Artificial Intelligence 357 (2026) 104543

Contents lists available at ScienceDirect

Artificial Intelligence

journal homepage: www.elsevier.com/locate/artint

Auction design with ex post ROI constraints

Hongtao Lv [a,b], Xiaohui Bei [c], Zhenzhe Zheng [d,*], Fan Wu [d,*]

[a] Joint SDU-NTU Centre for Artificial Intelligence Research (C-FAIR), Shandong University, Jinan, China
[b] School of Software, Shandong University, Jinan, China
[c] School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore, Singapore
[d] School of Computer Science, Shanghai Jiao Tong University, Shanghai, China

ARTICLE INFO

Keywords:
Return on investment (ROI)
Mechanism design
Myerson auction

ABSTRACT

Motivated by practical constraints in online advertising, we investigate single-parameter auction design for bidders with constraints on their Return On Investment (ROI) – a targeted minimum ratio between the obtained value and the payment. We focus on *ex post* ROI constraints, which require the ROI condition to be satisfied for every realized value profile. With ROI-constrained bidders, we first provide a full characterization of the allocation and payment rules of dominant-strategy incentive compatible (DSIC) auctions. In particular, we show that given any monotone allocation rule, the corresponding DSIC payment should be the Myerson payment with a *rebate* for each bidder to meet their ROI constraints. Furthermore, we also determine the optimal auction structure when the item is sold to a single bidder under a mild regularity condition. This structure entails a randomized allocation scheme and a first-price payment rule, which differs from the deterministic Myerson auction and previous works on ex ante ROI constraints. Finally, for multiple ROI-constrained bidders, we investigate simple auction design and, in particular, deterministic auction design and how they can be applied in sponsored search auctions.

1. Introduction

Online advertising auctions are a vital source of revenue for many IT companies. In recent years, with tens of millions of ad auctions being conducted in real-time each day, this large-scale and complex market has prompted modern online advertising platforms to develop auto-bidding services, which allow the advertisers to set up high-level marketing goals for their ad campaigns and then bid on behalf of the advertisers.

In these auto-bidding scenarios, advertisers' financial constraints such as budget and return on investment (ROI) constraints have become critical in auction design. While auctions for budget-constrained bidders have been extensively studied in the literature [1–3], research on auction design for bidders with ROI constraints is still in its nascent stage. The ROI constraints of advertisers require that the payment cannot be more than a certain fraction of the obtained advertising value. In other words, there is a targeted minimum ratio between the obtained value and the payment for an ROI-constrained bidder. Unlike budget constraints which set a hard limit on payment, ROI constraints establish a payment limit that is linearly related to the allocated value. Previous studies [4,5] have demonstrated that ROI constraints align better with real-world empirical evidence than budget constraints, and it is the aim of this paper to explore how to design auctions with good incentive and revenue guarantees for ROI-constrained bidders.

The existing literature on auction design for ROI-constrained bidders primarily focuses on *ex ante* ROI constraints, which requires an *expected* ROI with respect to the prior value distributions of bidders [4,6]. This approach is suitable for advertisers who participate

* Corresponding authors.
E-mail addresses: lht@sdu.edu.cn (H. Lv), xhbei@ntu.edu.sg (X. Bei), zhengzhenzhe@sjtu.edu.cn (Z. Zheng), fwu@cs.sjtu.edu.cn (F. Wu).

https://doi.org/10.1016/j.artint.2026.104543
Received 15 January 2024; Received in revised form 6 February 2026; Accepted 20 April 2026
Available online 24 April 2026
0004-3702/© 2026 Elsevier B.V. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

in a large number of auctions daily and are only concerned with their average spend per unit of value. However, in reality, most ad campaigns experience the “long-tail phenomenon” [7], which means they only receive dozens of or fewer clicks per day. Under these conditions, an auction with ex ante ROI guarantees may have a non-negligible probability of violating the ROI constraints of these ad campaigns over a day. Due to these reasons, in this work, we focus on the *ex post* or *hard* ROI constraints, which ensures that the auction respects the ROI constraints of bidders for any realized value profile. This is a stronger requirement compared to ex ante ROI constraints and addresses the limitation of current auction design methods.

1.1. Our results

In this work, we examine the design of truthful and optimal auction design for ex post ROI-constrained bidders. We inherit the setting from the classic single-parameter mechanism design and consider the values of bidders as private information and the targeted ROIs as public information. In this single-parameter environment, the ROI constraints can be integrated into the objective function (see Section 2 for details), resulting in a transformed utility model: $u_i = M_i v_i x_i - p_i$, where u_i represents the utility of bidder i, v_i is the value, x_i is the allocation quantity, and p_i is the payment. Here $M_i > 1$ is the targeted ROI ratio, which differentiates this model from the classical quasilinear utility model. We show that this new utility model is related to both the traditional *utility maximizer* model and the *value maximizer* models [8,9] emerging in recent years, which takes the allocation value as the objective without subtracting the payment.

We first study the characterizations of truthful auctions with ROI-constrained bidders. Compared to Myerson’s characterization of truthful auctions in the single-parameter environment, we show that the monotonicity requirement of the allocation rule remains true for ex post ROI-constrained bidders, but the unique payment rule in [10] should be modified by subtracting a max term, which can be interpreted as a “rebate” equal to the largest “violation” of the Myerson payment to the ROI constraint for all lower valuations. This is a full characterization that completely describes all truthful auctions with ROI-constrained bidders. This result can be proved using similar techniques from Myerson’s analysis. It can also be derived from the following alternative interpretation of the payment rule (we will explain this in detail later on): note that the ROI-constrained bidder assigns a weight $M_i > 1$ to her obtained value $v_i x_i$ from the allocation, so a naive approach is to charge the bidder M_i times the Myerson payment. However, to not violate the individual rationality (IR) condition, we must iteratively apply the Myerson payment increment (multiplied by M_i) in small intervals and “cap” or truncate the payment at the obtained value whenever necessary.

Next, we turn our focus to the optimal (*i.e.* revenue-maximizing) auction design. The additional max term in our payment rule poses a significant challenge to the optimal auction design, since it is unclear how this term can be incorporated into a modified virtual valuation function as seen in previous literature. Instead, we concentrate on the case of selling a single item to a single bidder. Our main result suggests that under a mild regularity assumption known as *decreasing marginal revenue* (DMR),[1] the optimal auction for selling to a single ex post ROI-constrained bidder employs a randomized scheme. More specifically, the allocation rule $x(\cdot)$ starts with a *first-price* interval, where the payment always matches the obtained value, until it reaches the highest allocation and $x(\cdot)$ becomes constant thereafter. This finding is in contrast to the classic Myerson auction [10] and previous results for bidders with ex ante ROI constraints, where the optimal auctions are always deterministic. It implies that similar to much literature on optimal mechanism design for multi-parameter settings and some single-parameter environments, including public budget [11] and risk aversion [12,13], a slight generalization of the inclusion of the M_i term can also lead to randomized optimal auctions.

We then turn our attention to multi-bidder scenarios. The randomized allocation for the single-bidder setting implies that any extension could be very technical; hence, we mainly focus on simple auctions and show that simple and approximately optimal auctions for traditional bidders established in previous literature also have approximation guarantees for ROI-constrained bidders. Furthermore, for deterministic auctions which are often preferred in real-world applications such as online advertising, it is revealed that, when there is a single item and multiple ROI-constrained bidders, the optimal deterministic auction is equivalent to the Myerson auction.

Finally, we apply our results to the specific context of online sponsored search auctions. In our characterization of truthful deterministic sponsored search auctions with ROI-constrained bidders, the payment could be determined by taking the minimum of a ladder-style payment and a critical-price payment. Furthermore, we show that when there is only one ROI-constrained bidder with finite valuation space, the optimal deterministic auction can be efficiently computed using a dynamic programming algorithm.

1.2. Related work

There are two main threads of studies of auctions with ROI-constrained bidders. The first thread investigates how the bidding strategies of the bidders are affected by the ROI constraints in classic VCG or generalized second price (GSP) auctions [14–20]. The second thread, which our paper follows, focuses on the *design* of auctions with ROI-constrained bidders, which is of practical interest to many online advertising platforms [4,6,16,21–26]. In this line of study, the most related work to ours is [4], which showed empirically that a fraction of the buyers in online advertising are indeed ROI-constrained. They also took the first step towards revenue-maximizing auction design for bidders with Bayesian incentive compatibility (BIC) and *ex ante* ROI constraints, which only require

[1] DMR requires the marginal revenue, $vf(v) + F(v) - 1$ to be non-decreasing in the *value space*. This is different from the usual definition of regularity which requires the same monotonicity but in the *quantile space*. Please see the related work section for a more detailed discussion of their differences and more related works.

the ROI conditions to be met in expectation and are strictly weaker than ex post ROI constraints. They proved that the allocation rule is very similar to Myerson auction with a modified virtual value function. The payment rule is related to the ROI constraint: if the ROI constraint is low to the extent that it is not binding, then the modified virtual values reduce to the classical virtual values; if the ROI constraint is moderate, then the mechanism operates with the modified virtual values; if the ROI constraint is high, then the mechanism needs to subsidize buyers to ensure their participation. It is worth noting that this mechanism is a deterministic one. Another recent work [6] considered the scenario where either the value or the ROI constraint is private information of the bidder. They also considered ex ante ROI, but they used a different constraint of ex ante IC. For the public ROI and private valuation setting, they found that although ex ante IC leads to a larger feasible set than BIC, the optimal mechanism is exactly the one in [4] as stated above. Unlike these works, we concentrate on *ex post* ROI constraints, which provide a hard ROI guarantee for bidders in every possible value realization. In [27,28], the authors considered ex post ROI constraints in multiple stages, and assumed that each bidder maintains a fixed bid multiplier among stages, which leads to completely different problems from ours.

One particular line of research focuses on requirements of truthfulness for ex post ROI constraints. Cavallo *et al.* studied the same utility function as ours in [22, Appendix A], and investigated the corresponding payment rules. The main difference is that they limited their focus on *deterministic* mechanisms for bidders with *identical* ROI constraints, while we consider a more general single-parameter setting in the randomized mechanism domain. They show that the truthful payment rule for deterministic sponsored search auction is a generalization of the classical VCG and GSP mechanism. Our results on deterministic auctions in Section 5 also generalize their results to non-identical ROI constraints. Li et al. [21] proposed a condition on the truthfulness of the ROI information, based on which they provided a mechanism framework using tools from control theory. They took the ROI constraints as private information, instead of the value, which leads to a substantially different problem from ours.

The DMR assumption used in our optimal auction characterization has been widely discussed in the literature on auction and dynamic pricing [29–32]. It means that the function $\psi(v) \triangleq vf(v) + F(v) - 1$ is non-decreasing, or equivalently $v \cdot (1 - F(v))$, which is the expected revenue of selling the item at price p, is concave, and this is where the name of this condition comes from. Intuitively, many commonly used distribution functions satisfy this assumption, *e.g.*, uniform distributions, and any distribution of finite support and monotone non-decreasing density. The DMR condition was first proposed in [29] for bidders with budget constraints. In [31], the authors found that the DMR condition is more natural in their setting than the traditionally used notion of *regularity* [10], since DMR precisely removes the requirement of ironing in the *value* space, instead of in the *quantile* space as in [10]. In [30], DMR was discussed comprehensively, and the authors showed that the optimal mechanism is deterministic under the DMR condition in a multi-unit setting with private demands. We refer the reader to their work for concrete examples and more discussion.

2. Preliminaries

We consider a general single-parameter auction environment, which consists of a seller and n bidders $\boldsymbol{N} = \{1, 2, \dots, n\}$. Each bidder i has a private valuation t_i per unit of the good. We represent x_i as the quantity of the allocated good to bidder i and p_i as the payment of bidder i. Without loss of generality, we assume the maximum possible allocation is $x_i^{\max} = 1$ and the good is indivisible, that is, x_i denotes the probability of bidder i receiving the good. Besides the allocated value, each bidder also has a return on investment (ROI) constraint M_i, as public information,[2] which specifies the minimum targeted ratio between her obtained value and the payment. We assume $1 < M_i < +\infty$ in this work. We note that the ROI constraint is considered in an ex post measure, *i.e.*, it requires that $\frac{t_i x_i}{p_i} \geq M_i$ strictly holds in the outcome of every auction. Note that the same model is also adopted in [22, Appendix A].

