Characterizing Off-Chain Influence Proof Transaction Fee Mechanisms

Aadityan Ganesh
Princeton University, NJ, USA

Clayton Thomas
Yale University, New Haven, CT, USA

S. Matthew Weinberg
Princeton University, NJ, USA

Abstract

Roughgarden [31] initiates the study of Transaction Fee Mechanisms (TFMs), and posits that the on-chain game of a "good" TFM should be on-chain simple (OnC-S), i.e., incentive compatible for both the users and the miner. Recent work of Ganesh, Thomas an Weinberg [21] posit that they should additionally be Off-Chain Influence-Proof (OffC-IP), which means that the miner cannot achieve any additional revenue by separately conducting an off-chain auction to determine on-chain inclusion. They observe that a cryptographic second-price auction satisfies both properties, but leave open the question of whether other mechanisms (such as those not dependent on cryptography) satisfy these properties.

In this paper, we characterize OffC-IP TFMs: They are those satisfying a *burn identity* relating the burn rule to the allocation rule. In particular, we show that auction is OffC-IP if and only if its (induced direct-revelation) allocation rule $\overline{X}(\cdot)$ and burn rule $\overline{B}(\cdot)$ (both of which take as input users' values $v_1, \ldots, v_n$) are truthful when viewing $\left(\overline{X}(\cdot), \overline{B}(\cdot)\right)$ as the allocation and pricing rule of a multi-item auction for a single additive buyer with values $\left(\varphi(v_1), \ldots, \varphi(v_n)\right)$ equal to the users' virtual values.

Building on this burn identity, we characterize OffC-IP and OnC-S TFMs that are deterministic and do not use cryptography: They are posted-price mechanisms with specially-tuned burns. As a corollary, we show that such TFMs can only exist with infinite supply and prior-dependence. However, we show that for *randomized* TFMs, there are additional OnC-S and OffC-IP auctions that do not use cryptography (even when there is finite supply, under prior-dependence with a bounded prior distribution). Holistically, our results show that while OffC-IP is a fairly stringent requirement, families of OffC-IP mechanisms can be found for a variety of settings.

2012 ACM Subject Classification Theory of computation → Algorithmic mechanism design; Applied computing → Online auctions

Keywords and phrases Transaction Fee Mechanism Design, Off-Chain Influence Proofness, Blockchain, Decentralized Finance, Simple Auctions

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.65

Related Version *Full Version*: `https://aadityanganesh.in/characterizing-off-chain-influence-proof-transaction-fee-mechanisms/`

Funding *Aadityan Ganesh*: Supported by an Ethereum Foundation Academic Grant.
S. Matthew Weinberg: Supported by an Ethereum Foundation Academic Grant and NSF CAREER Award CCF-1942497.

1 Introduction

In increasingly-important and high-demand blockchain applications, miners use Transaction Fee Mechanisms (TFMs) to allocate block space based on user-submitted bids. TFMs are subject to a host of novel auction design constraints in this distributed trustless setting.

Recent work of Ganesh, Thomas and Weinberg [21] highlights in particular the need for TFMs to be resistant to miners conducting an off-chain auctions in order to determine behavior on-chain, and call such a TFM off-chain influence-proof (OffC-IP). They show that EIP-1559 (the existing state-of-the-art TFM) is not OffC-IP, but that a cryptographic variant of a second-price auction satisfies OffC-IP along with other desirable simplicity properties. However, they leave open the question of precisely which auctions satisfy OffC-IP; for instance, whether there are any practical OffC-IP TFMs that do not use cryptography.

In this paper, we characterize OffC-IP TFMs, and delineate precisely how this property is compatible with other desirable properties considered by prior work. Our findings specify how the burn rule of the TFM is uniquely determined by its allocation rule, and highlight posted-price mechanisms as the *unique* deterministic TFMs satisfying our desirable desiderata while avoiding cryptography.

TFM Design Thus-Far

Originally, blockchains such as bitcoin used a first-price auction as their TFM, requiring strategic bidding from the users. Interest grew in simplifying the TFM for users, leading to a new proposed TFM for the Ethereum blockchain known as EIP-1559. Motivated by these concerns, Roughgarden [31, 32] studies the theory of TFM design, and proposes three desiderata for a desirable TFM. First, the on-chain TFM should be "simple for users" (OnC-US). That is, *assuming the miner honestly implements the proposed TFM*, it should be a dominant strategy for users to simply bid their value. Second, the on-chain TFM should be "simple for miners" (OnC-MS). That is, *assuming users simply bid their value*, the miner should have no incentive to deviate from the proposed TFM (e.g., censor users or include shill bids), even after seeing all users' bids. Third, the TFM should be "robust to collusion between the miner and users" (SCP). That is, the miner and users together cannot jointly profit by submitting anything other than their true bids.[1] Roughgarden [32] shows that EIP-1559 satisfies all three desirable properties.

EIP-1559 is essentially a "posted-price and burn" mechanism. The protocol exogenously sets a "base fee" p. Any user who wishes to include their transaction can pay p to do so, and their payment is *burned* (that is, the miner receives no revenue).[2] Prior analysis suggests EIP-1559 is a very strategically-simple mechanism. For instance, since there is no amount of censoring or fabricating bids that can leave the miner with a positive revenue, EIP-1559 is "simple for miners".

Despite EIP-1559's theoretical appeal and practical success, recent work of Ganesh, Thomas and Weinberg [21] observes the following challenge. Suppose the miner publicly threatens to censor all users who do not pay her an entry fee of \$5 off-chain.[3] So long as *any* user capitulates and pays this extra \$5 to be included, the miner increases her revenue by making this threat. They therefore propose that a TFM should be "off-chain influence-proof" (OffC-IP). That is, even if the miner were willing to act like a Bayesian monopolist off-chain, she should optimizes her expected revenue by simply following the prescribed TFM on-chain with no off-chain behavior. They formally observe that EIP-1559 is not OffC-IP.

---

[1] Multiple variants of "robust to collusions" have been proposed. See [9, 19, 21] for discussions.

[2] This describes EIP-1559 in the "ideal" case where the base fee is high enough to reduce demand below the maximum block size. In case the base fee is too low, EIP-1559 devolves into a first-price auction.

[3] In fact, in the real deployment of EIP-1559 in Ethereum, users are able to submit an optional on-chain tip to the miner, allowing miners to conduct this attack entirely on-chain.

Furthermore, Ganesh, Thomas and Weinberg [21] establish that no TFM can satisfy all four desirable properties (OnC-US, OnC-MS, OffC-IP, and SCP), even with cryptography. On the other hand, a cryptographic second-price auction[4] is OnC-US, OnC-MS, and OffC-IP; they argue that this strongly suggests that users can simply bid their values in this TFM, without worry. However, the full possibilities and limitations imposed by the OffC-IP constraint remain highly unclear. This is exactly the gap our paper fills.

Main Contribution I: A Burn Identity for OffC-IP TFMs. Our first main result characterizes OffC-IP TFMs by establishing a reduction to a canonical auction problem with a monopsonist and many items. To explain this result concretely, consider a TFM as a function that takes as input a single bid from each of n bidders $\vec{v} = (v_1, \ldots, v_n)$, and outputs an allocation vector $X_1(\vec{v}), \ldots, X_n(\vec{v})$, a payment vector $P_1(\vec{v}), \ldots, P_n(\vec{v})$ and a burn amount $\mathsf{Burn}(\vec{v})$. On input $\vec{v}$, each bidder i receives the item with probability $X_i(\vec{v})$, pays $P_i(\vec{v})$, and the miner receives a revenue $\sum_i P_i(\vec{v}) - \mathsf{Burn}(\vec{v})$.

Now, given that user values are drawn from a prior with virtual value function φ, define $X_i^{\varphi}(\vec{\beta}) := X_i(\varphi^{-1}(\beta_1), \ldots, \varphi^{-1}(\beta_n))$ to be the allocation rule that takes as input $\vec{\beta}$ and allocated to bidder i according to X_i *as if the submitted bids were* $\vec{v} = \varphi^{-1}(\vec{\beta})$ *instead of* $\vec{\beta}$. Now view $X^{\varphi}(\cdot)$ as a function that takes as input n values $\varphi(\vec{v}) = (\varphi(v_1), \ldots, \varphi(v_n))$ from a monopsonist with an additive valuation, awards that monopsonist the item i with probability $X_i^{\varphi}(\varphi(\vec{v})) = X_i(\vec{v})$ for all i, and charges her a payment $\mathsf{Burn}^{\varphi}(\varphi(\vec{v})) = \mathsf{Burn}(\vec{v})$. The monoposonist's utility is given by $\sum_i \varphi(v_i) X_i^{\varphi}(\varphi(v_i)) - \mathsf{Burn}^{\varphi}(\varphi(v_i))$, which in expectation (by Myerson's lemma), equals the net revenue received by the miner in equilibrium. Theorem 11 and Theorem 16 establish that $\big(X(\cdot), P(\cdot), \mathsf{Burn}(\cdot)\big)$ is OffC-IP as a TFM for n single-dimensional bidders if and only if $\big(X^{\varphi}(\cdot), \mathsf{Burn}(\cdot)\big)$ is DSIC as an auction for the multi-dimensional monopsonist, thus *fully* characterizing the allocation and burn rules of OffC-IP TFMs.

Main Contribution II: Posted Price Mechanisms are (nearly) the only plaintext OnC-S and OffC-IP TFMs. After establishing our burn identity for OffC-IP mechanisms, we apply this result to OnC-S (i.e., OnC-US and OnC-MS) mechanisms to investigate the limits of these desirable properties. In particular, we ask whether any such TFMs are *plaintext*, i.e., allowing users to submit un-encrypted bids without resorting to the strong use of cryptography employed by the second-price auction considered by Ganesh, Thomas and Weinberg [21].

On this front, we first show that any nontrivial plaintext OnC-S and OffC-IP mechanism must be *prior-dependent*. That is, the TFM must somehow "know" the same prior as the miner and set prices and burns based on this prior.[5] At a high level, this is because our burn identity shows that any OffC-IP mechanism is dependent on the virtual value φ, which in turn, depends on the prior. If the miner herself is the source of how the blockchain learns the prior, this input is provided by the miner after seeing the users' plain-text bids, and we show that this cannot be OnC-MS. Therefore, the protocol's knowledge of the prior must be independent of the miner.

[4] In a cryptographic second-price auction, all users submit encrypted bids. Then, without decrypting the bids, the miner chooses which bids to censor, whether to insert fake bids, and how to set a reserve. Then, all bids are decrypted.

[5] This may be an entirely realistic assumption in case the miner forms their prior exclusively from on-chain data, and an entirely unrealistic assumption in case the miner forms their prior primarily based on off-chain data (such as a tweet announcing an NFT drop). We leave as an important direction for future work to understand the fraction of blocks for which on-chain data suffices to form a reasonably accurate prior, and how the magnitude of inaccuracy impacts OffC-IP.

Given the above, we ask: are there any desirable prior-dependent mechanisms satisfying these properties? Under the additional assumption that there is infinite supply (i.e., any number of users can be included), there is a natural deterministic OnC-S and OffC-IP mechanism: The revenue-optimal posted-price mechanism that charges each included user the Myerson reserve p and transfers all charged prices to the miner (i.e., burning nothing). This mechanism is indeed OnC-S and OffC-IP. In fact, we show that the *only* deterministic plaintext TFMs that are OnC-S and OffC-IP are posted price auctions; specifically, they are generalizations of the revenue-optimal mechanism in which a specially-tuned fraction of each price p is burnt (depending on p and on the prior distribution). Since posted-price mechanisms are not feasible for a finite supply, this implies that there exists *no* deterministic plaintext TFM with finite supply that is OnC-S and OffC-IP.

