

MIT Open Access Articles

Auction Design for ROI-Constrained Buyers

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Golrezaei, Negin, Lobel, Ilan and Paes Leme, Renato. 2021. "Auction Design for ROI-Constrained Buyers." Proceedings of the Web Conference 2021.

As Published: 10.1145/3442381.3449841

Publisher: Association for Computing Machinery (ACM)

Persistent URL: <https://hdl.handle.net/1721.1/144165>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution 4.0 International license



Auction Design for ROI-Constrained Buyers

Negin Golrezaei
golrezae@mit.edu
Massachusetts Institute of
Technology (MIT)
Cambridge, MA, USA

Ilan Lobel
ilobel@stern.nyu.edu
New York University (NYU)
New York, NY, USA

Renato Paes Leme
renatoppl@google.com
Google Research,
New York, NY, USA

ABSTRACT

We combine theory and empirics to (i) show that some buyers in online advertising markets are financially constrained and (ii) demonstrate how to design auctions that take into account such financial constraints. We use data from a field experiment where reserve prices were randomized on Google’s advertising exchange (AdX). We find that, contrary to the predictions of classical auction theory, a significant set of buyers lowers their bids when reserve prices go up. We show that this behavior can be explained if we assume buyers have constraints on their minimum return on investment (ROI). We proceed to design auctions for ROI-constrained buyers. We show that optimal auctions for symmetric ROI-constrained buyers are either second-price auctions with reduced reserve prices or subsidized second-price auctions. For asymmetric buyers, the optimal auction involves a modification of virtual values. Going back to the data, we show that using ROI-aware optimal auctions can lead to large revenue gains and large welfare gains for buyers.

KEYWORDS

Return on Investment, Performance Buyers, Mechanism Design, Online Advertising

ACM Reference Format:

Negin Golrezaei, Ilan Lobel, and Renato Paes Leme. 2021. Auction Design for ROI-Constrained Buyers. In *Proceedings of the Web Conference 2021 (WWW ’21), April 19–23, 2021, Ljubljana, Slovenia*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3442381.3449841>

1 INTRODUCTION

Second-price auctions under financial constraints. Online advertisement is a large and growing business, having generated US revenues greater than 70 billion dollars in 2016 and 80 billion in 2017. Most online ad markets are based on second-price auctions. A key feature that makes second-price auctions so appealing to both market designers and market participants (buyers) is that they are truthful. That is, it is a dominant strategy for buyers to report their true valuations as their bids. This statement, however, is predicated on an important and underappreciated assumption: the truthfulness of second-price auctions depends on buyers having quasilinear preferences. A buyer’s utility is quasilinear if it can be computed by subtracting the amount paid from the expected value from winning the auction.

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW ’21, April 19–23, 2021, Ljubljana, Slovenia

© 2021 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-8312-7/21/04.

<https://doi.org/10.1145/3442381.3449841>

Quasilinearity is a stronger assumption than might appear at first sight. In particular, it implies the buyers are not financially constrained. When a buyer is financially constrained, existing theory shows that he should shade his bid in a second-price auction ([14] and [3]).¹ That is, a financially constrained buyer should submit a bid below his true value in such an auction.

Our empirical findings. In this paper, we show empirically that a fraction of the buyers in online ad markets behave as if they are financially constrained. We do this by using data from a field experiment that was run by Google on its advertising exchange (AdX). In the experiment, impressions were offered for sale via second-price auctions with randomized reserve prices. For a quasilinear buyer, there are two sensible answers to a reserve price being increased: to maintain or to raise his bid. However, we show that in practice many buyers *lower* their bids in response to higher reserve prices! While such shading behavior is inconsistent with quasilinear utilities, we show in Theorem 3.1 that this is in accordance with equilibrium behavior under financial constraints.

We detect shading behavior by focusing on a group that we call *performance buyers*, which are buyers that are focused on the short-term returns they gain from their ads (as opposed to brand buyers, who care more about the long-term visibility of their brands). Some categories of performance buyers include in-app advertisers and retargeting advertisers (advertisers that show consumers products they previously browsed). This finding is consistent with a performance buyer’s goal being to maximize returns from his ad purchases given his financial constraints.

Return on investment. The prior literature regarding auctions with financially constrained buyers typically assumes buyers have budgets ([15], [7], [16], and [19]). That is, the typical assumption is that buyers have finite resources and no access to capital markets. A more standard business approach to modeling financial constraints is to assume that firms see online ads as one of several available investment opportunities and that they avail themselves of the investments that generate the highest returns. With this formulation in mind, we propose a framework where buyers require a certain return on investment (ROI) in order to participate in the system. That is, each buyer has a target ROI, and he will not participate unless he obtains a certain return on his investment. As we show in Section 4, ROI-constrained behavior better matches our empirical evidence than budget constraints.

Auction design under ROI constraints. Given our empirical findings, we focus on the problem of how to optimize an auction given that buyers have ROI constraints. We consider first the case of symmetric buyers, with all buyers having the same target ROI. We show in Theorem 5.3 that an optimal auction for ROI-constrained

¹Throughout the paper, we use female pronouns to refer to the seller and male ones to refer to the buyers.

buyers takes one of three forms, depending on the value of the target ROI. The first form is the standard second-price auction with *the Myersonian reserve price* (optimal for low target ROIs), the second one is a second-price auction with a reduced reserve price (optimal for moderate target ROIs), and the third one is a second-price auction without reserve plus a participation subsidy for each buyer (optimal for high target ROIs). With asymmetric buyers and different target ROIs for the different buyers, the optimal auction is a bit more complex, as is true in classical auction theory. As we show in Theorem 5.5, the optimal auction for ROI-constrained buyers can be interpreted in terms of *modified virtual values*, where the modification is designed to make a buyer's ROI constraint binding when appropriate.

Counterfactual analysis. Combining our empirical data with our optimal auction design, we estimate the potential practical effects of accounting for ROI constraints when designing ad auctions. We use the optimal auction design for symmetric buyers due to its simplicity and practicality, and we assume multiple independent copies of real-world buyers are present. Accounting for ROI constraints leads the designer to reduce reserve prices. Therefore, using an ROI-aware auction not only increases revenues but also increases the welfare of the buyers.

1.1 Related Literature

Auctions with financially constrained buyers. The first paper to argue that auction designers should take into account the financial constraints of buyers was [15]. They considered a setting with symmetric buyers with identical budgets and showed that an all-pay auction with reserve is optimal. Over time, a long series of papers in economics and computer science [5, 7, 8, 10–12, 16, 17, 19] has improved this result and extended it in various directions. The most complete of these results is [19], which considers a multidimensional setting where both valuations and budgets are private information.

Shading in second-price auctions. A key idea that enables us to empirically identify financially constrained buyers is that such buyers shade their bids in second-price auctions. That is, they bid less than their true valuations in equilibrium. While we prove this result for our specific model, this idea has been previously introduced in the literature on repeated second-price auctions with budget-constrained buyers [3, 4, 9, 14].

Return on investment. Most papers on auctions with financially constrained buyers assume budgets, but there are a few that consider ROI instead, as we do. The first papers to do so are [20] and [6], which consider an ROI-based heuristic for deciding on how to bid across different keywords in search advertising. Both papers approach the problem from the dynamical systems view and find that such a heuristic leads to cyclic behavior. More recently, [22] studied auction design for buyers that do not care about payments as long as they are under budget or satisfy an ROI constraint. Their main result is that any monotone allocation can be implemented truthfully in such settings.

Empirical evidence. The only empirical work we could find on financial constraints in auctions is by [2]. Using search advertising data from Yahoo, they find that most buyers behave consistently with the ROI-constrained heuristic described in [6]. Our empirical

work differs from [2] in several ways. First, we use 2017 data from a display advertising exchange, instead of 2002–2003 data from search advertising. Display and search are roughly equally large markets, but they are vastly different in terms of how the markets operate. Second, our paper analyzes a setting with second-price auctions, while theirs considers first price auctions. [7] show that financially constrained buyers behave in very different ways in first and second-price auctions and that the revenue equivalency theorem does not hold in such settings. Third, perhaps most importantly, we use data from a field experiment, rather than observational data. Field experiments allow us to determine causality with greater certainty.

Performance buyers. Our paper focuses on a set that we call *performance buyers*, which are buyers that focus on short-term financial performance above long-term outcomes. This idea is borrowed from [13].

2 MODEL

We consider a model that consists of a seller, who owns an indivisible item, and n interested buyers. Each buyer $i \in [n]$ has a valuation for the item being sold that is represented by v_i , where $[n] = \{1, 2, \dots, n\}$. Each valuation v_i is independently drawn from a probability distribution $F_i : [\underline{v}, \bar{v}] \rightarrow [0, 1]$ for some $0 \leq \underline{v} < \bar{v} \leq \infty$, with a continuous probability density function f_i . While the distributions F_i 's are common knowledge, the valuation v_i is known only to buyer i . We denote the vector of valuations (v_1, v_2, \dots, v_n) by \mathbf{v} . Furthermore, we let \mathbf{v}_{-i} denote the vector of valuations of all buyers except for buyer i 's; that is, $\mathbf{v}_{-i} = (v_1, v_2, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$.

Let $\phi_i(v_i) = v_i - (1 - F_i(v_i))/f_i(v_i)$ be the virtual value of buyer i ([18]). We make the standard assumption that virtual values are non-decreasing. This assumption is satisfied by many commonly used distributions, including the normal, lognormal, uniform, and exponential distributions. Next, we define the space of the selling mechanisms. By the revelation principle, we focus without loss of generality on direct mechanisms, where buyers are asked to report their valuations (or types). A direct mechanism \mathcal{M} consists of a pair (q, p) , where $q_i : \mathbb{R}_+^n \rightarrow [0, 1]$ maps from the reported types to buyer i 's allocation probability, and $p_i : \mathbb{R}_+^n \rightarrow \mathbb{R}$ maps from the reported types to buyer i 's payment to the mechanism. In a typical mechanism design paper, once the revelation principle is invoked, only incentive-compatible (IC) mechanisms are considered. In our work, we also wish to consider the equilibrium outcome of some non-IC mechanisms. Therefore, we will allow for arbitrary strategy profiles in our analysis.

Buyers are allowed to play mixed strategies $\sigma_i : [\underline{v}, \bar{v}] \rightarrow \Delta([\underline{v}, \bar{v}])$, where $\Delta([\underline{v}, \bar{v}])$ represents the space of possible randomizations over possible reports. We denote strategy profiles by $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n)$ and strategy profiles that exclude player i by $\sigma_{-i} = (\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_n)$. We use $q_{i, \sigma_{-i}}(\hat{v}_i)$ and $p_{i, \sigma_{-i}}(\hat{v}_i)$ to respectively represent buyer i 's interim expected allocation and payment given a report of \hat{v}_i and assuming other buyers play according to strategy profile σ_{-i} ; that is, $q_{i, \sigma_{-i}}(\hat{v}_i) = \mathbb{E}[q_i(\hat{v}_i, \sigma_{-i}(\mathbf{v}_{-i}))]$ and $p_{i, \sigma_{-i}}(\hat{v}_i) = \mathbb{E}[p_i(\hat{v}_i, \sigma_{-i}(\mathbf{v}_{-i}))]$. Here, the expectation is with respect to (w.r.t.) valuation of other buyers \mathbf{v}_{-i} and any randomness in the strategy profile of other buyers; that is, σ_{-i} .