With the above definitions, the utility of bidder i is given by

$$u_i = \begin{cases} t_i x_i - p_i & \text{if } \frac{t_i x_i}{p_i} \geq M_i \\ -\infty & \text{otherwise.} \end{cases} \tag{1}$$

It is worth noting that this is the standard quasilinear utility model with the addition of the ROI constraint. We can further define

$$v_i = \frac{t_i}{M_i},$$

which could be interpreted as the maximum willingness-to-pay of the bidder i per unit of the good. Then, we can rewrite the utility function as

$$u_i = \begin{cases} M_i v_i x_i - p_i & \text{if } v_i x_i \geq p_i \\ -\infty & \text{otherwise.} \end{cases} \tag{2}$$

One can observe that, as M_i is a public constant, v_i and t_i are completely interchangeable. To avoid confusion, we use the term *value* to represent v_i, and *initial value* to represent t_i in the following discussion. Each value v_i is independently drawn from a probability distribution $F_i : [0, v_{\max}] \to [0, 1]$, with a continuous probability density function f_i. While the distributions F_i's are common knowledge, the exact value v_i is known only to the bidder i. We denote $\mathbf{v}$ as the value profile of all bidders, and $\mathbf{v}_{-i}$ as that of all bidders except bidder i.

In an auction, each bidder reports her value as b_i, which is not necessarily equal to v_i. We define $\mathbf{b}$ and $\mathbf{b}_{-i}$ similarly as the notations of $\mathbf{v}$ and $\mathbf{v}_{-i}$. Based on the reported bids, an auction mechanism consists of an allocation rule $x_i(b_i, \mathbf{b}_{-i})$, mapping the bid

[2] This setting is practical and prevalent in practice, *e.g.*, in online advertising, the targeted ROI typically remains the same over a certain period.

profile to the allocated quantity to each bidder i, and a payment rule $p_i(b_i, \mathbf{b}_{-i})$, mapping the bid profile to the payment for each bidder. We also use $u_i(b_i, v_i, \mathbf{b}_{-i})$ to represent the utility of bidder i who has value v_i and bid b_i. When clear from context, we will omit $\mathbf{b}_{-i}$ in the mappings. In the following discussion, we assume that the allocation rule $x_i(\cdot)$ is always right-differentiable, and there are finite non-differentiable points. When $x_i(\cdot)$ is non-continuous at v, let $x_i(v) = \lim_{z\to v^+} x_i(z)$.

We are interested in auctions that are *dominant-strategy incentive compatible* (DSIC) and *ex post individually rational* (ex post IR), where DSIC intuitively means that reporting the true value (*i.e.*, $b_i = v_i$) is no worse than reporting any other values, and ex post IR represents that submitting true values in the auction never suffers a loss.

Definition 1 (Dominant-Strategy Incentive Compatibility, DSIC). A mechanism is dominant-strategy incentive compatible if and only if

$$u_i(v_i, v_i, \mathbf{b}_{-i}) \geq u_i(b_i, v_i, \mathbf{b}_{-i}), \quad \forall b_i, \mathbf{b}_{-i}, i \in \boldsymbol{N}.$$

Definition 2 (Ex Post Individual Rationality, ex post IR). A mechanism is ex post individually rational if and only if

$$u_i(v_i, v_i, \mathbf{b}_{-i}) \geq 0, \quad \forall v_i, \mathbf{b}_{-i}, i \in \boldsymbol{N}.$$

In subsequent sections, we will omit $\mathbf{b}_{-i}$ when it does not cause ambiguity. For ease of notation, we use *truthfulness* to represent the properties of both DSIC and ex post IR in the following sections, and we also use IR to represent ex post IR when there is no ambiguity. In addition, for truthful auctions, we do not distinguish v_i and b_i hereinafter.

The revenue of a truthful auction is defined as

$$\mathsf{rev} = \mathbb{E}_{\mathbf{v}}\left[\sum_{i\in \boldsymbol{N}} p_i(v_i)\right].$$

The aim of this work is to characterize both truthful and revenue-maximizing (optimal) auctions with ex post ROI-constrained bidders.

3. Characterize the structure of DSIC auctions

In this section, we present characterizations of the DSIC auctions with ex post ROI constraints. These results generalize the classical Myerson's Lemma [10] for the traditional utility model (*i.e.*, $M_i = 1$), which states that in the single-parameter environment, a mechanism is DSIC if and only if its allocation rule is monotone and the payment scheme follows a unique rule.

Lemma 1 (Myerson's Lemma [10]). *For traditional bidders with $M_i = 1$, a single-parameter mechanism is DSIC if and only if:*

Monotone Allocation Rule the allocation rule is monotonically non-decreasing, i.e.*, $x_i(v) \leq x_i(v')$ for all $v < v'$ and bidder i;*

Unique Payment Rule for each monotonically non-decreasing allocation rule $x_i(\cdot)$, and $p_i(0) = 0$, the payments are given by

$$p_i(v) = v x_i(v) - \int_0^v x_i(z)\,\mathrm{d}z. \tag{3}$$

Clearly, these results cannot be directly applied to the ROI-constrained bidders, because the payment derived from Myerson's Lemma may violate the ROI constraints, or otherwise induce bidders to misreport a higher bid to achieve higher utility. The main result in this section is a complete characterization of the DSIC mechanisms with ROI-constrained bidders. We will see that the monotonicity condition for the allocation remains the same, but the payment rule needs to be modified appropriately to accommodate the ROI constraints. We will also demonstrate later in Lemma 2 that this characterization can be adapted to Bayesian incentive compatible (BIC) mechanisms.

Theorem 1 (Characterization). *For ex post ROI-constrained bidders, a single-parameter mechanism is DSIC if and only if:*

Monotone Allocation Rule the allocation rule is monotonically non-decreasing, i.e.*, $x_i(v) \leq x_i(v')$ for all $v < v'$ and bidder i;*

Unique Payment Rule for each monotonically non-decreasing allocation rule $x_i(\cdot)$, and $p_i(0) = 0$, the payments are given by

$$p_i(v) = M_i\widetilde{p}_i(v) - \max_{0\leq z\leq v}\{M_i\widetilde{p}_i(z) - z x_i(z)\}, \tag{4}$$

where $\widetilde{p}_i$ is the Myerson payment given in (3).

We can interpret this characterization from two perspectives: First, from the perspective of the initial value t_i of bidder i, it suggests that compared to the classic Myerson auction, an ROI-constrained bidder with value v will need to pay the initial Myerson payment $M_i\widetilde{p}_i(v)$ (recall that $t_i = M_i v_i$), minus a "rebate" which equals the largest "violation" of the Myerson payment to the ROI requirement when the bidder's valuation is no more than v.

Second, from the perspective of the value v_i, since the ROI-constrained bidder assigns a weight $M_i > 1$ to her obtained value from the allocation, but not to her payment, a naive application that charges the bidder M_i times of the Myerson payment may violate the IR constraint. Therefore, we need to iteratively apply the Myerson payment increment (multiplied by M_i) in small intervals and truncate the payment at the value whenever necessary. These two perspectives are mathematically equivalent, and we adopt the second perspective in the following for exposition convenience.

Next, before proving this theorem, some observations are immediate from this characterization. We defer the proof of these observations to Appendix A.1.

Proposition 1.

1. *We always have $p_i(v) \le v x_i(v)$ and $p_i(v) \le M_i \widetilde{p}_i(v)$ for any bidder i and value v in DSIC mechanisms.*
2. *The payment function $p_i(\cdot)$ is monotonically non-decreasing for any bidder i in DSIC mechanisms.*
3. *For the same allocation function $x_i(\cdot)$, we always have $p_i(v) \ge \widetilde{p}_i(v)$ for any valuation v,* i.e.*, the payment of an ROI-constrained bidder in a truthful mechanism is always no lower than that in Myerson auction.*

Now we proceed to prove Theorem 1. The proof consists of showing the following claims in sequence. It is not difficult to see that these three claims together imply Theorem 1.

1. For ex post ROI-constrainted bidders, if a mechanism is DSIC, then the allocation rule must be monotonically non-decreasing.
2. Any monotonically non-decreasing allocation rule $x_i(\cdot)$ with the payment rule given in (4) produces a DSIC mechanism.
3. Given any monotone allocation rule $x(\cdot)$, the payment rule $p(\cdot)$ such that (x, p) is DSIC, if exists, must be unique.

The analyses of steps (1) and (3) are very similar to the proof of the original Myerson's Lemma, and we defer the details to Appendices A.2 and A.3. Next, we prove step (2). When clear from context, we will drop the subscript i in $x_i(\cdot), p_i(\cdot), u_i(\cdot)$ and M_i as shorthand in the following proofs.

Proof of Step (2). Consider a bidder i with private valuation v and fix the other bids $\mathbf{b}_{-i}$. We examine the utilities of bidder i when she bids her true valuation and when she bids some different value $v' \neq v$. Consider two cases.

- When $v' < v$, we have $\max_{0\le z\le v}\{M\widetilde{p}(z) - zx(z)\} \ge \max_{0\le z\le v'}\{M\widetilde{p}(z) - zx(z)\}$, which implies
$$p(v) - p(v') \le M(\widetilde{p}(v) - \widetilde{p}(v')).$$
  This inequality effectively removes the max term in the payment formula (4) and reduces the problem to that with the Myerson payment. This allows us to apply the standard argument for the Myerson auction to show the DSIC property of our mechanism. We show the analysis below for completeness.
  We can compute the utility difference of bidder i when she bids v and v', and get
$$\begin{aligned} u(v,v) - u(v',v) &= (Mvx(v) - p(v)) - \big(Mvx(v') - p(v')\big) \\ &\ge M(vx(v) - \widetilde{p}(v)) - M(vx(v') - \widetilde{p}(v')) \\ &= M\int_0^v x(z)\,\mathrm{d}z - M\left(vx(v') - v'x(v') + \int_0^{v'} x(z)\,\mathrm{d}z\right) \\ &= M\left(\int_{v'}^{v} x(z)\,\mathrm{d}z - (v-v')x(v')\right) \\ &\ge M\left(\int_{v'}^{v} x(v')\,\mathrm{d}z - (v-v')x(v')\right) = 0, \end{aligned}$$
  where the second equality is by plugging in the Myerson payment formula (3). This means bidder i has no incentive to misreport her valuation v as v' in this case.
- When $v' > v$, we examine $\max_{0\le z\le v}\{M\widetilde{p}(z) - zx(z)\}$ and $\max_{0\le z\le v'}\{M\widetilde{p}(z) - zx(z)\}$. There are two possibilities:
  - If these two terms are equal, then we can apply the same argument as in the previous case (and also as in the Myerson auction analysis) to prove the DSIC property. We omit the details here.
  - If $\max_{0\le z\le v}\{M\widetilde{p}(z) - zx(z)\} < \max_{0\le z\le v'}\{M\widetilde{p}(z) - zx(z)\}$, this means $\arg\max_{0\le z\le v'}\{M\widetilde{p}(z) - zx(z)\} = v^* > v$. Then at valuation v', we should have
$$\begin{aligned} p(v') &= M\widetilde{p}(v') - (M\widetilde{p}(v^*) - v^*x(v^*)) \\ &= M\left(v'x(v') - v^*x(v^*) - \int_{v^*}^{v'} x(z)\,\mathrm{d}z\right) + v^*x(v^*) \\ (\text{replace } M \text{ by } 1) &\ge v'x(v') - \int_{v^*}^{v'} x(z)\,\mathrm{d}z \\ &\ge v'x(v') - \int_{v^*}^{v'} x(v')\,\mathrm{d}z \\ &= v'x(v') - (v'-v^*)x(v') \\ &= v^*x(v') > vx(v'). \end{aligned}$$

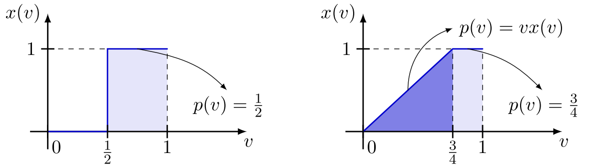

(a) The Myerson auction

(b) The optimal auction for ROI-constrained bidders

Fig. 1. The Myerson auction and the optimal auction for one bidder with uniform value distribution over $[0,1]$ and $M = 2$.

That is to say, when reporting v', the payment of bidder i will be greater than the value she obtains (which is $vx(v')$), therefore violating the IR condition. So bidder i has no incentive to misreport as v' in this case. $\square$

Furthermore, we demonstrate that this characterization can be directly adapted to *Bayesian incentive compatible* (BIC) mechanisms. A mechanism is called BIC if for every bidder i and valuation v_i, truthful bidding will lead to an optimal expected utility, where the expectation is with respect to the prior value distribution of others $\mathbf{v}_{-i}$. The rationale for this extension is anchored in the nature of the proof for Theorem 1, which notably does not rely on any terms associated with the values of other bidders. Consequently, it is feasible to adapt the allocation and payment terms delineated in Theorem 1 by considering them in expectation, relative to the prior value distribution of other bidders. It is important to note that the ex post ROI constraint imposes a hard feasibility requirement. This ensures that the ROI condition is satisfied for every realization, thereby precluding the $-\infty$ utility outcome and allowing a natural generalization to the Bayesian setting, thereby broadening the applicability of our findings. We defer the proof of this lemma to Appendix A.4.

