Towards Efficient Auction Design with ROI Constraints

Xinyu Tang
School of Software & C-FAIR, Shandong University
Jinan, China
xinyutang@mail.sdu.edu.cn

Hongtao Lv
School of Software & C-FAIR, Shandong University
Jinan, China
lht@sdu.edu.cn

Yingjie Gao
School of Software & C-FAIR, Shandong University
Jinan, China
yjgao@mail.sdu.edu.cn

Fan Wu
Department of Computer Science and Engineering, Shanghai Jiao Tong University
Shanghai, China
wu-fan@sjtu.edu.cn

Lei Liu
School of Software & C-FAIR, Shandong University
Shandong Research Institute of Industrial Technology
Jinan, China
l.liu@sdu.edu.cn

Lizhen Cui
School of Software & C-FAIR, Shandong University
Jinan, China
clz@sdu.edu.cn

ABSTRACT

Online advertising stands as a significant revenue source of the Internet. Recently, the trend among advertisers tilting towards the use of auto-bidding tools has heralded the emergence of a new model of bidders operating with constraints related to return on investment (ROI). However, most of the current research on ROI-constrained bidders in auction design only focuses on either the ROI constraints or values of bidders being private, while it is more practical to keep them both private in reality. Designing a truthful mechanism for bidders with both private values and ROI constraints introduces complexities because of the characteristics of designing mechanisms with multiple parameters. To remedy this, we divide bidders into binary classes: the traditional *utility maximizers* (UMs) who can be viewed as having an ROI constraint of 1, and the *ROI-constrained bidders* (RBs) who share a fixed ROI constraint denoted as γ. This framework retains the essence of multi-parameter mechanism design but transitions this into a more tractable form. Then we introduce a novel auction mechanism, cleverly combining the conventional VCG mechanism and an existing mechanism for public ROI-constrained bidders which is called Cavallo's mechanism. Our mechanism can achieve an approximation ratio of $\frac{3}{2}$ on social welfare. Additionally, we unearth new insights into the limitations posed by ROI constraints. When the ROI constraint γ exceeds 2, the lower bound of social welfare is $\frac{5}{4}$; when it falls below 2, the lower bound becomes $\frac{3+\gamma}{2+3\gamma-\gamma^2}$.

KEYWORDS

Return on investment; Mechanism design; Incentive compatibility

ACM Reference Format:
Xinyu Tang, Hongtao Lv, Yingjie Gao, Fan Wu, Lei Liu, and Lizhen Cui. 2024. Towards Efficient Auction Design with ROI Constraints. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 9 pages.

Hongtao Lv and Lei Liu are the corresponding authors.

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

1 INTRODUCTION

Online advertising auctions constitute a significant source of income for numerous tech companies, bringing in hundreds of billions of dollars each year. Recently, as daily ad auctions number in the tens of millions, happening in real-time, this vast and intricate marketplace has driven contemporary online advertising platforms to create auto-bidding services. These services allow advertisers to set broad marketing objectives for their campaigns and then automatically place bids on their behalf. In such auto-bidding setups, financial limitations of advertisers, such as budget or return on investment (ROI), have become key factors in auction design, and the traditional model of utility maximizers in game theory can not fit them very well. While auctions for bidders with budget limits have been extensively studied for more than two decades [3, 10, 14, 21, 27, 30], research on ROI-constrained bidders has only recently gained significant attention [6, 17, 18, 22–24]. The ROI constraints demand that the payment can not exceed a certain portion of the advertising value gained. Simply put, an ROI-constrained bidder aims for a specific minimum proportion between the value gained and the payment. This is distinct from budget constraints that set a fixed upper limit on payment. ROI constraints impose a payment ceiling that is proportionally linked to the value allocated. Empirical studies have shown that ROI constraints more accurately mirror real-world scenarios compared to budget constraints [2, 18], and the analysis of bidding strategies [17, 31, 32] and auction mechanism design [6, 11, 34] under ROI constraints have also been proposed.

However, most existing research on mechanism design primarily focuses on either value or ROI being private, while it is more practical for them both to be private in reality. It has been widely accepted in literature that the advertisers' values are private. The ROI constraints, symbolizing the minimum return on investment advertisers aim for, are shaped by their internal assessments in contrast with other marketing channels or their internal accounting, which makes these ratios remain undisclosed from the seller's viewpoint [6, 24]. Both value and ROI remaining private pose a challenge of multi-parameter mechanism design, historically recognized as a difficult problem in game theory [29].

To overcome the difficulty and offer insights for realistic applications, we simplify this model by examining binary ROI scenarios involving two distinct classes of bidders: the traditional *utility maximizers* (UMs) who can be viewed as having an ROI constraint of 1 (we will discuss this later) and the *ROI-constrained bidders* (RBs) who share a fixed ROI constraint denoted as γ. This adaptation, while preserving the essence of multi-parameter mechanism design, transforms the problem into a tractable problem. Intuitively, this adaption simplifies the misreporting space from a two-dimensional plane into two one-dimensional lines, thus reducing the complexity of multi-parameter mechanism design.

Based on this idea, we propose truthful and efficient mechanisms in this work. To warm up, we first consider the setting with public class information and propose a mechanism for BINary bidders with Public classes (BIN-Pub) and private values. BIN-Pub ranks advertisers by their bids and charges UMs using the Vickrey-Clarke-Grove (VCG) mechanism [12, 19, 33] and RBs via the mechanism proposed in [11], which we call *Cavallo's mechanism.* However, for private classes, we show that BIN-Pub may be "unfair" for RBs to some extent. Thus, we introduce a mechanism for BINary bidders with Private classes (BIN-Pri) and private values. Its core rationale is to compute the payment for each slot as the greater of the payment in VCG or Cavallo's mechanism, rather than pricing based on the class of the bidder itself. BIN-Pri first allocates RBs according to their values and subsequently identifies the position that maximizes the utilities for UMs. We show that this mechanism is truthful and ensures that the approximation ratio for social welfare is capped at a maximum of $\frac{3}{2}$. Particularly, we find that, when the ROI constraint γ exceeds 2, the lower bound of the ratio is $\frac{5}{4}$, while when it is less than 2, the lower bound becomes $\frac{3+\gamma}{2+3\gamma-\gamma^2}$.

The contributions of our study are threefold:

(1) In practice, it is more practical for bidders when both values and ROI constraints are private. In contrast to most existing studies where only value or ROI is private, we make one of the first attempts at mechanism design wherein both value and ROI remain private.

(2) We propose a novel mechanism for two classes of bidders with ROI constraints, named BIN-Pri, which guarantees economic properties and provides some interesting findings. BIN-Pri indicates that truthful mechanisms could give preference to bidders who have high ROI constraints, allocating them higher positions even when their bids are smaller than those from bidders with low ROI constraints.

(3) We established a lower bound in this both-private scenario, demonstrating that it is impossible to achieve social welfare with a better approximation ratio, let alone reach the optimal solution. This result advances the understanding of mechanism design for ROI-constrained bidders.

