Auto-bidding with Budget and ROI Constrained Buyers

Xiaodong Liu and **Weiran Shen***
Gaoling School of Artificial Intelligence, Renmin University of China
{xiaodong.liu, shenweiran}@ruc.edu.cn

Abstract

In online advertising markets, an increasing number of advertisers are adopting auto-bidders to buy advertising slots. This tool simplifies the process of optimizing bids based on various financial constraints.

In our study, we focus on second-price auctions where bidders have both private budget and private ROI (return on investment) constraints. We formulate the auto-bidding system design problem as a mathematical program and analyze the auto-bidders' bidding strategy under such constraints. We demonstrate that our design ensures truthfulness, i.e., among all pure and mixed strategies, always reporting the truthful budget and ROI is an optimal strategy for the bidders. Although the program is non-convex, we provide a fast algorithm to compute the optimal bidding strategy for the bidders based on our analysis. We also study the welfare and provide a lower bound for the PoA (price of anarchy). Moreover, we prove that if all bidders utilize our auto-bidding system, a Bayesian Nash equilibrium exists. We provide a sufficient condition under which the iterated best response process converges to such an equilibrium. Finally, we conduct extensive experiments to empirically evaluate the effectiveness of our design.

1 Introduction

Ever since the seminal works of the Vickrey-Clarke-Groves (VCG) auction [Vickrey, 1961; Clarke, 1971; Groves, 1973] and Myerson's optimal auction [Myerson, 1981], the auction design problem has been one of the central topics at the intersection of economics and computer science, leading to the development and implementation of various auction mechanisms across numerous fields [Cramton *et al.*, 2004; Aggarwal *et al.*, 2006; Varian, 2007; Edelman *et al.*, 2007].

One of the most successful applications of auction theory is online advertising, which is a large and still growing business. Online advertising has become the main source of revenue for many Internet companies, such as Meta, Google, and TikTok. According to the statistics by Statista [2022], 521 billion dollars have been spent on digital advertisements in 2021. Most online ad platforms adopt the second-price auction or its variants. The appeal of the second-price auction lies in its truthfulness, where buyers are incentivized to report their true values as bids. Additionally, the second-price auction is known for maximizing social welfare, ensuring that the buyer with the highest value ultimately secures the item or ad slot.

Although the second-price auction is truthful, many bidders are still constantly changing their bids on these platforms. This behavior can be attributed to the presence of various financial constraints faced by advertisers, including budget limitations and profitability considerations. Consequently, advertisers often find it challenging to bid their true values, as doing so may conflict with their constraints, such as exceeding their allocated budget. However, setting fine-grained bids for different auctions is a notoriously difficult task for advertisers, and only those large advertisers have the ability to fine-tune the bidding strategy. Conversely, smaller advertisers, who may be more constrained by financial limitations, often lack the capacity to optimize their bids effectively.

Auto-bidding systems have gained significant popularity among advertisers as they aim to simplify the bidding process in online advertising. These systems only require advertisers to provide high-level financial targets, such as their budget and target CPA (cost per acquisition), instead of setting separate bids for each auction. The auto-bidding system then leverages these targets to compute optimized bids and participate in auctions on behalf of the advertisers. Such auto-bidding systems optimize the advertisers' bids for them and allow them to focus more on high-level goals. Thus more and more advertisers are adopting auto-bidding systems.

The simplest auto-bidding strategy only considers the budget constraint for the bidders [Borgs *et al.*, 2007; Pai and Vohra, 2014; Conitzer *et al.*, 2018; Balseiro *et al.*, 2017]. They ensure that the buyers' actual payment does not exceed their budget. Another set of strategies considers the bidder's return on investment (ROI) constraints, which is also called the return on spend (ROS) constraint in some research [Golrezaei *et al.*, 2021; Balseiro *et al.*, 2021b]. Aggarwal *et al.* [2019] assume that the bidders' budget and ROI constraints are publicly known, and Balseiro *et al.* [2022] consider both constraints but assume that the buyers' budget is publicly known, while the buyers' ROI constraint is privately known.

*To whom correspondence should be addressed.

To the best of our knowledge, there is no existing work that considers both private budget and private ROI constraints. In our paper, we focus on this setting and assume the underlying auction mechanism is the second-price auction with no reserve prices.

1.1 Our Contributions

In this paper, we formulate the problem as a mathematical program and analyze the properties of the optimal solution. We then propose an efficient algorithm based on theoretical analysis to compute the optimal strategy. We also show that under our mechanism, reporting the truthful budget and ROI targets is the optimal strategy even if the buyers can adopt mixed strategies. Furthermore, we show that an equilibrium always exists if all bidders adopt our auto-bidding algorithm. We then give a sufficient condition under which the iterated best response process converges to such an equilibrium. We also analyze the social welfare of the auto-bidding strategy and provide a lower bound for the PoA (price of anarchy). In the end, we conduct extensive experiments based on both synthetic and realistic data sets to demonstrate the performance of our bidding strategy.

1.2 Related Works

Budget Constraint: In the online advertising literature, Balseiro and Gur [2019] propose adaptive pacing algorithms to control the buyer's payment within the budget. Balseiro *et al.* [2017] provide a detailed comparison of budget management strategies commonly used in practice. Conitzer *et al.* [2022] and Chen *et al.* [2021] study the computational complexity in the pacing equilibrium.

ROI Constraint: Aggarwal *et al.* [2019] study the setting where the bidders are value maximizers with ROI constraint. Golrezaei *et al.* [2021] analyze the setting where the bidders are utility maximizers with ROI constraint in second-price auctions with reserve prices. They show that the buyer's optimal strategy is to shade his bid, i.e., bid lower than his valuation. Such shading strategies can also be found in [Balseiro *et al.*, 2015; Balseiro and Gur, 2019; Conitzer *et al.*, 2017; Gummadi *et al.*, 2012].

Social Welfare: Aggarwal *et al.* [2019] show that the truthful mechanism can achieve at least half of the optimal social welfare when bidders have financial constraints. Balseiro *et al.* [2021a] and Deng *et al.* [2021] study how to use boost and reserve price to improve the social welfare when auto-bidders are value maximizers under ROI constraints. Mehta [2022] studies how to use randomized mechanism to improve social welfare under the auto-bidding setting. Balseiro *et al.* [2021b] and Balseiro *et al.* [2022] analyze the setting where the buyer's ROI target is private and the budget constraint is public. Their settings are all different from ours as we consider both private budget constraints and private ROI targets.

Bayesian Nash Equilibrium: Both Aggarwal *et al.* [2019] and our paper consider the existence of a Bayesian Nash equilibrium. However, Aggarwal *et al.* [2019] assume there is only one particular impression and the queries and slots form a continuum, while we make assumptions about bidders' value distributions.

2 Preliminaries

We consider the online advertising setting where there is a seller with an item for sale to n buyers. Let $[n] = \{1, 2, \cdots, n\}$ denote the set of buyers. Each buyer $i \in [n]$ has a value v_i for the item, drawn independently from a publicly known cumulative distribution function $F_i(v_i) : [0, \bar{v}] \mapsto [0, 1]$ for some $0 < \bar{v} < +\infty$. Assume $F_i(v_i)$ is differentiable and $f_i(v_i)$ is the corresponding probability density function. Let $\mathbf{v} = (v_1, v_2, \cdots, v_n)$ be the value profile, which contains the values of all buyers. Similarly, let $\mathbf{v}_{-i} = (v_1, \cdots, v_{i-1}, v_{i+1}, \cdots, v_n)$ denote the value profile of all buyers except buyer i. Though the value distribution $F_i(v_i), i \in [n]$ is publicly known, the realized value v_i is only known to buyer i. Let b_i be the bid of buyer i. Similarly, we define $\mathbf{b} = (b_1, b_2, \cdots, b_n)$ as the bid profile of all buyers and $\mathbf{b}_{-i} = (b_1, \cdots, b_{i-1}, b_{i+1}, \cdots, b_n)$ as the bid profile of all buyers except buyer i.

Assume that the seller uses the second-price auction as the base auction mechanism, which is widely used in the online advertising industry, especially in ad exchanges. In a second-price auction, the bidder with the highest bid wins the item and pays the second highest bid. Let $x(\mathbf{b}) = (x_1(\mathbf{b}), \ldots, x_n(\mathbf{b}))$ be the allocation rule and $p(\mathbf{b}) = (p_1(\mathbf{b}), \ldots, p_n(\mathbf{b}))$ the payment rule:

$$x_i(\mathbf{b}) = \begin{cases} 1 & \text{if } b_i \geq b_j, \forall j \\ 0 & \text{otherwise} \end{cases}, \tag{1}$$

$$p_i(\mathbf{b}) = \begin{cases} \max_{j \neq i} b_j & \text{if } b_i \geq b_j, \forall j \\ 0 & \text{otherwise} \end{cases}. \tag{2}$$

If there are more than one highest bids, we can break ties arbitrarily. Assume that buyer i's utility functions u_i are quasi-linear, i.e., $u_i(v_i) = v_i x(b_i, \mathbf{b}_{-i}) - p(b_i, \mathbf{b}_{-i})$.

For a buyer i, let $D_i = \max_{j \neq i} b_j$ be the highest bid of all other buyers, which is a random variable. Denote by $G_i(D_i)$ and $g_i(D_i)$ the cumulative distribution function and the probability density function of D_i, respectively. Since $F_i(v_i)$ is differentiable for all i, D_i also has a density function, which implies that the probability of $b_i = D_i$ is simply 0. Thus the tie-breaking rule does not affect the auction outcome.