Lemma 2 (Characterization of BIC Mechanisms). *For ex post ROI-constrained bidders, a mechanism (x,p) is Bayesian Incentive Compatible (BIC) if and only if for every bidder i:*

- *The interim allocation rule $X_i(v_i) \triangleq \mathbb{E}_{v_{-i}}[x_i(v_i, v_{-i})]$ is monotonically non-decreasing;*
- *The interim payment rule $P_i(v_i) \triangleq \mathbb{E}_{v_{-i}}[p_i(v_i, v_{-i})]$ satisfies:*

$$P_i(v_i) = M_i \tilde{P}_i(v_i) - \max_{0 \le z \le v_i} \left\{ M_i \tilde{P}_i(z) - z X_i(z) \right\},$$

where $\tilde{P}_i(v_i) = v_i X_i(v_i) - \int_0^{v_i} X_i(z) dz$ is the standard interim Myerson payment associated with the interim allocation $X_i(\cdot)$.

4. Optimal auction design for a single bidder

Having obtained the precise characterization of the allocation rule and payment function in the setting with ROI constraints, we now turn to the revenue maximization auction design. Recall that in the Myerson auction [10] and previous works in the ex ante ROI constraints setting [4,6], the revenue maximization problem is reduced to the problem of maximizing (modified) virtual welfare. Unfortunately, with ex post ROI constraints, the payment function characterization (4) involves an additional $\max$ term compared to the Myerson payment, and it is unclear how to incorporate this term into a modified virtual valuation formulation. We present the following simple example with a single bidder to demonstrate that, unlike the Myerson auction, the allocation that maximizes the virtual welfare may no longer be optimal with ROI-constrained bidders.

Example 1. Consider selling a single item to a single bidder with ROI constraint $M = 2$ and valuation for the item v following a uniform distribution $U[0,1]$. If we disregard the ROI constraint (*i.e.*, let $M = 1$), the virtual valuation of this bidder is $\phi(v) = 2v - 1$, and the optimal Myerson auction, as shown in Fig. 1a, sells the item at price $p = \frac{1}{2}$ with the expected revenue of $\frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}$.

However, with the ROI constraint $M = 2$ in presence, this allocation rule ($x(v) = 0$ when $v < 1/2$ and $x(v) = 1$ otherwise) is no longer optimal. As shown in Fig. 1b, the optimal auction, which will be proved in Theorem 2 later in this section, is a randomized auction with the allocation rule given by

$$x(v) = \begin{cases} \frac{4}{3}v & \text{if } v \le \frac{3}{4} \\ 1 & \text{if } v > \frac{3}{4}. \end{cases}$$

This allocation rule would generate an expected revenue of $\frac{3}{8}$, which is higher than $\frac{1}{4}$.

This example already highlights an important feature of the optimal auction with ROI constraints: the allocation and payment may be randomized, even in the simple setting with a single bidder and uniform value distribution. It also suggests that it is difficult to follow the Myerson auction regime and reduce the revenue maximization problem to a welfare maximization problem with some modified virtual valuation. It seems a very challenging problem to obtain a characterization for the optimal auction in this setting. Instead, in this section, we focus on the special case when the item is sold to a single bidder. As we will show in the following analysis, this is already a nontrivial and interesting problem to design an optimal auction for a single bidder.

First, we show that with a single ROI-constrained bidder, the $\max$ term in the payment formula (4) would always be 0, reducing the payment rule (4) to the standard Myerson payment (multiplied by a factor of M).

Lemma 3. *In the optimal auction with a single ROI-constrained bidder, let (x, p) be the revenue-maximizing auction, then we have $M\widetilde{p}(v) \le v \cdot x(v)$ for any value $v \in [0, v_{\max}]$. In other words, $\max_{0\le z\le v}\{M\widetilde{p}(z) - zx(z)\} = 0$ (which is achieved at $z = 0$) for all $v \in [0, v_{\max}]$, and the payment rule reduces to $p(v) = M\widetilde{p}(v)$.*

Proof. We prove this statement by contradiction. We first let $v^* = \min\{v \mid v \in \arg\max_{0\le z\le v_{\max}}\{M\widetilde{p}(z) - zx(z)\}\}$. Suppose the claim is not true, then we know that there must be a violation at v^*, and, for all $0 \le v < v^*$, the difference $d(v, v^*)$ between $M\widetilde{p}(v^*) - v^*x(v^*)$ and $M\widetilde{p}(v) - vx(v)$ is positive:

$$\begin{aligned} d(v,v^*) &= (M\widetilde{p}(v^*) - v^*x(v^*)) - (M\widetilde{p}(v) - vx(v)) \\ &= M\left(v^*x(v^*) - vx(v) - \int_v^{v^*} x(z)\,\mathrm{d}z\right) - (v^*x(v^*) - vx(v)) \\ &= (M-1)(v^*x(v^*) - vx(v)) - M\int_v^{v^*} x(z)\,\mathrm{d}z > 0. \end{aligned} \tag{5}$$

Given an allocation $x(\cdot)$, our plan is to construct a new monotone allocation rule, which, together with the corresponding payment rule, can generate a higher revenue. More specifically, we select a sufficiently small $\delta > 0$ such that letting $d(v, v^*) = 0$ by increasing the allocation $x(v)$ in the interval $(v^* - \delta, v^*)$ does not break the monotonicity of the allocation rule (we will discuss why this is possible later). Then, the new allocation rule $\bar{x}(\cdot)$ is defined as

$$\bar{x}(v) = \begin{cases} x(v) & \text{if } v \le v^* - \delta \quad \text{or} \quad v \ge v^* \\ \left(\frac{v}{v^*}\right)^{\frac{1}{M-1}} x(v^*) & \text{if } v \in (v^* - \delta, v^*). \end{cases}$$

This way, for all values $v \in (v^* - \delta, v^*)$, with respect to $\bar{x}(\cdot)$, we have

$$\begin{aligned} d(v,v^*) &= (M-1)(v^*\bar{x}(v^*) - v\bar{x}(v)) - M\int_v^{v^*} \bar{x}(z)\,\mathrm{d}z \\ &= (M-1)\left(v^*x(v^*) - v\left(\frac{v}{v^*}\right)^{\frac{1}{M-1}} x(v^*)\right) - M\int_v^{v^*} \left(\frac{z}{v^*}\right)^{\frac{1}{M-1}} x(v^*)\,\mathrm{d}z \\ &= x(v^*)(M-1)\left(v^* - v\left(\frac{v}{v^*}\right)^{\frac{1}{M-1}}\right) - x(v^*)M\int_v^{v^*} \left(\frac{z}{v^*}\right)^{\frac{1}{M-1}}\mathrm{d}z \\ &= x(v^*)(M-1)\left(\frac{1}{v^*}\right)^{\frac{1}{M-1}}\left(v^{*\frac{M}{M-1}} - v^{\frac{M}{M-1}}\right) \\ &\quad - x(v^*)M\left(\frac{1}{v^*}\right)^{\frac{1}{M-1}}\int_v^{v^*} z^{\frac{1}{M-1}}\,\mathrm{d}z \\ &= 0. \end{aligned} \tag{6}$$

Then, we prove that $\bar{x}(v) > x(v)$ for all values $v \in (v^* - \delta, v)$. Assume by contradiction that for δ defined above, there exists some $v \in (v^* - \delta, v^*)$ such that $\bar{x}(v) \le x(v)$. Because we have assumed that both $\bar{x}(\cdot)$ and $x(\cdot)$ are right-differentiable and only have finite non-differentiable points, this means there must exist some $\delta' > 0$, such that $\bar{x}(v) \le x(v)$ holds for all $v \in (v^* - \delta', v^*)$. Pick an arbitrary v in this interval. We then have

$$\begin{aligned} d(v,v^*) &= (M-1)(v^*x(v^*) - vx(v)) - M\int_v^{v^*} x(z)\,\mathrm{d}z \\ &\le (M-1)(v^*x(v^*) - v\bar{x}(v)) - M\int_v^{v^*} \bar{x}(z)\,\mathrm{d}z \\ &= (M-1)(v^*\bar{x}(v^*) - v\bar{x}(v)) - M\int_v^{v^*} \bar{x}(z)\,\mathrm{d}z \\ &= 0, \end{aligned}$$

where the first equality is by Eq. (5) and the last equality is by Eq. (6). This is a direct contradiction to our previous claim that $d(v, v^*) > 0$ for all $v < v^*$. Therefore, we obtain that $\bar{x}(\cdot)$ remains non-decreasing because we only modify x in the interval of $(v^* - \delta, v^*)$ and we have: (1) $\bar{x}(v) > x(v)$ for all values $v \in (v^* - \delta, v^*)$; (2) $\bar{x}(v)$ is increasing in $(v^* - \delta, v^*)$; and (3) $\bar{x}(v)$ is continuous at v^*.

Finally, we analyze the revenue generated by $\bar{x}(\cdot)$. Let $q(\cdot)$ be the corresponding payment rule of $\bar{x}(\cdot)$ derived from Theorem 1. We compare $q(v)$ and $p(v)$ at each value v.

- When $v \le v^* - \delta$, we have $q(v) = p(v)$. This is because the allocation remains unchanged in the interval $[0, v^* - \delta)$, and the payment of a bid at value v only depends on the allocation at interval $[0, v]$.
- When $v \in (v^* - \delta, v^*)$, we always have $v \in \arg\max_{0\le z\le v}\{M\widetilde{p}(z) - z\bar{x}(z)\}$, which means the payment of $q(v)$ reduces to the first price, and we have $q(v) = v\bar{x}(v) > vx(v) \ge p(v)$.

- When $v \geq v^*$, we again have $q(v) = p(v)$. This is because when $v \geq v^*$, we have $v^* \in \arg\max_{0 \leq z \leq v}\{M\widetilde{p}(z) - zx(z)\}$ for both (x, p) and $(\bar{x}, q)$. Then the payment formulation of (4) reduces to

$$p(v) = M\widetilde{p}(v) - (M\widetilde{p}(v^*) - v^*x(v^*)) = v^*x(v^*) + M(\widetilde{p}(v) - \widetilde{p}(v^*)).$$

  This payment only depends on the allocation at interval $[v^*, v]$, which again $x(\cdot)$ and $\bar{x}(\cdot)$ agree on.

Combining these cases together, we have $\int_v q(v)f(v)\,\mathrm{d}v > \int_v p(v)f(v)\,\mathrm{d}v$. That is, the new mechanism $(\bar{x}, q)$ has a higher revenue compared to (x, p). This is a contradiction, and this completes the proof. □

With Lemma 3 at hand, it seems with a single bidder, we are back to the classic Myerson regime, where the revenue maximization problem can be converted to a welfare maximization problem with respect to the virtual valuation. That is, recall from the Myerson's theorem [10], we have

$$\mathsf{rev} = M\int_0^{v_{\max}} \phi(v)x(v)f(v)\,\mathrm{d}v,$$

where $\phi(v) = \left(v - \frac{1-F(v)}{f(v)}\right)$ is the *virtual valuation*. However, we still have the additional constraint that $M\widetilde{p}(v) \leq v \cdot x(v)$ for every v. This restricts our allocation space and turns the problem into a constrained welfare maximization problem.

In the following lemma, we provide some further characterizations of the structure of the optimal auction with a single ROI-constrained bidder.

Lemma 4. *With a single ROI-constrained bidder, there always exists a revenue-maximizing auction such that for any valuation v, at least one of the following statements holds:*

- *the derivative of allocation rule at the valuation v exists, and $x'(v) = 0$;*
- *the payment follows the first-price rule,* i.e.*, $p(v) = vx(v)$.*

Proof. Assume that in some revenue-maximizing auction (x, p), there exists a value $\hat{v}$ such that the two conditions claimed in the lemma do not hold. That is, we have $p(\hat{v}) < \hat{v}x(\hat{v})$, and $x'(\hat{v}) > 0$ or $x'(\hat{v})$ does not exist. Since $x(\cdot)$ is right-differentiable with finite non-differentiable points, by Theorem 1, we have $p(\cdot)$ is also continuous at all but a finite number of points in its domains. This means that, by letting $\underline{v} = \hat{v}$ we can always find an interval $[\underline{v}, \bar{v}]$, such that $p(v) < vx(v)$ and $x_i'(v) > 0$ for all valuations $v \in [\underline{v}, \bar{v}]$.