If buyer i were quasilinear, his utility given type v_i and report \hat{v}_i would be given by $u_{i, \sigma_{-i}}(v_i, \hat{v}_i) = v_i \cdot q_{i, \sigma_{-i}}(\hat{v}_i) - p_{i, \sigma_{-i}}(\hat{v}_i)$,

assuming others were playing according to σ_{-i} . We call this expression buyer i 's *unadjusted utility*. We define buyer i 's ROI as his expected adjusted utility divided by his expected payment, assuming the expected payment is positive. Formally, we define buyer i 's ROI as follows: $\text{ROI}_{i,\sigma}$ is $\mathbb{E}[u_{i,\sigma_{-i}}(v_i, \sigma_i(v_i))]/\mathbb{E}[p_{i,\sigma_{-i}}(\sigma_i(v_i))]$ if $\mathbb{E}[p_{i,\sigma_{-i}}(\sigma_i(v_i))] > 0$ and ∞ otherwise. Here, the expectation is w.r.t. valuation of all the buyers, including buyer i , and any any randomness in the strategy profile of buyers. The ROI determines how much utility a buyer expects to get as a fraction of his cash outlay; that is, $\mathbb{E}[p_{i,\sigma_{-i}}(\sigma_i(v_i))]$. If the expected payment is non-positive, we let $\text{ROI}_{i,\sigma}$ be infinity, since financial constraints are not applicable in this case.

The ROI we compute is an ex-ante measure (it is not a function of the valuation v_i), since we assume a buyer will make a decision whether or not to participate in the auction before knowing his precise value for that auction. Our definition of ROI represents an empirical measure of ROI as perceived by a buyer over many auctions: ROI computes the predicted utility earned from an auction divided by the predicted monetary cost of participating in the auction. This notion of ROI is inspired by display advertising markets where buyers participate in millions of auctions each day. In these markets, buyers are less sensitive to their obtained ROI in a single auction but are instead chiefly concerned about their expected ROI across many auctions. Put differently, this definition is consistent with a buyer who has an outside option with a known aggregate ROI (for example, television ads or an alternative platform) and needs to plan his marketing campaign multiple times a year without knowing his valuation for every single auction.

Each buyer i is assumed to have a target ROI, denoted by $\gamma_i \geq 0$. Buyer i will only agree to participate in the auction if he expects his ROI to match or exceed γ_i . We can formalize this notion by letting buyer i 's *adjusted utility* given his valuation v_i and report \hat{v}_i be

$$U_{i,\sigma}(v_i, \hat{v}_i) = \begin{cases} u_{i,\sigma_{-i}}(v_i, \hat{v}_i) & \text{if } \text{ROI}_{i,\sigma} \geq \gamma_i; \\ -\infty & \text{if } \text{ROI}_{i,\sigma} < \gamma_i. \end{cases} \quad (1)$$

We assume that the ROIs of all buyers are common knowledge. This assumption greatly simplifies our problem, as making γ_i 's private would make the seller's mechanism design problem multidimensional. In Section 6, we show how a seller can design an experiment to try to estimate the buyers' target ROIs.

A strategy profile is a (Bayes Nash) equilibrium if no buyer can deviate to a better strategy given his type v . That is, a strategy profile σ is an equilibrium if $\mathbb{E}[U_{i,\sigma}(v, \sigma_i(v))] \geq \mathbb{E}[U_{i,(\sigma'_i, \sigma_{-i})}(v, \sigma'_i(v))]$ for all i and any deviation σ'_i , where the expectation is taken with respect to any randomness in the valuation of other buyers; that is, \mathbf{v}_{-i} , and their strategy profiles σ_{-i} . When clear from context, we will suppress the subscripts σ and σ_{-i} to lighten the notation.

3 ROI AND SECOND-PRICE AUCTIONS

Before we turn our attention to auction design, we first focus on the kind of auctions that are most commonly used in online advertising: second-price auctions with reserve (SPA). We will now analyze how ROI-constrained buyers behave when faced with second-price auctions. In Section 4, we will build on the theoretical analysis we develop here to argue that some buyers behave in practice as if they are ROI-constrained.

Second-price auctions are truthful mechanisms when buyers are utility maximizers. But, as we will show in this section, this is not the case when buyers have ROI constraints. With ROI constraints, buyers may have an incentive to shade their bids. With bid shading, buyers win fewer bids, but earn greater surplus on the bids they win.

We now focus on the equilibrium bidding strategy of a buyer i in an SPA. Buyer i wins the item in an SPA when his submitted bid is greater than the reserve price r and the highest competing bid $\max_{j \neq i} \{b_j\}$. That is, the buyer wins if his bid is greater than the highest competing bid $D = \max\{r, \max_{j \neq i} \{b_j\}\}$. Let $G(\cdot)$ be the distribution of the highest competing bid D . In an SPA, the distribution G of the highest competing bid is a sufficient statistic of σ_{-i} for buyer i .

The distributions of valuations being regular implies that the distribution of v_i is atomless for all i . Therefore, it is natural to assume that players will choose strategies that do not assign point masses to bids above the reserve price r . We say that G is well-behaved if the players other than i choose strategies σ_{-i} such that in the interval (r, ∞) , the distribution G is differentiable.

Consider a buyer i with target ROI γ . Let D be the random variable that represents the highest competing bid generated by a well-behaved G . Then, strategy $\sigma_i(\cdot) = b(\cdot)$ is a best response by buyer i if and only if $b(\cdot)$ solves

$$\max_{b(\cdot)} \mathbb{E}[(v - D)\mathbb{1}\{b(v) \geq D\}] \quad \text{s.t.} \quad \frac{\mathbb{E}[(v - D)\mathbb{1}\{b(v) \geq D\}]}{\mathbb{E}[D\mathbb{1}\{b(v) \geq D\}]} \geq \gamma. \quad (\text{Bid-SPA})$$

We now introduce the ROI function $\mathcal{R}(\cdot)$, which we will use in our shading theorem. For any $\beta \in [0, 1]$, let $\mathcal{R}(\beta) = \mathbb{E}[(v - D)\mathbb{1}\{v\beta \geq D\}]/\mathbb{E}[D\mathbb{1}\{v\beta \geq D\}]$ if $\mathbb{E}[D\mathbb{1}\{v\beta \geq D\}] > 0$, and ∞ otherwise. We point out that $\mathcal{R}(\cdot)$ is continuous almost everywhere, given a well-behaved G . Note that $\mathcal{R}(\beta) = \infty$ when $\beta \leq \frac{r}{\bar{v}}$ and positive otherwise. In addition, \mathcal{R} is discontinuous only at $\beta = \frac{r}{\bar{v}}$. If $\bar{v} = \infty$, then \mathcal{R} is continuous everywhere.

THEOREM 3.1. *Consider a buyer with a target ROI of γ . Assume the highest competing bid D is drawn from a well-behaved G . Let $\bar{\gamma} = \max_{\beta \in [0, 1]} \{\mathcal{R}(\beta)\}$ and $\underline{\gamma} = \mathcal{R}(1)$. Then, the buyer's best response is characterized by:*

- If $\gamma \leq \underline{\gamma}$, then the buyer bids truthfully, i.e., $b(v) = v$.
- If $\gamma \in (\underline{\gamma}, \bar{\gamma}]$, then buyer's bidding strategy is $b^*(v) = \beta v$ where $\beta \in (\frac{1}{1+\gamma}, 1)$ is the unique solution of $\mathcal{R}(\beta) = \gamma$.
- If $\gamma > \bar{\gamma}$, then the buyer does not participate in the SPA.

Theorem 3.1 shows that when the target ROI is low ($\gamma < \underline{\gamma}$), the buyer bids truthfully. This occurs when ROI is not a binding constraint. In this case, the buyer's ROI is γ . When the target ROI is moderate ($\gamma \in (\underline{\gamma}, \bar{\gamma}]$), the buyer shades his bid by a constant factor. In this scenario, the ROI constraint is binding. Finally, when the target ROI is too large and cannot be obtained under any shading rule, the buyer does not participate in the SPA.

PROPOSITION 3.2. *Consider buyer $i \in [n]$ bidding according to Theorem 3.1. Let the buyer's value be v and consider two possible reserve prices $r_L < r_H$, and assume that the buyer chooses to participate in the SPA under r_H . Assume $G(\cdot)$ above the higher reserve price*

r_H is the same in the two cases. Then, the buyer's bid will be strictly lower under r_H than under r_L unless the ROI constraint is not binding under the higher reserve price, in which case, the buyer's bid will be the same under both reserve prices.

In contrast, shading is *not* how a rational player in a standard auction model behaves. In a standard (non-ROI) auction model, there is no reason for a buyer to lower his bid in an SPA when the reserve price goes up. We will take advantage of this difference in behavior to empirically demonstrate that some real-world buyers behave as if they face ROI constraints.

4 EMPIRICAL EVIDENCE

In this section, we provide empirical evidence that demonstrates that a significant subset of the buyers in online advertising exchanges shade their bids in SPAs. Specifically, we empirically show that buyers shade their bids in SPAs when reserve prices are high. As argued in Section 3, shading is consistent with buyers being financially constrained but inconsistent with buyers being quasilinear. We also argue that, based on the data, ROI constraints seem to better capture buyer behavior than budget constraints.

The experiment. We use bids submitted to Google's advertising exchange (AdX) over the course of two days. We use data generated from an internal AdX online experiment, which we did not design or control. The experiment was run on a small fraction of AdX's traffic, but one that was randomly sampled and that still amounted to hundreds of millions of queries (auctions) per day. For each auction, we have access to all the submitted bids and the winner's payment. The auctions are all second-price auctions with reserve. In addition to the bids and payment, our dataset also includes the specific ad slot being sold in each auction.

The queries that are selected for the experiment are randomly partitioned into treatment and control groups. The only difference between the two groups is how the reserve prices are set. For a given query q , let $r_c(q)$ and $r_t(q)$ be the reserve price under the control and treatment groups, respectively. We are not at liberty to discuss the algorithm used to compute $r_c(\cdot)$ and $r_t(\cdot)$ for a given query. For our purposes, it is sufficient to know that $r_t(q) \geq r_c(q)$ for any query q . That is, if a query is assigned to the treatment group, rather than to the control group, the query will carry a higher reserve price (in Google's AdX, buyers are informed of reserve prices up front). On a typical query q , the fraction $(r_t(q) - r_c(q))/r_c(q)$ is within the range of 1%-4%, so these exogenous shocks to the reserve prices are fairly modest. Using this data, we seek to understand whether buyers respond to reserve price increases as predicted by classical theory (by maintaining their bids) or by the ROI constraint theory we presented in Section 3 (by shading their bids). Buyers may also interpret a higher reserve price as a signal that a given query is particularly valuable, and thus buyers would increase their bids in response to increased reserve prices. As we will show, this signaling behavior is not borne out in our dataset.