We consider the setting where all buyers have both budget and ROI constraints. Let B_i and γ_i be the budget and ROI constraints of buyer i. Formally, buyer i's bids b_i satisfies the budget constraint if:

$$\mathbf{E}[D_i \mathbb{I}\{b_i \geq D_i\}] \leq B_i,$$

where the function $\mathbb{I}(\cdot)$ is the indicator function, and the expectation is taken over both b_i and D_i. Since the seller uses a second-price auction, buyer i wins the auction only if $b_i \geq D_i$. And when buyer i wins, their payment is D_i. Thus the left-hand side of the above equation is simply their expected payment in the auction.

The ROI can be defined as follows.

Definition 1 (Return on Investment)**.** *Buyer i's return on investment is the ratio between his expected gain from the auction and his expected payment. Formally, assuming $\mathbf{E}[D_i \mathbb{I}(b_i \geq D_i)] > 0$, we define the buyer i's ROI as*

$$\frac{\mathbf{E}[(v_i - D_i)\mathbb{I}\{b_i \geq D_i\}]}{\mathbf{E}[D_i \mathbb{I}\{b_i \geq D_i\}]}.$$

Buyer i's bid satisfies the ROI constraint if:

$$\frac{\mathbf{E}[(v_i - D_i)\mathbb{I}\{b_i \geq D_i\}]}{\mathbf{E}[D_i \mathbb{I}\{b_i \geq D_i\}]} \geq \gamma_i.$$

A major reason for the buyer's ROI constraint is that the buyer has outside options. Suppose that the buyer can have a profit margin of 5% if they invest the money elsewhere. Then they will only participate in the auction if they can secure an ROI of at least 5% in the advertising platform. We assume that both B_i and γ_i are the private information of buyer i, and thus need to be reported to the platform.

Throughout the paper, we assume that the budget B_i has a lower bound $\underline{B} > 0$, as a 0 budget simply indicates that the buyer does not want to participate in the auctions. Also, we assume that the ROI γ_i has an upper bound $\bar{\gamma}$ with $\bar{\gamma} > 0$, since no platform cannot guarantee an arbitrarily high ROI.

Note that in both two constraints, we consider the expected utility and payment. This is because buyers participate in large volumes of online ad auctions each day, and focusing on the expected quantities makes more sense than considering only a single auction.

We follow literature convention (see, e.g., [Golrezaei *et al.*, 2021; Mehta, 2022]) and focus on shading strategies[1] for the bidders: $b_i = \beta_i v_i$. We restrict β_i to be a positive number (i.e., $\beta_i > 0$) for obvious reasons. We call β_i the *shading parameter*.

Since we assume

$$\mathbf{E}[D_i \mathbb{I}(b_i \geq D_i)] > 0,$$

the ROI constraint becomes:

$$\mathbf{E}[(v_i - D_i)\mathbb{I}\{b_i \geq D_i\}] \geq \gamma_i \, \mathbf{E}[D_i \mathbb{I}\{b_i \geq D_i\}].$$

Plugging in the bidding strategy $b_i = \beta_i v_i$, we obtain the following inequality for the ROI constraint:

$$\mathbf{E}\{[(1+\gamma_i)D_i - v_i]\mathbb{I}\{\beta_i v_i \geq D_i\} \leq 0.$$

In this paper, we first consider how to design an autobidding strategy (which is fully characterized by the parameter β_i) for each buyer given their reported budget and ROI. We assume that the buyers' utility is $-\infty$ if either constraint is violated, i.e., the constraints are hard constraints. With the above discussion, we can formulate the problem as a mathematical program based on B_i and γ_i:

$$\begin{aligned} \text{maximize:} \quad & \mathbf{E}[(v_i - D_i)\mathbb{I}\{\beta_i v_i \geq D_i\}] \\ \text{subject to:} \quad & \mathbf{E}[D_i \mathbb{I}\{\beta_i v_i \geq D_i\}] \leq B_i \\ & \mathbf{E}[((1+\gamma_i)D_i - v_i)\mathbb{I}\{\beta_i v_i \geq D_i\}] \leq 0 \end{aligned} \tag{3}$$

3 The Optimal Shading Parameter

If the buyer i does not have financial constraints, their optimal strategy is to use their valuations as the bids ($\beta_i = 1$) in a second-price auction. However, when the bidders have financial constraints, they may try to shade their bids and use a different parameter β_i. In this section, we first analyze the properties of Program (3) when buyer i has both budget and ROI constraints. Then we propose an efficient algorithm to solve this program based on the properties, although Program (3) is non-convex.

3.1 Properties of the Optimal Parameter

In this section, we derive important properties of Program (3) that will be useful for designing our algorithm. Due to space limit, we defer all the proofs to the Appendix.

We first analyze the feasible region of Program 3. For simplicity, define the bidder i's expected payment as

$$\begin{aligned} P_i(\beta_i) &= \mathbf{E}[D_i \mathbb{I}\{\beta_i v_i \geq D_i\}] \\ &= \int_0^{\bar{v}} \int_0^{\beta_i v_i} D_i g_i(D_i) \,\mathrm{d}D_i f_i(v_i) \,\mathrm{d}v_i. \end{aligned} \tag{4}$$

Lemma 1. *The buyer's expected payment $P_i(\beta_i)$ is monotone increasing with respect to β_i.*

According to Lemma 1, we must have $\beta_i \leq P_i^{-1}(B_i)$ in order to satisfy the bidder's budget constraint $P_i(\beta_i) \leq B_i$. Therefore, the budget constraint actually restricts β_i to the interval $(0, P_i^{-1}(B_i)]$.

Then we analyze the ROI constraint. Define

$$\begin{aligned} h_i(\beta_i) &= \mathbf{E}[((1+\gamma_i)D_i - v_i)\mathbb{I}\{\beta_i v_i \geq D_i\}] \\ &= \int_0^{\bar{v}} \int_0^{\beta_i v_i} [(1+\gamma_i)D_i - v_i] g_i(D_i) \,\mathrm{d}D_i f_i(v_i) \,\mathrm{d}v_i. \end{aligned} \tag{5}$$

The following lemma analyzes how $h_i(\beta_i)$ changes with respect to β_i.

Lemma 2. *Function $h_i(\beta_i)$ is monotone decreasing in $(0, \frac{1}{1+\gamma_i}]$, and monotone increasing in $(\frac{1}{1+\gamma_i}, +\infty)$.*

It is easy to see that $h_i(0) = 0$. Then Lemma 2 implies that there is at most one $\hat{\beta} > 0$ such that $h_i(\hat{\beta}) = 0$.

Lemma 3. *If there exists $\hat{\beta} > 0$ with $h_i(\hat{\beta}) = 0$, then the ROI constraint is equivalent to $\beta_i \leq \hat{\beta}$. Otherwise, we have $h_i(\beta_i) < 0, \forall \beta_i \in (0, +\infty)$.*

Lemma 3 says that the ROI constraint either requires β_i to be in the interval $(0, \hat{\beta}]$ or imposes no restrictions on β_i (i.e., β_i should be in the interval $(0, +\infty)$).

Combining Lemma 1 and 3, we can immediately get the following corollary:

Corollary 1. *The feasible region of Program (3) is an interval $(0, \bar{\beta}]$. If there exists $\hat{\beta} > 0$ with $h_i(\hat{\beta}) = 0$, $\bar{\beta} = \min\{P_i^{-1}(B_i), \hat{\beta}\}$. Otherwise, $\bar{\beta} = P_i^{-1}(B_i)$.*

Now we analyze the objective function. Write buyer i's expected utility as a function of β_i:

$$\begin{aligned} U_i(\beta_i) &= \mathbf{E}[(v_i - D_i)\mathbb{I}\{\beta_i v_i \geq D_i\}] \\ &= \int_0^{\bar{v}} \int_0^{\beta_i v_i} (v_i - D_i) g_i(D_i) \,\mathrm{d}D_i f_i(v_i) \,\mathrm{d}v_i \end{aligned} \tag{6}$$

Lemma 4. *The buyer's utility is monotone increasing with respect to β_i when $\beta_i \leq 1$, and monotone decreasing with respect to β_i when $\beta_i > 1$.*

[1] Although the strategy is called a shading strategy in our paper, we actually allow β_i to be larger than 1. However, later analyses show that the optimal β_i never exceeds 1, and hence a "shading" strategy.

Lemma 4 shows that the utility function is maximized when $\beta_i = 1$ if there is no restriction on β. However, we already know from Corollary 1 that the feasible region of Program (3) is an interval. The following result is straightforward from Lemma 4 and Corollary 1.

Corollary 2. *The optimal solution β_i to program* (3) *satisfies* $0 < \beta_i \leq 1$.

Since $\beta_i \leq 1$ and the expected payment of the buyer is monotone in β_i, the maximum possible payment is obtained when $\beta_i = 1$. Define the maximum possible payment as

$$B^H = P_i(1) = \mathbf{E}[D_i \mathbb{I}\{v_i \geq D_i\}]. \quad (7)$$

Therefore, we have the following corollary:

Corollary 3. *If the buyer's budget B_i satisfies $B_i \geq B^H$, the buyer's budget constraint will have no effect on the value of β_i, i.e., The budget constraint will always be satisfied.*

In this paper, we only consider the case $\underline{B} \leq B^H$, since otherwise, the budget constraint will always be satisfied.

Now we show that choosing a larger β_i yields a lower ROI.

Lemma 5. *The actual ROI of buyer i is monotone decreasing with respect to β_i. And as β_i approaches 0, the ROI approaches infinity.*

Define

$$\gamma^L = \frac{\mathbf{E}[(v_i - D_i)\mathbb{I}\{v_i \geq D_i\}]}{\mathbf{E}[D_i \mathbb{I}\{v_i \geq D_i\}]} \quad (8)$$

to be the ROI if buyer i bids their values, i.e., $\beta_i = 1$.