Recall that when $p(v) < vx(v)$ for all v, by Lemma 3, the payment reduces to $p(v) = M\widetilde{p}(v)$ and we can write the revenue of the mechanism as $\mathsf{rev} = M\int_0^{v_{\max}} \phi(z)x(z)f(z)\,\mathrm{d}z$. Next, we look at the virtual values $\phi(v)$ within this interval $[\underline{v}, \bar{v}]$. Our plan is to modify the allocation $x(v)$ in a subinterval of this interval based on the sign of $\phi(v)$ while maintaining the monotonicity of $x(\cdot)$ and $p(v) < vx(v)$ in the interval, and the expected revenue will (weakly) increase. We note that there always exists such a subinterval due to the right-hand differentiability of the allocation function.

We consider the following three cases:

- There exists an interval $[a, b] \subseteq [\underline{v}, \bar{v}]$ such that $\phi(v) > 0, \forall v \in [a, b]$. In this case, we define

$$\bar{x}(v) = \begin{cases} x(v) & \text{if } v \leq a \quad \text{or} \quad v \geq b \\ x(v) + \delta \cdot \dfrac{b - v}{b - a} & \text{if } v \in (a, b), \end{cases}$$

  where $\delta > 0$ is sufficiently small such that:
  1. $\bar{x}(\cdot)$ is still non-decreasing; and
  2. $p(v) < v\bar{x}(v)$ still holds in the interval $[\underline{v}, \bar{v}]$, which means the corresponding payment is still $p(v) = M\widetilde{p}(v)$ in this interval.

  Conditions (1) and (2) imply we can still write the expected revenue as

$$\mathsf{rev} = M\int_0^{v_{\max}} \phi(z)\bar{x}(z)f(z)\,\mathrm{d}z$$

  for the new mechanism with allocation $\bar{x}$. In the meanwhile, $\bar{x}(v)$ is point-wise larger than $x(v)$ at all values $v \in [a, b]$ where $\phi(v)$ is always positive, and outside this interval the two allocations remain the same. This means $\bar{x}(\cdot)$, together with its corresponding payment rule, would yield strictly higher revenue than the previous mechanism. A contradiction.
- There exists an interval $[a, b]$ in $[\underline{v}, \bar{v}]$ such that $\phi(v) < 0, \forall v \in [a, b]$. Similar to the first case, we define

$$\bar{x}(v) = \begin{cases} x(v) & \text{if } v \leq a \quad \text{or} \quad v \geq b \\ x(v) - \delta \cdot \dfrac{v - a}{b - a} & \text{if } v \in (a, b), \end{cases}$$

  where $\delta > 0$ is sufficiently small such that $\bar{x}(\cdot)$ is still non-decreasing and $p(v) < vx(v)$ still holds in the interval $[\underline{v}, \bar{v}]$. Using the same argument as in the previous case, we can again argue that $\bar{x}(\cdot)$ gives a higher revenue than $x(\cdot)$. Again a contradiction.
- If neither of the first two cases happens, we must have $\phi(v) = 0, \forall v \in [\underline{v}, \bar{v}]$. In this case, as long as the monotonicity of $x(\cdot)$ and $p(v) \leq vx(v)$ is maintained in the interval, any modification of $x(\cdot)$ would generate the same revenue. Therefore we can always find an allocation $\bar{x}(\cdot)$ that satisfies one of the required two conditions and have the same revenue as that of $x(\cdot)$. Therefore $\bar{x}(\cdot)$ still gives a revenue-maximizing auction.

This concludes the proof. □

Lemma 4 allows us to focus on auctions with a very specific structure: as long as the allocation is not constant, it always follows the first-price payment rule. In particular, combining with Lemma 3, it implies whenever $x'(v)$ exists and $x'(v) > 0$, we always have $p = vx(v) = M\widetilde{p}(v)$.

Next, we want to obtain a further characterization of the optimal auction with a single bidder. However, to do this would require us to make a mild assumption on the value distribution of the bidder, which is known as the *Decreasing Marginal Revenue (DMR)* condition by [29,30,32].

Definition 3 (Decreasing Marginal Revenue, DMR)**.** The value distribution of a bidder satisfies the condition of decreasing marginal revenue if and only if the function

$$\psi(v) \triangleq \phi(v) f(v) = v f(v) + F(v) - 1$$

is monotonically non-decreasing.

Note that $\psi(v)$ being non-decreasing is equivalent to the fact that $v \cdot (1 - F(v))$, which is the expected revenue of selling the item at price p, being concave, and this is where the name of this condition comes from. Intuitively, many commonly used distribution functions satisfy this assumption, *e.g.*, uniform distributions, and any distribution of finite support and monotone non-decreasing density.

For comparison, we also provide the widely used assumption of *regularity* proposed in [10].

Definition 4 (Regularity)**.** We assume the valuation distributions are regular, i.e., the virtual valuation functions $\phi_i(v_i)$ are monotonically non-decreasing.

The DMR condition is closely related to the regularity condition but they are incompatible.[3] We refer the reader to [30] for concrete examples and more discussion.

With the assumption of DMR, the optimal auction exhibits an even simpler structure than what is described in Lemma 4, namely that there exist only two intervals in the optimal auction: interval of $(0, D)$ with $x'(v) > 0$ and interval of $(D, v_{\max})$ with $x'(v) = 0$, where D is a threshold valuation between them. This leads to our main theorem in this section, which characterizes the optimal allocation rule and payment rule for a single ROI-constrained bidder. We also note that the DMR assumption of v_i is equivalent to the same assumption of t_i, and the proof is presented in Appendix A.5.

Theorem 2. *The optimal auction for a single ex post ROI-constrained bidder with a DMR value distribution over* $[0, v_{\max}]$ *is as follows:*

- *when $v < D$, the allocation is given by*

  $$x(v) = \left(\frac{v}{D}\right)^{\frac{1}{M-1}},$$

  and the payment follows the first-price rule, i.e.*,* $p(v) = vx(v)$*;*
- *when $v \geq D$, the allocation rule is $x(v) = 1$, and the payment is given by $p(v) = D$.*

Here D is a threshold valuation given by $\int_0^{D^*} \psi(v) v^{\frac{1}{M-1}}\, \mathrm{d}v = 0$.

This theorem provides an important insight that the optimal auction in the ROI-constrained setting is a randomized mechanism. Note that Myerson's optimal auction in the single-parameter setting is deterministic, but a decent body of work has shown that many generalizations to multi-parameter settings will lead to randomized optimal auctions [33,34]. The necessity of randomization is also known in some single-parameter environments, including public budget [11] and risk aversion [12,13]. Theorem 2 indicates that the slight generalization of ROI constraints of bidders will also lead to a randomized optimal auction. Moreover, we also remind the reader that DSIC and BIC are equivalent when there is a single agent.

We prove the theorem via the following steps. First, we derive the allocation of an optimal auction in an interval $(0, v^*)$ when $x'(v)$ is always positive in that interval. Then, we show in Lemma 5 that there exist only two intervals in the optimal auction: the interval of $(0, D)$ with $x'(v) > 0$ and the interval of $(D, v_{\max})$ with $x'(v) = 0$. Finally, we will compute the optimal threshold valuation D between these two intervals. The DMR assumption is used in the second and the last steps.

Proposition 2. *If for some $v^* \in (0, v_{\max}]$, we have $x'(v) > 0$ for all $v \in (0, v^*)$ in an optimal auction, then $x(\cdot)$ is continuous at v^*, and the allocation rule $x(v)$ for all $v \in [0, v^*]$ is given as:*

$$x(v) = \left(\frac{v}{v^*}\right)^{\frac{1}{M-1}} x(v^*).$$

Proof. We first assume $x(\cdot)$ is continuous at v^*, and we will prove later that, if it is discontinuous, we can improve the revenue without violating the DSIC property. By Lemmas 3 and 4, we get that for all valuations $v \in [0, v^*]$, $M\widetilde{p}(v) - vx(v) = 0$ always holds. That is,

$$M\left(vx(v) - \int_0^v x(z)\,\mathrm{d}z\right) - vx(v) = 0.$$

[3] The regularity condition is equivalent to the expected revenue being concave in the *quantile* space, while the DMR condition means the expected revenue is concave in the *value* space.

After transposition and derivation, this translates to

$$x(v) - (M-1)vx'(v) = 0.$$

By solving this differential equation with the value of $x(v^*)$ at valuation v^*, we can get

$$x(v) = \left(\frac{v}{v^*}\right)^{\frac{1}{M-1}} x(v^*), \quad \forall v \in [0, v^*]. \tag{7}$$

Next, if $x(\cdot)$ is discontinuous at v^*, we need to replace $x(v^*)$ in (7) with $x(v^{*-})$, *i.e.*, the left limit of $x(\cdot)$ at v^* (recall that we denote $x(v^*)$ as the right limit when it is discontinuous). Since $x(v^{*-}) < x(v^*)$, we can observe that directly using (7) as the allocation rule will increase the revenue without violating the DSIC property, which also makes $x(\cdot)$ continuous at v^*. This concludes the proof. □

Lemma 5. *For a single ROI-constrained bidder with DMR value distribution, if there exist intervals with $x'(v) = 0$ in an optimal auction, then there is exactly one such interval, and it appears in the highest value region.*

Proof. Assume by contradiction that there exist multiple intervals with $x'(v) = 0$ in an optimal auction. Pick $(\underline{v}, \bar{v})$ to be the first such interval. That is, we have $x'(v) > 0$ for all $v \in (0, \underline{v})$, $x'(v) = 0$ for all $v \in (\underline{v}, \bar{v})$, and $\bar{v} < v_{\max}$.

First, we must have $\underline{v} > 0$. That is, the allocation rule cannot start with a flat interval. To see why this is true, assume otherwise that $\underline{v} = 0$. We focus on a point $v' = \bar{v} + \delta$ in the next interval for a sufficiently small $\delta > 0$ such that $v' < M\bar{v}$. Then there are two cases: (1) $x'(v') = 0$ and (2) $x'(v')$ is strictly positive. In the first case, we will have $x(\bar{v}) > 0$, and the Myerson price at $\bar{v}$ will be $\widetilde{p}(\bar{v}) = \bar{v} \cdot x(\bar{v}) < M\widetilde{p}(\bar{v})$, which directly contradicts Lemma 3. In the second case, we look at the payment $p(v')$ at point v'. Note that since $x'(v) > 0$, by Lemmas 3 and 4, we should have $M\widetilde{p}(v') = v'x(v')$. But this cannot happen because

$$\begin{aligned} M\widetilde{p}(v') &= M\left(v'x(v') - \int_0^{v'} x(z)\,\mathrm{d}z\right) = M\left(v'x(v') - \int_{\bar{v}}^{v'} x(z)\,\mathrm{d}z\right) \\ &> M\left(v'x(v') - (v' - \bar{v})x(v')\right) = M\bar{v}x(v') > v'x(v'). \end{aligned}$$

Knowing $\underline{v} > 0$, by Proposition 2, we know $x(\cdot)$ is continuous at $\underline{v}$, and the allocation in $[0, \underline{v}]$ is given as:

$$x(v) = \left(\frac{v}{\underline{v}}\right)^{\frac{1}{M-1}} x(\underline{v}), \quad \forall v \in [0, \underline{v}].$$

Next, combined with $p(\bar{v}) = \bar{v}x(\bar{v})$, we have $M\widetilde{p}(\bar{v}) - \bar{v}x(\bar{v}) = M\widetilde{p}(\underline{v}) - \underline{v}x(\underline{v})$, that is,

$$x(\bar{v}) = \frac{M\bar{v} - \underline{v}}{(M-1)\bar{v}} \cdot x(\underline{v}). \tag{8}$$

In order to argue that allocation rule $x(\cdot)$ is not revenue-maximizing, we construct a new allocation rule as

$$\bar{x}(v) = \begin{cases} x(v) & \text{if } v \geq \bar{v} \\ \left(\frac{v}{\bar{v}}\right)^{\frac{1}{M-1}} x(\bar{v}) & \text{if } v \in [0, \bar{v}). \end{cases} \tag{9}$$