2 RELATED WORK

Two primary threads dominate the discussion regarding auction design with ROI-constrained bidders. The initial thread examines the impact of ROI constraints of bidding strategies in classical VCG or generalized second price (GSP) auctions [1, 4, 9, 17, 20, 31, 32]. The second thread, which our study aligns with, explores the *design* of auction mechanisms for ROI-constrained bidders [6, 8, 11, 17, 18, 22–25, 28, 34, 36]. Among these studies, Golrezaei *et al.* [18] empirically demonstrated that, in online advertising, a portion of bidders are indeed bound by ROI constraints. The existing studies on ROI in mechanism design can be divided into two distinct categories, namely *ex ante* ROI constraints [6, 18] and *ex post* ROI constraints [7, 13, 25]. Ex ante ROI constraints mandate that the ROI must be met on an expected basis, denoting the expected ROI based on the prior value distributions of bidders, while ex post ROI constraints offer a hard ROI constraint for advertisers across all potential value realizations. Intuitively, ex ante ROI constraints fit the advertisers who have adequate budgets and participate in a large number of auctions in a day, while ex post ROI constraints fit the advertisers who have limited budgets and get few clicks per day [25]. In this work, we follow the thread of ex post ROI constraints.

In the thread of studies on ex post ROI constraints in mechanism design, Cavallo *et al.* [11] first started the research by introducing the utility function of the RB model as used in our work and proposing a truthful mechanism, which we call Cavallo's mechanism in our study with the associated payment rules. Lv *et al.* [25] further conducted a study on revenue-maximizing mechanism design for a single bidder with their RB model. The authors of [22] and [36] considered multi-round ROI constraints and proposed corresponding conditions for truthful mechanisms. These studies take either the values or the ROI constraints as private, whereas our study places emphasis on the scenario where both values and ROI constraints are kept private.

In a series of studies concerning ROI [16, 26, 34, 35], the concept of *value maximizer* (VM), which only maintains the value term in the utility function, to some extent, can be viewed as a particular instance of a large ROI constraint scenario. Wilkens *et al.* [35] showed that when an advertiser's ROI is relatively high (exceeding 2 or 3), their behavioral pattern tends to closely mirror what is observed in the VM model. Based on this model, the work in [26] introduced a truthful mechanism for scenarios where UMs and VMs coexist and achieved a nearly tight approximation ratio on social welfare. Their study can be viewed as a special case of ours, where we adopt a more generalized model of RBs and provide a corresponding truthful and efficient mechanism.

Another closely relevant work to our research is [24], which also studied the multi-parameters private settings and developed approximately revenue-maximizing auction mechanisms within the context of budget and ROI constraints. They explored both fully and partially private settings regarding agent information, demonstrating constant approximation mechanisms for different agent demands and improved approximation ratios when only agents' target bars are private. However, there are significant distinctions between their work and ours. They concentrated on value maximizers and revenue maximization, while we focused on (ROI-constrained) utility maximizers and the optimization of social welfare.

3 PRELIMINARIES

We adopt a standard advertising auction model with K ad slots (or interchangeably, positions), labeled as $k \in \{1, 2, ..., K\}$, from bottom to up. Each slot k is associated with a click-through rate

(CTR) denoted as x_k. It is assumed that $x_K \geq x_{K-1} \geq ... \geq x_1 > 0$. Moreover, we introduce a virtual slot denoted as $k = 0$ with $x_0 = 0$, representing a slot even lower than the lowest slot. We consider a group of bidders represented as $\mathcal{N} = \{1, 2, ..., n\}$, labeled by i, and every individual bidder possesses a private value t_i for a click. We assume for simplicity that $n > K$ without affecting the generality of our model. Within the auction mechanism defined as M, the allocation outcome is represented as $\Pi = \{\pi_0, \pi_1, \pi_2, ..., \pi_K\}$, where π_k denotes the bidder allocated to the kth advertising slot. We also introduce the notation a_i to represent the slot index assigned to bidder i, *i.e.*, $a_i = k$ if $\pi_k = i$. Subsequently, a payment p_i is incurred by each bidder i per click. In our study, we differentiate between two advertiser classes: the conventional *utility maximizers* (UMs) and *ROI-constrained bidders* (RBs) bound by a return on investment (ROI) constraint, represented as γ. Here, γ denotes the value-to-payment ratio, signifying that for each auction outcome, the inequality $\frac{t_i}{p_i} \geq \gamma$ must be strictly satisfied, and we assume $1 < \gamma < +\infty$ in this work. We consider γ as a shared constant for all ROI-constrained bidders, but the classes of all bidders are private, which maintains the essence of a multi-parameter mechanism design problem. We adopt the *ex post* ROI constraints, ensuring that the auction outcomes meet ROI requirements for any value profile, thus providing a more stringent guarantee compared to *ex ante* constraints as mentioned in the preceding section. Given these premises, the utility for an ROI-constrained bidder i can be expressed as:

$$u_i = \begin{cases} t_i x_{a_i} - p_i x_{a_i} & \text{if } \frac{t_i}{p_i} \geq \gamma \\ -\infty & \text{otherwise.} \end{cases}$$

We can further define

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

representing the maximum willingness-to-pay of bidder i per click. To circumvent ambiguity, we utilize the notations *value* for v_i and *click revenue* for t_i. For coherence, we also provide a definition for the value of each utility maximizer $v_i = t_i$. Next, following the definitions and annotations in [11] and [26], we provide formal descriptions for UMs and RBs:

Definition 1 (Utility Maximizer, UM). *A utility maximizer i strategizes to maximize her utility* $u_i = v_i x_{a_i} - p_i x_{a_i}$.

Definition 2 (ROI-Constrained Bidder, RB). *An ROI-constrained bidder i strategizes to maximize her utility*

$$u_i = \begin{cases} \gamma v_i x_{a_i} - p_i x_{a_i} & \textit{if } v_i \geq p_i \\ -\infty & \textit{otherwise.} \end{cases} \tag{1}$$

One can learn from the definitions above that the models of UM and RB converge when $\gamma = 1$. For ease of presentation, we assume $v_i \neq v_j$ for each pair of bidders i and j.[1] Similar to the terms used in [26], we employ $\tau_i \in \{UM, RB\}$ to express the *class* of bidder i and $\theta_i = (v_i, \tau_i)$ as her *type*. We presume that a bidder may misreport both her value and class, which means she may report her type as $\hat{\theta}_i = \{\hat{v}_i, \hat{\tau}_i\}$ with $\hat{v}_i \neq v_i$ and/or $\hat{\tau}_i \neq \tau_i$. Moreover, θ signifies the type profiles of all bidders, with θ_{-i} representing that of all bidders excluding i. Using these notations, $p_i(\hat{\theta}_i, \hat{\theta}_{-i})$ and $u_i(\hat{\theta}_i|\theta_i, \theta_{-i})$ represent the payment and utility of bidder i when reporting her type as $\hat{\theta}_i$, respectively, given her type θ_i and the type profile of others θ_{-i}.

Following the classic research on advertising auction mechanisms, there are two primary requirements for mechanism design: *incentive compatibility* (IC) and *individual rationality* (IR).