Similarly, we have the following corollary:

Corollary 4. *When buyer i's ROI γ_i satisfies $\gamma_i \leq \gamma^L$, the buyer's ROI constraint can always be satisfied for all $\beta_i \in (0, 1]$.*

Let P_i denote the buyer's expected payment and Γ_i denote the buyer i's actual ROI after using the bidding strategy β_i obtained by solving the above program. Combining the above properties, we can get the theorem below.

Theorem 1. *Let B_i and γ_i be the reported budget constraint and the ROI constraint of buyer i. Define $B^H = \mathbf{E}[D_i \mathbb{I}\{v_i \geq D_i\}]$ and $\gamma^L = \frac{\mathbf{E}[v_i \mathbb{I}\{v_i \geq D_i\}]}{\mathbf{E}[D_i \mathbb{I}\{v_i \geq D_i\}]} - 1$. The relation between point (B_i, γ_i) and the expected payment and ROI point (P_i, Γ_i) is shown in Figure 1, where the solid curve are the set of all points that satisfy both the budget constraints and ROI constraints simultaneously, and*

- *if (B_i, γ_i) lies in area 1, the realized point (P_i, Γ_i) is the point right above (B_i, γ_i) in the solid curve and satisfies $P_i = B_i$;*
- *if (B_i, γ_i) lies in area 2, the realized point (P_i, Γ_i) is the point to the left of (B_i, γ_i) in the curve and satisfies $\Gamma_i = \gamma_i$;*
- *if (B_i, γ_i) lies in area 3, the realized point (P_i, Γ_i) is (B^H, γ^L).*

3.2 Algorithm

Since Program (3) is non-convex, it is not easy to directly find it solution. In this section, we present an indirect algorithm based on the theoretical analyses so far.

Theorem 2. *Algorithm 1 can correctly solve the Program (3).*

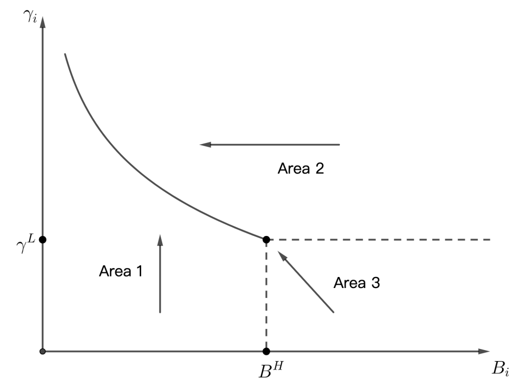

Figure 1: The relationship between (B_i, γ_i) and (P_i, Γ_i).

Algorithm 1: Finding the optimal β_i

Input: The buyer's budget constraint B_i and ROI target γ_i.
Output: The buyer's best bidding strategy β_i

```
Compute B^H and γ^L according to Equation (7) and (8);
if B_i ≥ B^H then
    β_i^1 ← 1;
else
    Use binary search to solve equation P_i(β_i^1) = B_i,
      where P_i(β_i) is defined in Function (4);
if γ_i ≤ γ^L then
    β_i^2 ← 1;
else
    Use binary search to solve equation h_i(β_i^2) = 0,
      where h_i(β_i) is defined in Function (5);
return β_i = min{β_i^1, β_i^2}
```

4 Strategic Issues

In the above analyses, we only study how to optimize bidder i's utility given their reported B_i and γ_i, but have not considered their strategic behaviors. The following results show that if the auto-bidding system optimizes the buyer's bids by solving Program (3), the buyer's best response is to report B_i and γ_i truthfully. We omit all the proofs due to the space limit.

Lemma 6. *Buyer i's optimal pure strategy is to report (B_i, γ_i) truthfully.*

Lemma 6 only considers deterministic strategies. Now we show that reporting the true B_i and γ_i is also the optimal strategy even if we consider mixed strategies.

Suppose buyer i uses a mixed strategy that leads to a distribution of possible β_i's. Thus the expected payment and utility becomes $\mathbf{E}_{\beta_i}[P_i(\beta_i)]$ and $\mathbf{E}_{\beta_i}[U_i(\beta_i)]$. Note that the expected ROI is $\frac{\mathbf{E}_{\beta_i}[U_i(\beta_i)]}{\mathbf{E}_{\beta_i}[P_i(\beta_i)]}$ rather than $\mathbf{E}_{\beta_i}[\Gamma_i(\beta_i)]$. Therefore, we cannot directly analyze how the expected ROI changes, but focus on $\mathbf{E}_{\beta_i}[P_i(\beta_i)]$ and $\mathbf{E}_{\beta_i}[U_i(\beta_i)]$ instead.

Lemma 7. *View $U_i(\beta_i)$ and $P_i(\beta_i)$ as a parametric equation, which induces an implicit function $U_i(P_i)$. The function*

$U_i(P_i)$ *is concave.*

Theorem 3. *Buyer i's optimal strategy is to report (B_i, γ_i) truthfully, even if mixed strategies are considered.*

Now we give an example here to illustrate our results.

Example 1. *Suppose there is a bidder i with v_i drawn from the uniform distribution $U[0, 10]$. Assume that D_i also follows the uniform distribution $U[0, 10]$.*

We first compute B^H and γ^L.

$$B^H = \mathbf{E}[D_i \mathbb{I}\{v_i \geq D_i\}] = \int_0^{10} \int_0^{v} D \frac{1}{10} \mathrm{d}D_i \frac{1}{10} \mathrm{d}v = \frac{5}{3},$$

$$\gamma^L = \frac{\mathbf{E}[v_i \mathbb{I}\{v_i \geq D_i\}]}{\mathbf{E}[D_i \mathbb{I}\{v_i \geq D_i\}]} - 1 = \frac{\int_0^{10} v \int_0^{v} \mathrm{d}D_i \, \mathrm{d}v}{\int_0^{10} \int_0^{v} D_i \, \mathrm{d}D_i \, \mathrm{d}v} - 1 = 1.$$

Now we compute the two thresholds β_i^1 and β_i^2. If $B_i \geq \frac{5}{3}$, we have $\beta_i^1 = 1$. Otherwise, if $B_i < \frac{5}{3}$, we can obtain β_i^1 by solving the equation $\mathbf{E}[D_i \mathbb{I}\{\beta^1 v_i \geq D_i\}] = B_i$. In this example, we can get the closed-form solution $\beta_i^1 = \sqrt{\frac{3}{5} B_i}$.

As for β_i^2, if $\gamma_i \leq 1$, we have $\beta_i^2 = 1$ and if $\gamma_i > 1$, the threshold is $\beta_i^2 = \frac{2}{1+\gamma_i}$, which is the solution to equation

$$\mathbf{E}\{[(1 + \gamma_i) D_i - v_i] \mathbb{I}(\beta v_i \geq D_i)\} = 0.$$

The optimal solution is $\beta_i = \min\{\beta_i^1, \beta_i^2\}$. We compute the payments and utilities for different β_i's. The results are shown in Figure 2, which confirms Lemma 5 and 7.

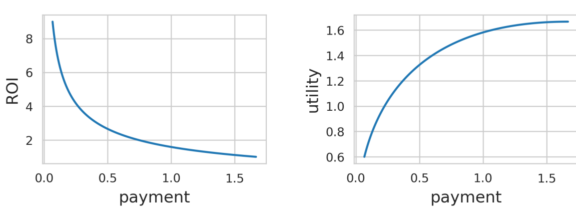

(a) ROI vs payment (b) Utility vs payment

Figure 2: The buyer's ROI and utility as a function the the payment.

5 Equilibrium Analysis

Until now, we have analyzed the behavior of a single bidder. The bidder will shade their bids to satisfy the budget and ROI constraints. In this section, we analyze the Bayesian game induced by the auto-bidding system. We show the existence of a Bayesian Nash equilibrium and provide a sufficient condition for the iterated best response process to converge to such an equilibrium. We also analyze the social welfare of the auto-bidding system. We make the following additional assumptions about $f_i(v_i)$ that hold for a wide range of value distributions:

1. For any i, $f_i(v_i)$ is bounded by $m \leq f_i(v_i) \leq M$;
2. There exits $\xi > 0$, such that $f_i(v_i)$ is ξ-Lipschitz for all i, i.e., $|f_i(v_i) - f_i(v_i')| \leq \xi |v_i - v_i'|$.

Since for any i and value $v_i \in [0, \bar{v}]$, $f_i(v_i) \leq M$, it follows that the function $F_i(v)$ is M-Lipschitz. Formally, for any value $v_i, v_i' \in [0, \bar{v}]$, we have

$$|F_i(v_i) - F_i(v_i')| \leq M |v_i - v_i'|, \ \forall i \in [n].$$

The shading parameters β_i depend on the bidders' budget and ROI constraints. Since both the budget and ROI constraints are bounded, we can provide a lower bound for β_i.

Let $\boldsymbol{\beta} = [\beta_1, \beta_2, \ldots, \beta_n]$ denote a shading profile, and $\boldsymbol{\beta}_{-i} = [\beta_1, \beta_2, \ldots, \beta_{i-1}, \beta_{i+1}, \ldots, \beta_n]$ the shading profile of all buyers except for bidder i. For each bidder i, $\boldsymbol{\beta}_{-i}$ will induce a probability density function $g_i(D_i, \boldsymbol{\beta}_{-i})$ over $D_i = \max_{j \neq i} b_j$, as the other bidders' bids depend on $\boldsymbol{\beta}_{-i}$.