That is, we replace the first price interval $[0, \underline{v}]$ and the flat interval $[\underline{v}, \bar{v}]$ in $x(\cdot)$ with a single first-price interval $[0, \bar{v}]$ in $\bar{x}(\cdot)$. Comparing the two allocation rules $x(\cdot)$ and $\bar{x}(\cdot)$, we first note from Lemmas 3 and 4 that $p(\bar{v}) = M\widetilde{p}(\bar{v}) = \bar{v}x(\bar{v}) = \bar{v}\bar{x}(\bar{v})$, which indicates that

$$\int_0^{\bar{v}} x(z)\,\mathrm{d}z = \int_0^{\bar{v}} \bar{x}(z)\,\mathrm{d}z, \tag{10}$$

because both sides equal to $\bar{v}x(\bar{v})(M-1)/M$. Next, we have for $v \in [0, \underline{v}]$,

$$\begin{aligned} \bar{x}(v) - x(v) &= \left(\frac{v}{\bar{v}}\right)^{\frac{1}{M-1}} x(\bar{v}) - \left(\frac{v}{\underline{v}}\right)^{\frac{1}{M-1}} x(\underline{v}) \\ &= v^{\frac{1}{M-1}} \cdot x(\underline{v}) \cdot \left(\left(\frac{1}{\underline{v}}\right)^{\frac{1}{M-1}} - \left(\frac{1}{\bar{v}}\right)^{\frac{1}{M-1}} \cdot \frac{M\bar{v} - \underline{v}}{(M-1)\bar{v}}\right). \end{aligned}$$

Since v only appears in the first term, $\bar{x}(v) - x(v)$ must be constantly positive or negative in $(0, \underline{v}]$, determined by the last term. Combined with (10) and $x'(v) = 0, \forall v \in (\underline{v}, \bar{v})$, we know it is constantly negative, *i.e.*, $\bar{x}(v) < x(v)$ for all $v \in (0, \underline{v}]$. Therefore, there exists a threshold $v^* \in (\underline{v}, \bar{v})$ such that $x(v) > \bar{x}(v)$ for all $v \in [0, v^*)$ and $x(v) \leq \bar{x}(v)$ for all $v \in [v^*, \bar{v})$. Combining this with Eq. (10), we have

$$\int_0^{v^*} (x(z) - \bar{x}(z))\,\mathrm{d}z = \int_{v^*}^{\bar{v}} (\bar{x}(z) - x(z))\,\mathrm{d}z > 0. \tag{11}$$

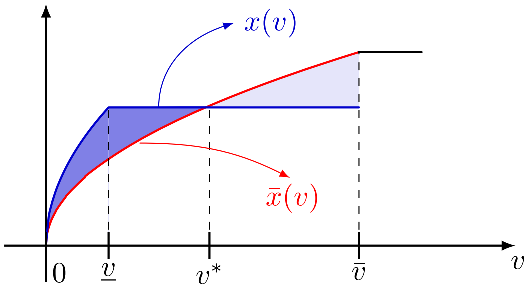

Fig. 2. An illustration for the proof of Lemma 5 with $M = 3$. The blue line denotes an allocation rule $x(\cdot)$ with $x'(v) = 0$ in $(\underline{v}, \bar{v})$ where $\bar{v} < v_{\text{max}}$, and $x(v)$ for $v \in [0, \underline{v})$ is computed by Proposition 2. The red line denotes our constructed allocation rule $\bar{x}(\cdot)$ as given in (9). The point v^* denotes the intersection of $x(\cdot)$ and $\bar{x}(\cdot)$. By Eq. (11) we have the areas of the two shadowed regions are the same. Furthermore, as $\psi(\cdot)$ is non-decreasing, we can conclude that $\bar{x}(\cdot)$ leads to a higher revenue than $x(\cdot)$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Next, for any allocation rule $x(\cdot)$, we denote

$$\mathsf{rev}_{[0,\bar{v}]}^{x(\cdot)} = M \int_0^{\bar{v}} \psi(z)x(z)\,\mathrm{d}z,$$

and we can compare the revenue generated from $x(\cdot)$ and $\bar{x}(\cdot)$ in the interval $[0, \bar{v}]$:

$$\mathsf{rev}_{[0,\bar{v}]}^{x(\cdot)} - \mathsf{rev}_{[0,\bar{v}]}^{\bar{x}(\cdot)} = M\left(\int_0^{v^*} \psi(z)(x(z) - \bar{x}(z))\,\mathrm{d}z - \int_{v^*}^{\bar{v}} \psi(z)(\bar{x}(z) - x(z))\,\mathrm{d}z\right).$$

Finally, we can see that this difference is always negative due to Eq. (11) and the fact that $\psi(\cdot)$ is non-decreasing by the DMR condition.[4] Fig. 2 demonstrates the idea of this argument.

Therefore, we conclude that $\bar{x}(\cdot)$ generates a higher revenue than $x(\cdot)$, contradicting the fact that (x, p) is an optimal auction. This completes the proof. □

We now proceed to complete the proof of Theorem 2.

Proof of Theorem 2. By Lemma 5, we know $x'(v) > 0$ for all valuations $v \in (0, D)$ and $x'(v) = 0$ for all valuations $v \in (D, v_{\text{max}})$. First, we have $x(D) = 1$, since otherwise, the revenue could be improved by setting $x(v) = 1$ for all $v \in [D, v_{\text{max}}]$. Next, we find the optimal threshold D that maximizes the overall revenue. By Proposition 2, the allocation rule for valuations $v \in [0, D]$ is

$$x(v) = \left(\frac{v}{D}\right)^{\frac{1}{M-1}} \cdot x(D) = \left(\frac{v}{D}\right)^{\frac{1}{M-1}}, \forall v \in [0, D].$$

And the overall revenue is

$$\mathsf{rev} = M\left(\int_0^D \psi(v)\left(\frac{v}{D}\right)^{\frac{1}{M-1}} \mathrm{d}v + \int_D^{v_{\text{max}}} \psi(v)\,\mathrm{d}v\right).$$

In order to maximize this revenue, we compute its derivative

$$\frac{\mathrm{d}\mathsf{rev}}{\mathrm{d}D} = -\frac{M}{M-1} \cdot D^{-\frac{M}{M-1}} \cdot \int_0^D \psi(v)v^{\frac{1}{M-1}}\,\mathrm{d}v. \tag{12}$$

Looking at this derivative, we see that the term $-\frac{M}{M-1} \cdot D^{-\frac{M}{M-1}}$ is always negative and $v^{\frac{1}{M-1}}$ is always positive. Furthermore, we have $\psi(0) = -1 < 0$ and $\psi(v_{\text{max}}) > 0$. By the DMR condition, $\psi(\cdot)$ is non-decreasing. Then, by integration by parts, we have

$$\begin{aligned}\int_0^{v_{\text{max}}} \psi(v)v^{\frac{1}{M-1}}\,\mathrm{d}v &= \int_0^{v_{\text{max}}} (vf(v) + F(v) - 1)v^{\frac{1}{M-1}}\,\mathrm{d}v \\ &= v(F(v) - 1) \cdot v^{\frac{1}{M-1}}\Big|_0^{v_{\text{max}}} + \frac{1}{M-1}\int_0^{v_{\text{max}}} (v(1 - F(v))v^{\frac{2-M}{M-1}}\,\mathrm{d}v.\end{aligned}$$

One can hence observe that the first term is 0 and the last term is constantly positive, that is, at $D = v_{\text{max}}$, the derivative is negative. Therefore, by letting $\int_0^{D^*} \psi(v)v^{\frac{1}{M-1}}\,\mathrm{d}v = 0$, the revenue will increase with D until D^* and then decrease. This indicates that the optimal solution is achieved at $D = D^*$, completing the proof. □

[4] We omit an ill-defined special case where $\psi(v) = 0, \forall v \in [0, \bar{v}]$, since in such case $\psi(v)$ will be infinity when $v \to 0$.

Next, to illustrate the necessity of the DMR condition for the characterization in Theorem 2, we construct a counterexample using a discrete bimodal distribution that violates the DMR condition. In this setting, the optimal mechanism deviates from the structure described in Theorem 2 and achieves strictly higher revenue.

Example 2. Consider a single bidder with an ROI constraint $M = 2$. The bidder's valuation v follows a discrete distribution with probability mass concentrated at two points: low value $v_L = 1$ with probability 0.5 and high value $v_H = 10$ with probability 0.5. This distribution violates the DMR condition due to the significant gap between v_L and v_H where the probability density is zero. We compare the revenue of the mechanism constrained by the structure in Theorem 2 against the true optimal mechanism.

Mechanism under Theorem 2 Structure: Theorem 2 states that under DMR, the optimal allocation rule $x(v)$ must be strictly increasing (specifically, $x(v) = (v/D)^{\frac{1}{M-1}}$) until it reaches saturation ($x(v) = 1$). With $M = 2$, this implies a linear allocation rule $x(v) = v/D$ for $v < D$. To extract maximum revenue from the high-value bidder, the optimal threshold is $D = 10$. Consequently, for the high type ($v_H = 10$): $x_H = 1$ and $p_H = 10$ (binding ROI); For the low type ($v_L = 1$): the structure forces $x_L = v_L/D = 0.1$, with payment $p_L = v_L x_L = 0.1$. The expected revenue is:

$$\mathsf{rev}_{\text{thm7}} = 0.5 \times 0.1 + 0.5 \times 10 = \mathbf{5.05}.$$

True Optimal Mechanism: Without the DMR condition, the optimal mechanism is not restricted to the strictly increasing form. We can improve revenue by increasing the allocation for the low type (x_L) while maintaining DSIC and IR. Let $x_H = 1$ and $p_H = 10$. We maximize x_L subject to the DSIC constraint of the high type:

$$u_H(v_H) \geq u_H(v_L) \implies M v_H x_H - p_H \geq M v_H x_L - p_L.$$

Substituting the values ($M = 2, v_H = 10, p_L = x_L$):

$$2 \cdot 10 \cdot 1 - 10 \geq 2 \cdot 10 \cdot x_L - x_L \implies 10 \geq 19 x_L \implies x_L \leq \frac{10}{19} \approx 0.526.$$

By setting $x_L = 10/19$ and $p_L = 10/19$, the expected revenue becomes:

$$\mathsf{rev}_{\text{opt}} = 0.5 \times \frac{10}{19} + 0.5 \times 10 \approx \mathbf{5.263}.$$

In this example, the true optimal mechanism yields a revenue of approximately 5.263, which is strictly higher than the 5.05 obtained under the restrictions of Theorem 2. This demonstrates that without the DMR condition, the optimal allocation rule may exhibit a "staircase" shape (flat intervals) rather than the strictly increasing behavior characterized in Theorem 2. Thus, the DMR condition is necessary for the result.

5. Auction design for multiple bidders

As we have discussed, the optimal auction with ROI-constrained bidders may involve fractional allocations in the simple setting of a single item being sold to a single bidder. In other words, when the item is indivisible, a probabilistic allocation is necessary to implement the optimal auction. On the one hand, this implies that any extension will probably be very technical; on the other hand, this is, unfortunately, often undesired in practical scenarios such as online advertising due to the volatility and uncertainty of the auction outcome. Therefore, for the purpose of practical applications, in this section, we investigate simple auctions and, in particular, deterministic auctions for multiple ROI-constrained bidders.

First, we would like to introduce an important observation, which shows that any "simple v.s. optimal" mechanisms for traditional bidders in previous literature [35,36] could also be generalized for ROI-constrained bidders in our setting, with a modified approximation ratio. In particular, this result for ROI-constrained bidders can be interpreted as an application of the general reduction framework for non-linear agents developed by [37].

Theorem 3. *In the setting with multiple traditional bidders with quasilinear utility functions, assume that a mechanism $\mathcal{M} = (x(\cdot), p(\cdot))$ is DSIC and IR, and it achieves an approximation ratio of R; then, in our adapted setting with multiple ROI-constrained bidders, there exists a corresponding mechanism $\widehat{\mathcal{M}}$ which is DSIC and IR, and it achieves an approximation ratio of no worse than $M^{\max} \cdot R$, where $M^{\max}$ denotes the maximal ROI constraint among all bidders.*

Proof. To make the proof clear, for the group of ROI-constrained bidders in an auction, we construct a corresponding group of traditional bidders by simply setting their ROI constraints as 1, and then we apply different mechanisms to these two groups. In the following, for each mechanism we consider, the notation of revenue rev is understood to be computed from the payments of the type of bidders for which the mechanism is designed (traditional bidders or ROI-constrained bidders).