Given the above, it is natural to ask whether there exist any *randomized* plain-text, OnC-S, and OffC-IP TFMs that are meaningfully distinct from posted-price mechanisms. Surprisingly, we show this is indeed the case. In particular, we construct a plaintext, OnC-S, and OffC-CIP "position auction", where the k^{th} highest bidder is awarded an item with some probability $x^{(k)}$, for any bounded regular distribution, even for a finite block capacity. While our constructed auction is fairly complex and likely impractical, it shows that the space of OffC-IP and OnC-S mechanisms has an intricate boundary. Additionally, our techniques in this construction extend to solve a related open problem in TFM design posted by Chung, Roughgarden and Shi [9] and Gafni and Yaish [20].[6]

Additional Discussions. Our results highlight the importance of posted-price mechanisms in TFM design. There are several variants of posted-price mechanisms.[7] Due to Ganesh, Thomas and Weinberg [21], none of these can be OnC-US, OnC-MS, OffC-IP, and SCP simultaneously. However, several variants satisfy three of the four desiderata; hence, we initiate a discussion of the key tradeoffs involved.

- **SCP or OffC-IP?** We address the tradeoff between OffC-IP and SCP: together with OnC-US and OnC-MS, it is possible to have one of them, but not both. OffC-IP is robust to miners acting as Bayesian monopolists, or colluding with users they do not trust. But, OffC-IP mechanisms are not robust to collusion between trusted parties who know each others' private values. SCP on the other hand is robust even to collusion between trusted parties who know each others' private values (i.e. even Coinbase cannot profit in EIP-1559 by having its stakers collude with its users). But, SCP is not robust to miners acting as Bayesian monopolists. Which property is more desirable in practice, and whether either admit meaningful approximate variants, remain intriguing open directions. See the full version for more details.
- **Prior-dependence, cryptography, or communication?** In the full version, we give additional discussion on the merits and drawbacks of different implementations of a posted-price variant that are OnC-US, OnC-MS, and OffC-IP. We discuss three such variants.

[6] Specifically, Chung, Roughgarden and Shi [9] and Gafni and Yaish [20] study the design of Global Strong Collusion Proof (GSCP) mechanisms: mechanisms that are resilient to collusion by the global coalition containing the miner and all users. They pose the question of designing OnC-S and GSCP mechanisms that also yield a positive revenue to the miner. We design such (randomized) mechanisms in the full version, which are also feasible for a finite block when the users' values are drawn from some bounded support.

[7] For example: is the payment burned or given to the miner? Is cryptography used? Can the protocol set prior-dependent reserves?

- First, the protocol can set the price. This may be possible using on-chain data, but under what conditions and for what fraction of the blocks?
- Second, users can submit encrypted bids that will surely decrypt after the miner has finalized the input. Striking modern advances in cryptography make this more realistic every day, but will intensive cryptography ever be able to match the throughput required?
- Third, and perhaps most-speculatively, miners can "commit" to a price in advance by, for example, posting it on a previous block. Can this cover realistic variations in demand over time, or do miners' prior change so quickly from block-to-block that a rigorous analysis of this sequential pricing game is necessary to understand its viability?

Finally, we remark on the role of our paper within the literature on TFM design. One way to frame the main question of this literature is: How will a block-building protocol devolve into the actual mechanism faced by users? OnC-MS stipulates that miners may manipulate the protocol by myopically best-responding; thus, for example, a second-price auction will devolve into a first-price auction (since the miner will insert fake bids just below the winner's). On the other hand, OffC-IP stipulates that the miner will act as a monopolist running whatever mechanism she wants to determine the inputs to the protocol; thus, for example, EIP-1559 will devolve into a mechanism where users must tip the miner to be included. More generally, under the lens of OffC-IP, the *only* binding constraint on the miner is the set of all possible allocations that are supported by the TFM and the burns corresponding to these allocations, which act as a production cost. Hence, we expect the miner to implement (by any means necessary) the revenue-optimal mechanism given these production costs, and the block-building process does not have the ability to determine the payment rule faced by users in the end. However, when the mechanism is OffC-IP, users can (at least) reliably expect the entire mechanism to be conducted on-chain. This motivates the search for a deep understanding on the burn rules supported by OffC-IP TFMs, so that other properties that enable better user experience, like OnC-MS, can be guaranteed in conjunction with OffC-IP. This is precisely what our paper provides.

1.1 Related Work

Our work contributes to the rapidly-expanding literature on transaction fee mechanism design and more broadly, on game-theoretic applications to decentralized platforms. While Lavi, Sattath and Zohar [28], Basu et al. [6] and Yao [35] were amongst the first to view inclusion of transactions in a block as a mechanism design problem, a research agenda around TFMs started building after Roughgarden [31, 32] established basic desiderata that a "good" TFM must satisfy, and analyzed EIP-1559, the TFM adopted by Ethereum, through the lens of these desiderata. However, Chunag and Shi [10] argue that simultaneously satisfying the three properties proposed by Roughgarden [31] – truthfulness for users and the miner and collusion-resistance – is impossible for any non-trivial mechanism that ever allocates some user with a positive probability. Various notions of collusion-resistance have been considered [10, 15] while Gafni and Yaish [19] and Chung, Roughgarden and Shi [9] investigate the relationship between these collusion-resistance desiderata. Further work has also considered relaxing the concept of truthfulness for users from DSIC to BIC [18, 7]. Shi, Chung and Wu [33, 34] consider expanding the space of TFMs by arming the mechanism with cryptography. While all the aforementioned works treat TFMs as single-shot mechanisms, [16, 4, 1] investigate the long-term dynamics of the repeated game that arises from running the TFM in every block. Bahrani, Garimidi and Roughgarden [5] are the first to revisit the desiderata for TFMs in blockchains with MEV.

Beyond TFMs, a wide body of literature has been growing around economic design in blockchains, starting with Babaioff et al. [3] and Eyal and Sirer [14]. More recent works range from designing credible auctions via cryptography [17, 13, 8, 22, 11] and persuasion mediated via blockchains [12]. There has also been growing literature on cryptographic and consensus protocols in the presence of rational agents [23], [26], [2], [25], [24], [27].

Our paper is a follow-up to Ganesh, Thomas and Weinberg [21], which define off-chain influence-proofness (OffC-IP). They construct a cryptographic OffC-IP and on-chain simple (OnC-S) TFM based on a second price auction, and prove that OffC-IP and OnC-S are incompatible with strong notions of collusion-proofness considered in previous papers. Besides fully characterizing the allocation and burn rules of any OffC-IP mechanism, the present paper advances beyond the results of Ganesh, Thomas and Weinberg [21] by (1) uncovering *plaintext* OffC-IP and OnC-S TFMs, thus removing the strong use of cryptography, and (2) providing formal impossibility results regarding OffC-IP and OnC-S TFMs (regardless of whether the TFM is collusion-proof).

1.2 Road-map

After giving preliminaries, our paper proceeds along the lines of our two main contributions as follows:

1. First, we characterize the allocation and burn rule of any OffC-IP TFM: they are exactly those satisfying the *burn identity* we define in Theorem 11.
   - Section 3 first provides a much-simpler warm-up by characterizing the burn rule of OffC-IP *posted-price* TFMs.
   - Section 4 presents the burn identity for general mechanisms.
2. Second, we examine settings under which OffC-IP and OnC-S TFMs exist, and show that posted-price mechanisms are the only deterministic plaintext OffC-IP and OnC-S TFMs.
   - In Section 5, we consider prior-independent TFMs, and show that in order to be OffC-IP and OnC-S, such TFMs require *both* cryptography and miner-advice (as leveraged by Ganesh, Thomas and Weinberg [21] to construct such a TFM).
   - Motivated by the above, in Section 6 we consider prior-dependent mechanisms. We prove our second main result in Theorem 17: that deterministic OffC-IP and OnC-S TFMs must *necessarily* be the posted-price mechanisms of Section 3. Then we examine various implications and limitations of this result.
     - In particular, Corollary 21 observes that deterministic OffC-IP and OnC-S TFMs thus *require infinite supply*, i.e., they do not exist with finite capacity.
     - In contrast, Subsection 6.3 shows that there exist *randomized* OffC-IP and OnC-S TFMs with finite supply, so long as the prior distribution of values is bounded.

2 Preliminaries: Model of Transaction Fee Mechanisms

We begin with review of transaction fee mechanisms. For preliminaries on Bayesian mechanism design generally, see the full version.

We follow Ganesh, Thomas and Weinberg [21], and consider a model of TFMs which allows the miner to perform both on-chain manipulations (i.e., dropping bids or submitting fake bids) and off-chain manipulations (i.e., conducting an arbitrary off-chain mechanism to determine behavior on-chain). In this model, the TFM is specified by a *block-building process* $\mathcal{B}$, which is an algorithm for constructing a block based on bids collected from all the users and an input from the miner termed the "miner advice" (for example, a reserve price).

Strategic play proceeds in either the *on-chain game* $\mathcal{C}$ or the *off-chain game* $\mathcal{D}$. In $\mathcal{C}$, the miner submits inputs to $\mathcal{B}$ by – as a function of the submitted bids – deciding on an advice, censoring any set of bids, and submitting her own fabricated bids. In $\mathcal{D}$, the miner decides on an arbitrary *off-chain mechanism* $\mathcal{M}_{\mathsf{off}}$ which solicits inputs from users and dictates the on-chain strategies of all agents in $\mathcal{C}$. In addition to the payments made in the on-chain game $\mathcal{C}$, the miner can also direct off-chain transfers through $\mathcal{M}_{\mathsf{off}}$ in $\mathcal{D}$. We omit the full details of this model of TFMs, and refer the reader to Ganesh, Thomas and Weinberg [21].

We adopt the notation $\big(X(\cdot), P(\cdot), \mathsf{Burn}(\cdot)\big)$ for specifying the block-building process $\mathcal{B}$, where $X_i(a_{\mathsf{mi}}, b_1, \ldots, b_n)$ and $P_i(a_{\mathsf{mi}}, b_1, \ldots, b_n)$ specify user i's allocation and payment, and $\mathsf{Burn}(a_{\mathsf{mi}}, b_1, \ldots, b_n)$ specifies the burn, for the miner's advice a_{mi} and bids $b_1, \ldots, b_n$. The on-chain revenue to the miner equals $\sum_{i=1}^{n} P_i(a_{\mathsf{mi}}, b_1, \ldots, b_n) - \mathsf{Burn}(a_{\mathsf{mi}}, b_1, \ldots, b_n)$. We summarize the main concepts and notation in Table 1.

Most of our paper focuses on the *plaintext* model in which bids are un-encrypted and visible to the miner, and hence her strategy can condition on the values of the bids. Ganesh, Thomas and Weinberg [21] also consider a *cryptographic* model of TFMs, which differs only in that users submit encrypted bids, and hence the miner's action must be independent of the values of the bid.

We restrict attention to TFMs in which the block-building process $\mathcal{B}$ is *anonymous*.[8] As in standard mechanism design, this means that each user is treated identically by the TFM, i.e., that if the users' $\{1, \ldots, n\}$ are relabeled with any permutation π, then the outcome is identical except that it is relabeled by π.

Before proceeding, we highlight in concrete terms the main "actions" that a miner can take in the off-chain game, which we need to exploit in our proofs. A miner can (1) censor bids, (2) fabricate and submit their own bids, and (3) ask bidders to submit some alternative bid. Since the miner conducts an entirely separate off-chain mechanism to decide how users behave on-chain, each of actions (1)-(3) can depend on the entire profile of users' messages in the off-chain game. We always assume the users play in a BNE in the off-chain game. Hence, by applying the revelation principle to the off-chain mechanism chosen by the miner, we might as well assume that the off-chain equilibrium is truthful, and the miner learns the values of all the users. It is notationally convenient to describe the total payments made by the users in the off-chain game (i.e, in the composite mechanism including both, the on-chain and off-chain components) as the off-chain payments.[9] Thus, the off-chain payments made by users must satisfy the Myerson's payment identity [29].