Buyers, ad slots, and shading factors. We did not collect data on all buyers but only on those that we term *performance buyers*. Performance buyers are advertisers that we believe seek immediate rewards from their campaigns, as opposed to long-term benefits such as, brand awareness ([13]). We focused on these buyers because we do not believe all buyers are likely to be responsive to short-term

financial constraints. Buyers focused on long-term brand awareness might not be able to even measure the financial return from their ad purchases. The selection of the buyers was not random: we selected 35 buyers from among the 200 largest buyers on AdX whose bidding behavior was indicative of a focus on performance. Typical examples of performance buyers in our dataset include retargeting advertisers, who show ads to consumers displaying products they previously shopped for (in an attempt to get consumers to purchase those products), and mobile app advertisers attempting to maximize app installs. Retargeting is a significant fraction of the online ad business. According to a survey by [1], 24% of large marketers spend at least 50% of their online ad budgets on retargeting.

Different ad slots have very different values to buyers. In order to avoid confounding, we focus on the top 40 ad slots for each buyer, and we estimate the shading factor for each ad slot. We then average the 40 values to compute our estimated shading factor for each buyer, on each day of the experiment. More formally, for each buyer and day, we let b_t^a and b_c^a be the average submitted bid for ad slot a under the control and treatment groups, respectively. The estimated shading factor for ad slot a is $s^a = b_t^a/b_c^a$, and the estimated shading factor of a buyer on a given day is simply the average over all 40 top ad slots.

The shading factors we estimate are with respect to control reserve prices. That is, they are not the exact same shading factors as discussed in Section 3, which are with respect to true buyer valuations. We discuss in a bit more detail the distinction between these two shading factors in Section 6. The shading factors we estimate here can be interpreted as price elasticity results. That is, they reflect how bids change as reserve prices change.

Experimental results. Figure 1 depicts the estimated shading factors of the 35 performance buyers for the two consecutive days of the experiment. The figure shows both point estimates and their 95% confidence intervals.




Figure 1: Average shading factor of performance buyers across two consecutive days.

A shading factor of 1 or greater would indicate no shading behavior, while a shading factor strictly below 1 indicates shading. All 70 point estimates are below 1. For 47 of the point estimates (67%), the 95% confidence interval does not include 1. That is, according to the 95% confidence intervals, we can establish shading on 2/3 of the buyer-day pairs. All 35 buyers have estimates below 1, according to their 95% confidence intervals, when we aggregate submitted bids across both days.

As the plots demonstrate, there is great consistency across the days: buyers who shaded their bids more on day 1 were also more likely to shade more on day 2. We let (x_i, y_i) be the day 1 and

day 2 point estimates of each buyer i 's shading factor, and we ran a linear regression on this data. We found the R^2 to be 0.88, indicating that most of the variability in the day 2 shading factor can be explained by the variability in the day 1 shading factor. In other words, the shading factors across two days are highly correlated. This consistency demonstrates that the shading factors are a product of the buyers' bidding strategies, not the random assignment of queries.

Budgets versus ROI. There are different ways to model financial constraints. The most common model in the literature is to assume buyers have budgets. In the business world, due to the availability of capital markets, companies tend to think in terms of ROI instead. Given an investment's ROI and prevailing interest rates, a financially constrained company can decide whether borrowing money to pursue that investment is a good idea.

We now make the case that the performance buyers in our dataset seem to be behaving more like they are ROI constrained than budget constrained. Whichever metric a buyer targets (budget or ROI) is likely to fluctuate less than the other metric. We argue that ROI is the financial metric that buyers target by showing that buyers' total spending (budget) fluctuates more than ROI.

We compare the budget of the 35 performance buyers and their ROIs across all their ad slots. In particular, we compare the volatility of the budget and of the ROI of the buyers over time. To compute the ROI of each buyer, we consider all the auctions in which the buyer wins. We are not able to observe buyers' true valuations, and as a result, their true ROIs cannot be computed. To overcome this obstacle, we report a proxy for to the buyer's ROI. This proxy is defined as the ratio of expected submitted bids divided by his expected payment in all the auctions in which the buyer wins. We note that if the buyer is truthful, this proxy is equal to buyer's ROI plus one. Further, when the buyer shades his bid by a factor of β , this proxy simply scales with the same factor. Thus, showing that the proxy is stable across multiple days implies that the buyer's ROI is stable. In the rest of this section, we refer to this proxy as the buyer's ROI.

To compare the volatility of the amount spent and of the ROI of the buyers, we compute the coefficient of variations of the daily amount spent and the estimated ROI across seven consecutive days. We use all the queries in these days, including those that did not go through the online experiment. We observe significantly higher volatility in the buyers' expenditures than their ROIs. The average coefficient of variations of the budget is 37% higher than that of the ROI. For 27 buyers, the coefficient of variation of the budget is higher than the coefficient of variation of the ROI, while the opposite is true for eight buyers.

To better visualize these fluctuations, we now focus on two buyers: buyers 15 and 23, and present their daily expenditures and ROIs across seven days. We first focus on one of the top ad slots for each of these buyers, and we compute their total spending and ROI; see Figures 2a and 2b. Then, in Figures 2c and 2d, we present the total spending and ROI of these buyers across all of his ad slots. Due to the sensitivity of the data, we do not reveal the absolute values of spending and ROI of the buyers, reporting normalized values instead. Figure 2 shows that the ROIs of the buyers are relatively stable across different days, while total spending fluctuates significantly. For instance, Figure 2a shows that the total spending of

buyer 1 in his top slot changes by almost 36% when we move from day one to day six. However, his ROI only changes by 4%. Similarly, in Figure 2c, we observe that the total spending of buyer 1 across all the ad slots changes by 85% when moving from day one to day four. But, the change in his ROI is less than 22%. These figures illustrate that the buyers do not have to spend a fixed amount each day but that the buyers do seem to attempt to maintain a constant ROI over time.

5 OPTIMAL MECHANISM DESIGN

Thus far, we have established that a second-price auction with the monopoly price as reserve is not the revenue-maximizing mechanism when buyers are ROI-constrained (see Example 1). In this section, we characterize the structure of the optimal mechanism when buyers are ROI-constrained.

By the revelation principle, we focus on direct incentive compatible and individually rational mechanisms. A direct mechanism \mathcal{M} consists of a pair (q, p) where $q_i : [\underline{v}, \bar{v}]^n \rightarrow [0, 1]$ is an allocation rule for buyer i and $p_i : [\underline{v}, \bar{v}]^n \rightarrow \mathbb{R}$ is his transfer to the mechanism. With a slight abuse of notation, we denote $\mathbb{E}_{v_{-i}}[q_i(v_i, v_{-i}) | v_i]$ by $q_i(v_i)$.

Definition 5.1. A direct mechanism is incentive compatible (IC) if for each buyer i , truth-telling is a best response given that all other buyers report truthfully, i.e., $v_i q_i(v_i, v_{-i}) - p_i(v_i, v_{-i}) \geq v_i q_i(\hat{v}_i, v_{-i}) - p_i(\hat{v}_i, v_{-i})$ for all v_i, \hat{v}_i and ROI constraints under truthful reporting are satisfied. That is, $\frac{\mathbb{E}[q_i(\mathbf{v})v_i]}{\mathbb{E}[p_i(\mathbf{v})]} \geq \gamma_i + 1, i \in [n]$.

For incentive compatible mechanisms, the definition above is equivalent to our definition in Section 2. To ensure that a mechanism is IC, we also need to guarantee that the buyers can reach their target ROI by bidding truthfully. The following lemma identifies necessary and sufficient under which a direct mechanism is IC.

PROPOSITION 5.2. A mechanism \mathcal{M} with allocation rule $q(\cdot)$ is IC if and only if for any $i = 1, 2, \dots, n$, $q_i(v_i) = \mathbb{E}[q_i(v_i, v_{-i})]$ is weakly increasing in v_i , the expected utility of buyer i with valuation v_i , denoted by $u_i(v_i)$, is given by $u_i(\underline{v}) + \int_{\underline{v}}^{v_i} q_i(x) dx$, and

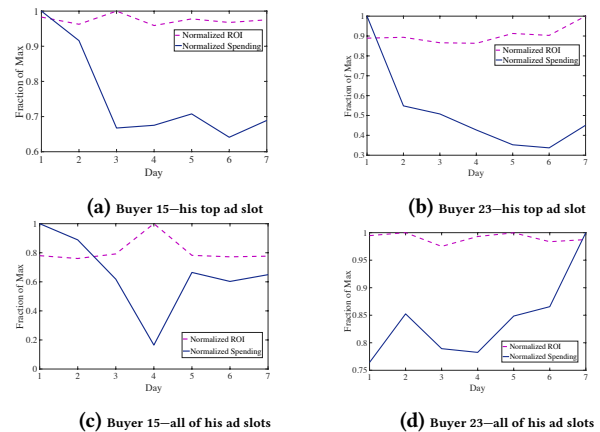


Figure 2: Normalized spending and ROI of buyers 15 and 23.

the ROI constraint is satisfied under truthful reporting. That is, $\frac{\mathbb{E}[u_i(v_i)]}{\mathbb{E}[v_i q_i(v_i) - u_i(v_i)]} \geq \gamma_i$ for $i \in [n]$.

The proof follows from the fact that buyers aim to maximize their utilities after their ROI constraints are satisfied. The proof follows Myerson's classical proof and is, thus, omitted.

An IC mechanism is individually rational (IR) if for each buyer i , the expected utility under the truthful reporting is non-negative, that is, $v_i q_i(v_i, \mathbf{v}_{-i}) - p_i(v_i, \mathbf{v}_{-i}) \geq 0$ for any $(v_i, \mathbf{v}_{-i}) \in [\underline{v}, \bar{v}]^n$. The IR constraints ensure that the buyers participate in the mechanism.

Having defined the IC and IR constraints, we now define the revenue-optimal mechanism. An IC and IR mechanism is revenue-optimal if it yields the highest revenue among all IC and IR mechanisms; that is, it maximizes $\mathbb{E}[\sum_{i=1}^n p_i(\mathbf{v})]$. Using the standard Myersonian technique, we have

$$\mathbb{E}[p_i(\mathbf{v})] = \mathbb{E}[\phi_i(v_i)q_i(\mathbf{v}) - u_i(\underline{v})]. \quad (2)$$

The seller's revenue optimization problem can thus be written as:

$$\begin{aligned} \max_{q_i(\cdot), u_i(\underline{v})} \quad & \mathbb{E}\left[\sum_{i=1}^n (\phi_i(v_i)q_i(\mathbf{v}) - u_i(\underline{v}))\right] \\ \text{s.t.} \quad & u_i(v_i) = u_i(\underline{v}) + \int_{\underline{v}}^{v_i} q_i(x)dx, \quad i \in [n], v_i \in [\underline{v}, \bar{v}] \\ & q_i(v_i) \text{ is weakly increasing in } v_i, \quad i \in [n] \\ & \mathbb{E}[v_i q_i(v_i)] - (\gamma_i + 1)\mathbb{E}[v_i q_i(v_i) - u_i(v_i)] \geq 0, \quad i \in [n] \\ & u_i(v_i) \geq 0, \quad i \in [n], v_i \in [\underline{v}, \bar{v}] \\ & \sum_{i=1}^n q_i(\mathbf{v}) \leq 1, \quad \mathbf{v} \in [\underline{v}, \bar{v}]^n, \end{aligned} \quad (\text{OPT})$$

The first three sets of constraints follow from Proposition 5.2, where we identify the necessary and sufficient conditions for incentive compatibility. In particular, the third set of constraints ensures that the buyers obtain their target ROI. The fourth set of constraints guarantees that the mechanism is IR. The last set of constraints ensures that the mechanism does not allocate the item to more than one buyer.