First, let the allocation rule of $\widehat{\mathcal{M}}$ be the same as $\mathcal{M}$, *i.e.*, $\hat{x}(\cdot) = x(\cdot)$, and its payment rule $\hat{p}(\cdot)$ be computed following Theorem 1. It is straightforward that $\widehat{\mathcal{M}}$ is DISC and IR because the allocation rule satisfies monotonicity for all bidders. Next, we denote the optimal mechanism for the group of ROI-constrained bidders as $\mathcal{M}_{\text{opt}} = (x_{\text{opt}}, p_{\text{opt}})$, and the optimal mechanism for the group of traditional bidders (*i.e.*, Myerson auction) as $\mathcal{M}_{\text{mye}} = (x_{\text{mye}}, p_{\text{mye}})$. Then we can define a mechanism $\bar{\mathcal{M}}$ for traditional bidders with an allocation rule $\bar{x}(\cdot) = x_{\text{opt}}(\cdot)$ and a corresponding payment rule $\bar{p}(\cdot)$ following Myerson's Lemma. Next, for the group of traditional bidders, we can observe the relation among their revenue: $\mathsf{rev}_{\bar{\mathcal{M}}} \leq \mathsf{rev}_{\mathcal{M}_{\text{mye}}} \leq R \cdot \mathsf{rev}_{\mathcal{M}}$. Also, we have that $\mathsf{rev}_{\mathcal{M}} \leq \mathsf{rev}_{\widehat{\mathcal{M}}}$ because $\mathcal{M}$ and $\widehat{\mathcal{M}}$ have the same allocation rule, and hence, the corresponding payment of $\mathcal{M}$ for each bidder is lower or equal to $\widehat{\mathcal{M}}$, given by the third argument of Proposition 1. Finally, the first argument of Proposition 1 tells us that the payment of a bidder i under $\mathcal{M}_{\text{opt}}$ is no more than M_i times that under $\bar{M}$. This leads to $\mathsf{rev}_{\mathcal{M}_{\text{opt}}} \le M^{\max} \cdot \mathsf{rev}_{\bar{\mathcal{M}}}$, where $M^{\max}$ denotes the maximal ROI constraint among all bidders. Combining all the above inequalities, we can conclude the proof. □

Based on this theorem, we can immediately have the following observation.

Proposition 3. *For ex post ROI-constrained bidders, simply adopting the allocation rule of Myerson auction regardless of the ROI constraints leads to an $M^{\max}$-approximately optimal mechanism.*

Moreover, simple mechanisms, such as VCG with monopoly reserves in previous literature [35], could also guarantee DSIC and IR, and achieve corresponding approximation ratios for ROI-constrained bidders [37]. Additionally, following the framework of [38,39], it might be possible to design an adaptive sequential mechanism that achieves a constant approximation ratio when selling a single item to multiple bidders. However, this framework requires the characterization of the optimal single-bidder mechanism under a given upper-bound on its expected allocation probability (Theorem 1 in [38]). This characterization is nontrivial in our setting, as it may no longer possess the simple two-interval structure. Therefore, we leave the design of such an approximate multi-bidder mechanism as our future work.

Next, we consider a particular class of simple mechanisms, *i.e.*, deterministic auctions. A deterministic auction is a mechanism where the allocation of items and the corresponding payments are exclusively determined by the bidders' submitted bids, with no element of randomness. We now present the optimal deterministic auctions for a single item and multiple ROI-constrained bidders. In this setting, the class of deterministic DSIC mechanisms for traditional agents is the same as the class of deterministic DSIC mechanisms for ROI-constrained agents. Also, we will show that bidders in this setting have the identical preference order of outcomes as that of the bidders with standard quasilinear utility functions. Therefore, the optimal deterministic auction for ex post ROI-constrained bidders has the same structure as the Myerson auction.

Theorem 4. *The optimal deterministic auction for a single item and multiple ex post ROI-constrained bidders with regular value distributions is the same as the Myerson auction [10]. That is, the item is allocated to the bidder i with the highest positive virtual value $\phi_i(v_i)$, and the payment is $\phi_i^{-1}(\max\{\phi_j(v_j), 0\})$ where j is the bidder with the second-highest virtual value.*

Proof. By Theorem 1, the allocation for a bidder must be monotone in her value v_i. Therefore, for a deterministic mechanism with a single item, there must be a threshold value v^*, below which nothing is allocated, *i.e.*, $x_i(v_i) = 0$, and above which the slot is allocated to her, *i.e.*, $x_i(v_i) = 1$. Thus, for any two feasible outcomes in any mechanism, denoted as $o_i = (x_i(v_i), p_i(v_i))$ and $\hat{o}_i = (\hat{x}_i(v_i), \hat{p}_i(v_i))$, we can show that the preference orders of a bidder with ROI-constrained utility function are identical to that with the quasilinear utility function. We discuss this argument in the following cases[5]:

- if $x_i(v_i) = \hat{x}_i(v_i) = 1$, or $x_i(v_i) = \hat{x}_i(v_i) = 0$, then the preference orders of the outcomes are both inverted from the order of payments for bidders with both utility functions;
- if $x_i(v_i) = 1$ and $\hat{x}_i(v_i) = 0$, then $u_i(o_i) \ge 0 \ge u_i(\hat{o}_i)$ for bidders with both utility functions;
- if $x_i(v_i) = 0$ and $\hat{x}_i(v_i) = 1$, then $u_i(o_i) \le 0 \le u_i(\hat{o}_i)$ for bidders with both utility functions.

In summary, the preference order is always identical for these two types of utility functions for deterministic auctions with a single item. In addition, the feasible regions of allocation and payment rule are also the same for both ROI-constrained advertisers and traditional advertisers. Thus, each truthful mechanism for traditional advertisers is also a truthful mechanism for ROI-constrained advertisers with the same revenue, and vice versa. On top of all the above arguments, the optimal auctions for ROI-constrained advertisers are identical to the celebrated Myerson auction. □

This theorem seems to suggest that the optimal deterministic auction design may reduce to the classic Myerson auction when bidders have ex post ROI constraints. However, it turns out not to be the case when we go beyond the single-item setting. Next, we demonstrate this point in the environment of online sponsored search auctions.

6. Sponsored search auctions

In this section, we investigate sponsored search auctions with S ad slots. A sponsored search auction can be seen as a special case of multi-item environments. In a sponsored search auction, each bidder bids a single value b_i that represents their "value per click" v_i, and the auction determines the allocation of these S slots to bidders and their corresponding payments p_i. Note that an ROI-constrained bidder still has an "initial value" $t_i = M_i \cdot v_i$ as in our previous definitions. Each slot s has a click-through rate (CTR) c_s denoting the probability of the slot being clicked after it is allocated to a bidder. With the notations, we can get the utility of a bidder i with a slot s allocated to her:

$$u_i = \begin{cases} M_i v_i c_s - p_i & \text{if } v_i c_s \ge p_i \\ -\infty & \text{otherwise.} \end{cases} \tag{13}$$

In a sponsored search, the CTR for each slot decreases from top to bottom, *i.e.* $c_1 < c_2 < \ldots < c_S$, and each bidder is allocated at most one slot. Without loss of generality, we assume $c_S = 1$ and $S < n$, and we also define a dummy slot 0 with $c_0 = 0$ for notation

[5] Note that we only need to discuss the cases where the IR property is satisfied.

convenience. For sponsored search auctions, a *deterministic* auction is a mechanism where the allocation of slots and the corresponding payments are exclusively determined by the bidders' submitted bids, with no element of randomness.

First, we show that our DSIC auction characterization of Theorem 1 reduces to the following form in the deterministic sponsored search auction setting.

Proposition 4. *In a DSIC deterministic sponsored search auction with S ad slots and ROI-constrained bidders, if we focus on an arbitrary bidder i and fix the other bids $\mathbf{b}_{-i}$, we have*

- *the allocation rule could be represented as a piecewise function*

$$x_i(\cdot) = \left\{ \left([0, v_{i,(1)}), c_0\right), \left([v_{i,(1)}, v_{i,(2)}), c_1\right), \left([v_{i,(2)}, v_{i,(3)}), c_2\right), \dots, \left([v_{i,(S)}, v_{\max}], c_S\right) \right\},$$

  where $v_{i,(1)} \le v_{i,(2)} \le \dots \le v_{i,(S)}$, and $\left([v_{i,(s)}, v_{i,(s+1)}), c_s\right)$ means that slot s is allocated to bidder i when she bids $b_i \in [v_{i,(s)}, v_{i,(s+1)})$;
- *the payment for slot s is given by*

$$p_{i,(s)} = \min\left\{p_{i,(s-1)} + M_i v_{i,(s)}(c_s - c_{s-1}), v_{i,(s)} c_s\right\}.$$

Proof. The proof of the piecewise allocation rule is straightforward by Theorem 1. We next provide proof of the payment rule. For notation simplicity, we omit the subscript i in the following. First, for all values $v \in [v_{(s-1)}, v_{(s)})$, we have

$$M\widetilde{p}(v) - vx(v) \le M\widetilde{p}(v_{(s-1)}) - v_{(s-1)}x(v_{(s-1)}), \tag{14}$$

because $vx(v)$ increases and $\widetilde{p}(v)$ remains unchanged when $v \in [v_{(s-1)}, v_{(s)})$ increases. Next, we denote

$$v_{(s^*)} = \underset{0 \le z \le v_{(s-1)}}{\arg\max}\, \{M\widetilde{p}_i(z) - zx(z)\},$$

then it suffices to discuss the following two cases:

- When $v_{(s^*)} = v_{(s-1)}$, combining with (14), we have that

$$\underset{0 \le z < v_{(s)}}{\arg\max}\{M\widetilde{p}(z) - zx(z)\} = \underset{0 \le z \le v_{(s-1)}}{\arg\max}\, \{M\widetilde{p}(z) - zx(z)\} = v_{(s-1)}.$$

  Therefore

$$\begin{aligned} p_{(s)} &= M\widetilde{p}(v_{(s)}) - \max_{z \in \{v_{(s)}, v_{(s-1)}\}} \{M\widetilde{p}(z) - zx(z)\} \\ &= \min\left\{p_{(s-1)} + Mv_{(s)}(c_s - c_{s-1}), v_{(s)}c_s\right\}. \end{aligned}$$

- When $v_{(s^*)} < v_{(s-1)}$, combining with (14), we have that

$$\underset{0 \le z < v_{(s)}}{\arg\max}\{M\widetilde{p}(z) - zx(z)\} = v_{(s^*)},$$

  and

$$p_{(s-1)} = M\widetilde{p}(v_{(s-1)}) - (M\widetilde{p}(v_{(s^*)}) - v_{(s^*)}x(v_{(s^*)})).$$

  Therefore

$$\begin{aligned} p_{(s)} &= M\widetilde{p}_i(v_{(s)}) - \max_{z \in \{v_{(s)}, v_{(s^*)}\}} \{M\widetilde{p}(z) - zx(z)\} \\ &= \min\left\{p_{(s-1)} + Mv_{(s)}(c_s - c_{s-1}), v_{(s)}c_s\right\}. \end{aligned}$$

  This concludes the proof. □

This proposition tells us that a DSIC deterministic auction will set allocation thresholds for each slot, and set the payment as the minimum of the following two terms: 1) the "ladder-style" payment computed by adding a ladder of payment (multiplied by M_i) to the price of the closest lower slot, and 2) the critical value for obtaining the slot. We note that this proposition was also shown in [22] but only for a special case where bidders all have identical ex post ROI constraints. Our proposition 4, which comes directly from our characterization Theorem 1, already generalizes their result and allows for diverse ROI constraints of bidders.

With this proposition at hand, we proceed to the optimal deterministic sponsored search auctions. Again, the min term in the payment rule in Proposition 4 poses critical challenges in the ROI-constrained setting, similar to the optimal randomized auction design in Section 4. Thus, we turn to the optimal deterministic auction design for a single ROI-constrained bidder. We first use a simple example to illustrate the structure of optimal deterministic auctions. For notation simplicity, we omit the subscript i in the following.

Example 3. Assume there are two slots with CTRs $c_1 = 0.5$ and $c_2 = 1.0$, and one bidder with uniform value distribution over $[0, 1]$ and $M = 2$. The optimal deterministic auction is illustrated in Fig. 3. When $v < \frac{2}{7}$, the allocated CTR is 0; when $\frac{2}{7} \le v < \frac{4}{7}$, the allocated CTR is $x(v) = c_1 = 0.5$, and the payment is $p(v) = \frac{1}{7}$; when $v \ge \frac{4}{7}$, the allocated CTR is $x(v) = c_2 = 1.0$, and the payment is $p(v) = \frac{4}{7}$. This optimal deterministic auction generates revenue of 0.286.

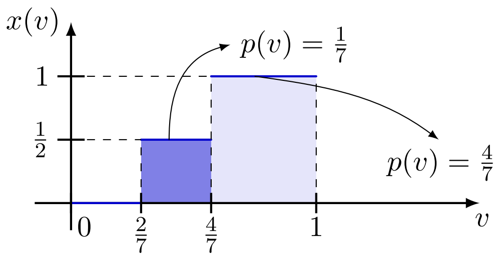

Fig. 3. The optimal deterministic auction for two slots with CTRs $c_1 = 0.5$ and $c_2 = 1.0$ and one bidder with uniform value distribution over $[0, 1]$ and $M = 2$.