Ganesh, Thomas and Weinberg [21] consider three core desiderata for TFM design: On-chain user simplicity, on-chain miner simplicity, and off-chain influence proofness, which each capture different ways in which agents (the users or the miner) will want to deviate from a specified strategy profile. We now briefly recall these definitions, and again refer the reader to Ganesh, Thomas and Weinberg [21] for full details.

- An on-chain miner strategy and user BNE $\sigma^{\mathcal{C}} = (s^{\mathcal{C}}_{\mathsf{mi}}, s^{\mathcal{C}}_{\mathsf{usr},1}, \ldots, s^{\mathcal{C}}_{\mathsf{usr},n})$ is *on-chain user simple* if all users bid their value in $\sigma^{\mathcal{C}}$, and moreover, the equilibrium is DSIC (i.e., each user best-responds by bidding his value for *any* bids of the other users).

[8] Anonymity was also implicitly assumed in the impossibility result in Ganesh, Thomas and Weinberg [21] (e.g., this assumption is needed to say that the outcome of the block-building process when a miner fabricates a bid of 0 is the same as when a user submits a bit of 0).

[9] For example, if the TFM is a posted-price mechanism where the users are expected to pay p on-chain and nothing off-chain, we say that $\mathcal{M}_{\mathsf{off}}$ charges an off-chain payment p from its users.

Table 1 Main elements of the model.

Concept	Notation	Brief Description
Block-building process	$\mathcal{B}(a_{\mathsf{mi}}, \vec{b})$	Algorithm which determines outcome based on miner advice a_{mi} and users' bids $\vec{b}$
Allocation, price, & burn rule of $\mathcal{B}$	(X, P, Burn)	$X_i(a_{\mathsf{mi}}, \vec{b})$ denotes i's allocation, and $P_i(a_{\mathsf{mi}}, \vec{b})$ denotes i's price. The miner earns $\sum_i P_i(a_{\mathsf{mi}}, \vec{b}) - \mathsf{Burn}(a_{\mathsf{mi}}, \vec{b})$.
On-chain game	$\mathcal{C}\big(s^{\mathcal{C}}_{\mathsf{mi}}, s^{\mathcal{C}}_{\mathsf{usr}}(\vec{v})\big)$	Game in which each user i bids $s^{\mathcal{C}}_{\mathsf{usr},i}(v_i)$, the miner sees their bids, then forwards some advice, fabricated bids, and subset of user bids to $\mathcal{B}$
User BNE in $\mathcal{C}$	$\sigma^{\mathcal{C}} = \big(s^{\mathcal{C}}_{\mathsf{usr},i}(v_i)\big)_{i=1}^{n}$	Strategy profile in $\mathcal{C}$ such that, given that the miner plays some fixed $s^{\mathcal{C}}_{\mathsf{mi}}$, the users are playing in a Bayes-Nash equilibrium
Off-chain mechanism	$\mathcal{M}_{\mathsf{off}}(\vec{u})$	An arbitrary mechanism, decided by the miner, that determines the users' and miner's strategy in the on-chain game
Off-chain game	$\mathcal{D}\big(\mathcal{M}_{\mathsf{off}}; s^{\mathcal{D},\mathcal{M}_{\mathsf{off}}}_{\mathsf{usr}}(\vec{v})\big)$	Game in which the miner commits to $\mathcal{M}_{\mathsf{off}}$, and users submit messages and payments to $\mathcal{M}_{\mathsf{off}}$ to determine users' and the miner's strategies in $\mathcal{C}$
User BNE in $\mathcal{D}$	$\sigma^{\mathcal{D}} = \big(s^{\mathcal{D},\mathcal{M}_{\mathsf{off}}}_{\mathsf{usr},i}(v_i)\big)_{i=1}^{n}$	Strategy profile in $\mathcal{D}$ such that, given that the miner commits to some $\mathcal{M}_{\mathsf{off}}$, the users are playing in a Bayes-Nash equilibrium
Revenue in $\mathcal{D}$	$\mathsf{Rev}^{\mathcal{D}}(\mathcal{M}_{\mathsf{off}}, \sigma^{\mathcal{D}}(\vec{v}))$	The miner's total reward in $\mathcal{D}$, i.e., the sum of the on- and off-chain payments.

- An on-chain miner strategy and user BNE $\sigma^{\mathcal{C}} = (s^{\mathcal{C}}_{\mathsf{mi}}, s^{\mathcal{C}}_{\mathsf{usr},1}, \ldots, s^{\mathcal{C}}_{\mathsf{usr},n})$ is *on-chain miner simple* if the miner is best-responding to the strategies of the users, and moreover, $s^{\mathcal{C}}_{\mathsf{mi}}$ never drops or fabricates any bids and always submits the same constant miner advice. In such a case, we say that the miner's strategy is *compliant*.
- An on-chain miner strategy and user BNE $\sigma^{\mathcal{C}} = (s^{\mathcal{C}}_{\mathsf{mi}}, s^{\mathcal{C}}_{\mathsf{usr},1}, \ldots, s^{\mathcal{C}}_{\mathsf{usr},n})$ is *off-chain influence proof* if the miner achieves her maximum revenue in $\sigma^{\mathcal{C}}$ over *all off-chain mechanisms* $\mathcal{M}_{\mathsf{off}}$ and all possible user BNE $\big(\tilde{s}^{\mathcal{D},\mathcal{M}_{\mathsf{off}}}_{\mathsf{usr},i}\big)_{i=1}^{n}$ under $\mathcal{M}_{\mathsf{off}}$.

We say that $\sigma^{\mathcal{C}}$ is *on-chain simple* if it satisfies both on-chain user and miner simplicity.

Ganesh, Thomas and Weinberg [21] also provide supplementary desiderata on collusion resistance, mirroring the definitions of Chung and Shi [10] and Chung, Roughgarden and Shi [9]:

- An on-chain miner strategy and user BNE $\sigma^{\mathcal{C}} = (s^{\mathcal{C}}_{\mathsf{mi}}, s^{\mathcal{C}}_{\mathsf{usr},1}, \ldots, s^{\mathcal{C}}_{\mathsf{usr},n})$ is 1-1-*strong collusion proof* if any coalition containing the miner and one other user cannot deviate to increase the coalition's joint utility.
- An on-chain miner strategy and user BNE $\sigma^{\mathcal{C}} = (s^{\mathcal{C}}_{\mathsf{mi}}, s^{\mathcal{C}}_{\mathsf{usr},1}, \ldots, s^{\mathcal{C}}_{\mathsf{usr},n})$ is *global strong collusion proof* if the global coalition containing the miner and all users cannot deviate to increase their joint utility.

The difference between coalitions and the off-chain mechanisms is the following. While forming a coalition, the agents in the coalition entirely stop being strategic with each other: the users truthfully reveal their values to the miner, and the coalition is purely interested in maximizing their joint utility as if they were a single entity. In contrast, in the off-cain mechanism the users and the miner continue to be strategic with each other (i.e., they must be playing in an equilibrium).

3 Warmup: Characterizing Off-Chain Influence Proof Posted-Price Mechanisms

In this section, we give exposition into our framework and our results by considering a simple sub-class of TFMs: posted-price mechanisms (with infinite supply, where the price and the burn are prior-dependent and determined by the blockchain). We give a short proof characterizing the burn rule such that a posted-price mechanism is off-chain influence proof. Our later sections use more involved arguments to generalize this characterization of the burn rule to all off-chain influence proof TFMs.

Formally, our goal is as follows. Consider a posted-price mechanism with infinite supply. The block-building process $\mathcal{B} = (X, P, \mathsf{Burn})$ does not receive any advice from the miner, and is determined by two fixed parameters: a price Q and a burn B. The mechanism includes all users with values larger than Q, charges each allocated user a payment Q and burns B per allocated user. Consider the on-chain equilibrium $\sigma^{\mathcal{C}}_{\mathsf{honest}}$ where the users bid their values truthfully and the miner does not fabricate or censor bids. We want to compute Q and B such that the equilibrium induced by $\sigma^{\mathcal{C}}_{\mathsf{honest}}$ in the off-chain game is off-chain influence proof.

▶ **Proposition 1.** *For a distribution $\mathcal{T}$ of user values with a continuous virtual value function φ, suppose that the block-building process posts a price Q and burns B per allocated user. Then, the equilibrium $\sigma^{\mathcal{C}}_{\mathsf{honest}}$ is off-chain influence proof if and only if $B = \varphi(Q)$.*

Proof. Suppose that the miner runs an off-chain mechanism $\mathcal{M}_{\mathsf{off}}$ and induces a BNE $\sigma^{\mathcal{D}} = (s^{\mathcal{D}}_{\mathsf{mi}}, s^{\mathcal{D}}_{\mathsf{usr},1}, \ldots, s^{\mathcal{D}}_{\mathsf{usr},n})$. When users have a value profile $\vec{v}$, let their on-chain allocation be $\overline{X}(\vec{v}) = X(s^{\mathcal{D}}_{\mathsf{mi}}(\vec{v}), s^{\mathcal{D}}_{\mathsf{usr},1}(v_1), \ldots, s^{\mathcal{D}}_{\mathsf{usr},n}(v_n))$. Let $\overline{P}(\vec{v})$ denote the equilibrium payments as dictated by Myerson's payment identity for the allocation rule $\overline{X}$.

Then, a Bayesian monopolist miner maximizes her net expected revenue equal to

$$\begin{aligned}
\mathop{\mathbb{E}}_{\vec{v}\sim\mathcal{T}^n}\left[\mathsf{Rev}^{\mathcal{D}}(\mathcal{M}_{\mathsf{off}}, \sigma^{\mathcal{D}}(\vec{v}))\right] &= \mathop{\mathbb{E}}_{\vec{v}\sim\mathcal{T}^n}\left[\sum_{i=1}^{n} \overline{P}_i(\vec{v}) - B \cdot \overline{X}_i(\vec{v})\right] \\
&\qquad \text{(By Myerson's lemma, [29])} \\
&= \mathop{\mathbb{E}}_{\vec{v}\sim\mathcal{T}^n}\left[\sum_{i=1}^{n} \varphi(v_i)\,\overline{X}_i(\vec{v}) - B \cdot \overline{X}_i(\vec{v})\right] \\
&= \mathop{\mathbb{E}}_{\vec{v}\sim\mathcal{T}^n}\left[\sum_{i=1}^{n} \Big(\varphi(v_i) - B\Big) \cdot \overline{X}_i(\vec{v})\right] \qquad (1)
\end{aligned}$$

The miner optimizes her expected revenue precisely by allocating all users with a virtual value at least B, which corresponds to setting a price Q such that $B = \varphi(Q)$. ◀

Proposition 1 shows that, with a carefully-tuned burn rule, any posted price auction yields an off-chain influence proof TFM. Additionally, it is not hard to see that these TFMs are on-chain simple too. This highlights the sharp difference between the TFMs we consider and those of prior work; for instance, Chung and Shi [10] show that there is *no* TFM with positive miner revenue that satisfies on-chain simplicity along with 1-1-strong collusion proofness.

Equation (1) gives intuition towards our burn identity in Section 4, which generalizes this result to apply to *all* TFMs (and not only posted-price mechanisms). Our reduction in Section 4 views the miner as a (single) bidder in a multi-item auction, where each item i corresponds to including bidder i, and relates the total burn in the TFM to the price charged in the multi-item auction. From here, we argue that the miner's objective is to maximize the expectation of $\sum_{i=1}^{n} \varphi(v_i)\, X_i$ minus the total burn, generalizing Equation (1).