Definition 3 (Incentive Compatibility, IC). *A mechanism is incentive-compatible if and only if*

$$u_i(\theta_i|\theta_i, \theta_{-i}) \geq u_i(\hat{\theta}_i|\theta_i, \theta_{-i}), \quad \forall \theta_i, \hat{\theta}_i, \theta_{-i}, i \in \mathcal{N}.$$

Definition 4 (Individual Rationality, IR). *A mechanism is individually rational if and only if*

$$p_i(\theta_i, \theta_{-i}) \leq v_i, \quad \forall \theta_i, \theta_{-i}, i \in \mathcal{N}.$$

In other words, IC ensures that no bidder in the mechanism will have an incentive to misreport her type, while IR guarantees that no bidder will ever have a negative utility when bidding truthfully. Note that IR implies $v_i \geq p_i$ for an RB i, which means that we can simplify (1) as $u_i = \gamma v_i x_{a_i} - p_i x_{a_i}$ under the IR requirement. We refer to a mechanism as being both IC and IR by the word "truthful". It is widely recognized that the VCG mechanism is truthful for UMs. Additionally, recent research [11] has shown that the following mechanism is truthful for RBs with an ROI constraint of γ, which we call the *Cavallo's mechanism* in our work:

Mechanism 1 (Cavallo's mechanism).

- ***Allocation****: In accordance with the order of bidders' values, slots are allotted from the highest to the lowest.*
- ***Payment****: For every bidder labeled as i, where* $1 \leq a_i \leq K$, *consider j to be the bidder allocated the closest slot below,* i.e., $a_j = a_i - 1$, *and the payment is* $p_i = \min\left\{v_j, \frac{1}{x_{a_i}}\left(p_j x_{a_j} + \gamma v_j (x_{a_i} - x_{a_j})\right)\right\}$.

One can observe that the payment of Cavallo's mechanism constitutes a recursive term, wherein the subsequent term p_i is contingent upon its antecedent term p_j, and this dependency continues iteratively through a chain of terms.

In this study, we aim to design an efficient mechanism, *i.e.*, to maximize the total social welfare of a truthful auction for the setting where both UMs and RBs have private classes and values. It is noteworthy that the traditional definition of social welfare, which is an aggregate of the utilities of all bidders combined with the seller's revenue (equivalent to summing up the click revenues obtained by all bidders), does not align well with the nature of RBs. We highlight that the utility function of bidders should be scale-free, that is, the expressions $u_i = \gamma v_i x_{a_i} - p_i x_{a_i}$ and $u_i = v_i x_{a_i} - \frac{1}{\gamma} p_i x_{a_i}$ should be mathematically equivalent, because they indicate the identical preference of the bidder. However, when optimizing the traditional concept of social welfare, these two expressions of utility functions are no longer equivalent. Therefore, building upon prior research in mechanism design tailored for budget-constrained bidders, we introduce the concept of *liquid social welfare* (LSW)[2] [13, 15, 26] to maintain the scale-free property.

[1] When $v_i = v_j$, we choose one bidder following the lexicographic order of their non-manipulable unique IDs, and then subtract a sufficiently small positive number ϵ from her value, making their values distinct and ordered (in cases with more than two bidders having identical values, we can do this similarly). Additionally, ϵ should be small enough such that, the allocation will not change by further decreasing it.

[2] Here, "liquid" means "transferable money" in the auction.

Definition 5 (Liquid Social Welfare, LSW). *In a mechanism, the liquid social welfare of an allocation result* Π *is defined as the total of the highest bids that all bidders are willing to pay for the allocation. In other words,*

$$Wel(\Pi) = \sum_{k=1}^{K} v_{\pi_k} x_k.$$

By applying the definition of LSW, it becomes straightforward to deduce that VCG achieves LSW optimality for UMs, while Cavallo's mechanism attains this optimality for RBs. Consequently, in the development of an LSW-optimal mechanism for binary bidders, a key requisite is *consistency*, which means, in special cases, it can be deduced to classical mechanisms. This implies that the mechanism should naturally evolve as an extension of existing mechanisms:

Definition 6 (Consistency). *A mechanism ensures consistency if its outcome aligns with that of the VCG mechanism when all are UMs, and it mirrors Cavallo's mechanism when all are RBs.*

An ideal auction mechanism is supposed to be IC, IR, consistent, and LSW-optimal. Given our emphasis on truthful mechanisms in this research, we treat θ_i and $\hat{\theta}_i$ as equivalent in instances devoid of ambiguity.

4 MECHANISM FOR BINARY BIDDERS WITH PUBLIC CLASSES

In this section, as a warm-up, we consider a fundamental scenario where the class information of bidders is public to all bidders and the auctioneer. This scenario falls under the umbrella of single-parameter mechanism design [29], which, in this case, simplifies the problem while providing insights for our subsequent multi-parameter mechanism design problem. The proposed mechanism for BINary bidders with Public classes is referred to as BIN-Pub.

As shown in Algorithm 1, BIN-Pub sorts all bidders by their respective values (Line 3). The top K bidders are grouped into a set N, with the $(K+1)$st bidder assigned to a virtual slot, labeled by 0, for pricing purposes (Line 4). Following this, BIN-Pub sequentially fills the slots using the sorted bidders (Lines 5-6). The algorithm then deduces the payments for each slot utilizing the VCG and Cavallo's mechanism separately (Lines 7-8). The final payment is calculated based on the class of each bidder (Lines 9-13).

Next, we prove BIN-Pub is IC, IR, consistent, and LSW-optimal. Due to space limitations, the proof is omitted and will be included in the full version of the paper.

Theorem 1. *If the classes of bidders are public, BIN-Pub is IC, IR, consistent, and LSW-optimal.*

5 MECHANISM FOR BINARY BIDDERS WITH PRIVATE CLASSES

In the previous section, we studied an optimal mechanism for UMs and RBs when their classes are public. However, the situation shifts when such class information is kept private, steering the problem into the realm of multi-parameter mechanism design—a domain known for its complexity [29]. We first scrutinize the viability of BIN-Pub in a setting where class information is private. We can observe that the payment heavily relies on the class of bidders.

Algorithm 1: BIN-Pub

Input: All bidders' type profile θ, the CTR x_k for each slot, and the ROI constraint γ.

Output: The allocation and payment outcome.

1 $p_i \leftarrow 0, \forall i \in \{1, ..., n\}$;

2 $p_{UM}^{(0)}$, $p_{RB}^{(0)} \leftarrow 0$, and $x_0 \leftarrow 0$;

3 Order bidders based on their values;

4 Define N as the group of the highest K bidders according to their values, and let π_0 represent the bidder with the $(K+1)$th highest value;

5 **for** *k from* 1 *to* $|N|$ **do**

6     $\pi_k \leftarrow$ the bidder with the kth lowest value in N;

7     $p_{UM}^{(k)} = \frac{1}{x_k}\left(p_{UM}^{(k-1)} x_{k-1} + v_{\pi_{k-1}}(x_k - x_{k-1})\right)$;

8     $p_{RB}^{(k)} = \min\left\{v_{\pi_{k-1}}, \frac{1}{x_k}\left(p_{RB}^{(k-1)} x_{k-1} + \gamma v_{\pi_{k-1}}(x_k - x_{k-1})\right)\right\}$;