5.1 Symmetric Buyers

In this subsection, we assume that the valuation of each buyer $i = 1, 2, \dots, n$ is independently drawn from distribution F with probability density f . We need a few definitions before presenting our result. Assume that the seller runs an SPA with the monopoly reserve $r^m = \phi^{-1}(0)$ and all the buyers bid truthfully. Let γ_L be the ROI of a truthful buyer under this mechanism. Precisely, $\gamma_L = \frac{\mathbb{E}[v q_i(\mathbf{v})]}{\mathbb{E}[p_i(\mathbf{v})]} - 1 = \frac{\int_{r^m}^{\bar{v}} v F^{n-1}(v) dF(v)}{\int_{r^m}^{\bar{v}} \phi(v) F^{n-1}(v) dF(v)} - 1$, where $F^{n-1}(v)$ is the probability that a buyer with valuation v has the highest valuation among n buyers. The numerator is the expected value the buyer gains from possibly winning the item, while the denominator corresponds to the expected virtual value from possibly winning the item. Similarly, let γ_H be the ROI of a buyer when the seller runs a SPA with zero reserve price and all buyers bid truthfully. That is, $\gamma_H = (\int_0^{\bar{v}} v F^{n-1}(v) dF(v)) / (\int_0^{\bar{v}} \phi(v) F^{n-1}(v) dF(v)) - 1$. For any

$\gamma \in [\gamma_L, \gamma_H]$, let r_γ be the value that solves the following equation:

$$\gamma = \frac{\int_{r_\gamma}^{\bar{v}} v F^{n-1}(v) dF(v)}{\int_{r_\gamma}^{\bar{v}} \phi(v) F^{n-1}(v) dF(v)} - 1. \quad (3)$$

The above definition implies that when all buyers are truthful and the seller runs an SPA with reserve r_γ , any truthful buyer earns ROI of γ , for $\gamma \in [\gamma_L, \gamma_H]$. As we will show in Theorem 5.3, the reserve price $r_\gamma \in [0, r^m]$ and is decreasing in γ . Note that $r_{\gamma_H} = 0$ and $r_{\gamma_L} = r^m$. For any $\gamma \in [\gamma_L, \gamma_H]$, let λ_γ be the value that solves $\bar{\phi}(r_\gamma, \lambda_\gamma, \gamma) = 0$, where

$$\bar{\phi}(x, \lambda, \gamma) = \lambda x + \phi(x)(1 - \lambda(1 + \gamma)). \quad (4)$$

The term $\bar{\phi}(r_\gamma, \lambda_\gamma, \gamma)$ corresponds to a modification of the virtual value function, which takes into account a Lagrangian relaxation of the ROI constraint. Under this interpretation, the optimal Lagrange multiplier of the ROI constraint corresponds to:

$$\lambda^* = \begin{cases} 0 & \text{if } \gamma \leq \gamma_L; \\ \lambda_\gamma & \text{if } \gamma \in (\gamma_L, \gamma_H); \\ 1/(1 + \gamma) & \text{if } \gamma \geq \gamma_H. \end{cases} \quad (5)$$

We are now ready to present the optimal mechanism.

THEOREM 5.3. *Consider a setting with symmetric buyers with regular value distributions. The following mechanism maximizes the seller's revenues: If $\gamma < \gamma_L$, then run an SPA with reserve price r^m ; If $\gamma \in [\gamma_L, \gamma_H]$, then run an SPA with reserve price r_γ , where r_γ is defined in Eq. (3); If $\gamma > \gamma_H$, then run an SPA with reserve price 0. In addition, provide a subsidy s to each buyer, where s solves*

$$\frac{\int_0^{\bar{v}} v F^{n-1}(v) f(v) dv}{-s + \int_0^{\bar{v}} \phi(v) F^{n-1}(v) f(v) dv} = (1 + \gamma). \quad (6)$$

In addition, the optimal revenue of the seller is equal to

$$\mathbb{E}[(\max_{i \in [n]} \bar{\phi}(v_i, \lambda^*, \gamma))^+].$$

Theorem 5.3 shows that when buyers are ROI-constrained, the seller needs to reduce the reserve price or even provides buyers with a positive subsidy in order to maximize revenue. The theorem has three parts. When the target ROI is low ($\gamma < \gamma_L$), the ROI constraint is not binding and the optimal mechanism is the standard SPA with the monopoly reserve r^m . When the target ROI is moderate ($\gamma \in [\gamma_L, \gamma_H]$), the optimal mechanism is again an SPA. However, the reserve price should decrease so that the ROI constraints of the buyers are satisfied. When the target ROI is high ($\gamma > \gamma_H$), reducing the reserve price alone is not sufficient to produce a solution that satisfies the ROI constraint. To make sure that the ROI constraints are satisfied and the buyers participate, the seller needs to run an SPA without a reserve price, and give each buyer a subsidy. The amount of subsidy is chosen so that the ROI constraints are binding. **Welfare and revenue implications.** In the moderate target ROI case, $\gamma \in [\gamma_L, \gamma_H]$, the seller's optimizes revenues by choosing a reserve price below the monopoly r^m . This implies that seller that ignores ROI constraints creates a situation that is worse off for everyone: both seller revenues and buyers' surpluses would go up if the seller reduces the reserve price below r^m .

For the high target ROI case, $\gamma > \gamma_H$, an intuitive answer might have been that if the seller cannot get the buyers to participate even

without a reserve price, there is nothing for the seller to do. Theorem 5.3 proves this is false. In fact, with the proper choice of subsidy, the seller optimizes revenues by ensuring participation even if the target ROI is sky high. The following corollary formalizes this notion. This corollary follows from the fact that $\bar{\phi}(v_i, \frac{1}{1+\gamma}, \gamma) = \frac{v_i}{1+\gamma}$ and the expected revenue of the seller in the optimal mechanism is equal to $\mathbb{E}[(\max_{i \in [n]} \bar{\phi}(v_i, \lambda^*, \gamma))^+]$. It shows that, under a high target ROI, the optimal revenue of the seller decreases linearly in γ . It also shows that even with a high target ROI, the seller obtains a constant fraction of the consumer surplus.

COROLLARY 5.4. *Assume the target ROI is high, i.e., $\gamma \geq \gamma_H$. Then, the optimal revenue of the seller is equal to $\frac{1}{1+\gamma} \mathbb{E}[\max_{i \in [n]} \{v_i\}]$.*

5.2 Asymmetric Buyers

We now present the optimal mechanism for the case when buyers are asymmetric with respect to their valuation distributions and target ROIs. We will show that the heterogeneity among buyers leads to distortion in the allocation rule of the optimal mechanism, in the sense that the item is not allocated to the buyer with the highest submitted bid. We further show that the allocation rule of the optimal mechanism remains distorted even if buyers are only heterogeneous in their target ROIs.

We first consider the standard case where the valuation distributions of the buyers are heterogeneous, and buyers are not ROI constrained. Without ROI constraints, it is optimal to sort buyers based on their virtual values, $\phi_i(v_i)$, and allocate the item to the buyer with the highest nonnegative virtual value [18]. We will show that when buyers are ROI constrained, the structure of the optimal mechanism bears some resemblance to the aforementioned mechanism. However, the mechanism needs to be modified so that the buyers' ROI constraints are satisfied. We start by presenting a parametrized mechanism $\mathcal{M}(\Lambda, \Gamma, S)$:

Mechanism $\mathcal{M}(\Lambda, \Gamma, S)$

Subsidies: Each buyer i receives a subsidy of $s_i \geq 0$ from the seller.

Bids: Each buyer i submits his bid b_i .

Modified virtual values: For each buyer i , his modified virtual value, denoted by $\bar{\phi}_i(b_i, \lambda_i, \gamma_i)$, is calculated according to $\bar{\phi}_i(x, \lambda_i, \gamma_i) = \lambda_i x + \phi_i(x)(1 - \lambda_i(1 + \gamma_i))$.

Winner determination: The item is allocated to a buyer i^* with the highest nonnegative modified virtual value; that is,

$$i^* = \arg \max_{i \in [n], \bar{\phi}_i(b_i, \lambda_i, \gamma_i) \geq 0} \bar{\phi}_i(b_i, \lambda_i, \gamma_i). \quad (7)$$

The item is not allocated if the modified virtual values of all the buyers are negative.

The winner's payment: Define $\bar{\phi}^{(2)} = (\max_{j \in [n], j \neq i^*} \{\bar{\phi}_j(b_j, \lambda_j, \gamma_j)\})^+$ as the second highest modified virtual value. Let $\theta_i(x, \lambda_i, \gamma_i)$ be the inverse of $\bar{\phi}_i(v_i, \lambda_i, \gamma_i)$ with respect to its first argument v_i . Then, buyer i^* pays $p_{i^*} = \theta_{i^*}(\bar{\phi}^{(2)}, \lambda_{i^*}, \gamma_{i^*})$, and for any other buyer $i \neq i^*$, the payment p_i is set to zero.

Mechanism $\mathcal{M}(\Lambda, \Gamma, S)$, which can be described by Lagrangian multipliers $\Lambda = (\lambda_1, \dots, \lambda_n)$, target ROIs $\Gamma = (\gamma_1, \dots, \gamma_n)$, and subsidy levels $S = (s_1, \dots, s_n)$. Note that mechanism $\mathcal{M}(\Lambda, \Gamma, S)$ sorts buyers based on their modified virtual value and allocates the

item to the buyer with the highest nonnegative modified virtual value. When $(1 - \lambda_i(1 + \gamma_i)) \geq 0$ and distribution F_i is regular, then the modified virtual value, $\bar{\phi}_i(x, \lambda_i, \gamma_i)$, is increasing in its first argument. The monotonicity of the modified virtual value implies that the allocation probability of buyers increases as they submit higher bids.

The mechanism $\mathcal{M}(\Lambda, \Gamma, S)$ is similar to Myerson's optimal non-ROI-constrained mechanism, with the key distinction being that modified virtual values $\bar{\phi}_i(v_i, \lambda_i, \gamma_i)$ replace virtual values $\phi_i(v_i)$. This change in the allocation is designed to ensure the target ROIs of the buyers are met. In order to meet the target ROIs, each buyer is assigned a Lagrangian multiplier $\lambda_i \in [0, 1/(1 + \gamma_i)]$. As the target ROI of a buyer increases, the mechanism will assign the buyer a larger Lagrangian multiplier. If the target ROI cannot be achieved even at $\lambda_i = \frac{1}{1 + \gamma_i}$, then the mechanism needs to provide buyer i a subsidy s_i .

In mechanism $\mathcal{M}(\Lambda, \Gamma, S)$, buyer i does not get the item when his value is less than his personalized reserve price $r_i = \theta_i(0, \lambda_i, \gamma_i)$. In addition, when buyer i wins the item, his payment is at least r_i . In fact, the payment of the winning buyer is the minimum bid that he needs to submit in order to win the item.

