIPS/eip-1559.

[6] Ruggiero Cavallo. 2006. Optimal Decision-Making With Minimal Waste: Strategyproof Redistribution of VCG Payments. In *Proc. of the 5th Int. Joint Conf. on Autonomous Agents and Multi Agent Systems (AAMAS'06)*. Hakodate, Japan, 882–889. http://econcs.seas.harvard.edu/files/econcs/files/cavallo-redis.pdf

[7] Ruggiero Cavallo. 2008. Efficiency and Redistribution in Dynamic Mechanism Design. In *Proc. 9th ACM Conf. on Electronic Commerce (EC'08)*. Chicago, IL, 220–229. http://econcs.seas.harvard.edu/files/econcs/files/cavallo-ec08.pdf

[8] Hao Chung and Elaine Shi. 2023. Foundations of transaction fee mechanism design. In *ACM-SIAM SODA*. 3856–3899.

[9] Edward Clarke. 1971. Multipart pricing of public goods. *Public Choice* 11, 1 (1971), 17–33.

[10] Sankarshan Damle, Manisha Padala, and Sujit Gujar. 2024. Designing Redistribution Mechanisms for Reducing Transaction Fees in Blockchains. arXiv:2401.13262 [cs.GT]

[11] Sourav Das, Vinith Krishnan, Irene Miriam Isaac, and Ling Ren. 2022. Spurt: Scalable distributed randomness beacon with transparent setup. In *IEEE S&P*.

[12] Geoffroy de Clippel, Victor Naroditskiy, Maria Polukarov, Amy Greenwald, and Nicholas R. Jennings. 2014. Destroy to save. *Games and Economic Behavior* 86, C (2014), 392–404.

[13] Paul Dütting, Zhe Feng, Harikrishna Narasimhan, David Parkes, and Sai Srivatsa Ravindranath. 2019. Optimal auctions through deep learning. In *ICML*. PMLR, 1706–1715.

[14] Boi Faltings. 2005. A Budget-balanced, Incentive-compatible Scheme for Social Choice. In *Agent-Mediated Electronic Commerce VI. Theories for and Engineering of Distributed Mechanisms and Systems: AAMAS 2004 Workshop, AMEC 2004*. 30–43.

[15] Matheus V. X. Ferreira, Daniel J. Moroz, David C. Parkes, and Mitchell Stern. 2021. Dynamic posted-price mechanisms for the blockchain transaction-fee market. In *ACM Conference on Advances in Financial Technologies (AFT)*. 86–99.

[16] Theodore Groves. 1973. Incentives in Teams. *Econometrica* 41, 4 (1973), 617–31.

[17] Sujit Gujar and Y. Narahari. 2011. Redistribution Mechanisms for Assignment of Heterogeneous Objects. *J. Artif. Int. Res.* 41, 2 (May 2011), 24.

[18] Mingyu Guo. 2012. Worst-case Optimal Redistribution of VCG Payments in Heterogeneous-item Auctions with Unit Demand. In *AAMAS*. 745–752.

[19] Mingyu Guo and Vincent Conitzer. 2007. Worst-case Optimal Redistribution of VCG Payments. In *ACM EC*. 30–39.

[20] Mingyu Guo and Vincent Conitzer. 2008. Better Redistribution with Inefficient Allocation in Multi-unit Auctions with Unit Demand. In *ACM EC*. 210–219.

[21] Mingyu Guo and Vincent Conitzer. 2009. Worst-case optimal redistribution of VCG payments in multi-unit auctions. *Games and Economic Behavior* 67, 1 (2009), 69–98.

[22] Jason D. Hartline and Tim Roughgarden. 2008. Optimal Mechanism Design and Money Burning. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing* (Victoria, British Columbia, Canada) *(STOC '08)*. ACM, New York, NY, USA, 75–84. https://doi.org/10.1145/1374376.1374390

[23] Mahimna Kelkar, Fan Zhang, Steven Goldfeder, and Ari Juels. 2020. Order-fairness for byzantine consensus. In *Annual International Cryptology Conference (CRYPTO)*. 451–480.

[24] Klaus Kursawe. 2020. Wendy, the good little fairness widget: Achieving order fairness for blockchains. In *ACM Conference on Advances in Financial Technologies (AFT)*. 25–36.

[25] Padala Manisha, CV Jawahar, and Sujit Gujar. 2018. Learning Optimal Redistribution Mechanisms Through Neural Networks. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. 345–353.

[26] Eric Maskin, J. J. Laffont, and J. J. Laffont. 1979. *A Differential Approach to Expected Utility Maximizing Mechanisms*. North Holland, 289-308.

[27] Johnnatan Messias, Mohamed Alzayat, Balakrishnan Chandrasekaran, and Krishna P Gummadi. 2020. On Blockchain Commit Times: An analysis of how miners choose Bitcoin transactions. In *The Second International Workshop on Smart Data for Blockchain and Distributed Ledger (SDBD2020)*.

[28] Herve Moulin. 2009. Almost budget-balanced VCG mechanisms to assign multiple objects. *Journal of Economic Theory* 144, 1 (2009), 96–119.

[29] Roger B Myerson and Mark A Satterthwaite. 1983. Efficient mechanisms for bilateral trading. *Journal of economic theory* 29, 2 (1983), 265–281.

[30] Satoshi Nakamoto. 2008. Bitcoin: A peer-to-peer electronic cash system. *Decentralized Business Review* (2008), 21260.

[31] Ariel Orda and Ori Rottenstreich. 2021. Enforcing fairness in blockchain transaction ordering. *Peer-to-peer Networking and Applications* 14, 6 (2021), 3660–3673.

[32] David C. Parkes, Jayant R. Kalagnanam, and Marta Eso. 2001. Achieving Budget-Balance with Vickrey-Based Payment Schemes in Exchanges. In *Proc. 17th International Joint Conference on Artificial Intelligence (IJCAI'01)*. 1161–1168. http://econcs.seas.harvard.edu/files/econcs/files/combexch01.pdf

[33] Ryan Porter and Yoav Shoham. 2003. On cheating in sealed-bid auctions. In *Proceedings of the 4th ACM conference on Electronic commerce*. 76–84.

[34] Tim Roughgarden. 2021. Transaction Fee Mechanism Design. In *ACM Conference on Economics and Computation (ACM EC)*. 792.

[35] Shoeb Siddiqui, Ganesh Vanahalli, and Sujit Gujar. 2020. BitcoinF: Achieving Fairness For Bitcoin In Transaction Fee Only Model. In *AAMAS*. 2008–2010.

[36] Corwin Smith. 2023. PROOF-OF-STAKE REWARDS AND PENALTIES. ethereum.org/en/developers/docs/consensus-mechanisms/pos/rewards-and-penalties/. [Online].

[37] Andrea Tacchetti, DJ Strouse, Marta Garnelo, Thore Graepel, and Yoram Bachrach. 2019. A neural architecture for designing truthful and efficient auctions. *arXiv preprint arXiv:1907.05181* (2019).

[38] W. Vickrey. 1961. Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance* 16, 1 (March 1961), 8–37.

[39] YCHARTS. 2022. Bitcoin Transactions Per Day. ycharts.com/indicators/bitcoin_transactions_per_day.

[40] YCHARTS. 2022. Ethereum Transactions Per Day. ycharts.com/indicators/ethereum_transactions_per_day.

[41] Zishuo Zhao, Xi Chen, and Yuan Zhou. 2022. Bayesian-Nash-Incentive-Compatible Mechanism for Blockchain Transaction Fee Allocation. In *Crypto Economics Security Conference (CESC)*.