9 **for** *k from* 1 *to* $|N|$ **do**

10   **if** π_k *is a UM* **then**

11         $p_{\pi_k} \leftarrow p_{UM}^{(k)}$;

12   **else**

13         $p_{\pi_k} \leftarrow p_{RB}^{(k)}$;

14 Return π_k and p_{π_k} for each bidder.

Therefore, one can verify that, if an RB misreports her class as a UM but reports her value truthfully, she could potentially benefit from a reduced payment, without altering her allocation. This highlights an inherent bias in BIN-Pub against RBs, and such a bias increases the opportunity for strategic manipulation when their classes are private. Following our previous discussion, in the context of private classes, the payment should be independent of her class. This implies that the payment for a bidder, be it from the UM or RB class, should be invariant if she secures a specific slot, assuming other allocation parameters remain consistent. To this end, we introduce a nuanced payment rule that is anchored more to the slot than to the individual bidder. This rule melds components from both VCG and Cavallo's mechanism frameworks and is influenced by the classes and values of the bidders positioned below the slot. Consequently, we posit that payment at the slot should be the greater of two elements: 1) the *VCG-style payment*, originating from the closest lower UM; 2) the *Cavallo-style payment*, originating from the closest lower RB. To expound, for any slot indexed by $k \geq 1$, if i_R represents the closest RB below k at position k_R, and i_U represents the closest UM below k at position k_U, then the payment at slot k can be formulated as:

$$p^{(k)} = \max\left\{\hat{p}_U^{(k)}, \hat{p}_R^{(k)}\right\}, \tag{2}$$

where

$$\hat{p}_U^{(k)} = \frac{1}{x_k}\left(p^{(k_U)} x_{k_U} + v_{i_U}(x_k - x_{k_U})\right), \tag{3}$$

and

$$\hat{p}_R^{(k)} = \min\left\{v_{i_R}, \frac{1}{x_k}\left(p^{(k_R)} x_{k_R} + \gamma v_{i_R}(x_k - x_{k_R})\right)\right\}. \tag{4}$$

Algorithm 2: BIN-Pri

Input: The type profile θ of all bidders, the CTR x_k for all slots, the ROI constraint γ.
Output: The allocation and payment outcome.

1 $p_i \leftarrow 0, \forall i \in \{1, ..., n\}$, $p^{(0)} \leftarrow 0$ and $x_0 \leftarrow 0$;
2 Order bidders based on their values, and define N as the set of top K bidders;
3 Set i as $(K+1)$st highest bidder, and $\pi_0 \leftarrow i$;
4 Define S as the set of all UMs in N and T as all RBs in N;
`// Slot Allocation to RBs in` T.
5 **if** T *is non-empty* **then**
6   **for** $k = 1$ *to* $|T|$ **do**
7         $\pi_k \leftarrow$ RB with kth lowest value from T;
8   Update payment $p^{(k)}$ for slots $1 \leq k \leq |T| + 1$ using (2);
`// Iterative Slot Allocation to UMs.`
9 **while** S *is non-empty* **do**
10   Select i with the lowest value in S;
11   Set $\bar{k} \leftarrow K - |S| + 1$;
12   Find k^i to maximize $x_{k^i}(v_i - p^{(k^i)})$ for $1 \leq k^i \leq \bar{k}$;
13   **if** $k^i \neq \bar{k}$ **then**
14   `// Shift existing bidders by one slot up.`
        for $k = \bar{k}$ *to* $k^i + 1$ *descending* **do**
15             $\pi_k \leftarrow \pi_{k-1}$;
16   Assign $\pi_{k^i} \leftarrow i$;
17   Update $p^{(k)}$ for slots $k^i + 1 \leq k \leq \bar{k} + 1$ using (2);
18   Remove i from S;
19 Set payment for each slot: $p_{\pi_k} \leftarrow p^{(k)}$ where $k \in \{1, ..., K\}$;
20 Return π_k and p_i for each slot k and bidder i.

Here, we use $p^{(k)}$ to represent the (interim) payment at slot k. In the absence of an RB or UM below slot k, we designate a payment of 0 for that corresponding term in (2).

In light of this payment rule and inspired by [26], we propose our mechanism for BINary bidders with Private classes as BIN-Pri. Algorithm 2 specifies BIN-Pri in detail. At its heart, BIN-Pri first allocates slots to all RBs and then, iteratively, allocates UMs to the slots that optimize their utilities. Within Algorithm 2, the bidders with top K values are represented as a set N (Line 2). A virtual slot, denoted as 0, is assigned to the bidder ranking $(K+1)$st in value, as the baseline for pricing (Line 3). Here, S represents the set of all UMs in N, while T corresponds to the set of all RBs (Line 4). Sequentially, the lowest $|T|$ slots (excluding the virtual slot) are occupied by the RBs in T (Lines 5-7). After these allocations, the payment for each slot, ranging from 1 to $|T| + 1$, is determined using (2), based on the allocated RBs (Line 8). Given that UMs in S have not been allocated slots at this step, the variable $\hat{p}_U^{(k)}$ defaults to 0, simplifying the price determination. Subsequent to this price calculation, the unallocated UM with the lowest value is represented by i, and an optimal slot k^i is identified for her, that is, the slot with the highest utility; if there is a tie, the lower slot is preferred (Lines 10-12). Crucially, there exist two potential slot options: 1) $k^i = \bar{k} = K - |S| + 1$, indicating the slot just above the previously assigned bidders (with K being the sum of the sizes of T and S initially, and S gets updated as the set of all remaining UMs); and 2) $k^i < \bar{k}$, indicating an already occupied slot. In the former scenario, bidder i is immediately allocated slot k^i. In the latter, all bidders located at slot k^i or above are moved up a slot to make room for bidder i (Lines 13-16). Prices of slots ranking above k^i are then revised based on the value of the newly placed UM (Line 17). Once this is done, UM i is removed from set S, and this entire process repeats until all UMs in S secure a slot (Line 18). In the end, if a bidder is allocated to slot k, her pay-per-click is defined as the price of the slot; if not, she incurs no charge (Line 19).

Slot	CTR
4	0.4
3	0.3
2	0.2
1	0.1
0	0

Bidder	Class	Value
A	RB	3
B	RB	4
C	RB	5
D	UM	6
E	UM	7

Figure 1: The illustration of four slots with their CTRs, and five bidders with their classes and values in Example 1.

To elucidate the idea of BIN-Pri, we now provide an example. Notably, it is quite intriguing to observe from the example that, RBs with lower values, might be allocated to higher slots in comparison to UMs that have higher values.