We need to define some terms we will use to describe the optimal mechanism. Let

$$\begin{aligned} \Lambda^* &= (\lambda_1^*, \lambda_2^*, \dots, \lambda_n^*) \\ &= \arg \min_{\lambda_i \in [0, \frac{1}{1+\gamma_i}], i \in [n]} \mathbb{E} \left[\left(\max_{i \in [n]} \{\bar{\phi}_i(v_i, \lambda_i, \gamma_i)\} \right)^+ \right] \end{aligned} \quad (8)$$

be the optimal Lagrangian multipliers of the ROI constraints. By duality theory, the above optimization problem is convex for any set of functions $\bar{\phi}_i$ and can be computed via standard convex optimization algorithms. Furthermore, let $S^* = (s_1^*, \dots, s_n^*)$ be defined as

$$s_i^* = \begin{cases} \mathbb{E} \left[(-v_i + (\gamma_i + 1)\phi_i(v_i))q_i(\mathbf{v}) \right] / (1 + \gamma_i) & \text{if } \lambda_i^* = 1/(1 + \gamma_i); \\ 0 & \text{if } \lambda_i^* < 1/(1 + \gamma_i). \end{cases} \quad (9)$$

Here, $q_i(\mathbf{v}) = 1$ if $i = \arg \max_{j \in [n]} \bar{\phi}_j(v_j, \lambda_j^*, \gamma_j)$ and $\bar{\phi}_i(v_i, \lambda_i^*, \gamma_i) \geq 0$. Otherwise, $q_i(\mathbf{v}) = 0$.

THEOREM 5.5. *Consider a setting with asymmetric buyers with regular value distributions. The mechanism $\mathcal{M}(\Lambda^*, \Gamma, S^*)$ maximizes the seller's revenues, where Λ^* and S^* are respectively defined in equations (8) and (9). In addition, the optimal revenue of the seller is equal to $\mathbb{E} \left[(\max_{i \in [n]} \bar{\phi}_i(v_i, \lambda_i^*, \gamma_i))^+ \right]$.*

The intuition for this result is similar to the one for symmetric buyers. Under a low target ROI, the ROI constraint is not binding, and $\lambda_i^* = 0$. In this case, the modified virtual values revert to the virtual value. Under a moderate target ROI, the virtual value is modified to make the ROI constraint feasible. Under a high target ROI, the virtual value is maximally modified, with $\lambda_i^* = 1/(1 + \gamma_i)$, but this is not sufficient to ensure buyer participation. Therefore, a buyer subsidy is also necessary in this final case. As in the symmetric case, it is always better for the seller to subsidize a buyer than to have the buyer not participate in the auction, regardless of the buyer's target ROI.

6 COUNTERFACTUAL ANALYSIS

In this section, we combine the theoretical results on optimal auction design under ROI constraints from Section 5 with the data from Section 4 to estimate how much is being left on the table due to not accounting for ROI constraints.

The optimal auctions for asymmetric buyers are complex and not very practical. This is true not only of our ROI-constrained optimal auctions (see Subsection 5.2), but also of Myerson’s auction for asymmetric buyers. We therefore focus on the ROI-constrained optimal auctions for symmetric buyers from Subsection 5.1. Optimizing an ROI-constrained symmetric auction requires only choosing a reserve price or a subsidy level (if the optimal reserve price is zero).

We use the two buyers for whom we estimated the bid distributions for their top ad slots—buyers 15 and 23—and construct optimal auctions for these buyers. To maintain the auction symmetry, we do not combine the buyers into a single auction. That is, we create an auction for a setting with multiple versions of buyer 15, and a different auction for a setting with multiple versions of buyer 23.

Estimating value distributions. We are not able to directly observe buyer value distributions and target ROIs. All that we can observe is buyer bidding under treatment and control reserve prices, and create estimates based on those. We call the empirical distributions of submitted bids $\hat{F}_t(x)$ and $\hat{F}_c(x)$ under treatment and control, respectively, $\hat{F}_t(x) = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \mathbb{1}\{b_t \leq x\}$, and $\hat{F}_c(x) = \frac{1}{|\mathcal{C}|} \sum_{t \in \mathcal{C}} \mathbb{1}\{b_t \leq x\}$, where \mathcal{T} and \mathcal{C} represent the sets of bids under treatment and control.

Based on the theory we developed, we know that buyers are more likely to be ROI-constrained when reserve prices are high, which corresponds to the treatment group. Therefore, the empirical control distribution is likely to be more similar to the true value distribution. We therefore assume $\hat{F}_c(\cdot)$ is the true value distribution for the purposes of this counterfactual analysis. The auctions in our dataset vary widely in quality and reserve prices. In order to maintain some quality consistency, we will use only auctions where the reserve price lies within one standard deviation of its mean, i.e., the reserve price of query q satisfies $r_c(q) \in [\mathbb{E}[r_c] - \sigma(r_c), \mathbb{E}[r_c] + \sigma(r_c)]$.

Estimating shading factors and target ROIs. Before estimating target ROIs, we estimate the shading factor used by the buyers. We estimate shading factors via structural estimation. We assume buyers are shading their bids as prescribed by Theorem 3.1 and that $\hat{F}_c(v)$ is their true value distribution. If a buyer with value distribution $\hat{F}_c(v)$ uses a shading factor β under treatment prices, we should observe bids distributed according to $\hat{F}_t(v/\beta)$ under treatment prices. We therefore compute our estimate $\hat{\beta}$ of the shading factor via the following regression: $\hat{\beta} = \arg \min_y \sum_n (\hat{F}_t(b_n) - \hat{F}_c(b_n/y))^2$. We find estimated shading factors of 0.93 and 0.88 for the top ad slots of buyers 15 and 23, with standard errors of 0.012 and 0.015.

Our theory predicts that the buyer is more likely to be ROI-constrained under high prices. We thus determine a buyer’s target ROI by assuming the ROI constraint is binding under treatment prices. We can therefore compute our estimated target ROI $\hat{\gamma}$ by computing $\hat{\gamma} = \frac{\sum_{t \in \hat{\mathcal{T}}} w_t (b_t / \hat{\beta} - p_t)}{\sum_{t \in \hat{\mathcal{T}}} w_t p_t}$, where p_t is the payment of the buyer in auction t and $w_t \in \{0, 1\}$ is one if the buyer wins auction

t , and is zero otherwise. Here, $\hat{\mathcal{T}}$ is the set of all the auctions in the treatment group that have a reserve price in the range $[\mathbb{E}[r_t] - \sigma(r_t), \mathbb{E}[r_t] + \sigma(r_t)]$. For competitive reasons, we do not disclose the estimated $\hat{\gamma}$ of the buyers.

Note that we are probably undermeasuring the true target ROI using our technique. That is, $\hat{\gamma}$ is possibly higher than true γ since our structural estimation procedure assumes the ROI constraint is not binding under control prices. If we were able to measure true target ROIs, the measurements would most likely be more pronounced, and the implications for revenue and welfare would be even more stark.

Counterfactual analysis results. We now consider a setting with $n = 1, \dots, 8$ buyers having valuation distribution $\hat{F}_c(\cdot)$ and target ROI $\hat{\gamma}$. We do this exercise twice, the first time with data from buyer 15, and the second time with data from buyer 23.

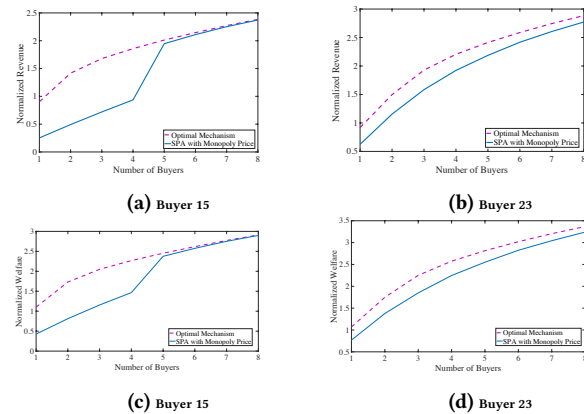


Figure 3: Normalized revenue and welfare of the optimal mechanism and the SPA with monopoly reserve price as a function of the number of buyers, assuming multiple copies of buyers 15 and 23 are present.

For these set of problem parameters, when the number of buyers is small, we find that the optimal mechanism according to Theorem 5.3 is an SPA with a positive reserve price which is below the monopoly price. In contrast, when the number of buyers is large, the optimal mechanism is an SPA with no reserve and a positive subsidy. This is so because by increasing the number of buyers, competition among buyers gets more severe and the ROIs of the buyers decrease. In Figure 3, we plot the normalized revenue and welfare of the optimal mechanism. To normalize revenue (welfare), we set the revenue (welfare) of Myerson’s optimal non-ROI-constrained mechanism with one buyer to be equal to one. As a benchmark, we also plot the revenue and welfare under an SPA with monopoly price, where buyers are assumed to shade according to Theorem 3.1.

In Figure 3a, we see that the revenue gain from taking ROI into account can be as low as 0.5% ($n = 8$) to as high as over 200% ($n = 1$). These numbers are based on buyers who behave like buyer 15. For this case, the average and the median revenue gain are 85% and 51%, respectively. The range of outcomes is not as wide for the buyer 23 model (Figure 3b). In this scenario, revenue gains range from 4% ($n = 8$) to 46% ($n = 1$), with mean 17% and median 12%. In Figure

3c, we see that the welfare gains for the buyer 15 also span a wide range, starting from 0.4% (n=8) going up to 150% (n=1), with mean and median 51% and 29%, respectively. In Figure 3d, we see the welfare gains for the buyer 23 model, which ranges from 4% and 39%, with mean 16% and median 12%. Overall, the revenue gains and welfare improvement are smaller when the market is thick, i.e., the number of buyers is large. Interestingly, the revenue and welfare gains are highly correlated. The situations where the ROI-aware auction performs much better than the standard auction are settings where ROI constraints are causing buyers to significantly shade their bids. In such situations, a reduction in the reserve price can simultaneously greatly improve revenue and generate a very large surplus for buyers.