We first prove the following lemmas, which will be useful for later arguments.

Lemma 8. *Given any $\boldsymbol{\beta}_{-i}$, $g_i(D_i, \boldsymbol{\beta}_{-i})$ satisfies the following:*

$$(n-1) m^{n-1} D_i^{n-1} \leq D_i g_i(D_i, \boldsymbol{\beta}_{-i}) \leq (n-1) M \bar{v}.$$

With the above lemma, we can bound β_i from below. The intuition is that although a small enough β_i may satisfy both the budget and the ROI constraints, it also lowers the bidder's utility according to Lemma 4.

Lemma 9. *Define*

$$\underline{\beta} = \max\left\{ \frac{2\underline{B}}{(n-1) M^2 \bar{v}^3}, \frac{1}{1+\bar{\gamma}} \right\}.$$

If each bidder i's financial constraints satisfy $B_i \geq \underline{B}$ and $\gamma_i \leq \bar{\gamma}$, then for all bidders, we have $\beta_i \geq \underline{\beta}, \forall i \in [n]$.

5.1 Existence of Bayesian Nash Equilibrium

In this section, we prove that if all the bidders use our auto-bidding system, then there exists a Bayesian Nash equilibrium in the Bayesian game induced by our mechanism. We also show that the iterated best response process converges to such an equilibrium under certain technical conditions.

Given any shading profile $\boldsymbol{\beta}$, for each bidder i, $\boldsymbol{\beta}_{-i}$ will induce a probability distribution $g_i(D_i, \boldsymbol{\beta}_{-i})$ over $D_i = \max_{j \neq i} b_j$, Let β_i' be the optimal solution to Program (3) by using $g_i(D_i, \boldsymbol{\beta}_{-i})$, i.e., β_i' can be viewed as a "best response" to $\boldsymbol{\beta}_{-i}$. Define function $X : [\underline{\beta}, 1]^n \mapsto [\underline{\beta}, 1]^n$ such that $X_i(\boldsymbol{\beta}) = \beta_i'$. Note that we assume that the bidders only can report $B_i \in [\underline{B}, +\infty)$ and $\gamma_i \in (0, \bar{\gamma}]$, the program (3) always has solutions and its solution always lies in $[\underline{\beta}, 1]$. Therefore, the function X is well-defined. It is clear that an equilibrium exists if the function X has a fixed point. To show this, we first prove that a small change in β_k with $k \neq i$ will not affect $g_i(D_i, \boldsymbol{\beta}_{-i})$ too much for any i (Lemma 10). Then we show that a slight perturbation in $g_i(D_i, \boldsymbol{\beta}_{-i})$ does not result in a sudden change in bidder i's response (Lemma 11). In the end, we show that even if all bidders change their strategies simultaneously, the change of function X can still be bounded (Lemma 12), which indicates that X is continuous. And with Brouwer's fixed point theorem, we know that X has a fixed point (Theorem 4).

Starting from any initial $\beta^{(0)}$, we define an iteration process as follows:

$$\boldsymbol{\beta}^{(t+1)} = X(\boldsymbol{\beta}^{(t)}), t = 0, 1, 2, \ldots$$

We perform convergence analysis of the above process. We show that under certain conditions, X becomes a contraction, and the convergence immediately follows (Theorem 5).

Lemma 10. *Given any $\boldsymbol{\beta}$, if the parameter β_k is changed to $\beta_k + \epsilon_k$, then for any $i \neq k$, the probability density function $g_i(D_i, \boldsymbol{\beta}_{-i})$ changes at most $\frac{(M^2+\underline{\beta}\xi)D_i}{\underline{\beta}^3}|\epsilon_k|$.*

Lemma 11. *Given any $\boldsymbol{\beta}$, if the parameter β_k for bidder k is changed to $\beta_k + \epsilon_k$, then for any i, $X_i(\boldsymbol{\beta})$ changes at most $\frac{(M^3+M\underline{\beta}\xi)(n+1)|\epsilon_k|}{12(n-1)\underline{\beta}^{n+2}m^n\bar{v}^{n-3}}$.*

Lemma 12. *Let $\boldsymbol{\epsilon} = (\epsilon_1, \epsilon_2, \dots, \epsilon_n)$. For any $\boldsymbol{\beta}$, function $X(\boldsymbol{\beta})$ satisfies the following:*

$$\|X(\boldsymbol{\beta}) - X(\boldsymbol{\beta} + \boldsymbol{\epsilon})\|_1 \leq \frac{(M^3 + M\underline{\beta}\xi)(n+1)\|\epsilon\|_1}{12(n-1)\underline{\beta}^{n+2}m^n\bar{v}^{n-3}},$$

where $\|\cdot\|_1$ is the 1-norm operator.

The above 3 lemmas show that the function X is continuous. Now we are ready to apply Brouwer's fixed point theorem to prove the existence of an equilibrium

Theorem 4. *For any problem instance satisfying the 3 assumptions specified at the beginning of this section, there exists a Bayesian Nash equilibrium $\boldsymbol{\beta}$, such that for each i, β_i is the best response to $\boldsymbol{\beta}_{-i}$.*

Corollary 5. *If*

$$\frac{(M^3 + M\underline{\beta}\xi)(n+1)}{12(n-1)\underline{\beta}^{n+2}m^n\bar{v}^{n-3}} < 1,$$

starting from any $\boldsymbol{\beta}^0$, the following iteration process converges to an equilibrium:

$$\boldsymbol{\beta}^{t+1} = X(\boldsymbol{\beta}^t), \forall t = 0, 1, 2, \dots$$

5.2 Welfare Analysis

In this section, we analyze the social welfare of the auto-bidding system. We show that our mechanism always achieves a certain fraction of the optimal welfare.

Definition 2. *(Social Welfare) Given any shading parameter profile $\boldsymbol{\beta}$, the social welfare is defined as*

$$SW(\boldsymbol{\beta}) = \mathbf{E}_{\mathbf{v}\sim F(\mathbf{v})}\left[\sum_{i=1}^{n} v_i x_i(\boldsymbol{b})\right],$$

where $F(\mathbf{v}) = \prod_i F_i(v_i)$ is the joint value distribution and $b_i = \beta_i v_i$.

Definition 3 (Price of Anarchy)**.** *Given any instance described by $F(\mathbf{v})$, Let $EQ(F)$ be the set of all the equilibria of the instance. The price of anarchy is defined as the worst ratio between the optimal social welfare and the social welfare in the worst equilibrium among all instances satisfying assumptions described in Section 5:*

$$PoA = \min_F \frac{\min_{\boldsymbol{\beta}\in EQ(F)} SW(\boldsymbol{\beta})}{\max_{\boldsymbol{\beta}} SW(\boldsymbol{\beta})}.$$

The optimal welfare is simply the welfare of the second-price auction, which can be achieved by setting $\beta_i = 1$ for all i. The following result shows that the price of anarchy of our mechanism is at least $\underline{\beta}$, i.e., our mechanism always achieves at least $\underline{\beta}$ fraction of the optimal welfare.

Theorem 5. *The price of anarchy of our mechanism is at least $\underline{\beta}$, i.e., $PoA \geq \underline{\beta}$.*

6 Experiments

In this section, we conduct experiments based on an open data set and report the experiment results.

We first consider a relatively simple case. Suppose that there are only five i.i.d. buyers. The value of each buyer is drawn from the uniform distribution $U[0, 10]$. Assuming that only one buyer has financial constraints and other four buyers simply bid their true values. The buyer's budget constraint is drawn randomly from a uniform distribution $U[0, 3]$ and the ROI target is drawn randomly from a uniform distribution $U[0, 5]$. We randomly sample 10000 budget constraint and ROI target pairs. For each pair, we simulate 1,000,000 auctions and compute the average payment and the utility. The results are shown in Figure 3 and Figure 4.

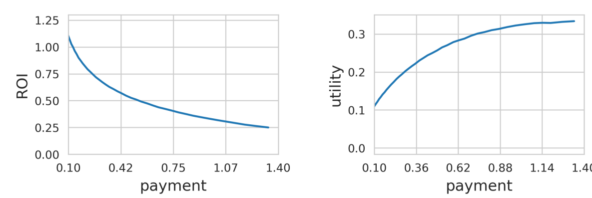

(a) ROI vs payment (b) Utility vs payment

Figure 3: The buyer's expected ROI, utility and payment for different budget and ROI constraints.

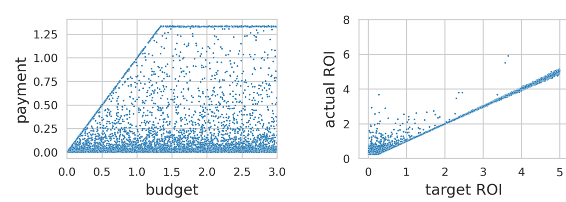

(a) Payment vs budget constraint (b) Actual ROI vs target ROI

Figure 4: Comparison between the buyer's actual payment and ROI and the corresponding constraints.

In Figure 3(a), the blue curve show how the actual ROI and payment changes. Note that although we have shown that the ROI approaches infinity as the payment goes to 0, the actual ROI can never reach infinity in the experiments.

Figure 3(b) shows that the buyer's utility function is indeed concave with respect to the payment. A mixed strategy leads to a point in the figure that is a convex combination of different points in the curve. Since the curve is convex, a mixed strategy cannot benefit the buyer. The curve also means that the effect of investing more money eventually diminishes.

Figure 4(a) shows that the buyer's actual payment is always lower than the buyer's budget satisfying the buyer's budget constraint. Similarly, Figure 4(b) shows that for the buyer's actual ROI is always higher than the buyer's ROI target, satisfying the buyer's ROI constraint.