Example 1. *Consider a scenario with four slots and five bidders, with their CTRs or types illustrated in Fig. 1. We set $\gamma = 1.2$ in this example. In BIN-Pri, we initially position bidder A, who has the fifth highest value, in the virtual slot 0. Subsequently, the RBs B and C occupy slots 1 and 2. This allocation leads to the payment for slots 1, 2, and 3 as $p^{(1)} = 3$, $p^{(2)} = 3.9$, and $p^{(3)} = 4.6$, respectively. Given these prices, the utilities for bidder D at each slot are computed as 0.3, 0.42, and 0.42 for slots 1, 2, and 3. Slot 2, yielding the maximum utility, is assigned to bidder D, moving bidder C to slot 3. Using (2), the updated prices for slots 3 and 4 are calculated as $p^{(3)} = 4.6$ and $p^{(4)} = 4.95$. Subsequently, bidder E's utilities at each slot are determined to be 0.4, 0.62, 0.72, and 0.82 for slots 1, 2, 3, and 4, hence, slot 4 is allocated to bidder E. The bidders' payments per click are determined by the respective prices of their assigned slots.*

5.1 Game Theoretical Properties

Prior to delving into the game-theoretical aspects of BIN-Pri, we introduce the concept of *marginal payment increase* [5] as it applies to a pair of ad slots. This notion measures the cost-benefit trade-off a bidder encounters when vying for a higher slot. Using this foundational concept, we introduce a series of lemmas that shed light on the core tenets governing BIN-Pri. These lemmas are instrumental in substantiating the proofs for IC and IR. Due to space constraints, we have omitted some proofs.

Definition 7. *For two slots $k > k'$, the marginal payment increase is defined as*

$$\Delta(k', k) \triangleq \frac{p_{\pi_k} x_k - p_{\pi_{k'}} x_{k'}}{x_k - x_{k'}}.$$

Lemma 1. *At the round after allocating a UM i to slot $a_i = k$, it follows that $\frac{p^{(k+1)} x_{k+1} - p^{(k)} x_k}{x_{k+1} - x_k} = v_i$.*

Lemma 2. *For two bidders i and j, if $v_i > v_j$ and $\tau_i = \tau_j$,* i.e., *they are both UMs or both RBs, then we have $a_i > a_j$.*

Corollary 1. *After allocating a UM i to a slot k, the allocations and payments of slot k and below will not change throughout the algorithm process.*

This corollary can be derived from Lemma 2, that is, when a UM i is allocated to position k, the next UM with higher values than i will definitely be at a higher position than k, so the prices at position k and below remain unchanged after the round of allocating i. Henceforth, in the following discussion, we will not distinguish between the notations of $p^{(k)}$ and p_{π_k} when there is no ambiguity. For notation simplicity, we use $p^{(k)}$ to denote the final price of slot k after the entire algorithm process if it is not specified.

Lemma 3. *For two bidders i and j, if bidder i is an RB and j is a UM, and $a_i < a_j$, then we have $v_i < v_j$.*

Lemma 4. *If a UM i is allocated a slot $a_i = k$, then we have $\Delta(k, k^*) \geq v_i, \forall k^* > k$.*

Lemma 5. *For two bidders i and j, if bidder i is an RB and j is a UM, and $a_i > a_j$, then we have $\gamma v_i \geq v_j$.*

Proof. Through Lemma 2, we only need to examine the proof for the bidder i nearest to j. We denote the closest UM below j as $\hat{j}$. Given the allocation rule of BIN-Pri, when selecting an optimal slot for bidder j, slot a_j emerges as the one offering maximal utility, resulting in

$$p^{(a_i)} x_{a_i} \geq p^{(a_j)} x_{a_j} + v_j (x_{a_i} - x_{a_j}) \tag{5}$$

From here, we can outline two potential cases of the payment $p^{(a_i)}$ in the round of allocating j.

Case 1: If $p^{(a_i)}$ comes from (4), combining with (5), we can deduce $\gamma v_i \geq v_j$.

Case 2: If $p^{(a_i)}$ comes from (3), by Lemma 4, we obtain

$$p^{(a_j)} x_{a_j} - p^{(a_{\hat{j}})} x_{a_{\hat{j}}} \geq v_{\hat{j}} (x_{a_j} - x_{a_{\hat{j}}}).$$

Combining with (5) and Lemma 2, we deduce

$$\begin{aligned} p^{(a_i)} x_{a_i} &\geq p^{(a_j)} x_{a_j} + v_j (x_{a_i} - x_{a_j}) \\ &\geq p^{(a_{\hat{j}})} x_{a_{\hat{j}}} + v_j (x_{a_i} - x_{a_j}) + v_{\hat{j}} (x_{a_j} - x_{a_{\hat{j}}}) \\ &> p^{(a_{\hat{j}})} x_{a_{\hat{j}}} + v_{\hat{j}} (x_{a_i} - x_{a_{\hat{j}}}), \end{aligned}$$

which means that $p^{(a_i)}$ can not come from (3). This leads to a contradiction and completes the proof. □

Corollary 2. *BIN-Pri is a consistent mechanism.*

The BIN-Pri algorithm naturally gives rise to this corollary: when all bidders are UMs, the outcome aligns with the VCG mechanism; when all bidders are RBs, as per Lemma 2, the outcome corresponds to Cavallo's mechanism. This corollary reinforces BIN-Pri as a generalization of both the VCG and Cavallo's mechanism, in alignment with our theoretical anticipations.

Building on the preceding discussions, we are well-positioned to conclude that BIN-Pri satisfies both IR and IC.

Theorem 2. *BIN-Pri is individually rational.*

Theorem 3. *BIN-Pri is incentive-compatible.*

Proof. We discuss the proof for UMs and RBs separately.

Case 1: We first consider a UM i. Define j and $\hat{j}$ as the closest UMs below and above bidder i under truthful reporting, respectively (in the absence of such UMs, we can introduce virtual UMs at slots 0 and $K + 1$). We denote their slots in the truthful scenario as a_j and $a_{\hat{j}}$. If bidder i misreports her type, their allocated slots become a'_j and $a'_{\hat{j}}$, respectively. We now examine the position of bidder i when she misreports her type, denoted by a'_i and we classify the scenario into three possibilities:

- a'_i lies between a'_j and $a'_{\hat{j}}$;
- a'_i is positioned higher than $a'_{\hat{j}}$;
- a'_i is positioned lower than a'_j.

For the first possibility, until bidder j is allocated, the outcome mirrors the truthful reporting scenario. Bidder i would have the choice between a_i and a'_i at the same price as in truthful reporting. Since she has a preference for a_i, it can be deduced that misreporting her type does not provide her with a more favorable outcome.

For the second possibility, by invoking Lemma 4, we observe that the utility bidder i derives at a'_i is always less or equivalent to the utility at slot $a'_{\hat{j}}$ when she misreports her type as $(v_{\hat{j}} - \epsilon, UM)$. Here, ϵ is a sufficiently small positive number. Bidder i displays a preference for a_i over the outcomes resulting from misreporting $(v_{\hat{j}} - \epsilon, UM)$, which falls into the first possibility.

For the third possibility, we define $\tilde{j}$ as the closest UM above a'_i in the untruthful scenario. The allocation for all slots below a'_i is consistent with the truthful setting. We define $a_{\tilde{j}}$ as the allocated slot of $\tilde{j}$ when i bids truthfully. If $a_{\tilde{j}}$ does not equal a'_i, then $\tilde{j}$ would have previously faced a choice between a'_i and $a_{\tilde{j}}$ and shown preference for $a_{\tilde{j}}$ during its allocation round under i's truthful reporting. As $v_i > v_{\tilde{j}}$, bidder i would also prefer $a_{\tilde{j}}$ and could obtain it by misreporting as $(v_{\tilde{j}} - \epsilon, UM)$ with a sufficiently small ϵ. Leveraging Lemma 1, we deduce that bidder i would prefer the slot just above $\tilde{j}$, denoted as $a_{\tilde{j}} + 1$. If $j > \tilde{j}$, this procedure can be repeated iteratively until $a_{\tilde{j}} + 1 > a'_j$. Conclusively, as it can be reduced to the first case, bidder i consistently prefers a_i.