REFERENCES

- [1] AdRoll. 2015. State of the Industry: A close look at retargeting and the programmatic marketer, Survey responses from 1,000 marketers. https://www.iab.com/wp-content/uploads/2015/07/US_AdRoll_State_of_the_Industry.pdf (2015).
- [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*. ACM, 1–9.
- [3] Santiago Balseiro, Omar Besbes, and Gabriel Y Weintraub. 2015. Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science* 61, 4 (2015), 864–884.
- [4] Santiago Balseiro and Yonatan Gur. 2017. Learning in Repeated Auctions with Budgets: Regret Minimization and Equilibrium. (2017).
- [5] Sayan Bhattacharya, Vincent Conitzer, Kamesh Munagala, and Lirong Xia. 2010. Incentive compatible budget elicitation in multi-unit auctions. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 554–572.
- [6] 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*. ACM, 531–540.
- [7] Yeon-Koo Che and Ian Gale. 1998. Standard auctions with financially constrained bidders. *The Review of Economic Studies* 65, 1 (1998), 1–21.
- [8] Yeon-Koo Che and Ian Gale. 2000. The optimal mechanism for selling to a budget-constrained buyer. *Journal of Economic Theory* 92, 2 (2000), 198–233.
- [9] Vincent Conitzer, Christian Kroer, Eric Sodomka, and Nicolas E Stier-Moses. 2017. Multiplicative Pacing Equilibria in Auction Markets. *arXiv preprint arXiv:1706.07151* (2017).
- [10] Nikhil R Devanur, Bach Q Ha, and Jason D Hartline. 2013. Prior-free auctions for budgeted agents. In *Proceedings of the 14th ACM Conference on Electronic Commerce*. ACM, 287–304.
- [11] Shahar Dobzinski and Renato Paes Leme. 2014. Efficiency guarantees in auctions with budgets. In *International Colloquium on Automata, Languages, and Programming*. Springer, 392–404.
- [12] Gagan Goel, Vahab Mirrokni, and Renato Paes Leme. 2014. Clinching auctions beyond hard budget constraints. In *Proceedings of the 15th ACM Conference on Economics and Computation*. ACM, 167–184.
- [13] Negin Golrezaei, Max Lin, Vahab Mirrokni, and Hamid Nazerzadeh. 2017. Boosted Second Price Auctions: Revenue Optimization for Heterogeneous Bidders. (2017).
- [14] Ramakrishna Gummadi, Peter Key, and Alexandre Proutiere. 2012. Repeated auctions under budget constraints: Optimal bidding strategies and equilibria. In *the Eighth Ad Auction Workshop*.
- [15] Jean-Jacques Laffont and Jacques Robert. 1996. Optimal auction with financially constrained buyers. *Economics Letters* 52, 2 (1996), 181–186.
- [16] Alexey Malakhov and Rakesh V Vohra. 2008. Optimal auctions for asymmetrically budget constrained bidders. *Review of Economic Design* 12, 4 (2008), 245.
- [17] Eric S Maskin. 2000. Auctions, development, and privatization: Efficient auctions with liquidity-constrained buyers. *European Economic Review* 44, 4-6 (2000), 667–681.
- [18] Roger B Myerson. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
- [19] Malleesh M Pai and Rakesh Vohra. 2014. Optimal auctions with financially constrained buyers. *Journal of Economic Theory* 150 (2014), 383–425.
- [20] Boleslaw K Szymanski and Juong-Sik Lee. 2006. Impact of ROI on bidding and revenue in sponsored search advertisement auctions. In *Second Workshop on Sponsored Search Auctions*.
- [21] Stanislav Uryasev and Panos M Pardalos. 2013. *Stochastic optimization: algorithms and applications*. Vol. 54. Springer Science & Business Media.
- [22] Christopher A Wilkens, Ruggiero Cavallo, and Rad Niazadeh. 2016. Mechanism design for value maximizers. *arXiv preprint arXiv:1607.04362* (2016).

7 APPENDICES: PROOFS OF MAIN RESULTS

The proof of lemmas in this section is presented in Section 8.

7.1 Proof of Theorem 3.1

The proof is divided into three segments. In the first and second segments, we assume that the target ROI is low and moderate, respectively. In the last segment, we assume that the target ROI is high.

Low target ROI ($\gamma \leq \underline{\gamma}$) Here we need to show that for any $\gamma \leq \underline{\gamma}$, the buyer bids truthfully. Recall that $\underline{\gamma}$ is the ROI of the buyer in the SPA with the monopoly reserve when the buyer bids truthfully. When the target ROI is less than $\underline{\gamma}$, the ROI constraint is non-binding under truthful bidding. Thus, the utility maximizer buyer bids truthfully.

Moderate target ROI ($\gamma \in (\underline{\gamma}, \bar{\gamma})$) Here, we construct an upper bound on Problem Bid-SPA by dualizing the ROI constraint. We then show that the bidding strategy presented in Theorem 3.1 achieves the upper bound and is, thus, optimal.

We start with by adding a constraint to Problem Bid-SPA:

$$\begin{aligned} \max_{b(\cdot)} \mathbb{E}[(v - D)\mathbb{1}\{b(v) \geq D\}] \\ \text{s.t. } \mathbb{E}[(v - (1 + \gamma)D)\mathbb{1}\{b(v) \geq D\}] &\geq 0, \\ \mathbb{E}[v\mathbb{1}\{b(v) \geq D\}] &> 0. \end{aligned} \quad (\text{Modified})$$

The second constraint in Problem Modified ensures that we do not pick a bidding strategy under which the buyer does not win at all. This ensures that the buyer’s ROI is finite. We show in Lemma 7.2 that at the optimal bidding strategy, this constraint is satisfied if the target ROI is moderate.

By the weak duality theorem, for any $\lambda \geq 0$, the optimal solution of Problem Modified is upper bounded by $\max_{b(\cdot)} \mathcal{L}(b, \lambda)$ where

$$\begin{aligned} \mathcal{L}(b, \lambda) &= \mathbb{E}[(v - D)\mathbb{1}\{b(v) \geq D\}] + \lambda (\mathbb{E}[(v - (1 + \gamma)D)\mathbb{1}\{b(v) \geq D\}]) \\ &= (1 + \lambda) \mathbb{E} \left[\left(v - \frac{(\lambda(1 + \gamma) + 1)D}{(1 + \lambda)} \right) \times \mathbb{1}\{b(v) \geq D\} \right]. \end{aligned} \quad (10)$$

The second equation implies that the effective payment of the buyer is not D ; it is $\frac{(\lambda(1 + \gamma) + 1)D}{(1 + \lambda)}$ instead. Therefore, to maximize $\mathcal{L}(b, \lambda)$, the buyer needs to shade his bid by a factor of $\frac{(1 + \lambda)}{(\lambda(1 + \gamma) + 1)}$, i.e., $b(v) = v \frac{(1 + \lambda)}{(\lambda(1 + \gamma) + 1)} = v\beta(\lambda, \gamma)$, or alternatively, $\beta(\lambda, \gamma) = \frac{(1 + \lambda)}{(\lambda(1 + \gamma) + 1)}$. Precisely, (10) can be written as

$$\mathcal{L}(b, \lambda) = (1 + \lambda) \mathbb{E} \left[\left(v - \frac{D}{\beta(\lambda, \gamma)} \right) \times \mathbb{1} \left\{ v \geq \frac{D}{\beta(\lambda, \gamma)} \right\} \right]. \quad (11)$$

To complete this part of the proof, we show that there exists a unique λ under which the ROI constraint is binding given that $b(v) = v\beta(\lambda, \gamma)$. We use the following definition in the rest of the proof:

$$\tilde{\mathcal{R}}(\beta) = \mathbb{E}[(v - D(1 + \gamma))\mathbb{1}\{v\beta \geq D\}]. \quad (12)$$

Assuming $\mathcal{R}(\beta)$ is finite, $\mathcal{R}(\beta) = \gamma$ if and only if $\tilde{\mathcal{R}}(\beta) = 0$. We can also have $\tilde{\mathcal{R}}(\beta) = 0$ when $\mathcal{R}(\beta) = \infty$, which correspond to values of β around zero. We say β is an “isolated root” of $\tilde{\mathcal{R}}(\cdot)$ if $\tilde{\mathcal{R}}(\beta) = 0$ and

there exists an open ball around β such that $\mathcal{R}(x) \neq 0$ for every x in the open ball. When β is an isolated root of $\tilde{\mathcal{R}}(\cdot)$, the ROI constraint is binding at the bidding strategy $b(v) = \beta v$, and thus $\mathcal{R}(\beta)$ is finite and $\mathcal{R}(\beta) = \gamma$.

In the following, we show that $\tilde{\mathcal{R}}(\cdot)$ has a unique isolated root. Let us denote this solution by β_0 . We can then choose λ^* as the unique solution of $\beta(\lambda^*, \gamma) = \beta_0$ since $\beta(\lambda, \gamma)$ is increasing in λ . Considering the fact that bidding strategy $b(v) = \beta_0 v$ maximizes $\mathcal{L}(b, \lambda^*)$ and the ROI constraint is binding at $b(v) = \beta_0 v$, we can conclude that this bidding strategy is optimal.

We complete our result by proving that $\tilde{\mathcal{R}}(\cdot)$ has a unique isolated root via a sequence of lemmas.

LEMMA 7.1. $\tilde{\mathcal{R}}(\beta)$ is non-decreasing in β when $\beta \leq \frac{1}{1+\gamma}$.

Lemma 7.1 shows that $\tilde{\mathcal{R}}(\beta)$ grows as β increases from 0 to $\frac{1}{1+\gamma}$. Then, considering the fact that $\tilde{\mathcal{R}}(0) = 0$, we have $\tilde{\mathcal{R}}(\frac{1}{1+\gamma}) \geq 0$. Thus, if we can show that $\tilde{\mathcal{R}}(\frac{1}{1+\gamma}) > 0$, we can conclude that there exists at most one $\beta \in (\frac{1}{1+\gamma}, 1)$ that solves $\tilde{\mathcal{R}}(\beta) = 0$. This follows from the fact that $\gamma > \underline{\gamma}$, and as a result $\tilde{\mathcal{R}}(1) < 0$.

The next lemma shows that $\tilde{\mathcal{R}}(\frac{1}{1+\gamma})$ is strictly positive. In particular, it shows that $\frac{1}{1+\gamma} > \frac{r}{\bar{v}}$, which implies that $\tilde{\mathcal{R}}(\frac{1}{1+\gamma}) \neq 0$ and, thus, $\tilde{\mathcal{R}}(\frac{1}{1+\gamma})$ is strictly positive.

LEMMA 7.2. Let $\gamma \leq \bar{\gamma}$. Then, $\frac{1}{1+\gamma} > \frac{r}{\bar{v}}$ and $\tilde{\mathcal{R}}(\frac{1}{1+\gamma}) > 0$.

So far, we have established that there exists at most one $\beta \in (\frac{1}{1+\gamma}, 1)$ that solves $\tilde{\mathcal{R}}(\beta) = 0$. The following lemma shows that such a solution is unique.

LEMMA 7.3. There exists a unique solution $\beta_0 \in (\frac{1}{1+\gamma}, 1)$ that solves $\tilde{\mathcal{R}}(\beta_0) = 0$.

Lemma 7.3 shows that there $\tilde{\mathcal{R}}(\cdot)$ has a unique isolated root in the range of $(\frac{1}{1+\gamma}, 1)$. By Lemma 7.2, we have $\beta_0 > \frac{1}{1+\gamma} > \frac{r}{\bar{v}}$, which implies that $\beta_0 \bar{v} > r$, and as a result, the second constraint in Problem Modified is satisfied; that is, $\mathbb{E}[v \mathbb{1}\{\beta_0 v \geq D\}] > 0$.

High target ROI ($\gamma > \bar{\gamma}$) Here, we first show that when $\gamma > \bar{\gamma}$, then the buyer does not participate. To this aim, we show that the upper bound of Problem Modified, defined below, is $-\infty$:

$$\min_{\lambda \geq 0} \max_{b(\cdot)} \mathcal{L}(b, \lambda). \tag{13}$$

This implies that Problem Modified is infeasible. To see why assume, contrary to our claim, that Problem Modified has a feasible solution. Then, considering the fact the objective function of Problem Modified at any feasible solution is positive, the upper bound of Problem Modified cannot be $-\infty$. This contradicts that Problem is feasible.