Then we conduct experiments based on the open data set iPinYou [Liao *et al.*, 2014]. The iPinYou data set is published by a major demand-side platform (DSP) and contains auction logs of 24 days in 3 auction seasons. The data set includes 78 million bid records with 24 million impressions. In our experiment, we select one day's impression log data in the 2nd season to estimate the competitor's highest bid's distribution.

Our chosen data contains five advertisers. We select one of them and use the "paying price" column as the competitor's highest bid. The advertiser participates in different auction platforms such as Goolge, Alibaba, and Tencent. We focus on the data from one platform and plot histogram of the competitor's highest bids in Figure 5(a). We fit the data to a log-normal distribution and plot the PDF of the fitted distribution as the orange curve.

Unfortunately, in the iPinYou dataset, the platform use a special strategy which always places a very high bid to win as many auctions as possible to collect data about the competitors' highest bids. Therefore, there is no data about the valuation of the advertiser in the DSP. To conduct the experiments, we assume that there are 3 major bidders [Shen *et al.*, 2020] in the auction platform, and that they have i.i.d. valuations and bid their values. Also, we assume that our chosen buyer also follows the same value distribution, which we show in Figure 5(b).

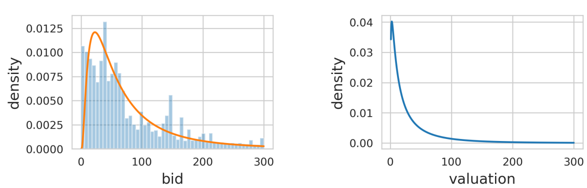

(a) Competitors' value distribution. (b) Buyer i's value distribution.

Figure 5: Value distributions obtained from the iPinYou data set.

Similar to the previous experiment, the budget constraint is drawn randomly from uniform distribution $U[0, 300]$ and the ROI target is drawn randomly from distribution $U[0, 5]$. We randomly sample 10000 budget constraint and ROI target pairs and simulate $1,000,000$ auctions for each pair to compute the expected payment and utility. The results are shown in Figure 6 and Figure 7.

The ROI is quite significant when the buyer's available budget is small, meaning that a small investment can lead to a substantial profit. In Figure 7(b), we can see a flat line when the target ROI is relatively small. This is because, according to Lemma 3, there is a lower bound γ^L for the actual ROI. These points are also the highest points in Figure 7(a), since when the ROI reaches the lower bound, the payment also hits the upper bound B^H.

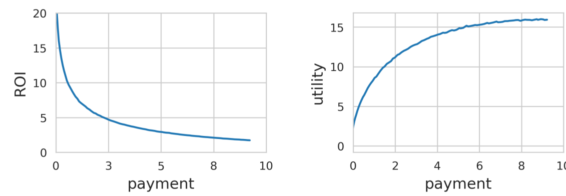

(a) ROI vs payment (b) Utility vs payment

Figure 6: The buyer's ROI and utility with payment.

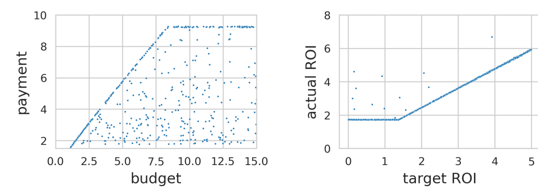

(a) Payment vs budget constraint. (b) Actual ROI vs target ROI.

Figure 7: The buyer's actual payment and ROI, and the corresponding constraints.

Figure 8 shows how the welfare changes with respect to β in both experiments. Since there is only one buyer with financial constraints in our experiments, Figure 8 shows how much it can affect the welfare by changing β for only one buyer. Note that the optimal welfare is simply the point with $\beta = 1$. Even when $\beta = 0$, our auto-bidding strategy achieves about 96% and 84% of the optimal welfare, In the experiments with the synthetic data set and the iPinYou data set, respectively. Interestingly, in Figure 8(a), the welfare drops quickly when β decreases, while in Figure 8(b), the welfare ratio is still above 90% for $\beta = 0.2$.

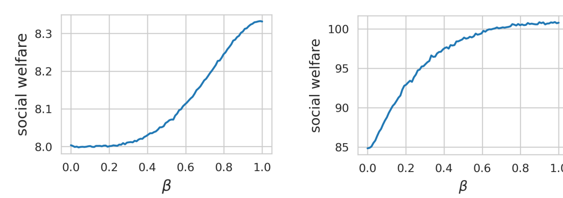

(a) Welfare of the experiments with synthetic data set. (b) Welfare of the experiments with iPinYou data set.

Figure 8: The relationship between the social welfare with β.

7 Conclusion

In this paper, we consider the auto-bidding problem in second-price auctions with the buyers having both budget and ROI constraints. We consider the buyer's bidding strategy and show that reporting the truthful budget and ROI targets is the optimal strategy even if they consider mixed strategies. We also show that the equilibrium exists under some assumptions and give a sufficient condition to find the equilibrium using iterated best response.

Acknowledgements

We thank all the anonymous reviewers for their helpful comments. This work was supported in part by the National Natural Science Foundation of China (No. 62106273 & No. 72192805), the Fundamental Research Funds for the Central Universities, and the Research Funds of Renmin University of China (22XNKJ16).

References

[Aggarwal *et al.*, 2006] Gagan Aggarwal, Ashish Goel, and Rajeev Motwani. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 1–7, 2006.

[Aggarwal *et al.*, 2019] Gagan Aggarwal, Ashwinkumar Badanidiyuru, and Aranyak Mehta. Autobidding with constraints. In *International Conference on Web and Internet Economics*, pages 17–30. Springer, 2019.

[Balseiro and Gur, 2019] Santiago R Balseiro and Yonatan Gur. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, 65(9):3952–3968, 2019.

[Balseiro *et al.*, 2015] Santiago R Balseiro, Omar Besbes, and Gabriel Y Weintraub. Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science*, 61(4):864–884, 2015.

[Balseiro *et al.*, 2017] Santiago Balseiro, Anthony Kim, Mohammad Mahdian, and Vahab Mirrokni. Budget management strategies in repeated auctions. In *Proceedings of the 26th International Conference on World Wide Web*, pages 15–23, 2017.

[Balseiro *et al.*, 2021a] Santiago Balseiro, Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. Robust auction design in the auto-bidding world. *Advances in Neural Information Processing Systems*, 34:17777–17788, 2021.

[Balseiro *et al.*, 2021b] Santiago R Balseiro, Yuan Deng, Jieming Mao, Vahab S Mirrokni, and Song Zuo. The landscape of auto-bidding auctions: Value versus utility maximization. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 132–133, 2021.

[Balseiro *et al.*, 2022] Santiago Balseiro, Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. Optimal mechanisms for value maximizers with budget constraints via target clipping. *Available at SSRN*, 2022.

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

[Brouwer, 1911] Luitzen Egbertus Jan Brouwer. Über abbildung von mannigfaltigkeiten. *Mathematische annalen*, 71(1):97–115, 1911.

[Chen *et al.*, 2021] Xi Chen, Christian Kroer, and Rachitesh Kumar. The complexity of pacing for second-price auctions. *arXiv preprint arXiv:2103.13969*, 2021.

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

[Conitzer *et al.*, 2017] Vincent Conitzer, Christian Kroer, Eric Sodomka, and Nicolas E Stier-Moses. Multiplicative pacing equilibria in auction markets. *arXiv preprint arXiv:1706.07151*, 2017.

[Conitzer *et al.*, 2018] Vincent Conitzer, Christian Kroer, Debmalya Panigrahi, Okke Schrijvers, Eric Sodomka, Nicolas E Stier-Moses, and Chris Wilkens. Pacing equilibrium in first-price auction markets. *arXiv preprint arXiv:1811.07166*, 2018.

[Conitzer *et al.*, 2022] Vincent Conitzer, Christian Kroer, Eric Sodomka, and Nicolas E Stier-Moses. Multiplicative pacing equilibria in auction markets. *Operations Research*, 70(2):963–989, 2022.

[Cramton *et al.*, 2004] Peter Cramton, Yoav Shoham, Richard Steinberg, et al. Combinatorial auctions. Technical report, University of Maryland, Department of Economics-Peter Cramton, 2004.

[Deng *et al.*, 2021] Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. Towards efficient auctions in an auto-bidding world. In *Proceedings of the Web Conference 2021*, pages 3965–3973, 2021.

[Edelman *et al.*, 2007] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review*, 97(1):242–259, 2007.

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

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

[Gummadi *et al.*, 2012] Ramakrishna Gummadi, Peter Key, and Alexandre Proutiere. Repeated auctions under budget constraints: Optimal bidding strategies and equilibria. In *the Eighth Ad Auction Workshop*, volume 4, 2012.

[Liao *et al.*, 2014] Hairen Liao, Lingxiao Peng, Zhenchuan Liu, and Xuehua Shen. ipinyou global rtb bidding algorithm competition dataset. In *Proceedings of the Eighth International Workshop on Data Mining for Online Advertising*, pages 1–6, 2014.

[Mehta, 2022] Aranyak Mehta. Auction design in an autobidding setting: Randomization improves efficiency beyond vcg. In *Proceedings of the ACM Web Conference 2022*, pages 173–181, 2022.

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

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

[Shen *et al.*, 2020] Weiran Shen, Binghui Peng, Hanpeng Liu, Michael Zhang, Ruohan Qian, Yan Hong, Zhi Guo, Zongyao Ding, Pengjun Lu, and Pingzhong Tang. Reinforcement mechanism design: With applications to dynamic pricing in sponsored search auctions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2236–2243, 2020.