Case 2: We now consider an RB i. If i is still assigned to the same position through misreporting her types, and given that other bidders are truthful, the utility will not change, and hence the bidder has no incentive to do so. Therefore, we only need to consider that i can attain either a higher or lower position k by misreporting her class or value. We use $\tilde{p}$ to denote the payment when bidding untruthfully and a_i to the allocated slot when bidding truthfully. Then we categorize the situations into four possibilities.

1). We assume that bidder i gets to a higher position k than her truthful position by misreporting her type, and k is only higher than UMs above i in truthful scenarios, but not than such RBs.

We assume that i^* is the UM at k^* that is closest to i above when bidding truthfully. We now prove that after i reaches a position k higher than i^* through misreporting, the IR property is violated, and hence i has no incentives to do so. After i misreporting her type, UM i^* might be allocated to $k^{*'}$. Because UMs always satisfy IR, we can obtain $v_{i^*} \geq \tilde{p}^{(k^{*'})}$. Through Lemma 4, we further deduce

$$\tilde{p}^{(k)} x_k \geq \tilde{p}^{(k^{*'})} x_{k^{*'}} + v_{i^*}(x_k - x_{k^{*'}}) \geq v_{i^*} x_k.$$

Combining with Lemma 3, we deduce $\tilde{p}^{(k)} \geq v_{i^*} > v_i$, which means violating the IR requirement.

2). We assume that bidder i gets to a higher position k than her truthful position by misreporting her type, and k is higher than at least an RB above i in truthful scenarios.

If the payment at k takes the former term of (4) in one step of the chain of terms, it would contradict IR. Therefore, we only need to consider the payment taking the latter term of (4) in every step of the chain. Therefore, combining (2) and (4), we can obtain

$$\tilde{p}^{(k)} x_k \geq p^{(a_i)} x_{a_i} + \gamma v_{\pi_{a_{i+1}}}(x_{a_{i+1}} - x_{a_i}) + \ldots + \gamma v_{\pi_{k-1}}(x_k - x_{k-1}).$$

We assume u_i as the truthful utility and $u_{i,k}$ as the utility when bidding untruthfully. The utility difference is non-negative:

$$u_i - u_{i,k} = \gamma v_i x_{a_i} - p^{(a_i)} x_{a_i} - (\gamma v_i x_k - \tilde{p}^{(k)} x_k)$$
$$= \gamma v_i (x_{a_i} - x_k) + \tilde{p}^{(k)} x_k - p^{(a_i)} x_{a_i} \geq 0.$$

Hence, misreporting her type does not increase the utility for i.

3). We assume that bidder i gets to a lower position k than her truthful position by misreporting her type, and k is only lower than RBs who are below i when i bids truthfully, but not than such UMs. The new payment has two possibilities.

- We consider the scenario that the payment for i is only from (4), *i.e.*, the recursive chain of calculating the payment only uses the (4). No matter what she bids, the price at position k remains the same. The price $p^{(a_i)}$ is derived from $p^{(k)}$, leading to the following inequality:

  $$p^{(a_i)} x_{a_i} \leq p^{(k)} x_k + \gamma v_{\pi_k}(x_{k+1} - x_k) + \ldots + \gamma v_{i-1}(x_{a_i} - x_{a_{i-1}}).$$

  From this, we represent u_i as the utility when bidding truthfully and $u_{i,k}$ as the utility when bidding untruthfully, then the difference in these utilities becomes

  $$u_i - u_{i,k} = \gamma v_i x_{a_i} - p^{(a_i)} x_{a_i} - (\gamma v_i x_k - p^{(k)} x_k)$$
  $$= \gamma v_i (x_{a_i} - x_k) + p^{(k)} x_k - p^{(a_i)} x_{a_i} \geq 0.$$

- We consider the senario that the payment for i is not only from (4). We assume an RB $\bar{i}$ is allocated to the slot $\overline{k}$, which is below a_i, with the VCG-style payment from the UM $\hat{i}$ allocated at $\hat{k}$ (if such RB does not exist, it can be reduced to the former possibility). From k to a_i, $\overline{k}$ is the highest position such that (3) is taken by the max function in (2), *i.e.*, $p^{(\overline{k})} x_{\overline{k}} = p^{(\hat{k})} x_{\hat{k}} + v_{\hat{i}}(x_{\overline{k}} - x_{\hat{k}})$, while others all take the term (4). From the previous possibility, we can obtain that the utility of i at $\overline{k}$ is not as high as being truthful at position a_i. Then we aim to validate the inequality $(\gamma v_i - p^{(\overline{k})}) x_{\overline{k}} \geq (\gamma v_i - p^{(k)}) x_k$, which further translates to

  $$p^{(k)} x_k \geq p^{(\hat{k})} x_{\hat{k}} + v_{\hat{i}}(x_{\overline{k}} - x_{\hat{k}}) + \gamma v_i (x_k - x_{\overline{k}}).$$

  Combining $p^{(k)} x_k \geq p^{(\hat{k})} x_{\hat{k}} + v_{\hat{i}}(x_k - x_{\hat{k}})$ and $\gamma v_i \geq v_{\hat{i}}$ drawing upon Lemma 5, we can complete this proof.

4). We assume that bidder i gets to a lower position k than her truthful position by misreporting her type, and k is lower than at least one UM who is below i when i bids truthfully.

Under truthful conditions, the closest UM advertiser j lower than bidder i is located at position $\overline{k}$. As demonstrated in the proof of Lemma 5, we have $\gamma v_i \geq v_j$. Consequently, when bidder i misreports her type and moves to position $\overline{k}$, her utility at that position will surpass any position below $\overline{k}$. Moreover, through Lemma 1, the utility at position $\overline{k}$ is inferior to that of the position $\overline{k}+1$. The utility of the position $\overline{k}+1$ is also less than what is realized under truthful bidding, reducing our analysis to the last situation.

This concludes the proof. □

Slot	CTR
2	$\frac{2}{\gamma}$
1	$\frac{1}{\gamma}$
0	0

Bidder	Class	Value
A	RB	ϵ
B	RB	$2+\epsilon$
C	UM or RB	4 or 1

Figure 2: The counter-example for $\gamma \geq 2$ in Theorem 5.

5.2 Approximation Ratio on LSW

We next show the upper and lower bounds on LSW, and again, we refer the reader to our full version for detailed proof.

Theorem 4. *BIN-Pri achieves an approximation ratio of no more than $\frac{3}{2}$ on LSW,* i.e., *$Wel_{our} \geq \frac{2}{3} Wel_{OPT}$, where Wel_{our} denotes the LSW of our mechanism, and $Wel_{OPT} = max_{\Pi} Wel(\Pi)$.*

Theorem 5. *No mechanism that is IC, IR, and consistent can guarantee an approximation ratio on LSW which is lower than*

$$\begin{cases} \frac{5}{4} & \text{when } \gamma \geq 2 \\ \frac{3+\gamma}{2+3\gamma-\gamma^2} & \text{when } 1 < \gamma < 2. \end{cases}$$

Proof. To substantiate this theorem, we employ two counter-examples.