Furthermore, Equation (1) can be used to find the revenue-optimal equilibrium of a posted price TFM when the mechanism is not off-chain influence proof. Specifically:

- When the burn B satisfies $\varphi(Q) > B$, the miner will want to include users with virtual values in the range $[B, \varphi(Q)]$ which otherwise are not included in the on-chain equilibrium $\sigma^{\mathcal{C}}_{\mathsf{honest}}$. The miner can ask the users with virtual values in the said interval to amplify their bids to Q and get included in the block. Since all users with a value at least $\varphi^{-1}(B)$ are now included, the miner has to charge each user a payment $\varphi^{-1}(B)$ in equilibrium. Thus, the miner issues a rebate $Q - \varphi^{-1}(B)$ to all users since they would have paid Q on-chain in order to get included.
- On the other hand, when $\varphi(Q) < B$, the miner would not want to include the users with virtual values $[\varphi(Q), B]$. Thus, the miner sets up an off-chain mechanism where it charges an entry fee $\varphi^{-1}(B) - Q$ on top of the Q charged on chain.

4 Myerson-in-Range Mechanisms and a Multi-Item Monopsonist Lens

In this section, we give a reduction relating off-chain influence proof mechanisms to mechanisms in the classical multi-item single-buyer environment. This allows us to derive what we call the *burn identity*, which relates the burn to the allocation rule (in pretty much the same way that the *payment identity* relates the payments to the allocation rule).

For an arbitrary block-building process $\mathcal{B}$, we would like to check whether a given equilibrium $\sigma^{\mathcal{D}}$ in the off-chain game is off-chain influence proof, which entails showing that $\sigma^{\mathcal{D}}$ maximizes revenue over the set of all off-chain mechanisms $\mathcal{M}_{\mathsf{off}}$ and all possible BNE supported by the mechanism $\mathcal{M}_{\mathsf{off}}$. We prove three easy-to-verify sufficient conditions for the equilibrium $\sigma^{\mathcal{D}}$ to be off-chain influence proof which will, in turn, be useful in deriving the burn identity.

4.1 Myerson-in-Range Mechanisms and Off-Chain Influence Proofness

As a first step towards deriving the burn identity, we re-interpret the revenue optimization problem faced by the miner in the off-chain game in terms of the equilibrium burn and the virtual values of the allocated users.

For a revenue optimal equilibrium $\sigma^{\mathcal{D}}$ and a corresponding mechanism $\mathcal{M}_{\mathsf{off}}$ in the off-chain game, we can apply the revelation principle to $\sigma^{\mathcal{D}}$ to unpack the underlying truth-telling mechanism. We say that the truth-telling mechanism $(\overline{X}, \overline{P}, \overline{B})$ corresponding to a revenue optimal equilibrium $\sigma^{\mathcal{D}}$ is *Myerson-in-Range*.

▶ **Definition 2** (Direct-revelation mechanism). *For a block-building process $\mathcal{B} = (X, P, \mathsf{Burn})$, an off-chain mechanism $\mathcal{M}_{\mathsf{off}}$, and a user BNE $\sigma^{\mathcal{D}} = (s^{\mathcal{D}}_{\mathsf{mi}}, s^{\mathcal{D}}_{\mathsf{usr},1}, \ldots, s^{\mathcal{D}}_{\mathsf{usr},n})$ with respect to $\mathcal{M}_{\mathsf{off}}$, the* direct-revelation mechanism *$(\overline{X}, \overline{P}, \overline{B})$ corresponding to $\mathcal{M}_{\mathsf{off}}$ and $\sigma^{\mathcal{D}}$ is obtained by applying the revelation principle; i.e.,*

$$(\overline{X}(\vec{v}), \overline{B}(\vec{v})) = \left(X(\sigma^{\mathcal{D}}(\vec{v})), \mathsf{Burn}(\sigma^{\mathcal{D}}(\vec{v}))\right)$$

and $\overline{P}$ from Myerson's payment identity for the allocation rule $\overline{X}$. We will say that $\overline{X}, \overline{P}$ and $\overline{B}$ are the value allocation, payment and burn rules respectively.

▶ **Definition 3** (Myerson-in-Range). *Fix a block-building process $\mathcal{B}$ and a distribution of user values $\mathcal{T}$. Let $\mathcal{M}_{\mathsf{off}}$ and $\sigma^{\mathcal{D}} = (s^{\mathcal{D}}_{\mathsf{mi}}, s^{\mathcal{D}}_{\mathsf{usr},1}, \ldots, s^{\mathcal{D}}_{\mathsf{usr},n})$ respectively be an off-chain mechanism and a user BNE with respect to $\mathcal{M}_{\mathsf{off}}$ that maximizes the miner's expected net revenue*

$\mathbb{E}_{\vec{v}\sim\mathcal{T}^n}\left[\mathsf{Rev}^{\mathcal{D}}(\mathcal{M}_{\mathsf{off}}, \sigma^{\mathcal{D}}(\vec{v}))\right]$ *for any number of users* n*. We say that* $(\overline{X}, \overline{P}, \overline{B})$ *is* Myerson-in-Range *for* $\mathcal{B}$ *and the distribution* $\mathcal{T}$ *if* $(\overline{X}, \overline{P}, \overline{B})$ *is the direct-revelation mechanism corresponding to such an* $\mathcal{M}_{\mathsf{off}}$ *and* $\sigma^{\mathcal{D}}$*.*

It immediately follows from the definition that an off-chain equilibrium $\sigma^{\mathcal{D}}$ with a trivial off-chain component is off-chain influence proof if and only if the corresponding direct-revelation mechanism is Myerson-in-Range.

▶ **Lemma 4.** *A BNE* $(s^{\mathcal{D}}_{\mathsf{mi}}, s^{\mathcal{D}}_{\mathsf{usr},1}, \ldots, s^{\mathcal{D}}_{\mathsf{usr},n})$ *with a trivial off-chain component is off-chain influence proof for the block-building process* $\mathcal{B} = (X, P, \mathsf{Burn})$ *if and only if its direct-revelation mechanism* $(\overline{X}, \overline{P}, \overline{B})$ *is Myerson-in-Range.*

It is notationally cumbersome to keep mentioning that the mechanism and the corresponding equilibrium satisfies some property (for eg. off-chain influence proofness). We abuse notation to instead say that the corresponding direct-revelation mechanism $(\overline{X}, \overline{P}, \overline{B})$ satisfies the same property.

In the remainder of this section, we characterize all Myerson-in-Range mechanisms for a given block-building process $\mathcal{B}$. We argue that the direct-revelation mechanism $(\overline{X}, \overline{P}, \overline{B})$ is Myerson-in-Range if and only if it optimizes the miner's expected *virtual utility* – the difference between her expected *virtual surplus* $\mathbb{E}_{\vec{v}\sim\mathcal{T}^n}\left[\sum_{i=1}^n \varphi(v_i)\,\overline{X}(\vec{v})\right]$ and the burn $\mathbb{E}_{\vec{v}\sim\mathcal{T}^n}\left[\overline{B}(\vec{v})\right]$ – for all $n \in \mathbb{N}$. In detail, for an equilibrium $\sigma^{\mathcal{D}}$ in the off-chain game with a corresponding direct-revelation mechanism $(\overline{X}, \overline{P}, \overline{B})$, the miner's net expected revenue equals $\mathbb{E}_{\vec{v}\sim\mathcal{T}^n}\left[\sum_{i=1}^n \overline{P}_i(\vec{v}) - \overline{B}(\vec{v})\right]$. Since $\overline{P}$ is the value payment rule of a BNE, we can apply Myerson's lemma [29] to replace the expected payments by the expected virtual surplus in the expected net revenue. In other words, we can pretend that the miner receives a surplus equal to $\sum_{i=1}^n \varphi(v_i)\,\overline{X}_i(\vec{v})$ upon allocating users with values $\vec{v}$, but is charged $\overline{B}(\vec{v})$ via burns. Therefore, the miner optimizes her expected virtual utility $\mathbb{E}_{\vec{v}\sim\mathcal{T}^n}\left[\sum_{i=1}^n \varphi(v_i)\,\overline{X}(\vec{v}) - \overline{B}(\vec{v})\right]$ over all possible BNE $\sigma^{\mathcal{D}}$ in the off-chain game.

In fact, for a regular distribution $\mathcal{T}$, the miner can pointwise optimize her virtual utility for $\vec{v} \sim \mathcal{T}^n$. The value allocation and burn rules $(\overline{X}, \overline{B})$ can take any value from the set of all *feasible outcomes* $\mathcal{F}_{\mathcal{B}}$ of $\mathcal{B} = (X, P, \mathsf{Burn})$, i.e,

$$\mathcal{F}_{\mathcal{B}} = \left\{\mathcal{B}(a_{\mathsf{mi}}; \vec{b})\right\}_{a_{\mathsf{mi}}, \vec{b}}.$$

For any allocation rule $\overline{X}$ supported on $\mathcal{F}_{\mathcal{B}}$, by Myerson's payment identity, $\overline{X}$ is the value allocation rule of some BNE $\sigma^{\mathcal{D}}$ if and only if $\overline{X}_i$ is monotone non-decreasing in v_i for all $1 \le i \le n$ ($\overline{X}_i$ is the allocation rule for user i). For a regular distribution $\mathcal{T}$, it is fairly straightforward to see that there exists

$$\overline{X}(\vec{v}) = \arg \max_{X:(X,B)\in\mathcal{F}_{\mathcal{B}}} \sum_{i=1}^n \varphi(v_i) X_i - B$$

that is monotone non-decreasing in user values. For that matter, any variant of $(\overline{X}, \overline{B})$ that is monotone, and pointwise optimizes for the virtual utility except with probability zero constitutes a Myerson-in-Range mechanism $(\overline{X}, \overline{P}, \overline{B})$ for a suitably chosen value payment rule $\overline{P}$. The miner's expected virtual utility is maximized even if there exists a set of value profiles with a probability measure zero for which the virtual utility optimal outcome is not implemented on-chain.

We summarize our discussion in the theorem below.

▶ **Theorem 5.** *Let $\mathcal{T}$ be a regular distribution. Then, the direct-revelation mechanism $(\overline{X}, \overline{P}, \overline{B})$ is Myerson-in-Range for the block-building process $\mathcal{B}$ if and only if (a) $\overline{X}$ is monotone non-decreasing for all $\vec{v} \in \mathsf{supp}(\mathcal{T}^n)$ and $n \in \mathbb{N}$, (b) $\overline{P}$ satisfies Myerson's payment identity and (c) except with probability zero,*

$$\big(\overline{X}(\vec{v}), \overline{B}(\vec{v})\big) = \arg \max_{(X,B)\in\mathcal{F}_{\mathcal{B}}} \sum_{i=1}^{n} \varphi(v_i) X_i - B$$

for $\vec{v} \sim \mathcal{T}^n$.[10]

4.2 Three Canonical Properties of Myerson-in-Range Mechanisms

In our second step towards obtaining the burn identity, we will identify three intuitive properties satisfied by Myerson-in-Range mechanisms.

The first property (labeled *optimal for n users*) is an immediate special case of our characterization of Myerson-in-Range mechanisms. While Theorem 5 suggests that $\overline{X}(\vec{v})$ pointwise optimizes $\sum_{i=1}^{n} \varphi(v_i)X - B$ almost everywhere for $(X, B) \in \mathcal{F}_{\mathcal{B}}$, the first condition specifically observes that $\varphi(\vec{v})\,\overline{X}(\vec{v}) - \overline{B}(\vec{v}) \geq \varphi(\vec{v})\,\overline{X}(\vec{w}) - \overline{B}(\vec{w})$ with probability 1 for $\vec{v} \sim \mathcal{T}^n$ and all $\vec{w} \in \mathsf{supp}(\mathcal{T}^n)$.