[Statista, 2022] Statista. Digital ad spend worldwide 2026. https://www.statista.com/statistics/237974/online-advertising-spending-worldwide/, 2022. Accessed: 2023-05-25.

[Varian, 2007] Hal R Varian. Position auctions. *international Journal of industrial Organization*, 25(6):1163–1178, 2007.

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

Appendix

A Omitted Proofs in Section 3

A.1 Proof of Lemma 1

Proof. Taking the derivative of $P_i(\beta_i)$ w.r.t. β_i, we get:

$$P_i'(\beta_i) = \int_0^{\bar{v}} \beta_i v_i^2 g_i(\beta_i v_i) f_i(v_i)\, \mathrm{d}_i v_i, \tag{9}$$

which is clearly positive. □

A.2 Proof of Lemma 2

Proof. Taking the derivative of $h_i(\beta_i)$ w.r.t. β yields:

$$h_i'(\beta_i) = \int_0^{\bar{v}} v_i^2[(1+\gamma_i)\beta_i - 1] g_i(\beta_i v_i) f_i(v_i)\, \mathrm{d}v_i$$

Therefore, when $\beta_i \in \left(0, \frac{1}{1+\gamma_i}\right]$, we have $h_i'(\beta_i) < 0$, and $h_i(\beta_i)$ is monotone decreasing. And when $\beta_i > \frac{1}{1+\gamma_i}$, we have $h_i'(\beta_i) > 0$, and $h_i(\beta_i)$ is monotone increasing. □

A.3 Proof of Lemma 3

Proof. It is easy to check that $h_i(0) = 0$. According to Lemma 2, we have $h_i(\frac{1}{1+\gamma_i}) < h_i(0) = 0$. Lemma 2 also says that $h_i(\beta_i)$ is monotone increasing when $\beta_i > \frac{1}{1+\gamma_i}$. This implies that there is at most one point $\hat{\beta} > 0$ such that $h_i(\hat{\beta}) = 0$.

If such $\hat{\beta}$ exists, then we clearly have that $h_i(\beta) \le 0, \forall \beta \in (0, \hat{\beta}]$ and $h_i(\beta) > 0, \forall \beta > \hat{\beta}$. In this case, the ROI constraint is equivalent to $\beta \le \hat{\beta}$.

If such $\hat{\beta}$ does not exist, then it must be the case that $h_i(\beta) < 0, \forall \beta > 0$. In this case, all $\beta > 0$ satisfies $h_i(\beta) < 0$.

□

A.4 Proof of Lemma 4

Proof. The derivative of $U_i(\beta_i)$ is:

$$U_i'(\beta_i) = (1-\beta_i) \int_0^{\bar{v}} v_i^2 g_i(\beta_i v_i) f_i(v_i)\, \mathrm{d}v_i, \tag{10}$$

Then we can get that when $\beta_i \le 1$, $U_i'(\beta_i) \ge 0$ and when $\beta_i > 1$, $U_i'(\beta_i) < 0$. □

A.5 Proof of Lemma 5

Proof. The ROI of buyer i is:

$$\frac{\mathbf{E}[(v_i - D_i)\mathbb{I}(\beta_i v_i \ge D_i)]}{\mathbf{E}[D_i \mathbb{I}(\beta_i v_i \ge D_i)]} = \frac{\mathbf{E}[v_i \mathbb{I}(\beta v_i \ge D_i)]}{\mathbf{E}[D_i \mathbb{I}(\beta_i v_i \ge D_i)]} - 1. \tag{11}$$

Define $\eta_i(\beta_i) = \frac{\mathbf{E}[v_i \mathbb{I}(\beta v_i \ge D_i)]}{\mathbf{E}[D_i \mathbb{I}(\beta v_i \ge D_i)]} = \frac{U_i(\beta_i) + P_i(\beta_i)}{P_i(\beta_i)}$. Consider the derivative of $\eta_i(\beta_i)$:

$$\eta_i'(\beta_i) = \frac{[U_i'(\beta_i) + P_i'(\beta_i)]P_i(\beta_i) - [U_i(\beta_i) + P_i(\beta_i)]P_i'(\beta_i)}{P_i^2(\beta_i)}. \tag{12}$$

The denominator is clearly positive. Plugging in Equation (9) and (10), the numerator can be written as:

$$\left[\int_0^{\bar{v}} \int_0^{\beta_i v_i} (D_i - v_i) g_i(D_i)\, \mathrm{d}D_i f_i(v_i)\, \mathrm{d}v_i\right] \left[\int_0^{\bar{v}} \beta_i v_i^2 g_i(\beta_i v_i) f_i(v_i)\, \mathrm{d}v_i\right], \tag{13}$$

where the second term is clearly positive, while the first term is non-positive since inside the integral, $D_i \le \beta_i v_i \le v_i$.

If β_i approaches 0, both $U_i(\beta_i)$ and $P_i(\beta_i)$ goes to 0. Thus the ROI becomes:

$$\lim_{\beta_i \to 0} \Gamma_i(\beta_i) = \lim_{\beta_i \to 0} \frac{U_i'(\beta_i)}{P_i'(\beta_i)} = \lim_{\beta_i \to 0} \frac{1-\beta_i}{\beta_i} = +\infty.$$

□

A.6 Proof of Theorem 1

Proof. Clearly, the solid curve can also be defined by $C = \{(P_i(\beta_i), \Gamma_i(\beta_i)) \mid \beta_i \in (0,1]\}$. Since $P_i(\beta_i)$ is an increasing function, while $\Gamma_i(\beta_i)$ is decreasing, it follows that Γ_i is decreasing with respect to P_i.

Let $\Sigma(B_i, \gamma_i) = \{(B, \gamma) \mid B \le B_i, \gamma \ge \gamma_i\}$ be the area satisfying the two constraint. Therefore, for any (B_i, γ_i), we need to choose a β_i such that the actual $(P_i(\beta_i), \Gamma_i(\beta_i))$ lies in $\Sigma(B_i, \gamma_i) \cap C$. And since $U_i(\beta_i)$ is an increasing function, we need to choose the maximum β_i possible, which must be the rightmost point in $\Sigma(B_i, \gamma_i) \cap C$, as $P_i(\beta_i)$ is also increasing in β_i.

We discuss the point (B_i, γ_i) case by case as follows:

1. If (B_i, γ_i) lies in area 3, the rightmost point in $\Sigma(B_i, \gamma_i) \cap C$ is clearly (B^H, γ^L);
2. If (B_i, γ_i) lies in area 2, the rightmost point in $\Sigma(B_i, \gamma_i) \cap C$ is at the lower edge of $\Sigma(B_i, \gamma_i)$. Thus the point satisfies $\Gamma_i = \gamma_i$;
3. If (B_i, γ_i) lies in area 2, the rightmost point in $\Sigma(B_i, \gamma_i) \cap C$ is at the right edge of $\Sigma(B_i, \gamma_i)$. Thus the point satisfies $P_i = B_i$;

□

A.7 Proof of Theorem 2

Proof. According to Lemma 1 and 3, both the budget and ROI constraints set two upper bounds β_i^1 and β_i^2 for β_i (β_i^2 can be $+\infty$ though). Combining Corollary 2, we know that the optimal solution to Program 3 is $\min\{\beta_i^1, \beta_i^2, 1\} = \min\{\min\{\beta_i^1, 1\}, \min\{\beta_i^2, 1\}\}$.

For the budget constraint, we know from Lemma 1 that $P_i(\beta_i)$ is monotone increasing w.r.t. β_i. This implies:

- If $B_i \ge P_i(1) = B^H$, then $\beta_i^1 \ge 1$ and $\min\{\beta_i^1, 1\} = 1$;
- If $B_i < P_i(1) = B^H$, there is exactly one β_i^1 such that $P_i(\beta_i^1) = B_i$. In this case, we can use the binary search algorithm to solve $P_i(\beta_i^1) = B_i$ since $P_i(\beta_i)$ is monotone.

As for the ROI constraint, Lemma 5 says that the ROI is monotone decreasing w.r.t. β_i. Similar analysis shows that $\min\{\beta_i^2, 1\}$ either equals 1 if $\gamma_i \le \gamma^L$ or can be found using the binary search algorithm. □

B Omitted Proofs in Section 4

B.1 Proof of Lemma 6

Proof. We prove the result by contradiction. Suppose reporting (B_i', γ_i') gives buyer i a strictly higher utility. The auto-bidding system solves program 3 according to (B_i', γ_i'), and gives an optimal solution β'. If the actual payment $P(\beta')$ and ROI $\Gamma(\beta')$ violate buyer i's true constraints, their utility becomes $-\infty$ as these constraints are hard constraints, which contradicts to the assumption that reporting (B_i', γ_i') yields a higher utility. If $P(\beta')$ and $\Gamma(\beta')$ do not violate buyer i's true constraints, then β' is also a feasible solution to the program 3 with the true (B_i, γ_i). Therefore, the optimal solution to the program 3 with the true (B_i, γ_i) should give the buyer a weakly higher utility, also a contradiction. □

B.2 Proof of Lemma 7

Proof. The derivative of U_i with respect to P_i is:

$$\frac{\mathrm{d}U_i}{\mathrm{d}P_i} = \frac{U_i'(\beta_i)}{P_i'(\beta_i)} = \frac{1}{\beta_i} - 1, \tag{14}$$

where the second equation is obtained by combining Equation (9) and (10). Therefore,

$$\frac{\mathrm{d}^2U_i}{\mathrm{d}P_i^2} = -\frac{1}{\beta_i^2}\frac{\mathrm{d}\beta_i}{\mathrm{d}P_i} = -\frac{1}{\beta_i^2 P_i'(\beta_i)}. \tag{15}$$