When $\gamma \geq 2$, we consider the situation presented in Fig. 2. We provide two slots with CTRs of $\frac{2}{\gamma}$ and $\frac{1}{\gamma}$, in addition to a virtual slot with the CTR of 0. We focus on three bidders A, B, and C. The respective types for A and B are defined as (ϵ, RB) and $(2+\epsilon, RB)$, where ϵ is a sufficiently small positive number. We then examine the type of bidder C across four distinct cases: **Case 1:** $(4, UM)$; **Case 2:** $(4, RB)$; **Case 3:** $(1, UM)$; **Case 4:** $(1, RB)$. We suppose that the approximation ratio, denoted as R, on LSW in $\widetilde{M}$ is strictly below $\frac{5}{4}$, then the allocation outcomes must be as following:

- In **Cases 1** and **2**, bidders A, B, and C are allocated slots in the order of 0, 1, and 2.

- In **Cases 3** and **4**, bidders A, B, and C are allocated slots in the order of 0, 2, and 1.

One can verify that the approximation ratio of any alternative allocations unlike the above will be higher than R. For example, in **Cases 1** and **2**, if bidders A, B, and C are allocated slots 0, 2, and 1, respectively, the approximation ratio is $\frac{10+\epsilon}{8+2\epsilon}$ which would be larger than R because we can always find an ϵ which is sufficiently small.

From earlier discussions, it is evident that the payments for bidder C in both **Case 1** and **Case 2** should be equal, and we denote this as p_h. Otherwise, C may misreport her class. A similar assertion can be made for **Case 3 and 4** with the payment represented as p_l. Moving to the first case, if C misreports her value as 1, the resulting allocation would mirror **Case 3**. By the IC constraint, her utility when bidding untruthfully should not exceed the utility from truthfully declaring a value of 4. Therefore, the following relation is derived

$$\frac{2}{\gamma} \cdot (4 - p_h) \geq \frac{1}{\gamma} \cdot (4 - p_l),$$

which, upon simplification, leads to

$$2p_h - p_l \leq 4. \tag{6}$$

A critical observation is that in both **Case 2 and 4**, all bidders are RBs. The consistency requirement allows us to employ Cavallo's mechanism to compute the payments. Specifically, we get $p_h = 2+\epsilon$ and $p_l = \epsilon$. Substituting these values, we find

$$2p_h - p_l = 4 + \epsilon > 4,$$

which starkly contradicts (6), completing our proof.

When $1 < \gamma < 2$, we consider the situation presented in Fig. 3. This situation is similar to the above but is slightly more complicated. We still have two slots, associated with CTRs of 1 and $\frac{1}{\gamma}$, respectively, alongside a virtual slot with a CTR of 0. We concentrate on three bidders: A, B, and C. Bidders A and B have types defined as (ϵ, RB) and $((2+\epsilon)(3-\gamma), RB)$, respectively, with ϵ being a sufficiently small positive number. We discuss the type of bidder C into four distinct cases: **Case 1:** $(4, UM)$; **Case 2:** $(4, RB)$; **Case 3:** $(1, UM)$; **Case 4:** $(1, RB)$. We assume the approximation ratio of $\widetilde{M}$ on LSW is less than $\frac{3+\gamma}{2+3\gamma-\gamma^2}$, then the allocation results will be certainly as following:

- In **Cases 1** and **2**, bidders A, B, and C are allocated slots 0, 1, and 2, in that order.
- In **Cases 3** and **4**, bidders A, B, and C are allocated slots 0, 2, and 1, in that order.

Again, one can verify that any alternative to this allocation pattern will be higher than the defined approximation ratio. Building upon prior discussions, we discern that the payments associated with bidder C for **Cases 1** and **2** must coincide and be represented as p_h. Otherwise, C might misreport her class. This logic analogously applies to **Cases 3** and **4**, where the payments are symbolized as p_l. Transitioning to **Case 1**, if C misreports her declared value to 1, it mirrors the allocation and payment of **Case 3**. The IC constraint asserts that her utility from this should not supersede the utility derived from a truthful declaration of 4. This provides

$$1 \cdot (4 - p_h) \geq \frac{1}{\gamma} \cdot (4 - p_l),$$

Slot	CTR
2	1
1	$\frac{1}{\gamma}$
0	0

Bidder	Class	Value
A	RB	ϵ
B	RB	$(2+\epsilon)(3-\gamma)$
C	UM or RB	4 or 1

Figure 3: The counter-example for $1 < \gamma < 2$ in Theorem 5.

which translates to

$$\gamma p_h - p_l \leq 4\gamma - 4. \tag{7}$$

Based on the consistency constraint in **Case 2** and **Case 4**, Cavallo's mechanism determines the payments. We get

$$p_h = \min\left\{(2+\epsilon)(3-\gamma), \frac{\epsilon}{\gamma} + (\gamma-1)(2+\epsilon)(3-\gamma)\right\}$$

and $p_l = \epsilon$. We now prove that, however, no matter which term in the p_h is taken, this yields stark contradiction to (7). If $p_h = (2+\epsilon)(3-\gamma)$, combining (7) and $1 < \gamma < 2$, we have that

$$2\gamma \leq (2+\epsilon)\gamma - \epsilon < (2+\epsilon)\gamma(3-\gamma) - \epsilon \leq 4\gamma - 4.$$

We can further obtain $\gamma > 2$, which goes against our premise $1 < \gamma < 2$. If $p_h = \frac{\epsilon}{\gamma} + (\gamma-1)(2+\epsilon)(3-\gamma)$, combining with (7), we get $(2+\epsilon)\gamma(3-\gamma) \leq 4$, which also goes against our premise $1 < \gamma < 2$. This completes the proof. □

6 CONCLUSION

In this study, we make one of the first attempts at mechanism design that treats both values and ROI constraints as private, which is a practical setting in reality. We introduced two distinctive mechanisms: BIN-Pub for the setting with public class information and BIN-Pri for the setting with private class information. Both mechanisms ensure IC, IR, consistency, and an approximately optimal LSW. This research paves the way for subsequent investigations into private ROI constraints and brings forth a handful of unresolved problems. A salient problem is to consider two classes of bidders with two distinct and arbitrary ROI constraints (instead of fixing the ROI constraint of one class as 1 as in our model). Further, we aim to generalize this mechanism to accommodate bidders with multiple classes and a diverse spectrum of ROI constraints.

ACKNOWLEDGEMENTS

This work was supported in part by China NSF grant No. 62220106004, 62132018, and 62302267, in part by NSFShandong No. ZR2021LZH006, ZR202211150156, in part by Taishan Scholars Program, in part by Project 2023M732063 funded by China Postdoctoral Science Foundation, in part by Alibaba Group through Alibaba Innovation Research Program, and in part by Tencent Rhino Bird Key Research Project. 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.