Since Problem Modified is infeasible, for any $\beta \in [0, 1]$, we have $\mathbb{E}[(v - (1 + \gamma)D) \mathbb{1}\{\beta v \geq D\}] < 0$. This implies that the Lagrangian multiplier λ that minimizes Eq. (13) is ∞ ; see the definition of $\mathcal{L}(b, \lambda)$ in Eq. (10). Since the objective function of Eq. (13) is $-\infty$, we conclude that the buyer does not participate in the auction.

7.2 Proof of Proposition 3.2

Assume that under the higher reserve price r_H , the buyer participates in the SPA and his ROI constraint is binding. If the buyer's ROI

constraint under reserve price r_L is non-binding, then the result holds. This follows from Theorem 3.1, where we show the buyer's shading factor is one when the ROI constraint is non-binding. Thus, we will focus on the case where the buyer's ROI constraint is also binding under the SPA with reserve price r_L .

By Theorem 3.1, the shading factor under the SPA with reserve price $r \in \{r_L, r_H\}$ is the unique solution of $\tilde{\mathcal{R}}(\beta) = 0$, see Eq. (12), where

$$\begin{aligned} \tilde{\mathcal{R}}(\beta) &= \mathbb{E}[(v - D(1 + \gamma)) \mathbb{1}\{v\beta \geq D\}] = \mathbb{E}_v \left[\int_r^{v\beta} (v - x(1 + \gamma)) d\hat{G}(x) \right] \\ &= \int_{\frac{r}{\beta}}^{\bar{v}} \int_r^{v\beta} (y - x(1 + \gamma)) d\hat{G}(x) dF(y) \quad \text{for } r \in \{r_L, r_H\}, \end{aligned}$$

where \hat{G} is the distribution of $\max_{j \neq i} \{b_j\}$. Note that $\tilde{\mathcal{R}}(\beta)$ implicitly depends on the reserve price r . From the proof of Theorem 3.1, we know that the unique solution of $\tilde{\mathcal{R}}(\beta) = 0$ in the range $(\frac{1}{1+\gamma}, 1)$ determines the equilibrium shading behavior. Since the right-hand side of the equation above is decreasing in r , the shading factor is greater under r_H than r_L .

7.3 Proof of Theorem 5.3

The proof is divided into three parts, each corresponding to low, moderate and high target ROIs.

Low target ROI ($\gamma < \gamma_L$) Recall that γ_L is the ROI of a buyer when an SPA with the monopoly reserve is run and buyers bid truthfully. This implies that for any $\gamma < \gamma_L$, the optimal mechanism is an SPA with the monopoly reserve price r^m . To see why, note that when $\gamma < \gamma_L$ and the seller runs the SPA with reserve of r^m , the ROI constraints are not binding under truthful bidding, and as a result, buyers' weakly dominant strategy is to bid truthfully.

Moderate target ROI ($\gamma \in [\gamma_L, \gamma_H]$). To characterize the optimal mechanism, we first construct an upper bound on Problem OPT. To construct an upper bound, we dualize the ROI constraint using a particular Lagrangian multiplier. We then show that the revenue of the SPA with reserve r_γ matches the upper bound. This shows that this mechanism is optimal. To do so, we use the following three lemmas.

LEMMA 7.4 (ROI DECREASES IN RESERVE PRICE). Assume that the seller runs an SPA with reserve $r \in [0, r^m]$ and buyers bid truthfully. Then, the ROI of a buyer, which is between γ_L and γ_H , is decreasing in r .

Lemma 7.4 shows that for any $\gamma \in [\gamma_L, \gamma_H]$, there exists r_γ that solves Eq. (3). In other words, for any moderate target ROI, there exists a reserve price r_γ that binds the ROI constraints. In the next lemma, we use r_γ to construct a Lagrange multiplier.

LEMMA 7.5 (OPTIMAL LAGRANGE MULTIPLIER). Let $\bar{\phi}(x, \lambda, \gamma) = \lambda x + \phi(x)(1 - \lambda(1 + \gamma))$. Then, there exists $\lambda^* \in [0, \frac{1}{1+\gamma}]$ that solves $\bar{\phi}(r_\gamma, \lambda^*, \gamma) = 0$.

We use λ^* to construct our upper bound. By weak duality, Problem OPT is upper bounded by:

$$\max_{(q_i(\cdot), u_i(\underline{v})) \in \mathcal{T}_i, i \in [n]} \mathbb{E} \left[\sum_{i=1}^n \phi(v_i) q_i(\mathbf{v}) - u_i(\underline{v}) + \lambda^* \left(-\gamma v_i q_i(v_i) + (\gamma + 1) u_i(v_i) \right) \right], \quad (14)$$

where $u_i(v_i) = u_i(\underline{v}) + \int_{\underline{v}}^{v_i} q_i(x) dx$ and the feasible region (ignoring the constraints $\sum_{i=1}^n q_i(\mathbf{v}) \leq 1$ for any $\mathbf{v} \in [\underline{v}, \bar{v}]^n$) is given by $\mathcal{T}_i = \{(q_i(\cdot), u_i(\underline{v})) \mid q_i'(v_i) \geq 0, u_i(\underline{v}) \geq 0\}$. We know from Myerson [18] that $u_i(v_i) = u_i(\underline{v}) + \int_{\underline{v}}^{v_i} q_i(x) dx$ and $\mathbb{E}[u_i(v_i)] = \mathbb{E}[\alpha(v_i) q_i(v_i)] + u_i(\underline{v})$, where $\alpha(v_i) = \frac{1-F(v_i)}{f(v_i)}$. Therefore, Eq. (14) can be rewritten as

$$\begin{aligned} & \max_{(q_i(\cdot), u_i(\underline{v})) \in \mathcal{T}_i, i \in [n]} \mathbb{E} \left[\sum_{i=1}^n q_i(\mathbf{v}) \left(v_i(1 - \lambda^*) - \alpha(v_i)(1 - \lambda^*(1 + \gamma)) \right) \right] \\ & - \sum_{i=1}^n u_i(\underline{v})(1 - \lambda^*(\gamma + 1)) \\ & = \max_{(q_i(\cdot), u_i(\underline{v})) \in \mathcal{T}_i, i \in [n]} \mathbb{E} \left[\sum_{i=1}^n q_i(\mathbf{v}) \bar{\phi}(v_i, \lambda^*, \gamma) \right] \\ & - \sum_{i=1}^n u_i(\underline{v})(1 - \lambda^*(\gamma + 1)). \end{aligned} \quad (15)$$

By Lemma 7.5, $\lambda^* \leq \frac{1}{1+\gamma}$. This implies that to maximize the above expression, we need to set $u_i(\underline{v})$ to zero. Furthermore, the item needs to get allocated to a buyer with the highest non-negative modified virtual value $\bar{\phi}(v_i, \lambda^*, \gamma)$. Observe that the modified virtual value, $\bar{\phi}(v_i, \lambda^*, \gamma)$, is increasing in v_i when the distribution F is regular. The monotonicity of $\bar{\phi}(v_i, \lambda^*, \gamma)$ follows from the fact that $\phi(v_i)$ is increasing and $\lambda^* \leq \frac{1}{1+\gamma}$. The fact that $\bar{\phi}$ is monotone implies that the item should be allocated to a buyer with the highest submitted bid/valuation as long as his valuation is greater than r_γ . Recall that by Lemma 7.5, r_γ is the unique solution of $\{x : \bar{\phi}(x, \lambda^*, \gamma) = 0\}$.

Now, consider an SPA auction with reserve price r_γ . This mechanism is a feasible solution of Problem OPT as the ROI constraint is binding in this mechanism; see Lemma 7.4. Furthermore, the revenue of this mechanism matches the upper bound given in Eq. (14). This implies that this mechanism is optimal.

High target ROI ($\gamma > \gamma_H$) Similar to the previous part, we first construct an upper bound on Problem OPT by dualizing the ROI constraint. We then show that the revenue of the SPA with no reserve and subsidy of s matches the upper bound.

Let $\lambda^* = \frac{1}{1+\gamma}$ be the Lagrangian multiplier of the ROI constraints. Then, the upper bound, given in Eq. (15), can be written as $\max_{(q_i(\cdot), u_i(\underline{v})) \in \mathcal{T}_i, i \in [n]} \mathbb{E} \left[\sum_{i=1}^n q_i(\mathbf{v}) \frac{v_i}{1+\gamma} \right]$. The above optimization problem is not a function of $u_i(\underline{v})$. Furthermore, the above expression gets maximized if we allocate the item to the buyer with the highest submitted bid. To complete the proof, we need to show that there is a feasible mechanism that obtains this upper bound.

Consider an SPA with no reserve and positive subsidy ($s = u_i(\underline{v})$) that solves Eq. (6). For this mechanism, the seller's revenue, which is equal to $n \int_0^{\bar{v}} \phi(v) F^{n-1}(v) f(v) dx - ns$, needs to be equal

to the upper bound $\max_{(q_i(\cdot), u_i(\underline{v})) \in \mathcal{T}_i, i \in [n]} \mathbb{E} \left[\sum_{i=1}^n q_i(\mathbf{v}) \frac{v_i}{1+\gamma} \right] = \frac{n}{1+\gamma} \int_0^{\bar{v}} v F^{n-1}(v) f(v) dv$. By the definition of the subsidy s given in Eq. (6), we conclude that the revenue of the seller in this mechanism matches its upper bound, and thus, the mechanism is optimal.

7.4 Proof of Theorem 5.5

By dualizing the ROI constraints and in particular, by Eq. (15) in the proof of Theorem 5.3, Problem OPT is upper bounded by

$$\begin{aligned} & \min_{\lambda_i \geq 0} \max_{(q_i(\cdot), u_i(\underline{v})) \in \mathcal{T}_i, i \in [n]} \mathbb{E} \left[\sum_{i=1}^n q_i(\mathbf{v}) \bar{\phi}_i(v_i, \lambda_i, \gamma_i) \right] \\ & - \sum_{i=1}^n u_i(\underline{v})(1 - \lambda_i(\gamma_i + 1)) \\ & = \min_{\lambda_i \geq 0} \max_{(q_i(\cdot), u_i(\underline{v})) \in \mathcal{T}_i, i \in [n]} \mathcal{L}(q, u(\underline{v}), \Lambda), \end{aligned}$$

where $\mathcal{L}(q, u(\underline{v}), \Lambda) = \mathbb{E} \left[\sum_{i=1}^n q_i(\mathbf{v}) \bar{\phi}_i(v_i, \lambda_i, \gamma_i) \right] - \sum_{i=1}^n u_i(\underline{v})(1 - \lambda_i(\gamma_i + 1))$ and with a slight abuse of notation, $q = (q_1(\cdot), \dots, q_n(\cdot))$ and $u(\underline{v}) = (u_1(\underline{v}), \dots, u_n(\underline{v}))$. We note that the outer optimization is always convex even if the inner one is non-convex. This is the case because $\mathcal{L}(q, u(\underline{v}), \Lambda)$ is linear in λ_i 's. Thus, the outer optimization problem can be solved polynomially using convex stochastic optimization techniques; see for example Uryasev and Pardalos [21]. Also observe that when $\lambda_i > \frac{1}{1+\gamma_i}$, then we can set $u_i(\underline{v})$ to ∞ in order to maximize $\mathcal{L}(q, u(\underline{v}), \Lambda)$. This implies that at any $\lambda_i > \frac{1}{1+\gamma_i}$, $\max_{(q_i(\cdot), u_i(\underline{v})) \in \mathcal{T}_i, i \in [n]} \mathcal{L}(q, u(\underline{v}), \Lambda) = \infty$. Therefore, we have