Since $P_i'(\beta_i) \geq 0$ according to Lemma 1, we have that $\frac{\mathrm{d}^2U_i}{\mathrm{d}P_i^2} \leq 0$, which implies that the function $U_i(P_i)$ is concave. □

B.3 Proof of Theorem 3

Proof. For each possible report (B_i', γ_i'), the auto-bidding system computes a β_i' by solving the corresponding program. Thus a mixed strategy results in a distribution over all possible $\beta_i \in (0, 1]$, hence also a distribution over possible payments P_i. Let B_i and γ_i be the true budget and ROI target of buyer i, and suppose that buyer i uses a mixed strategy that satisfies the two constraints. We show that there exists a pure strategy that give them a high utility. Since the mixed strategy satisfies the constraints, we have:

$$\mathbf{E}_{\beta_i}[P_i(\beta_i)] \leq B_i \quad \text{and} \quad \frac{\mathbf{E}_{\beta_i}[U_i(\beta_i)]}{\mathbf{E}_{\beta_i}[P_i(\beta_i)]} \geq \gamma_i. \tag{16}$$

Suppose β_i^* satisfies $P_i(\beta_i^*) = \mathbf{E}_{\beta_i}[P_i(\beta_i^*)]$. We claim that any strategy that leads to β_i^* is a better strategy than the mixed strategy. β_i^* clearly satisfies the budget constraint. As for the ROI constraint, we have:

$$\frac{\mathbf{E}_{\beta_i}[U_i(\beta_i)]}{\mathbf{E}_{\beta_i}[P_i(\beta_i)]} = \frac{\mathbf{E}_{P_i}[U_i(P_i)]}{P_i(\beta_i^*)} \leq \frac{U_i(\mathbf{E}_{P_i}[P_i])}{P_i(\beta_i^*)}, \tag{17}$$

which implies that the ROI is higher than γ_i. Meanwhile, the utility $U_i(\mathbf{E}_{P_i}[P_i])$ is also higher than that of the mixed strategy. □

B.4 Proof of Lemma 8

Proof. The cumulative distribution of D_i is

$$G_i(D_i, \beta_{-i}) = \prod_{j \neq i} F_j\left(\frac{D_i}{\beta_j}\right).$$

The derivative of $G_i(D_i, \boldsymbol{\beta}_{-i})$ with respect to D_i is

$$g_i(D_i, \boldsymbol{\beta}_{-i}) = \sum_{k \neq i}\left[\frac{1}{\beta_k} f_k\left(\frac{D_i}{\beta_k}\right) \prod_{j \neq k, j \neq i} F_j\left(\frac{D_i}{\beta_j}\right)\right].$$

For the lower bound, we have:

$$\begin{aligned}
&D_i g_i(D_i, \boldsymbol{\beta}_{-i}) \\
=&D_i \sum_{k \neq i}\left[\frac{1}{\beta_k} f_k\left(\frac{D_i}{\beta_k}\right) \prod_{j \neq k, j \neq i} F_j\left(\frac{D_i}{\beta_j}\right)\right] \\
\geq&(n-1)m^{n-1}D_i^{n-1}.
\end{aligned}$$

Now we prove the upper bound. For each buyer i, $v_i \leq \bar{v}$. Then we have the following two cases:

- If $D_i \geq \beta_k \bar{v}$, we have $f_k\left(\frac{D_i}{\beta_k}\right) = 0$;
- If $D_i < \beta_k \bar{v}$, we have $f_k\left(\frac{D_i}{\beta_k}\right) \leq M$.

In any case, $\frac{D_i}{\beta_k} f_k\left(\frac{D_i}{\beta_k}\right) \leq \bar{v}M$ holds. Therefore,

$$\begin{aligned}
&D_i g_i(D_i, \boldsymbol{\beta}_{-i}) \\
=&D_i \sum_{k \neq i}\left[\frac{1}{\beta_k} f_k\left(\frac{D_i}{\beta_k}\right) \prod_{j \neq k, j \neq i} F_j\left(\frac{D_i}{\beta_j}\right)\right] \\
\leq&\sum_{k \neq i} \frac{D_i}{\beta_k} f_k\left(\frac{D_i}{\beta_k}\right) \\
\leq&(n-1)M\bar{v},
\end{aligned}$$

□

B.5 Proof of Lemma 9

Proof. According to Lemma 1, we must have $\beta_i \leq P_i^{-1}(B_i)$. Since $P_i(\beta_i)$ is monotone increasing, we have that $P_i^{-1}(B_i)$ is also increasing. According to Lemma 4, we need to set β_i to be as large as possible in the interval$[0, 1]$. Therefore, the optimal strategy for buyer i is $\beta_i = \min\{P_i^{-1}(B_i), 1\} = \min\{P_i^{-1}(B_i), P_i^{-1}(B^H)\} \geq P_i^{-1}(\underline{B})$. It follows that β_i should always satisfy $P_i(\beta_i) \geq \underline{B}$. Or equivalently,

$$\int_0^{\bar{v}} \int_0^{\beta_i v_i} D_i g_i(D_i) \mathrm{d}D_i f_i(v_i) \mathrm{d}v_i \geq \underline{B}.$$

This means that $\beta_i \geq \frac{2\underline{B}}{(n-1)M^2\bar{v}^3}$, since otherwise, we have:

$$\begin{aligned}
&\int_0^{\bar{v}} \int_0^{\beta_i v_i} D_i g_i(D_i) \mathrm{d}D_i f_i(v_i) \mathrm{d}v_i \\
\leq&\int_0^{\bar{v}} \int_0^{\beta_i v_i} (n-1)M\bar{v} \mathrm{d}D_i M \mathrm{d}v_i \\
=&\frac{1}{2}(n-1)M^2\bar{v}^3\beta_i \\
<&\underline{B},
\end{aligned}$$

where the first inequality is from Lemma 8.

To satisfy the ROI constraint, according to Lemma 2, we have that $\beta_i \geq \frac{1}{1+\gamma_i} \geq \frac{1}{1+\bar{\gamma}}$.

Combining the two constraints, we obtain:

$$\beta_i \geq \max\left\{\frac{2\underline{B}}{(n-1)M^2\bar{v}^3}, \frac{1}{1+\bar{\gamma}}\right\} = \underline{\beta}.$$

□

C Omitted Proofs in Section 5

C.1 Proof of Lemma 10

Proof. Let β_{-i}^k be the strategy where β_k is changed to $\beta_k + \epsilon_k$. Therefore, we have:

$$\begin{aligned}
&\left|g_i(D_i, \beta_{-i}) - g_i(D_i, \beta_{-i}^k)\right| \\
=&\sum_{l\neq i,k}\left[\frac{1}{\beta_l} f_l\left(\frac{D_i}{\beta_l}\right) F_k\left(\frac{D_i}{\beta_k+\epsilon_k}\right)\prod_{j\neq l,i,k} F_j\left(\frac{D_i}{\beta_j}\right)\right] \\
&+\frac{1}{\beta_k+\epsilon_k} f_k\left(\frac{D_i}{\beta_k+\epsilon_k}\right)\prod_{j\neq i,k} F_j\left(\frac{D_i}{\beta_j}\right) \\
&-\sum_{l\neq i}\left[\frac{1}{\beta_l} f_l\left(\frac{D_i}{\beta_l}\right)\prod_{j\neq i,l} F_j\left(\frac{D_i}{\beta_j}\right)\right] \\
=&\sum_{l\neq i,k}\left\{\frac{1}{\beta_l} f_l\left(\frac{D_i}{\beta_l}\right)\left[F_k\left(\frac{D_i}{\beta_k+\epsilon_k}\right)\right.\right. \\
&\left.\left.-F_k\left(\frac{D_i}{\beta_k}\right)\right]\prod_{j\neq l,i,k} F_j\left(\frac{D_i}{\beta_j}\right)\right\} \\
&+\left[\frac{1}{\beta_k+\epsilon_k} f_k\left(\frac{D_i}{\beta_k+\epsilon_k}\right)\right. \\
&\left.-\frac{1}{\beta_k} f_k\left(\frac{D_i}{\beta_k}\right)\right]\prod_{j\neq i,k} F_j\left(\frac{D_i}{\beta_j}\right) \\
\leq& M\frac{1}{\underline{\beta}}\cdot M\left|\frac{D_i}{\beta_k+\epsilon_k} - \frac{D_i}{\beta_k}\right| + \xi\left|\frac{D_i}{\beta_k+\epsilon_k} - \frac{D_i}{\beta_k}\right| \\
=&\frac{(M^2+\underline{\beta}\xi)D_i}{\underline{\beta}\beta_k(\beta_k+\epsilon_k)}|\epsilon_k| \\
\leq&\frac{(M^2+\underline{\beta}\xi)D_i}{\underline{\beta}^3}|\epsilon_k|.
\end{aligned}$$

□

C.2 Proof of Lemma 11

Proof. From Theorem 1, we know that if bidder i's financial constraints are (B_i, γ_i), we can translate the bidder's ROI constraint to a budget constraint $\hat{B}_i$. Then the bidder's constraints can be represented by a budget constraint $B_i^* = \min\{B_i, \hat{B}_i\}$. Let β_i represent the optimal strategy before bidder k changes their strategy and $\hat{\beta}_i$ the optimal strategy after bidder k changes. Then we have:

$$\mathbf{E}[D_i\mathbb{I}\{\beta_i v_i \geq D_i\}] = B_i^*$$
$$\mathbf{E}[D_i\mathbb{I}\{\hat{\beta}_i v_i \geq D_i\}] = B_i^*$$