The second property (labeled *Negative φ's are suboptimal*) observes that including users with value below the monopoly reserve is suboptimal for a virtual utility maximizing miner. Suppose an allocation rule includes users with a negative virtual value with probability greater than zero. Then, the miner can increase her virtual utility by first censoring such a user's bid and fabricating an identical bid of her own. Since the same set of bids are submitted to the anonymous TFM, except for un-allocating the user with the negative virtual value, all other users are allocated the same and the burn remains unchanged too, thereby increasing the miner's expected virtual utility.

Finally, the third property (labeled *No censoring or fabricating*) connects the allocation rule for n users with the allocation rule for $\hat{n}$ users. It suggests that the miner's virtual utility should remain invariant to censoring or fabricating bids with a non-positive virtual value.

▶ **Theorem 6.** *Let $\mathcal{T}$ be a regular distribution with a continuous virtual value function φ and a probability mass function such that $\Pr_{v\sim\mathcal{T}}[\varphi(v) \leq 0] > 0$. Further, let $\mathcal{T}_{\varphi\leq 0}$ be the distribution $\mathcal{T}$ conditioned on the virtual value of a random draw having a non-positive virtual value. Then, a direct-revelation mechanism $(\overline{X}, \overline{P}, \overline{B})$ is Myerson-in-Range only if for any number n of users:*

(A) (Optimal for n users.) *For a value profile $\vec{v} \sim \mathcal{T}^n$ and $\vec{w} \in \mathsf{supp}(\mathcal{T}^n)$,*

$$\sum_{i=1}^{n} \varphi(v_i)\,\overline{X}_i(\vec{v}) - \overline{B}(\vec{v}) \geq \sum_{i=1}^{n} \varphi(v_i)\,\overline{X}_i(\vec{w}) - \overline{B}(\vec{w})$$

with probability one.

(B) (Negative φ's are suboptimal.) *For $\vec{v} \sim \mathcal{T}^n$, $\overline{X}_i(\vec{v}) = 0$ almost surely whenever $\varphi(v_i) < 0$.*

[10] The term "Myerson-in-Range" is derived from "Maximal-in-Range" which originated in the combinatorial auctions literature. A mechanism is said to be maximal over some range $\mathcal{F}_{\mathcal{B}}$ of feasible allocations if it implements the surplus optimal allocation over $\mathcal{F}_{\mathcal{B}}$ (typically, $\mathcal{F}_{\mathcal{B}}$ is a strict subset of the set of all feasible allocations). Similarly, a Myerson-in-Range mechanism maximizes the miner's virtual utility while restricting allocations to those permitted by the set $\mathcal{F}_{\mathcal{B}}$, which could possibly pack much less users than the capacity of the block.

(C) (No censoring or fabricating.) *For* $\vec{v} = (v_1, \ldots, v_n, v_{n+1}, \ldots, v_{n+t}) \sim \mathcal{T}^n \times \mathcal{T}^t_{\varphi \leq 0}$ *and* $\vec{w} = (v_1, \ldots, v_n)$,

$$\sum_{i=1}^{n} \varphi(v_i)\overline{X}_i(\vec{v}) - \overline{B}(\vec{v}) = \sum_{i=1}^{n} \varphi(v_i)\overline{X}_i(\vec{w}) - \overline{B}(\vec{w})$$

almost surely.

We defer the proof of Theorem 6 to the full version. For the remainder of our discussions, we solely focus on the class of regular distributions $\mathcal{T}$ meeting the requirements of Theorem 6, i.e., having a continuous virtual value function and a positive probability of a random draw having a non-positive virtual value. Most natural distributions such as the uniform distribution, the normal distribution and the exponential distribution satisfy this condition. We say that a distribution is *smooth* if it satisfies the requirements from Theorem 6.

4.3 The Burn Identity for Myerson-in-Range Mechanisms

In this section, we derive a closed-form expression for the burn rule of a Myerson-in-Range mechanism. At a high level, we will reduce calculating the burn rule of a Myerson-in-Range mechanism to computing equilibrium payments in multi-item single-buyer environments.

We show that the virtual utility maximization problem faced by the miner is similar to the utility maximization problem faced by a monopsonist in multi-item settings. In Theorem 5, we argued that a Bayesian miner chooses an outcome from the menu of user allocation probabilities and burns specified by $\mathcal{F}_{\mathcal{B}}$ in the off-chain game to optimize her virtual utility. In other words, the miner "buys" an outcome supported by $\mathcal{F}_{\mathcal{B}}$ by making a "payment" equal to the burn corresponding to the outcome. In return, she derives a "surplus" equal to the virtual surplus of the allocation. Concretely, the items in the multi-item single-buyer environment are analogous to users in the TFM environment and the monopsonist's values for the items correspond to the virtual values of the users and are drawn iid from $\mathcal{T}^{\varphi}$, the distribution of virtual values when the value is drawn from $\mathcal{T}$.

From the reduction discussed above, computing the equilibrium allocation and burn of a Myerson-in-Range mechanism corresponds to computing the equilibrium allocation and payments in the analogous multi-item single-buyer environment. Since we apply the revelation principle in the TFM environment, the value allocation and burn rules $(\overline{X}, \overline{B})$ will correspond to DSIC mechanisms for the monopsonist. However, there are some subtle differences between the two settings. In the multi-item single-buyer environment, the monopsonist should be disincentivized to deviate from truth-telling for all value profiles over the items. However, in the TFM environment, an off-chain equilibrium and the corresponding direct-revelation mechanism is Myerson-in-Range even if the miner would like to deviate from the equilibrium behaviour with probability zero. We will address the discrepancy in the definitions between the two environments by arguing that the allocation and the burn rules of a Myerson-in-Range mechanism can be *smoothened* over some measure zero collection of value profiles so that the miner pointwise optimizes her virtual utility for all value profiles $\vec{v}$ in the smoothened mechanism.[11]

[11] Note that the smoothened mechanism might not necessarily be feasible, i.e, be supported over $\mathcal{F}_{\mathcal{B}}$. However, we will use the smoothened mechanism to derive a burn identity which holds almost surely for the original mechanism.

▶ **Theorem 7.** *Let the users' values be drawn from a smooth regular distribution $\mathcal{T}$. Suppose that the direct-revelation mechanism $(\overline{X}, \overline{P}, \overline{B})$ is Myerson-in-Range for some block-building process $\mathcal{B}$. Then, there exists an allocation rule $\widetilde{X}$ and a burn rule $\widetilde{B}$ such that $(\widetilde{X}, \widetilde{B})$ is identical to $(\overline{X}, \overline{B})$ almost surely for any number n of users and for any profile of values $\vec{v}, \vec{w} \in \mathsf{supp}(\mathcal{T}^n)$,*

$$\sum_{i=1}^{n} \varphi(v_i)\, \widetilde{X}_i(\vec{v}) - \widetilde{B}(\vec{v}) \geq \sum_{i=1}^{n} \varphi(v_i)\, \widetilde{X}_i(\vec{w}) - \widetilde{B}(\vec{w}).$$

We will say that such a pair $(\widetilde{X}, \widetilde{B})$ is a *monopsonist smoothening* of $(\overline{X}, \overline{B})$.

▶ **Definition 8** (Monopsonist smoothening)**.** *An allocation and burn rule $(\widetilde{X}, \widetilde{B})$ is a* monopsonist smoothening *of a Myerson-in-Range mechanism $(\overline{X}, \overline{B})$ for a distribution $\mathcal{T}$ and a block-building process $\mathcal{B}$ if (a) $(\widetilde{X}, \widetilde{B})$ is identical to $(\overline{X}, \overline{B})$ except on a set of probability measure zero for any number n of users and (b) $\sum_{i=1}^{n} \varphi(v_i)\, \widetilde{X}_i(\vec{v}) - \widetilde{B}(\vec{v}) \geq \sum_{i=1}^{n} \varphi(v_i)\, \widetilde{X}_i(\vec{w}) - \widetilde{B}(\vec{w})$ for all $\vec{v}, \vec{w} \in \mathsf{supp}(\mathcal{T}^n)$.*

We defer the proof of Theorem 7 to the full version.

The monopsonist smoothening enables the application of results from existing literature on multi-item auctions to transaction fee mechanism design. For instance, for a given utility function $U(\vec{v})$ mapping the buyer's value profile to her utility[12], Rochet [30] argues that the allocation rule and payment rule of the DSIC mechanism that induces the utility function U is essentially unique.

▶ **Theorem 9** (Rochet, [30])**.** *A function $U : \mathbb{R}^n \to \mathbb{R}$ is the utility function of some DSIC mechanism in an n-item single-buyer environment if and only if U is convex and non-decreasing. The allocation rule of such a DSIC mechanism is given by $\overline{X}(\vec{v}) = \nabla U(\vec{v})$ and the payment charged to the buyer equals $P(\vec{v}) = \vec{v}^{\intercal} \nabla U(\vec{v}) - U(\vec{v})$.*

Note that when U is convex, ∇U is defined almost everywhere, and the allocation and payment rules $\overline{X}$ and $\overline{P}$ are uniquely defined at all such points with a well-defined gradient. At values where ∇U is not well-defined, setting $\overline{X}$ to be any sub-derivative of U satisfies dominant-strategy incentive compatibility.

Applying Theorem 7 in conjunction with Theorem 9 for a virtual utility optimizing miner yields the following.

▶ **Corollary 10.** *Consider a Myerson-in-Range mechanism $(\overline{X}, \overline{B})$ for a block-building process $\mathcal{B}$ and a smooth regular distribution $\mathcal{T}$. If,*

$$\overline{U}(\varphi(\vec{v})) = \sum_{i=1}^{n} \varphi(v_i)\, \overline{X}_i(\vec{v}) - \overline{B}(\vec{v}) \geq \sum_{i=1}^{n} \varphi(v_i)\, \overline{X}_i(\vec{w}) - \overline{B}(\vec{w})$$

almost surely for all $\vec{v} \sim \mathcal{T}^n$ and $\vec{w} \sim \mathsf{supp}(\mathcal{T}^n)$, then, there exists a function U convex and non-decreasing in virtual values $\varphi(\vec{v})$ such that

$$U(\varphi(\vec{v})) = \overline{U}(\varphi(\vec{v})), \qquad \overline{X}(\vec{v}) = \nabla^{\varphi} U(\varphi(\vec{v})) \quad \text{and} \quad \overline{B}(\vec{v}) = \sum_{i=1}^{n} \varphi(v_i)\, \nabla_i^{\varphi} U(\varphi(\vec{v})) - U(\varphi(\vec{v}))$$

with probability 1 for $\vec{v} \in \mathsf{supp}(\mathcal{T}^n)$, where ∇^{φ} is the gradient with respect to the virtual values $\varphi(\vec{v})$.

[12] Such a utility function can be derived even for non-truth-telling equilibria via the revelation principle (see Rochet , [30], for example).

We say that the function U is the *smoothened virtual utility function* of the mechanism $(\overline{X}, \overline{P}, \overline{B})$.