REFERENCES

[1] Gagan Aggarwal, Ashwinkumar Badanidiyuru, and Aranyak Mehta. 2019. Autobidding with constraints. In *Proceedings of the International Conference on Web and Internet Economics*. 17–30.

[2] Jason Auerbach, Joel Galenson, and Mukund Sundararajan. 2008. 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*. 1–9.

[3] Lawrence M Ausubel. 2004. An efficient ascending-bid auction for multiple objects. *American Economic Review* 94, 5 (2004), 1452–1475.

[4] Moshe Babaioff, Richard Cole, Jason Hartline, Nicole Immorlica, and Brendan Lucier. 2021. Non-Quasi-Linear Agents in Quasi-Linear Mechanisms. *Leibniz international proceedings in informatics* 185 (2021).

[5] Yoram Bachrach, Sofia Ceppi, Ian A Kash, Peter Key, and Mohammad Reza Khani. 2016. Mechanism design for mixed bidders. In *Proceedings of the 25th International Conference on World Wide Web*. 215–225.

[6] Santiago Balseiro, Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. 2021. The Landscape of Auto-Bidding Auctions: Value Versus Utility Maximization. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 132–133.

[7] Santiago Balseiro, Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. 2021. Robust Auction Design in the Auto-bidding World. *Advances in Neural Information Processing Systems* 34 (2021).

[8] Santiago Balseiro, Negin Golrezaei, Vahab Mirrokni, and Sadra Yazdanbod. 2019. A Black-Box Reduction in Mechanism Design with Private Cost of Capital. *Available at SSRN 3341782* (2019).

[9] Christian Borgs, Jennifer Chayes, Nicole Immorlica, Kamal Jain, Omid Etesami, and Mohammad Mahdian. 2007. Dynamics of bid optimization in online advertisement auctions. In *Proceedings of the 16th international conference on World Wide Web*. 531–540.

[10] Christian Borgs, Jennifer Chayes, Nicole Immorlica, Mohammad Mahdian, and Amin Saberi. 2005. Multi-unit auctions with budget-constrained bidders. In *Proceedings of the 6th ACM Conference on Electronic Commerce*. 44–51.

[11] Ruggiero Cavallo, Prabhakar Krishnamurthy, Maxim Sviridenko, and Christopher A Wilkens. 2017. Sponsored search auctions with rich ads. In *Proceedings of the 26th International Conference on World Wide Web*. 43–51.

[12] Edward H Clarke. 1971. Multipart pricing of public goods. *Public choice* (1971), 17–33.

[13] Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. 2021. Towards Efficient Auctions in an Auto-bidding World. In *Proceedings of the Web Conference 2021*. 3965–3973.

[14] Shahar Dobzinski, Ron Lavi, and Noam Nisan. 2012. Multi-unit auctions with budget limits. *Games and Economic Behavior* 74, 2 (2012), 486–503.

[15] Shahar Dobzinski and Renato Paes Leme. 2014. Efficiency guarantees in auctions with budgets. In *Proceedings of the International Colloquium on Automata, Languages, and Programming*. 392–404.

[16] Salman Fadaei and Martin Bichler. 2017. Truthfulness with value-maximizing bidders: On the limits of approximation in combinatorial markets. *European Journal of Operational Research* 260, 2 (2017), 767–777.

[17] Negin Golrezaei, Patrick Jaillet, Jason Cheuk Nam Liang, and Vahab Mirrokni. 2021. Bidding and Pricing in Budget and ROI Constrained Markets. *arXiv preprint arXiv:2107.07725* (2021).

[18] Negin Golrezaei, Ilan Lobel, and Renato Paes Leme. 2021. Auction design for ROI-constrained buyers. In *Proceedings of the Web Conference 2021*. 3941–3952.

[19] Theodore Groves. 1973. Incentives in teams. *Econometrica: Journal of the Econometric Society* (1973), 617–631.

[20] Benjamin Heymann. 2019. Cost per action constrained auctions. In *Proceedings of the 14th Workshop on the Economics of Networks, Systems and Computation*. 1–8.

[21] Jean-Jacques Laffont and Jacques Robert. 1996. Optimal auction with financially constrained buyers. *Economics Letters* 52, 2 (1996), 181–186.

[22] Bin Li, Xiao Yang, Daren Sun, Zhi Ji, Zhen Jiang, Cong Han, and Dong Hao. 2020. Incentive Mechanism Design for ROI-constrained Auto-bidding. *arXiv preprint arXiv:2012.02652* (2020).

[23] Juncheng Li and Pingzhong Tang. 2022. Auto-bidding Equilibrium in ROI-Constrained Online Advertising Markets. *arXiv preprint arXiv:2210.06107* (2022).

[24] Pinyan Lu, Chenyang Xu, and Ruilong Zhang. 2023. Auction Design for Value Maximizers with Budget and Return-on-Spend Constraints. In *Proceedings of the International Conference on Web and Internet Economics*. 474–491.

[25] Hongtao Lv, Xiaohui Bei, Zhenzhe Zheng, and Fan Wu. 2023. Auction Design for Bidders with Ex Post ROI Constraints. In *Proceedings of the International Conference on Web and Internet Economics*. 492–508.

[26] Hongtao Lv, Zhilin Zhang, Zhenzhe Zheng, Jinghan Liu, Chuan Yu, Lei Liu, Lizhen Cui, and Fan Wu. 2023. Utility maximizer or value maximizer: mechanism design for mixed bidders in online advertising. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 5789–5796.

[27] Alexey Malakhov and Rakesh V Vohra. 2008. Optimal auctions for asymmetrically budget constrained bidders. *Review of Economic Design* 12, 4 (2008), 245–257.

[28] Aranyak Mehta. 2022. Auction design in an auto-bidding setting: Randomization improves efficiency beyond VCG. In *Proceedings of the ACM Web Conference 2022*. 173–181.

[29] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. 2008. *Algorithmic Game Theory*. Cambridge University Press.

[30] Mallesh M Pai and Rakesh Vohra. 2014. Optimal auctions with financially constrained buyers. *Journal of Economic Theory* 150 (2014), 383–425.

[31] Boleslaw K Szymanski and Juong-Sik Lee. 2006. Impact of ROI on bidding and revenue in sponsored search advertisement auctions. In *Proceedings of the Second Workshop on Sponsored Search Auctions*, Vol. 1.

[32] Erik Tillberg, Peter Marbach, and Ravi Mazumdar. 2020. Optimal Bidding Strategies for Online Ad Auctions with Overlapping Targeting Criteria. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 4, 2, 1–55.

[33] William Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance* 16, 1 (1961), 8–37.

[34] Christopher A Wilkens, Ruggiero Cavallo, and Rad Niazadeh. 2017. GSP: the cinderella of mechanism design. In *Proceedings of the 26th International Conference on World Wide Web*. 25–32.

[35] Christopher A Wilkens, Ruggiero Cavallo, Rad Niazadeh, and Samuel Taggart. 2016. Mechanism design for value maximizers. *arXiv preprint arXiv:1607.04362* (2016).

[36] Yidan Xing, Zhilin Zhang, Zhenzhe Zheng, Chuan Yu, Jian Xu, Fan Wu, and Guihai Chen. 2023. Truthful auctions for automated bidding in online advertising. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*. 2915–2922.