We note that the Myerson auction and the optimal randomized auction for this two-slot setting are the same as the ones illustrated in Fig. 1, because the bidder is allocated to at most one slot in sponsored search auctions, and the highest slot has a CTR of 1.0. One can observe from this example that in contrast to the Myerson auction where the bidder is allocated either the highest slot or nothing, the optimal deterministic auction for an ROI-constrained bidder offers options of buying each slot from 1 to S. This phenomenon is also known as the *second-degree price discrimination* from [40]. Moreover, this example also demonstrates that the revenue of the optimal deterministic auction suffers a non-negligible loss compared with the optimal randomized auction (0.286 v.s. 0.375).

Next, we show that one can compute the optimal deterministic auction in polynomial time when the value space is discrete.

Theorem 5. *The optimal deterministic sponsored search auction with multiple slots and a single bidder can be computed in polynomial time.*

Proof. We will present a dynamic programming algorithm for computing the optimal auction. Let the value distribution $F(\cdot)$ of the bidder be finitely supported on a set $\{v^{(1)}, v^{(2)}, \dots, v^{(d)}\}$, with $v^{(1)} < v^{(2)} < \dots < v^{(d)}$. Denote $OPT(s, j)$ as the maximum revenue generated from only selling to values at most $v^{(j)}$, and only selling slots no higher than slot s. Denote $p_{opt}(s, j)$ as the price of the slot s in the mechanism of $OPT(s, j)$, which is updated synchronously with $OPT(s, j)$. We get the following recurrence relation:

$$OPT(s, j) = \max_{k \le j} \left\{ OPT(s-1, k) + p_{opt}(s-1, k)\left(F(v^{(j)}) - F(v^{(k)})\right) \right\},$$

If we denote k^* as the optimal k in the above equation, then $p_{opt}(s, j)$ can be updated as

$$p_{opt}(s, j) = \min\left(p_{opt}(s-1, k^*) + M v^{(k^*)}\left(c_s - c_{s-1}\right),\ v^{(k^*)} c_s \right).$$

The starting values are set as $OPT(0, j) = 0$ and $p_{opt}(0, j) = 0$ with $c_0 = 0$. Finally, $\max_{1 \le j \le d} OPT(S, j)$ gives the optimal revenue and its corresponding allocation rule can be computed via traceback. □

7. Conclusion

In this paper, we discuss optimal auction design for bidders who have ex post ROI constraints. We provide characterizations for DSIC auctions and optimal auctions with a single bidder in this setting. We show that the optimal auction may entail a randomized allocation scheme in the setting with ROI-constrained bidders. We also provide several results on simple auction design for multiple ex post ROI-constrained bidders.

There are several important open questions left in this model. The first and foremost one is to characterize the optimal auction (or approximately optimal auction) with a single item and multiple ROI-constrained bidders. As we have discussed in the paper, this would require us to go beyond the virtual welfare maximization regime in the Myerson auction setting. We believe such characterization could shed light on the mechanism design for bidders with non-quasilinear utility functions and provide useful insights for practical applications such as online advertising. Another direction is to study ex post ROI constraints when the target ratio M_i is also private information of bidder i. This brings the problem to the domain of multidimensional mechanism design, which is often challenging in the mechanism design literature. Here, one possible approach is also to identify conditions under which some simple and deterministic auctions are optimal or close to optimal.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used GPT-4 in order to improve language. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

CRediT authorship contribution statement

Hongtao Lv: Writing - original draft, Methodology, Investigation, Conceptualization; **Xiaohui Bei:** Writing - review & editing, Validation, Supervision, Investigation; **Zhenzhe Zheng:** Writing - review & editing, Supervision, Funding acquisition, Conceptualization; **Fan Wu:** Writing - review & editing, Supervision, Funding acquisition.

Data availability

No data was used for the research described in the article.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We thank Hu Fu, Yiqing Wang, Yidan Xing, Xiangyu Liu and anonymous reviewers for their insightful and helpful suggestions. This work was supported in part by China NSF grant No. 62302267, 62432007, 62322206, 62132018, 62272307, in part by the NTU SPMS Collaborative Research Award, in part by the Natural Science Foundation of Shandong (No. ZR2023QF083), in part by Taishan Scholars Program. The opinions, findings, conclusions, and recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agencies or the government.

Appendix A. Omitted proofs

A.1. Proof of Proposition 1

Proof. First, we have $M_i\widetilde{p}_i(0) - 0 \cdot x_i(0) = 0$ since $x_i(0) = 0$ and $p_i(0) = 0$. Then we can observe that

$$\max_{0\le z\le v}\{M_i\widetilde{p}_i(z) - zx_i(z)\} \ge M_i\widetilde{p}_i(0) - 0 \cdot x_i(0) = 0.$$

Therefore, $p_i(v) \le M_i\widetilde{p}_i(v)$. Similarly, we have

$$\max_{0\le z\le v}\{M_i\widetilde{p}_i(z) - zx_i(z)\} \ge M_i\widetilde{p}_i(v) - vx_i(v),$$

hence

$$p_i(v) \le M_i\widetilde{p}_i(v) - (M_i\widetilde{p}_i(v) - vx_i(v)) = vx_i(v).$$

For the second statement, we prove that for any values $v < v' \in [0, v_{\max}]$, $p_i(v) \le p_i(v')$. we consider the following two cases:

- When
$$\max_{0\le z\le v}\{M_i\widetilde{p}_i(z) - zx_i(z)\} = \max_{0\le z\le v'}\{M_i\widetilde{p}_i(z) - zx_i(z)\},$$
we have
$$p_i(v') - p_i(v) = M_i(\widetilde{p}_i(v') - \widetilde{p}_i(v)) \ge 0.$$
- When
$$\max_{0\le z\le v}\{M_i\widetilde{p}_i(z) - zx_i(z)\} < \max_{0\le z\le v'}\{M_i\widetilde{p}_i(z) - zx_i(z)\},$$
we denote
$$v^* = \arg\max_{0\le z\le v'}\{M_i\widetilde{p}_i(z) - zx_i(z)\}.$$
Then we have that
$$p_i(v^*) = M_i\widetilde{p}_i(v^*) - (M_i\widetilde{p}_i(v^*) - v^*x_i(v^*)) = v^*x_i(v^*),$$
and
$$\max_{0\le z\le v^*}\{M_i\widetilde{p}_i(z) - zx_i(z)\} = \max_{0\le z\le v'}\{M_i\widetilde{p}_i(z) - zx_i(z)\}.$$
By the analysis of the above case, we have
$$p_i(v') \ge p_i(v^*) = v^*x_i(v^*) > vx_i(v) \ge p_i(v).$$

For the third statement, we again denote

$$v^* = \arg\max_{0\le z\le v}\{M_i\widetilde{p}_i(z) - zx_i(z)\}.$$

Then we obtain the payment at valuation v^*:

$$p_i(v^*) = v^*x_i(v^*) \ge \widetilde{p}_i(v^*).$$

Moreover, we can observe that

$$p_i(v) - p_i(v^*) = M_i(\widetilde{p}_i(v) - \widetilde{p}_i(v^*)).$$

Combining the above arguments, we have $p_i(v) \geq \widetilde{p}_i(v)$. □

A.2. Proof of monotone allocation rule (step 1 of Theorem 1)

Lemma 6 (Monotone Allocation Rule). *For ex post ROI-constrained bidders, if a mechanism is DSIC, then the allocation rule is monotonically non-decreasing, that is, $x_i(v) \leq x_i(v')$ for all $v < v'$ and bidder i.*

Before the proof, we first present elementary inequalities derived from the property of truthfulness for ex post ROI-constrained bidders.

Given values v, v' with $v < v'$, by the truthfulness requirement, when a bidder i with value v misreports v', we have

$$M_i v x_i(v') - p_i(v') \leq M_i v x_i(v) - p_i(v), \tag{A.1}$$

or

$$p_i(v') > v x_i(v'). \tag{A.2}$$

Similarly, considering bidder i with value v' misreports v, we obtain that

$$M_i v' x_i(v) - p_i(v) \leq M_i v' x_i(v') - p_i(v'), \tag{A.3}$$

or

$$p_i(v) > v' x_i(v). \tag{A.4}$$

One can observe that (A.4) contradicts the IR requirement since $p_i(v) \leq v x_i(v) < v' x_i(v)$. Hence, in a truthful mechanism, (A.3) must hold, and at least one of (A.1) and (A.2) holds. On the basis of this elementary analysis, we now prove Lemma 6.

Proof. We prove the monotonicity by contradiction. Let $v < v'$, we assume $x_i(v) > x_i(v')$ in a truthful mechanism. Then we prove that (A.1) and (A.3) are not compatible. By (A.3), we get

$$p_i(v) - p_i(v') \geq M_i v' \cdot (x_i(v) - x_i(v')). \tag{A.5}$$

Similarly, by (A.1), we obtain

$$p_i(v) - p_i(v') \leq M_i v \cdot (x_i(v) - x_i(v')). \tag{A.6}$$

Combining (A.5) and (A.6), we have $M_i v' \cdot (x_i(v) - x_i(v')) \leq M_i v \cdot (x_i(v) - x_i(v'))$. However, since $x_i(v) > x_i(v')$ and $v' > v$, we can derive a contradiction.

Next, we continue to prove that (A.2) and (A.3) are not compatible. By the IR property, we have

$$p_i(v) \leq v x_i(v). \tag{A.7}$$

Combining (A.2), (A.3), and (A.7), we obtain

$$M_i v' x_i(v) - v x_i(v) < M_i v' x_i(v') - v x_i(v'),$$

that is,

$$(M_i v' - v) \cdot x_i(v) < (M_i v' - v) \cdot x_i(v'). \tag{A.8}$$

As we have $x_i(v) > x_i(v')$ and $v < v'$, (A.8) does not hold. In conclusion, we obtain that a mechanism could not be truthful if $x_i(v) > x_i(v')$, and hence, the allocation function $x_i(v)$ must be monotonously non-decreasing. □

A.3. Proof of uniqueness of payment (step 3 of Theorem 1)

Lemma 7 (Uniquessness of Payment). *Given any monotonically non-decreasing allocation rule $x_i(\cdot)$ and $p_i(0) = 0$ for every bidder i, there exists a unique payment rule $p_i(\cdot)$ such that (x, p) is DSIC.*

Proof. For any bidder i and any $v_i \in [0, v_{\max}]$. We consider two valuations v_i and $w_i = v_i + \delta$ for some sufficiently small $\delta > 0$. Following the payment sandwich inequality techniques used in the original Myerson's proof, we have

$$u_i(v_i, v_i) \geq u_i(w_i, v_i) \text{ and } u_i(w_i, w_i) \geq u_i(w_i, v_i).$$

This gives us two inequalities that can together sandwich the payment $p_i(v)$.

If we follow Myerson's original analysis and replace each utility with the value minus the payment for the bidder, these two inequalities would give us

$$M_i v_i (x_i(w_i) - x_i(v_i)) \leq p_i(w_i) - p_i(v_i) \leq M_i w_i (x_i(w_i) - x_i(v_i)).$$

In the limit, as δ approaches 0, we can divide each side by δ and get

$$\frac{\mathrm{d}p_i(v_i)}{\mathrm{d}v_i} = M_i v_i \frac{\mathrm{d}x_i(v_i)}{\mathrm{d}v_i}. \tag{A.9}$$

The uniqueness of the payment rule then follows by integrating (A.9) from 0 to each value v. However, this argument has a problem in the ROI-constrained setting: because of the extra factor M_i in the payment function due to the ROI condition, it is possible that at some point the payment grows to be larger than the bidder's obtained value, therefore violating the IR condition.

To deal with this issue, we first note that $p_i(v_i) \le v_i \cdot x_i(v_i) \le w_i \cdot x_i(v_i)$. This means when bidder i has value w_i and misreports her value as v_i, her payment will never violate the IR condition. This means we only need to discuss the case when the bidder i with value v_i misreports to w_i. There are two cases to discuss.