Corollary 10 expresses condition A from Theorem 6 in terms of the smoothened virtual utility. We can also similarly rewrite conditions B and C. Condition B states that the mechanism should almost never allocate a user with a virtual value smaller than zero. Thus, the gradient $\nabla_i^{\varphi} U(\varphi(\vec{v})) = \overline{X}_i(\vec{v}) = 0$ almost surely whenever $\varphi(v_i) < 0$. Since U is convex in $\varphi(\vec{v})$, $\nabla_i^{\varphi} U(\varphi(\vec{v}))$ is increasing in $\varphi(\vec{v})$ and thus, $\nabla_i^{\varphi} U(\varphi(\vec{v})) = 0$ for all v_i with a strictly negative virtual value.

Condition C suggests that fabricating or censoring a bid with a non-positive virtual value should almost never change the miner's virtual utility. Thus, for value profiles $\vec{v} = (v_1, \ldots, v_n, v_{n+1}, \ldots, v_{n+t})$ and $\vec{w} = (v_1, \ldots, v_n)$ such that $\varphi(v_{n+i}) \leq 0$ for all $1 \leq i \leq t$, the smoothened virtual utility function at $\vec{v}$ and $\vec{w}$ must satisfy $U(\vec{v}) = U(\vec{w})$.

We summarize the discussions above to state the burn identity.

▶ **Theorem 11** (Burn identity). *Let $(\overline{X}, \overline{B})$ be the value allocation and burn rule of a Myerson-in-Range mechanism for a smooth regular distribution $\mathcal{T}$. Then, there exists a family of smoothened virtual utility functions $\left(U^n\right)_{n \in \mathbb{N}}$, $U^n : \mathbb{R}^n \to \mathbb{R}$, such that for all $n \in \mathbb{N}$,*

1. *U^n is convex and non-decreasing as a function of the virtual values of the bids,*
2. *$\nabla_i^{\varphi} U^n(\varphi(\vec{v})) = 0$ whenever $\varphi(v_i) < 0$,*
3. *$U^n(\varphi(\vec{v})) = U^{n+1}(\varphi(\vec{w}))$ for all $\vec{v}, \vec{w}$ such that $\vec{v} = (v_1, \ldots, v_n, v_{n+1}, \ldots, v_{n+t}) \in \mathsf{supp}(\mathcal{T}^{n+t})$ and $\vec{w} = (v_1, \ldots, v_n)$ for $\varphi(v_{n+i}) \leq 0$ for all $1 \leq i \leq t$, and,*
4. *With probability 1 for $\vec{v} \sim \mathcal{T}^n$,*

$$\overline{X}(\vec{v}) = \nabla^{\varphi} U^n(\varphi(\vec{v})) \text{ and } \overline{B}(\vec{v}) = \sum_{i=1}^{n} \varphi(v_i) \nabla_i^{\varphi} U(\varphi(\vec{v})) - U(\varphi(\vec{v})).$$

Finally, we will combine the burn identity with Myerson's payment identity to derive a DSIC payment rule in terms of the smoothened virtual utility function. From the payment identity,

$$\overline{P}_i(\vec{v}) = v_i \overline{X}_i(\vec{v}) - \int_0^{v_i} \overline{X}_i(v, \vec{v}_{-i})\, dv.$$

The smoothened virtual utility function determines the allocation rule everywhere except at a measure zero set of points. Further, since the distribution $\mathcal{T}$ is regular, all points between the infimum and supremum of $\mathsf{supp}(\mathcal{T})$ have a positive probability density and as a consequence, even when $\overline{X}$ differs from $\nabla^{\varphi} U$ at a measure zero set of points, $\int_0^{v_i} \overline{X}_i(v, \vec{v}_{-i})\, dv = \int_0^{v_i} \nabla_i^{\varphi} U(\varphi(v), \varphi(\vec{v}_{-i}))\, dv$, almost everywhere. Thus, the payment rule of a Myerson-in-Range mechanism with a smoothened utility function U equals

$$\overline{P}_i(\vec{v}) = v_i \nabla_i^{\varphi} U(\varphi(\vec{v})) - \int_0^{v_i} \nabla_i^{\varphi} U(\varphi(v), \varphi(\vec{v}_{-i}))\, dv$$

with probability 1.

To conclude our discussion on Myerson-in-Range mechanisms, by Lemma 4, observe that a direct-revelation mechanism is off-chain influence proof only if it has a trivial off-chain component and its burn rule satisfies the burn identity.

The reduction from TFMs to multi-item auctions characterizing off-chain influence proof mechanisms can be extended to also characterize global strong collusion proof mechanisms. In the full version, we use our characterization to design a mechanism that is on-chain simple and global strong collusion proof which yields a positive revenue to the miner.[13]

[13] [9, 19] show that no such deterministic mechanism exists that yields positive miner revenue. We construct a randomized mechanism with positive miner revenue. Our mechanism applies for both infinite and bounded block size, although for a finite block, our mechanism requires the prior distribution to have a known upper bound (and it remains an open question whether such mechanisms exist without a known upper bound).

5 Prior-Independent On-Chain Miner Simple and Off-Chain Influence Proof Mechanisms Require Both Cryptography and Miner Advice

We now use the burn identity derived in Section 4 to reason about desirable TFMs in various settings. Going forward, we focus on TFMs that satisfy all the desirable properties highlighted by Ganesh, Thomas and Weinberg [21], and call these mechanisms *simple to participate*.

▶ **Definition 12.** *A direct-revelation mechanism* $(\overline{X}, \overline{P}, \overline{B})$ *is* simple to participate *if it is on-chain user simple, on-chain miner simple, and off-chain influence proof.*

In this section, we consider prior independent mechanisms, i.e., ones where the block-building process $\mathcal{B}$ has no knowledge of the distribution of values $\mathcal{T}$. Ganesh, Thomas and Weinberg [21] show that there exists a cryptographic variant of the second-price auction that is prior-independent and simple to participate for all regular distributions $\mathcal{T}$. However, their auction requires two features: heavy-duty cryptography, and allowing the miner to inform the mechanism through an advice (namely, allowing the miner to set a reserve). We show in this section that *both* of these features are necessary to get a nontrivial prior independent simple-to-participate TFM.

First, we show that the only (possibly cryptographic) prior-independent mechanism that does not rely on any advice from the miner and is simple to participate for a sufficiently rich class of distributions is the trivial mechanism that allocates users with zero probability irrespective of the distribution. In order to satisfy condition B of Theorem 6, an off-chain influence proof mechanism should be able to discriminate between users with values smaller than the monopoly reserve and the rest only using information available to the block-building process $\mathcal{B}$, so as to allocate the former with zero probability. However, $\mathcal{B}$ is prior-independent and is therefore agnostic to the monopoly reserve of $\mathcal{T}$.[14]

▶ **Theorem 13.** *Consider the class* T *of smooth regular distributions. Then, an on-chain user simple and off-chain influence proof mechanism that is both prior-independent and advice-independent for all* $\mathcal{T} \in \mathsf{T}$ *must be the trivial mechanism that never allocates any user irrespective of the block capacity.*

We prove Theorem 13 in the full version.

Second, we also rule out plain-text mechanisms that are simple to participate for a sufficiently rich class of distributions. At a high level, suppose there exists a distribution $\mathcal{T}$ and a value profile $\vec{v}$ for which the miner is awarded a revenue larger than the block reward from building an empty block. For another distribution $\hat{\mathcal{T}}$ with a monopoly reserve much larger than all values in $\vec{v}$, by condition B from Theorem 6, all users must receive no allocation when their value profile equals $\vec{v}$. However, the miner can deviate from the equilibrium behaviour, submit the advice $a^{\mathcal{T}}_{\mathsf{mi}}$ corresponding to $\mathcal{T}$ and trick the mechanism into believing that the users' values are drawn from $\mathcal{T}$. The miner will then receive a revenue larger than the block reward, contradicting on-chain miner simplicity.

[14] The reason we require on-chain user simplicity for Theorem 13 is subtle. Suppose the mechanism is not on-chain user simple and the users adopt different strategies for different distributions. The mechanism can potentially infer the value distribution based on the bids submitted by the users. However, when the mechanism is on-chain user simple, there is truly no information regarding the distribution that can be inferred from the users' behaviour. The users will always bid their values in equilibrium. We leave as an open problem whether the same impossibility can be proved purely via off-chain influence proofness, without having to rely on on-chain user simplicity.

▶ **Theorem 14.** *Consider the class* T *of smooth regular distributions* $\mathcal{T}$ *with no point masses, i.e, there exists no value* $\hat{v}$ *such that* $\Pr_{v\sim\mathcal{T}}[v = \hat{v}] > 0$*. Then, a simple-to-participate plain-text prior-independent mechanism for all* $\mathcal{T} \in \mathsf{T}$ *must be the trivial mechanism that almost never allocates any user irrespective of the block capacity or the distribution* $\mathcal{T}$*.*

We defer the proof to the full version.

6 Prior-Dependent, On-Chain User and Miner Simple, Off-Chain Influence Proof, Plain-Text Mechanisms

In this section, we consider the case of prior-dependent mechanisms, i.e., we assume the block-building process $\mathcal{B}$ knows the distribution $\mathcal{T}$ of user values. We focus on plain-text TFMs. While the two main desiderata of Section 5 are no longer relevant for this section,[15] other considerations emerge from the analysis as important. Specifically, we explore whether simple-to-participate mechanisms exist even for finite blocks, whether randomization is necessary for a TFM and finally, whether the distribution can have an arbitrary unbounded support. While we uncover additional TFMs in some settings, our main finding is that *posted-price mechanisms* (for an infinite block) as discussed in Section 3 are the only mechanisms that are deterministic, plain-text and simple to participate.

Throughout this section, we find it convenient to rank users $(1), \dots (n)$ in descending order of bids. For notational convenience, we will denote the quantities corresponding to the i^{th} largest bid by (i) (for example, $v^{(i)}$, $\overline{X}^{(i)} = \nabla^{\varphi}_{(i)} U^n$, etc).

6.1 A Sufficient Condition for Prior-Dependent On-Chain User Simple and Off-Chain Influence Proof Mechanisms

Before constructing simple-to-participate mechanisms, it would be useful to have a sufficient condition that ensures that a candidate mechanism $(\overline{X}, \overline{P}, \overline{B})$ is indeed on-chain user simple and off-chain influence proof. We will begin with such a sufficient condition very similar to the three properties in Theorem 6.

The three properties in Theorem 6 do not characterize Myerson-in-Range mechanisms due to the following reason – the three conditions do not automatically guarantee that, for some block-building process $\mathcal{B} = (X, P, \mathsf{Burn})$, $\varphi(\vec{v})\,\overline{X}(\vec{v}) - \overline{B}(\vec{v}) \geq \varphi(\vec{v})\,X - B$ for all $(X, B) \in \mathcal{F}_{\mathcal{B}}$. It only ensures that $\varphi(\vec{v})\,\overline{X}(\vec{v}) - \overline{B}(\vec{v}) \geq \varphi(\vec{v})\,X - B$ for $(X, B) = (\overline{X}(\vec{w}), \overline{B}(\vec{w}))$ for some value profile $\vec{w}$. However, suppose that the mechanism is on-chain user simple and $(\overline{X}(\vec{v}), \overline{P}(\vec{v}), \overline{B}(\vec{v})) = (X(\vec{v}), P(\vec{v}), \mathsf{Burn}(\vec{v}))$ for all $\vec{v}$.[16] Then, $\varphi(\vec{v})\,\overline{X}(\vec{v}) - \overline{B}(\vec{v}) \geq \varphi(\vec{v})\,X - B$ for all $(X, B) \in \mathcal{F}_{\mathcal{B}}$. In this section, we argue ensuring that the direct-revelation mechanism realizes all feasible outcomes in $\mathcal{F}_{\mathcal{B}}$, along with the three conditions from Theorem 6, is sufficient for the mechanism to be Myerson-in-Range.