It follows that

$$\begin{aligned}
&\int_0^{\bar{v}}\int_0^{\beta_i v_i} D_i g_i(D_i, \beta_{-i})\mathrm{d}D_i f_i(v_i)\mathrm{d}v_i \\
=&\int_0^{\bar{v}}\int_0^{\hat{\beta}_i v_i} D_i g_i(D_i, \beta_{-i}^k)\mathrm{d}D_i f_i(v_i)\mathrm{d}v_i.
\end{aligned} \tag{18}$$

Combining Lemma 8 yields:

$$\begin{aligned}
&\int_0^{\bar{v}}\int_0^{\beta_i v_i} D_i g_i(D_i, \beta_{-i})\mathrm{d}D_i f_i(v_i)\mathrm{d}v_i \\
\leq&\int_0^{\bar{v}}\int_0^{\beta_i v_i}\left[D_i g_i(D_i, \beta_{-i}^k)\right. \\
&\left.+\frac{(M^2+\underline{\beta}\xi)|\epsilon|D_i^2}{\underline{\beta}^3}\right]\mathrm{d}D_i f_i(v_i)\mathrm{d}v_i
\end{aligned} \tag{19}$$

$$\begin{aligned}
&\int_0^{\bar{v}}\int_0^{\beta_i v_i} D_i g_i(D_i, \beta_{-i})\mathrm{d}D_i f_i(v_i)\mathrm{d}v_i \\
\geq&\int_0^{\bar{v}}\int_0^{\beta_i v_i}\left[D_i g_i(D_i, \beta_{-i}^k)\right. \\
&\left.-\frac{(M^2+\underline{\beta}\xi)|\epsilon|D_i^2}{\underline{\beta}^3}\right]\mathrm{d}D_i f_i(v_i)\mathrm{d}v_i
\end{aligned} \tag{20}$$

Combine (18), (19) and (20), and we get

$$\begin{aligned}
&\left\|\int_0^{\bar{v}}\int_0^{\hat{\beta}_i v_i} D_i g_i(D_i, \beta_{-i}^k)\mathrm{d}D_i f_i(v_i)\mathrm{d}v_i\right. \\
&\left.-\int_0^{\bar{v}}\int_0^{\beta_i v_i} D_i g_i(D_i, \beta_{-i}^k)\mathrm{d}D_i f_i(v_i)\mathrm{d}v_i\right\| \\
=&\left\|\int_0^{\bar{v}}\int_{\beta_i v_i}^{\hat{\beta}_i v_i} D_i g_i(D_i, \beta_{-i}^k)\mathrm{d}D_i f_i(v_i)\mathrm{d}v_i\right\| \\
\leq&\frac{(M^2+\underline{\beta}\xi)|\epsilon|}{3\underline{\beta}^3}\mathbf{E}[v_i^3]
\end{aligned}$$

Also we can get that

$$\begin{aligned}
&\left\|\int_0^{\bar{v}}\int_{\beta_i v_i}^{\hat{\beta}_i v_i} D_i g_i(D_i, \beta_{-i}^k)\mathrm{d}D_i f_i(v_i)\mathrm{d}v_i\right\| \\
\geq&\int_0^{\bar{v}}\int_{\beta_i v_i}^{\hat{\beta}_i v_i}(n-1)m^{n-1}D_i^{n-1}\mathrm{d}D_i f_i(v_i)\mathrm{d}v_i \\
=&\frac{n-1}{n}m^{n-1}\mathbf{E}[v_i^n]\left\|\hat{\beta}_i^n - \beta_i^n\right\| \\
=&\frac{n-1}{n}m^{n-1}\mathbf{E}[v_i^n]\left\|(\hat{\beta}_i-\beta_i)\sum_{k=0}^{n-1}\hat{\beta}_i^{n-1-k}\beta_i^k\right\| \\
\geq&(n-1)m^{n-1}\mathbf{E}[v_i^n]\underline{\beta}^{n-1}\left\|\hat{\beta}_i-\beta_i\right\|
\end{aligned}$$

Then we can get

$$\|\hat{\beta}_i - \beta_i\| \le \frac{(ML + \underline{\beta}\xi)\,\mathbf{E}[v_i^3]|\epsilon|}{3(n-1)\underline{\beta}^{n+2}\,\mathbf{E}[v_i^n]m^{n-1}} \le \frac{(M^3 + M\underline{\beta}\xi)(n+1)|\epsilon|}{12(n-1)\underline{\beta}^{n+2}m^n\bar{v}^{n-3}}$$

□

C.3 Proof of Lemma 12

Proof. Let e_i denote the n-dimensional vector where the i-th element is 1 and others are 0. Note that

$$\begin{aligned}&\|X(\boldsymbol{\beta}) - X(\boldsymbol{\beta}+\boldsymbol{\epsilon})\|_1 \\ =&\left\|\sum_{l=0}^{n-1}\left[X\left(\boldsymbol{\beta}+\sum_{i=1}^{l}\epsilon_i e_i\right) - X\left(\boldsymbol{\beta}+\sum_{i=1}^{l+1}\epsilon_i e_i\right)\right]\right\|_1 \\ \le&\sum_{l=0}^{n-1}\left\|X\left(\boldsymbol{\beta}+\sum_{i=1}^{l}\epsilon_i e_i\right) - X\left(\boldsymbol{\beta}+\sum_{i=1}^{l+1}\epsilon_i e_i\right)\right\|_1.\end{aligned}$$

According to Lemma 11, we have that for each $0 \le l \le n-1$,

$$\begin{aligned}&\left\|X\left(\boldsymbol{\beta}+\sum_{i=1}^{l}\epsilon_i e_i\right) - X\left(\boldsymbol{\beta}+\sum_{i=1}^{l+1}\epsilon_i e_i\right)\right\|_1 \\ \le&\frac{(M^3 + M\underline{\beta}\xi)(n+1)|\epsilon_{l+1}|}{12(n-1)\underline{\beta}^{n+2}m^n\bar{v}^{n-3}}.\end{aligned}$$

Therefore,

$$\begin{aligned}&\|X(\boldsymbol{\beta}) - X(\boldsymbol{\beta}+\boldsymbol{\epsilon})\|_1 \\ \le&\sum_{l=0}^{n-1}\left\|X\left(\boldsymbol{\beta}+\sum_{i=1}^{l}\epsilon_i e_i\right) - X\left(\boldsymbol{\beta}+\sum_{i=1}^{l+1}\epsilon_i e_i\right)\right\|_1 \\ \le&\sum_{l=0}^{n-1}\frac{(M^3 + M\underline{\beta}\xi)(n+1)|\epsilon_{l+1}|}{12(n-1)\underline{\beta}^{n+2}m^n\bar{v}^{n-3}} \\ =&\frac{(M^3 + M\underline{\beta}\xi)(n+1)\|\boldsymbol{\epsilon}\|_1}{12(n-1)\underline{\beta}^{n+2}m^n\bar{v}^{n-3}}.\end{aligned}$$

□

C.4 Proof of Theorem 4

Proof. Consider function $X : [\underline{\beta}, 1]^n \mapsto [\underline{\beta}, 1]^n$. It is clear that the domain of function X is compact. And according to Lemma 12, function X is continuous since

$$\begin{aligned}&\lim_{\boldsymbol{\epsilon}\to 0}\|X(\boldsymbol{\beta}) - X(\boldsymbol{\beta}+\boldsymbol{\epsilon})\| \\ \le&\lim_{\boldsymbol{\epsilon}\to 0}\frac{(M^3 + M\underline{\beta}\xi)(n+1)\|\boldsymbol{\epsilon}\|}{12(n-1)\underline{\beta}^{n+2}m^n\bar{v}^{n-3}} = 0\end{aligned}$$

It follows from the Brouwer fixed point theorem ([Brouwer, 1911]) that a fixed point $\boldsymbol{\beta}^*$ of the function X exists. Such a fixed point is an equilibrium as for each i, β_i^* is the best response to $\boldsymbol{\beta}_{-i}^*$ by definition of X. □

C.5 Proof of Corollary 5

Proof. If $\frac{(M^3+M\underline{\beta}\xi)(n+1)}{12(n-1)\underline{\beta}^{n+2}m^n\bar{v}^{n-3}} < 1, \forall i \in [n]$, it is straightforward that function X becomes a contraction. Thus the process $\boldsymbol{\beta}^{t+1} = X(\boldsymbol{\beta}^t)$ converges to a fixed point $\boldsymbol{\beta}^*$, which is an equilibrium according to Theorem 4. □

C.6 Proof of Theorem 5

Proof. For any realized value profile $(v_1, v_2, \ldots, v_n)$, we can, without loss of generality, assume that $v_1 \ge v_2 \ge \cdots \ge v_n$. In this case, the optimal welfare is simply v_1.

Let $(\beta_1, \beta_2, \ldots, \beta_n)$ be any equilibrium profile $\boldsymbol{\beta}$ of all buyers. From Lemma 9, we can get that for any buyer $i \in [n], \beta_i \ge \underline{\beta}$.

If buyer 1 is the buyer with the largest $\beta_i v_i$, then the welfare of the auto-bidding system is still v_1. If buyer j is the buyer with the largest $\beta_i v_i$, then we have:

$$\beta_j v_j \ge \beta_1 v_1.$$

This means:

$$v_j \ge \frac{\beta_1}{\beta_j}v_1 \ge \frac{\underline{\beta}}{1}v_i = \underline{\beta}v_1,$$

which implies that the welfare is at least $\underline{\beta}$ times the optimal welfare. Combining the two cases and taking expectation over all possible value profiles completes the proof. □