- If there exists a neighborhood of v_i, such that by applying (A.9) to $p_i(v_i)$ to get $p_i(w_i)$, we have $p_i(w_i) \le v_i \cdot x_i(w_i)$. This means in this small neighborhood, the other direction of misreporting will also not be affected by the IR condition. Myerson's analysis can go through, and the derivative of $p_i(\cdot)$ remains fixed and unique at point v_i.
- If for any small neighborhood of v_i, applying (A.9) gives $p_i(w_i) > v_i \cdot x_i(w_i)$. This means bidder i with value v_i would not want to misreport her bid as w_i because it would violate the IR condition for her. Here we notice $v_i x_i(w_i) < p_i(w_i) \le w_i x_i(w_i)$ must hold. When δ is sufficiently small, it implies that we must have $p_i(w_i) = w_i x_i(w_i)$. This suggests that even though we can no longer apply Eq. (A.9) in this case. $p_i(w_i)$ is still a unique and fixed value.

Combining the two cases together, we know that at each point $v \in [0, v_{\max}]$, $p_i(v)$ is always unique, therefore proving this claim. □

A.4. Proof of Lemma 2

Proof. Recall that a mechanism is BIC if truthful reporting maximizes a bidder's expected utility, where the expectation is taken over the prior distributions of other bidders' valuations v_{-i}. Let $U_i(b_i, v_i)$ denote the expected utility of bidder i with true value v_i who reports b_i. Assuming the ROI constraints are satisfied, the utility function is given by:

$$\begin{aligned} U_i(b_i, v_i) &= \mathbb{E}_{v_{-i}}\left[M_i v_i x_i(b_i, v_{-i}) - p_i(b_i, v_{-i})\right] \\ &= M_i v_i \mathbb{E}_{v_{-i}}[x_i(b_i, v_{-i})] - \mathbb{E}_{v_{-i}}[p_i(b_i, v_{-i})] \\ &= M_i v_i X_i(b_i) - P_i(b_i). \end{aligned}$$

We observe that this expected utility function $U_i(b_i, v_i)$ has the exact same mathematical structure as the ex post utility function u_i defined in Eq. (13), with the deterministic allocation x_i replaced by the interim allocation X_i, and the deterministic payment p_i replaced by the interim payment P_i. Specifically, the problem of maximizing $M_i v_i X_i(b_i) - P_i(b_i)$ is isomorphic to the problem of maximizing $M_i v_i x_i(b_i) - p_i(b_i)$ solved in Theorem 1. Since Theorem 1 provides the necessary and sufficient conditions for the latter maximization to be achieved at $b_i = v_i$ (i.e., DSIC), the same logic applies directly to the interim functions for BIC. Thus, by applying the characterization from Theorem 1 to the interim functions $X_i(\cdot)$ and $P_i(\cdot)$, we conclude that the mechanism is BIC if and only if $X_i(\cdot)$ is monotone and $P_i(\cdot)$ follows the payment rule specified in the lemma. □

A.5. Proof of the DMR assumption on t_i

Proposition 5. *If the normalized valuation v_i satisfies the Decreasing Marginal Revenue (DMR) condition, then the initial valuation $t_i = M_i v_i$ (where $M_i > 1$ is a constant) also satisfies the DMR condition.*

Proof. Let $F_V(\cdot)$ and $f_V(\cdot)$ denote the cumulative distribution function (CDF) and probability density function (PDF) of the normalized valuation v_i, respectively. The DMR condition requires that the function $\psi_V(v)$ is monotonically non-decreasing in v, where:

$$\psi_V(v) \triangleq v f_V(v) + F_V(v) - 1.$$

Consider the initial valuation $t_i = M_i v_i$. Let $F_T(\cdot)$ and $f_T(\cdot)$ denote the CDF and PDF of t_i. The relationship between the distributions of t_i and v_i is given by:

$$F_T(t) = \mathbb{P}(T \le t) = \mathbb{P}(M_i V \le t) = \mathbb{P}\left(V \le \frac{t}{M_i}\right) = F_V\left(\frac{t}{M_i}\right).$$

By differentiating with respect to t, we obtain the relationship between the densities:

$$f_T(t) = \frac{d}{dt} F_T(t) = f_V\left(\frac{t}{M_i}\right) \cdot \frac{1}{M_i}.$$

We now examine the function $\psi_T(t)$ for the initial valuation t:

$$\psi_T(t) \triangleq t f_T(t) + F_T(t) - 1.$$

Substituting the expressions for $F_T(t)$ and $f_T(t)$ derived above:

$$\begin{aligned}\psi_T(t) &= t \cdot \left[f_V\left(\frac{t}{M_i}\right) \cdot \frac{1}{M_i} \right] + F_V\left(\frac{t}{M_i}\right) - 1 \\ &= \left(\frac{t}{M_i}\right) f_V\left(\frac{t}{M_i}\right) + F_V\left(\frac{t}{M_i}\right) - 1.\end{aligned}$$

Observing that $\frac{t}{M_i} = v$, we have:

$$\psi_T(t) = \psi_V\left(\frac{t}{M_i}\right).$$

Since M_i is a positive constant, $\frac{t}{M_i}$ is strictly increasing in t. Given that $\psi_V(\cdot)$ is non-decreasing by assumption, it follows that $\psi_T(t)$ is also non-decreasing in t. Thus, the initial valuation t_i satisfies the DMR condition. □

References

[1] A. Malakhov, R.V. Vohra, Optimal auctions for asymmetrically budget constrained bidders, Rev. Econ. Des. 12 (4) (2008) 245–257.
[2] C. Borgs, J. Chayes, N. Immorlica, M. Mahdian, A. Saberi, Multi-unit auctions with budget-constrained bidders, in: Proceedings of the 6th ACM Conference on Electronic Commerce, 2005, pp. 44–51.
[3] J.-J. Laffont, J. Robert, Optimal auction with financially constrained buyers, Econ. Lett. 52 (2) (1996) 181–186.
[4] N. Golrezaei, I. Lobel, R. Paes Leme, Auction design for ROI-constrained buyers, in: Proceedings of the Web Conference 2021, 2021, pp. 3941–3952.
[5] J. Auerbach, J. Galenson, M. Sundararajan, An empirical analysis of return on investment maximization in sponsored search auctions, in: Proceedings of the 2nd International Workshop on Data Mining and Audience Intelligence for Advertising, 2008, pp. 1–9.
[6] S.R. Balseiro, Y. Deng, J. Mao, V.S. Mirrokni, S. Zuo, The landscape of auto-bidding auctions: value versus utility maximization, in: Proceedings of the 22nd ACM Conference on Economics and Computation, 2021, pp. 132–133.
[7] E. Brynjolfsson, Y. Hu, D. Simester, Goodbye pareto principle, hello long tail: the effect of search costs on the concentration of product sales, Manag. Sci. 57 (8) (2011) 1373–1386.
[8] S. Fadaei, M. Bichler, Truthfulness and approximation with value-maximizing bidders, in: International Symposium on Algorithmic Game Theory, Springer, 2016, pp. 235–246.
[9] P. Lu, C. Xu, R. Zhang, Auction design for value maximizers with budget and return-on-spend constraints, arXiv preprint arXiv:2307.04302 (2023)
[10] R.B. Myerson, Optimal auction design, Math. Oper. Res. 6 (1) (1981) 58–73.
[11] S. Chawla, D.L. Malec, A. Malekian, Bayesian mechanism design for budget-constrained agents, in: Proceedings of the 12th ACM Conference on Electronic Commerce, 2011, pp. 253–262.
[12] H. Fu, J. Hartline, D. Hoy, Prior-independent auctions for risk-averse agents, in: Proceedings of the Fourteenth ACM Conference on Electronic Commerce, 2013, pp. 471–488.
[13] S. Chawla, K. Goldner, J.B. Miller, E. Pountourakis, Revenue maximization with an uncertainty-averse buyer, in: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2018, pp. 2050–2068.
[14] B.K. Szymanski, J.-S. Lee, Impact of ROI on bidding and revenue in sponsored search advertisement auctions, in: Second Workshop on Sponsored Search Auctions, 2006 1–8.
[15] C. Borgs, J. Chayes, N. Immorlica, K. Jain, O. Etesami, M. Mahdian, Dynamics of bid optimization in online advertisement auctions, in: Proceedings of the 16th International Conference on World Wide Web, 2007, pp. 531–540.
[16] N. Golrezaei, P. Jaillet, J.C.N. Liang, V. Mirrokni, Bidding and pricing in budget and ROI constrained markets, arXiv preprint arXiv:2107.07725 (2021)
[17] E. Tillberg, P. Marbach, R. Mazumdar, Optimal bidding strategies for online ad auctions with overlapping targeting criteria, 4, ACM New York, NY, USA, 2020 1–55.
[18] G. Aggarwal, A. Badanidiyuru, A. Mehta, Autobidding with constraints, in: International Conference on Web and Internet Economics, Springer, 2019, pp. 17–30.
[19] B. Heymann, Cost per action constrained auctions, in: Proceedings of the 14th Workshop on the Economics of Networks, Systems and Computation, 2019, pp. 1–8.
[20] M. Babaioff, R. Cole, J. Hartline, N. Immorlica, B. Lucier, Non-quasi-linear agents in quasi-linear mechanisms (Extended abstract), Leibniz Int. Proc. Inform. 185 (2021) 1.
[21] B. Li, X. Yang, D. Sun, Z. Ji, Z. Jiang, C. Han, D. Hao, Incentive mechanism design for ROI-constrained auto-bidding, arXiv preprint arXiv:2012.02652 (2020)
[22] R. Cavallo, P. Krishnamurthy, M. Sviridenko, C.A. Wilkens, Sponsored search auctions with rich ads, in: Proceedings of the 26th International Conference on World Wide Web, 2017, pp. 43–51.
[23] C.A. Wilkens, R. Cavallo, R. Niazadeh, GSP: The cinderella of mechanism design, in: Proceedings of the 26th International Conference on World Wide Web, 2017, pp. 25–32.
[24] J. Li, P. Tang, Auto-bidding equilibrium in ROI-constrained online advertising markets, arXiv preprint arXiv:2210.06107 (2022)
[25] S. Balseiro, N. Golrezaei, V. Mirrokni, S. Yazdanbod, A black-box reduction in mechanism design with private cost of capital, Available at SSRN 3341782 (2019)
[26] A. Mehta, Auction design in an auto-bidding setting: randomization improves efficiency beyond VCG, in: Proceedings of the ACM Web Conference 2022, 2022, pp. 173–181.
[27] Y. Deng, J. Mao, V. Mirrokni, S. Zuo, Towards efficient auctions in an auto-bidding world, in: Proceedings of the Web Conference 2021, 2021, pp. 3965–3973.
[28] S. Balseiro, Y. Deng, J. Mao, V. Mirrokni, S. Zuo, Robust auction design in the auto-bidding world, Adv. Neural Inf. Process. Syst. 34 (2021), pp. 17777–17788.
[29] Y.-K. Che, I. Gale, The optimal mechanism for selling to a budget-constrained buyer, J. Econ. Theory 92 (2) (2000), pp. 198–233.
[30] N.R. Devanur, N. Haghpanah, C.-A. Psomas, Optimal multi-unit mechanisms with private demands, in: Proceedings of the 2017 ACM Conference on Economics and Computation, 2017, pp. 41–42.
[31] A. Fiat, K. Goldner, A.R. Karlin, E. Koutsoupias, The fedex problem, in: Proceedings of the 2016 ACM Conference on Economics and Computation, 2016, pp. 21–22.
[32] M. Babaioff, S. Dughmi, R. Kleinberg, A. Slivkins, Dynamic pricing with limited supply, ACM Trans. Econ. Comput. 3 (1) (2015) 1–26.
[33] G. Pavlov, Optimal mechanism for selling two goods, BE J. Theor. Econ. 11 (1) (2011).
[34] S. Hart, P.J. Reny, Maximal revenue with multiple goods: nonmonotonicity and other observations, Theor. Econ. 10 (3) (2015) 893–922.
[35] J.D. Hartline, T. Roughgarden, Simple versus optimal mechanisms, in: Proceedings of the 10th ACM Conference on Electronic Commerce, 2009, pp. 225–234.
[36] C. Daskalakis, G. Pierrakos, Simple, optimal and efficient auctions, in: International Workshop on Internet and Network Economics, Springer, 2011, pp. 109–121.
[37] Y. Feng, J.D. Hartline, Y. Li, Simple mechanisms for non-linear agents, in: Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2023, pp. 3802–3816.

[38] S. Alaei, A. Makhdoumi, A. Malekian, Revenue maximization under unknown private values with non-obligatory inspection, in: Proceedings of the 22nd ACM Conference on Economics and Computation, 2021, pp. 27–28.
[39] Y.A. Gonczarowski, N. Immorlica, Y. Li, B. Lucier, Revenue maximization for buyers with costly participation, in: Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2024, pp. 41–73.
[40] C.A. Wilkens, R. Cavallo, R. Niazadeh, S. Taggart, Mechanism design for value maximizers, arXiv preprint arXiv:1607.04362 (2016)