[15] Section 5 considers whether the TFM needs to use miner advice, and whether it needs to use cryptography. A prior-dependent TFM does not require miner advice, since (for example) the block-building process itself can now know the right reserve price to set based on the distribution $\mathcal{T}$. Similarly, the simple-to-participate cryptographic variant of the second-price auction discussed by Ganesh, Thomas and Weinberg [21] for the prior-independent setting can easily be modified for the prior-dependent setting – by simulating the miner's advice for the distribution $\mathcal{T}$. So, we focus solely on the plain-text case.

[16] Remember that the mechanism is prior-dependent and we are assuming that the miner does not submit any advice to the block-building process (X, P, Burn). Thus, X, P and Burn take as input a profile of bids and output the allocations, payments and the burn respectively.

▶ **Lemma 15.** *Let $\mathcal{T}$ be a smooth regular distribution. For a prior-dependent block-building process $\mathcal{B} = (X, P, \mathsf{Burn})$, suppose a direct-revelation mechanism $(\overline{X}, \overline{P}, \overline{B})$ satisfies $(\overline{X}, \overline{P}, \overline{B}) = (X, P, \mathsf{Burn})$. Further, let $(\overline{X}, \overline{P}, \overline{B})$ satisfy the three conditions from Theorem 6 for any number n of users and for all value profiles belonging to $\mathsf{supp}(\mathcal{T}^n)$.*

(A) *(Optimal for n users.) For all $\vec{v}, \vec{w} \in \mathsf{supp}(\mathcal{T}^n)$, we have*

$$\sum_{i=1}^{n} \varphi(v_i)\, \overline{X}_i(\vec{v}) - \overline{B}(\vec{v}) \geq \sum_{i=1}^{n} \varphi(v_i)\, \overline{X}_i(\vec{w}) - \overline{B}(\vec{w}).$$

(B) *(Negative φ's are suboptimal.) For $\vec{v} \in \mathsf{supp}(\mathcal{T}^n)$, $\overline{X}_i(\vec{v}) = 0$ whenever $\varphi(v_i) < 0$.*

(C) *(No censoring or fabricating.) For $\vec{v} = (v_1, \ldots, v_n, v_{n+1}, \ldots, v_{n+t}) \in \mathsf{supp}(\mathcal{T}^n) \times \mathcal{T}^t_{\varphi \leq 0}$ and $\vec{w} = (v_1, \ldots, v_n)$,*

$$\sum_{i=1}^{n} \varphi(v_i)\overline{X}_i(\vec{v}) - \overline{B}(\vec{v}) = \sum_{i=1}^{n} \varphi(v_i)\overline{X}_i(\vec{w}) - \overline{B}(\vec{w}).$$

Then, $(\overline{X}, \overline{P}, \overline{B})$ is Myerson-in-Range.

We defer the proof to the full version.

Note that the direct-revelation mechanism $(\overline{X}, \overline{P}, \overline{B})$ satisfying the requirements in Lemma 15 is on-chain user simple and off-chain influence proof. For on-chain user simplicity, observe that the block-building process takes in bids and already implements the payment rule corresponding to the DSIC direct-revelation mechanism. For off-chain influence proofness, once again, equilibrium payments are charged from the users by the payment rule $P = \overline{P}$ on-chain, and thus, the miner will not have to make any extraneous off-chain transfers to steer users towards the underlying BNE. Thus, the Myerson-in-Range mechanism $(\overline{X}, \overline{P}, \overline{B})$ has a trivial off-chain component and by Lemma 4, the mechanism is off-chain influence proof.

Remember that the payment identity in the multi-item environment (Theorem 9) is both necessary and sufficient for a mechanism to be DSIC for a monopsonist. Applying Theorem 9 for a virtual utility maximizing miner, we get a partial converse to Theorem 11.

▶ **Theorem 16** (Partial converse to the burn identity)**.** *For a prior-dependent block-building process $\mathcal{B} = (X, P, \mathsf{Burn})$ and a direct-revelation mechanism $(\overline{X}, \overline{P}, \overline{B}) = (X, P, \mathsf{Burn})$, suppose there exists a family of virtual utility functions $\left(U^n\right)_{n \in \mathbb{N}}$ such that*

$$\overline{X}(\vec{v}) = \nabla^{\varphi} U^n(\varphi(\vec{v})), \quad \overline{P}_i(\vec{v}) = v_i \nabla^{\varphi}_i U(\varphi(\vec{v})) - \int_0^{v_i} \nabla^{\varphi}_i U(\varphi(v), \varphi(\vec{v}_{-i}))\, dv$$

$$\text{and} \quad \overline{B}(\vec{v}) = \sum_{i=1}^{n} \varphi(v_i)\, \nabla^{\varphi}_i U(\varphi(\vec{v})) - U(\varphi(\vec{v}))$$

satisfying

1. *U^n is convex and non-decreasing as a function of the virtual values of the bids,*
2. *$\nabla^{\varphi}_i U^n(\varphi(\vec{v})) = 0$ whenever $\varphi(v_i) < 0$, and,*
3. *$U^n(\varphi(\vec{v})) = U^{n+t}(\varphi(\vec{v}), \varphi(\hat{v}))$ for all $\hat{v} = (\hat{v}_{n+1}, \ldots, \hat{v}_{\hat{n}})$ such that $\varphi(\hat{v}_{n+i}) \leq 0$.*

for all $n \in \mathbb{N}$. Then, $(\overline{X}, \overline{P}, \overline{B})$ is on-chain user simple and off-chain influence proof.

6.2 Deterministic Simple-to-Participate Mechanisms

In this section, we argue that the only deterministic plain-text mechanism that is simple to participate is the class of all uniform posted-price mechanisms.[17]

▶ **Theorem 17.** *Let $\mathcal{T}$ be a smooth regular distribution. Then, a deterministic simple-to-participate mechanism allocates all bids above some threshold p at least the monopoly reserve, charges p from each allocated user and burns $\varphi(p) \geq 0$ per included bid apart from awarding a block reward to the miner independent of the contents of the created block.*

As a first step towards proving Theorem 17, we will show that a deterministic off-chain influence proof mechanism has a simple burn rule parametrized only by the number of users included in the block. Concretely, we will argue that there exists a sequence of burns $\left(\Phi_t\right)_{t\in\mathbb{N}}$ such that Φ_t is burnt whenever the miner builds a block containing t users. As a sanity check, note that the block reward equals the net revenue from building an empty block, which in turn equals $-\Phi_0$.

▶ **Lemma 18.** *For a smooth regular distribution $\mathcal{T}$ and any deterministic off-chain influence proof mechanism, there exists a sequence of burns $\left(\Phi_t\right)_{t\in\mathbb{N}}$ such that for a value profile $\vec{v} = (v^{(1)}, \ldots, v^{(n)})$, the direct-revelation mechanism $(\overline{X}, \overline{P}, \overline{B})$ almost surely includes users $(1), \ldots, (t)$ for*

$$t = \arg\max_{\hat{t}} \sum_{i=1}^{\hat{t}} \varphi(v^{(i)}) - \Phi_{\hat{t}}$$

and burns Φ_t for including exactly t users.

We defer the proof to the full version.

For the remainder of the section, we will show that $\Phi_t = t\,(\Phi_1 - \Phi_0) + \Phi$ for simple-to-participate deterministic mechanisms. In other words, the *marginal burn* $\Delta\Phi_t := \Phi_t - \Phi_{t-1}$, i.e, the burn for including the t^{th} user conditioned on including $t-1$ users, must be a constant $\Delta\Phi$. Proving a constant marginal burn will conclude the proof of Theorem 17 for $\varphi(p) = \Delta\Phi$ and $p = \varphi^{-1}(\Delta\Phi)$.

To show that the marginal burn $\Delta\Phi_t$ is a constant, we will first argue that the marginal burn is non-increasing in t (Lemma 19), and then, for a non-increasing sequence $\left(\Delta\Phi_t\right)_{t\in\mathbb{N}}$, we will show that they cannot decrease either (Lemma 20).

▶ **Lemma 19.** *Let $\mathcal{T}$ be a smooth regular distribution. Then, a deterministic TFM with marginal burns $\left(\Delta\Phi_t\right)_{t\in\mathbb{N}}$ that is simple to participate must satisfy $\Delta\Phi_t \geq \Delta\Phi_{t+1}$.*

We defer the proof to the full version.

▶ **Lemma 20.** *Let $\mathcal{T}$ be a smooth regular distribution. A deterministic simple-to-participate TFM with a sequence of non-increasing marginal burns $\left(\Delta\Phi_t\right)_{t\in\mathbb{N}}$ must satisfy $\Delta\Phi_t = \Delta\Phi$ for all $t \in \mathbb{N}$.*

[17] Remember that, by definition, we require our TFMs to be anonymous. Strictly speaking, a deterministic, anonymous, virtual utility maximizing mechanism is not possible for a block with a finite capacity. For example, for a block of capacity 1, if there are two users with equal values, both larger than the monopoly reserve, the mechanism either has to randomize or has to break ties in the favour of one bidder. However, we will consider mechanisms that are deterministic and anonymous up to tie-breaking. Our impossibility results are agnostic to how ties are broken, and do not rely on the edge-cases requiring tie-breaking.

We defer the formal proof to the full version.

For a block with a finite capacity, if the number of users with values more than the threshold p from Theorem 17 is larger than the capacity of the block, the mechanism cannot allocate all the users feasibly. Thus, there exists no deterministic mechanism that is simple to participate for a finite block.

▶ **Corollary 21.** *When the block has a finite capacity, there exists no non-trivial deterministic simple-to-participate mechanism for a smooth regular distribution $\mathcal{T}$.*

6.3 Position Auctions with Burn

Position auctions are mechanisms where the allocation probability for a user does not depend on the exact bid submitted by the user, but only on the rank of the user's bid amongst all other submitted bids. We consider position auctions where Φ_t is burnt to allocate t users.

▶ **Definition 22** (Position auctions)**.** *For a regular distribution $\mathcal{T}$, a* position auction *is given by a sequence of non-increasing allocation probabilities $\left(x^{(i)}\right)_{i\in\mathbb{N}}$ and* marginal burn per unit allocation *$\left(\Delta\Phi_i\right)_{i\in\mathbb{N}}$. Conditioned on allocating t users, the users are allocated with probabilities $x^{(1)},\ldots,x^{(t)}$ in descending order of their bids and $\Phi_t = \sum_{i=1}^{t}\Delta\Phi_i\, x^{(i)} + \Phi_0$ is burnt. For $\vec{v}\sim\mathcal{T}^n$, the mechanism almost surely maximizes the miner's virtual utility*

$$\max_t \sum_{i=1}^{t}\left(\varphi(v^{(i)}) - \Delta\Phi_i\right) x^{(i)} - \Phi_0 = \max_t \sum_{i=1}^{t}\varphi(v^{(i)})\, x^{(i)} - \Phi_t.$$

Similar to deterministic mechanisms, we will argue that the marginal burn per unit allocation $\Delta\Phi_t$ must be a constant $\Delta\Phi$, independent of the number of allocated users t.

▶ **Lemma 23.** *Let $\mathcal{T}$ be a smooth regular distribution. Then, the marginal burn per unit allocation $\left(\Delta\Phi_i\right)_{i\in\mathbb{N}}$ of a simple-to-participate position auction is a constant sequence $\left(\Delta\Phi\right)_{i\in\mathbb{N}}$.*

The proof is extremely similar to Lemma 19 and Lemma 20, and defer the proof to the full version.

In Theorem 24, we show that there cannot exist simple-to-participate position auctions that are not uniform posted-price mechanisms for unbounded distributions, even for a block with infinite capacity. On the other hand, Theorem 25 shows the existence of simple-to-participate position auctions beyond posted-price mechanisms for bounded smooth regular distributions and finite blocks, and further, characterizes all such mechanisms.

▶ **Theorem 24.** *For an unbounded smooth regular distribution $\mathcal{T}$, consider a position auction given by the allocation rule $\left(x^{(i)}\right)_{i\in\mathbb{N}}$ and a marginal burn per unit allocation $\Delta\Phi$. If the mechanism is simple-to-participate, then, $x^{(i)} = x$ for a constant $0\le x\le 1$ for all $i\in\mathbb{N}$.*

▶ **Theorem 25.** *For any bounded smooth regular distribution $\mathcal{T}$ and a block with a finite capacity Ω, a position auction given by the allocation probabilities $\left(x^{(t)}\right)_{t\in\mathbb{N}}$ and marginal burn per unit allocation $\Delta\Phi$ is on-chain miner simple if and only if (and thereby simple-to-participate if and only if) $\sum_{t=1}^{\infty} x^{(t)} \le \Omega$ and*

$$t\,(x^{(t)} - x^{(t+1)})\,(\sup\mathcal{T} - \varphi^{-1}(\Delta\Phi)) < \Delta\Phi\, x^{(t+1)} \tag{2}$$

for all $t\in\mathbb{N}$.

We defer the proof to the full version, where we also construct a position auction that satisfies the conditions in Theorem 25 for a finite block.

6.4 Generalized Position Auctions

So far, we have argued that the posted-price mechanism with a suitable burn is simple-to-participate for an infinite block. In contrast, when the distribution is bounded, we showed that there exists position auctions with an allocation rule meaningfully different from that of the posted-price mechanism and are simple-to-participate even for a finite block. In this section, we propose *generalized position auctions* that are also meaningfully different from the posted-price mechanism, and are simple-to-participate for infinite blocks and value distributions with an unbounded support. However, we show generalized position auctions cannot be simple to participate for finite blocks and thus, we will have to search in a much larger class of mechanisms to construct simple-to-participate mechanisms when the block capacity is bounded.

At a high level, the allocation rule for a user in a position auction depends only its rank when all bids are arranged in descending order. This can be extended to a generalized position auction by choosing a class of single-agent allocation rules $\left(x^{(t)}\right)_{t\in\mathbb{N}}$, where user (i) with the i^{th} largest bid and value $v^{(i)}$ is allocated with probability $x^{(i)}(v^{(i)})$.

▶ **Definition 26** (Generalized position auctions). *A* generalized position auction *is given by a sequence of single-agent allocation rules $\left(x^{(t)}\right)_{t\in\mathbb{N}}$, where $x^{(t)}:\mathbb{R}_{\geq 0}\to[0,1]$. For a value profile $\vec{v}$, user (i) receives an allocation $x^{(i)}(v^{(i)})$. The payment and the burn rules are chosen to satisfy the payment and burn identities (Theorem 11) respectively.*

Next, we will construct simple-to-participate generalized position auctions for an infinite block. The strategy to show that the mechanism is indeed simple to participate is similar to Theorem 25.

▶ **Theorem 27.** *Let $\mathcal{T}$ be a smooth regular distribution with an unbounded virtual value function. Suppose that an on-chain user simple and off-chain influence proof generalized position auction with allocation rule $\left(x^{(t)}\right)_{t\in\mathbb{N}}$ satisfies*

$$t\left(x^{(t)}(w)-x^{(t+1)}(w)\right)\geq(t+1)\left(x^{(t+1)}(w)-x^{(t+2)}(w)\right) \tag{3}$$

for all $t\in\mathbb{N}$, and

$$t\times\int_{\varphi^{-1}(0)}^{w}\left(x^{(t)}(z)-x^{(t+1)}(z)\right)dz\leq\varphi(w)\,x^{(t+1)}(w)-\int_{0}^{\varphi(w)}x^{(t+1)}(z)\,d\varphi(z) \tag{4}$$

for all $t\in\mathbb{N}$ and $w\geq\varphi^{-1}(0)$. Then, the mechanism is also on-chain miner simple.

We defer the proof to the full version. In the full version, we construct a simple-to-participate generalized position auction for smooth regular distributions.

We conclude our discussion by arguing that the only simple-to-participate generalized position auction for a finite block and an unbounded distribution is the trivial mechanism.

▶ **Theorem 28.** *Let $\mathcal{T}$ be a smooth regular distribution. Then, the only simple-to-participate generalized position auction for a finite block when user values are drawn from $\mathcal{T}$ is the trivial auction that almost never allocates any user.*

We leave as an open question whether there exist simple-to-participate mechanisms for unbounded distributions for blocks with a finite capacity. In the full version, we also discuss the challenges of designing an inference procedure to learn the prior $\mathcal{T}$ while being resistant to censorship and fabricated bids and leave open the question of designing such a procedure.

References

1 Guillermo Angeris, Theo Diamandis, and Ciamac Moallemi. Multidimensional blockchain fees are (essentially) optimal. *arXiv preprint arXiv:2402.08661*, 2024. doi:10.48550/arXiv.2402.08661.

2 Gilad Asharov, Ran Canetti, and Carmit Hazay. Towards a game theoretic view of secure computation. In *Eurocrypt*, pages 426–445, 2011. doi:10.1007/978-3-642-20465-4_24.

3 Moshe Babaioff, Shahar Dobzinski, Sigal Oren, and Aviv Zohar. On bitcoin and red balloons. In *Proceedings of the 13th ACM conference on electronic commerce*, pages 56–73, 2012. doi:10.1145/2229012.2229022.

4 Moshe Babaioff and Noam Nisan. On the optimality of eip-1559 for patient bidders (draft–comments welcome), 2024.

5 Maryam Bahrani, Pranav Garimidi, and Tim Roughgarden. Transaction fee mechanism design in a post-mev world. *Cryptology ePrint Archive*, 2024.

6 Soumya Basu, David Easley, Maureen O'Hara, and Emin Gün Sirer. Towards a functional fee market for cryptocurrencies. *CoRR*, abs/1901.06830, 2019. arXiv:1901.06830.

7 Xi Chen, David Simchi-Levi, Zishuo Zhao, and Yuan Zhou. Bayesian mechanism design for blockchain transaction fee allocation. *arXiv preprint arXiv:2209.13099*, 2024.

8 Tarun Chitra, Matheus V. X. Ferreira, and Kshitij Kulkarni. Credible, optimal auctions via blockchains. Cryptology ePrint Archive, Paper 2023/114, 2023. URL: https://eprint.iacr.org/2023/114.

9 Hao Chung, Tim Roughgarden, and Elaine Shi. Collusion-resilience in transaction fee mechanism design. *EC 2024*, 2024. arXiv preprint arXiv:2402.09321. doi:10.48550/arXiv.2402.09321.

10 Hao Chung and Elaine Shi. Foundations of transaction fee mechanism design. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3856–3899, 2023. doi:10.1137/1.9781611977554.CH150.

11 Hao Chung, Ke Wu, and Elaine Shi. Foundations of platform-assisted auctions. *arXiv preprint arXiv:2501.03141*, 2025.

12 Kimon Drakopoulos, Irene Lo, and Justin Mulvany. Blockchain mediated persuasion. In *Proceedings of the 24th ACM Conference on Economics and Computation*, EC '23, 2023.

13 Meryem Essaidi, Matheus VX Ferreira, and S Matthew Weinberg. Credible, strategyproof, optimal, and bounded expected-round single-item auctions for all distributions. *arXiv preprint arXiv:2205.14758*, 2022.

14 Ittay Eyal and Emin Gün Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *Financial Cryptography and Data Security*, pages 436–454. Springer, 2014. doi:10.1007/978-3-662-45472-5_28.

15 Matheus VX Ferreira, Yotam Gafni, and Max Resnick. Incentive-compatible collusion-resistance via posted prices. *arXiv preprint arXiv:2412.20853*, 2024.

16 Matheus VX Ferreira, Daniel J Moroz, David C Parkes, and Mitchell Stern. Dynamic posted-price mechanisms for the blockchain transaction-fee market. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, pages 86–99, 2021.

17 Matheus VX Ferreira and S Matthew Weinberg. Credible, truthful, and two-round (optimal) auctions via cryptographic commitments. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 683–712, 2020.

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

19 Yotam Gafni and Aviv Yaish. Barriers to collusion-resistant transaction fee mechanisms. *EC 2024*, 2024. arXiv preprint arXiv:2402.08564. doi:10.48550/arXiv.2402.08564.

20 Yotam Gafni and Aviv Yaish. Discrete and bayesian transaction fee mechanisms. In *The International Conference on Mathematical Research for Blockchain Economy*, pages 145–171. Springer, 2024. doi:10.1007/978-3-031-68974-1_8.

21 Aadityan Ganesh, Clayton Thomas, and S Matthew Weinberg. Revisiting the primitives of transaction fee mechanism design. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 703–703, 2024.

22 Aadityan Ganesh and Qianfan Zhang. Truthful, credible, and optimal auctions for matroids via blockchains and commitments. In *Proceedings of the 26th ACM Conference on Economics and Computation*, pages 923–943, 2025. doi:10.1145/3736252.3742652.

23 Juan Garay, Jonathan Katz, Ueli Maurer, Björn Tackmann, and Vassilis Zikas. Rational protocol design: Cryptography against incentive-driven adversaries. In *FOCS*, pages 648–657, 2013. doi:10.1109/FOCS.2013.75.

24 Tiantian Gong, Aniket Kate, Hemanta K. Maji, and Hai H. Nguyen. Disincentivize collusion in verifiable secret sharing. Cryptology ePrint Archive, Paper 2025/446, 2025. URL: https://eprint.iacr.org/2025/446.

25 Adam Groce, Jonathan Katz, Aishwarya Thiruvengadam, and Vassilis Zikas. Byzantine agreement with a rational adversary. In *ICALP*, pages 561–572, 2012. doi:10.1007/978-3-642-31585-5_50.

26 Joseph Halpern and Vanessa Teague. Rational secret sharing and multiparty computation: extended abstract. In *STOC*, pages 623–632, 2004. doi:10.1145/1007352.1007447.

27 Mahimna Kelkar, Aadityan Ganesh, Aditi Partap, Joseph Bonneau, and S Matthew Weinberg. Breaking omertà: On threshold cryptography, smart collusion, and whistleblowing. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security*, pages 3505–3519, 2025. doi:10.1145/3719027.3765087.

28 Ron Lavi, Or Sattath, and Aviv Zohar. Redesigning bitcoin's fee market. In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, pages 2950–2956, 2019. doi:10.1145/3308558.3313454.

29 Roger B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, 1981. doi:10.1287/MOOR.6.1.58.

30 Jean-Charles Rochet. The taxation principle and multi-time hamilton-jacobi equations. *Journal of Mathematical Economics*, 14(2):113–128, 1985.

31 Tim Roughgarden. Transaction fee mechanism design for the ethereum blockchain: An economic analysis of eip-1559. *arXiv preprint arXiv:2012.00854*, 2020. arXiv:2012.00854.

32 Tim Roughgarden. Transaction fee mechanism design. *ACM SIGecom Exchanges*, 19(1):52–55, 2021. Full version at https://arxiv.org/abs/2106.01340.

33 Elaine Shi, Hao Chung, and Ke Wu. What can cryptography do for decentralized mechanism design. *Innovations in Theoretical Computer Science (ITCS)*, 2023.

34 Ke Wu, Elaine Shi, and Hao Chung. Maximizing miner revenue in transaction fee mechanism design. *ITCS*, 2024.

35 Andrew Chi-Chih Yao. An incentive analysis of some bitcoin fee designs. *CoRR*, abs/1811.02351, 2018. arXiv:1811.02351.