$$\begin{aligned} & \min_{\lambda_i \in [0, \frac{1}{1+\gamma_i}]} \max_{(q_i(\cdot), u_i(\underline{v})) \in \mathcal{T}_i, i \in [n]} \mathcal{L}(q, u(\underline{v}), \Lambda) \quad (16) \\ & = \min_{\lambda_i \in [0, \frac{1}{1+\gamma_i}]} \max_{(q_i(\cdot), u_i(\underline{v})) \in \mathcal{T}_i, i \in [n]} \mathbb{E} \left[\sum_{i=1}^n q_i(\mathbf{v}) \bar{\phi}_i(v_i, \lambda_i, \gamma_i) \right], \quad (17) \end{aligned}$$

where the inequality follows because $\lambda_i \in [0, \frac{1}{1+\gamma_i}]$. In fact, when $\lambda_i \in [0, \frac{1}{1+\gamma_i}]$, to maximize $\mathcal{L}(q, u(\underline{v}), \Lambda)$, we need to set $u_i(\underline{v})$ to zero, and when $\lambda_i = \frac{1}{1+\gamma_i}$, $\mathcal{L}(q, u(\underline{v}), \Lambda)$ does not depend on $u_i(\underline{v})$. Then, to maximize $\mathbb{E} \left[\sum_{i=1}^n q_i(\mathbf{v}) \bar{\phi}_i(v_i, \lambda_i, \gamma_i) \right]$, the item needs to be allocated to the buyer with the highest non-negative modified virtual value, $\bar{\phi}_i(v_i, \lambda_i, \gamma_i)$. Thus, Eq. (17) can be further simplified as follows

$$\min_{\lambda_i \in [0, \frac{1}{1+\gamma_i}]} \mathbb{E} \left[\left(\max_{i \in [n]} \bar{\phi}_i(v_i, \lambda_i, \gamma_i) \right)^+ \right] = \mathbb{E} \left[\left(\max_{i \in [n]} \bar{\phi}_i(v_i, \lambda_i^*, \gamma_i) \right)^+ \right], \quad (18)$$

where the equality follows from definition of λ_i^* , given in Eq. (8). Let $q_i^*(\mathbf{v}) = 1$ when buyer i has the highest modified virtual value, i.e., when $\bar{\phi}_i(v_i, \lambda_i^*, \gamma_i) \geq \max_{j \neq i} (\bar{\phi}_j(v_j, \lambda_j^*, \gamma_j))^+$, and zero otherwise. The next lemma establishes an important property of λ_i^* .

LEMMA 7.6 (COMPLEMENTARY SLACKNESS). *Let $(q^*, u^*(\underline{v}))$ be defined by $(q^*, u^*(\underline{v})) = \arg \max_{(q_i(\cdot), u_i(\underline{v})) \in \mathcal{T}_i, i \in [n]} \mathcal{L}(q, u(\underline{v}), \lambda^*)$. Then, for any $i \in [n]$, we have $\lambda_i^* \left(\mathbb{E} \left[(v_i - (y_i + 1) \phi_i(v_i)) q_i^*(\mathbf{v}) \right] + (1 + \gamma_i) u_i^*(\underline{v}) \right) = 0$.*

Lemma 7.6 shows that the optimal Lagrangian multiplier, λ_i^* , satisfies complementary slackness. That is, for each buyer i , either $\lambda_i^* = 0$, or $\mathbb{E}\left[(v_i - (\gamma_i + 1)\phi_i(v_i))q_i^*(\mathbf{v})\right] + (1 + \gamma_i)u_i^*(v) = 0$. We note that the latter implies that the ROI constraint of buyer i is binding.

So far, we have fully characterized an upper bound on Problem OPT. As our last step, we need to show that Mechanism $\mathcal{M}(\Lambda^*, \Gamma, S^*)$ obtains the upper bound. Mechanism $\mathcal{M}(\Lambda^*, \Gamma, S^*)$ allocates the item to the buyer with the highest non-negative modified virtual value $\bar{\phi}_i(v_i, \lambda_i^*, \gamma_i)$. Thus, to show that this mechanism is optimal, it suffices to show that $\lambda_i^* \left(\mathbb{E}\left[(v_i - (\gamma_i + 1)\phi_i(v_i))q_i^*(\mathbf{v})\right] + (1 + \gamma_i)u_i^*(v) \right) = 0$. Note that the above expression is the extra term that is added to the objective function of Problem OPT to construct the upper bound. That is, this term is the Lagrange multiplier times the ROI of the constraint of buyer i . The above equation holds by Lemma 7.6 and by observing the fact that $s_i^* = u_i^*(v)$. To see the latter note that both s_i^* and $u_i^*(v)$ are zero where $\lambda_i^* < \frac{1}{\gamma_i + 1}$, and they both solve $\mathbb{E}\left[(v_i - (\gamma_i + 1)\phi_i(v_i))q_i^*(\mathbf{v})\right] + (1 + \gamma_i)y = 0$ when $\lambda_i^* = \frac{1}{\gamma_i + 1}$; see the definition of subsidy s_i^* .

8 ADDITIONAL PROOFS

8.1 Proof of Lemma 7.1

Let $G(\cdot)$ and $g(\cdot)$ be the probability distribution and probability density of D . For any $\beta \leq \frac{r}{\bar{v}}$, $\tilde{\mathcal{R}}(\beta) = 0$. Now assume that $\beta > \frac{r}{\bar{v}}$. Then, by definition, $\tilde{\mathcal{R}}(\beta) = \int_r^{\bar{v}} vG(v\beta)f(v)dv - (1 + \gamma) \int_r^{\bar{v}} \int_r^{v\beta} yg(y)dyf(v)dv$. By taking derivative w.r.t. β , we get $\mathcal{R}'(\beta) = \frac{r^2}{\beta^3}G(r)f(\frac{r}{\beta}) + (1 - (1 + \gamma)\beta) \int_{\frac{r}{\beta}}^{\bar{v}} v^2g(v\beta)f(v)dv$. This expression is positive when $\beta \leq \frac{1}{(1 + \gamma)}$.

8.2 Proof of Lemma 7.2

Assume, contrary to our claim, that $\frac{1}{1 + \gamma} \leq \frac{r}{\bar{v}}$. This implies that $\gamma \geq \frac{\bar{v} - r}{r}$. We will show that this equation cannot hold. Since $\gamma \in (\underline{\gamma}, \bar{\gamma})$, there exists β such that $\mathcal{R}(\beta) = \gamma$. That is, $\frac{\mathbb{E}[(v - D)\mathbb{1}\{v\beta \geq D\}]}{\mathbb{E}[D\mathbb{1}\{v\beta \geq D\}]} = \gamma$. In addition, by our assumption, we have $\gamma \geq \frac{\bar{v} - r}{r}$. That is,

$$\frac{\mathbb{E}[(v - D)\mathbb{1}\{v\beta \geq D\}]}{\mathbb{E}[D\mathbb{1}\{v\beta \geq D\}]} \geq \frac{\bar{v} - r}{r}.$$

But, this cannot be satisfied because $v \leq \bar{v}$ and $D \geq r$.

8.3 Proof of Lemma 7.3

Contrary to our claim, suppose that there are $\beta_1 < \beta_2$ that solves $\tilde{\mathcal{R}}(\beta_1) = \tilde{\mathcal{R}}(\beta_2) = 0$. Let λ_1 and λ_2 solve $\beta_1 = \frac{(1 + \lambda_1)}{(\lambda_1(1 + \gamma) + 1)}$ and $\beta_2 = \frac{(1 + \lambda_2)}{(\lambda_2(1 + \gamma) + 1)}$. Then, by weak duality, the maximum utility obtained by the buyer is less than or equal to $\mathbb{E}[(v - D)\mathbb{1}\{\beta_i v \geq D\}] + \lambda_i (\mathbb{E}[v\mathbb{1}\{\beta_i v \geq D\}] - (1 + \gamma)\mathbb{E}[D\mathbb{1}\{\beta_i v \geq D\}])$ for $i = 1, 2$. Since $\tilde{\mathcal{R}}(\beta_i) = 0$, the maximum utility obtained by the buyer is less than or equal to $\mathbb{E}[(v - D)\mathbb{1}\{\beta_i v \geq D\}]$ for $i = 1, 2$. Now assume that the buyer shades his bids by a factor of β_2 . In this case, his utility is

given by $\mathbb{E}[(v - D)\mathbb{1}\{\beta_2 v \geq D\}]$. By weak duality again, the utility of the buyer should be less than $\mathbb{E}[(v - D)\mathbb{1}\{\beta_1 v \geq D\}]$; that is, $\mathbb{E}[(v - D)\mathbb{1}\{\beta_2 v \geq D\}] < \mathbb{E}[(v - D)\mathbb{1}\{\beta_1 v \geq D\}]$. But, this cannot happen because $\beta_2 > \beta_1$.

8.4 Proof of Lemma 7.4

Consider an SPA with reserve price r . Using the standard Myersonian technique, if all buyers bid truthfully, the ROI of a buyer is given by $\frac{\int_r^{\bar{v}} vF^{n-1}(v)dF(v)}{\int_r^{\bar{v}} \phi(v)F^{n-1}(v)dF(v)} - 1$. In the following, we show that when $r < r^m$, the ROI of the buyer decreases in the reserve price r . Then, the result follows because by definition, the ROI under an SPA with the monopoly reserve is γ_L and the ROI under an SPA with no reserve is γ_H .

Let $V = \int_r^{\bar{v}} vF^{n-1}(v)dF(v)$ and $P = \int_r^{\bar{v}} \phi(v)F^{n-1}(v)dF(v)$. By taking derivative of the ROI w.r.t. r , we get $\frac{F^{n-1}(r)f(r)}{P^2}(-rP + \phi(r)V) \leq 0$, where the inequality holds because $r \leq r^m$ and as a result, $\phi(r) \leq \phi(r^m) = 0$.

8.5 Proof of Lemma 7.5

We first argue that for any $\lambda > \frac{1}{\gamma + 1}$, $\bar{\phi}(r_\gamma, \lambda, \gamma) > 0$. This implies that $\bar{\phi}(r_\gamma, \lambda, \gamma)$ can be zero only when $\lambda \leq \frac{1}{1 + \gamma}$. Then, considering the fact that $\bar{\phi}(r_\gamma, 0, \gamma) = r_\gamma - \alpha(r_\gamma) = \phi(r_\gamma) \leq 0$, there exists $\lambda \in [0, \frac{1}{1 + \gamma}]$ that solves $\bar{\phi}(r_\gamma, \lambda, \gamma) = 0$. To see why $\bar{\phi}(r_\gamma, 0, \gamma) \leq 0$ note that by Lemma 7.4, $r_\gamma \leq r^m$. Suppose that $\lambda > \frac{1}{\gamma + 1}$. By definition, $\bar{\phi}(r_\gamma, \lambda, \gamma) = \lambda x + \phi(x)(1 - \lambda(1 + \gamma)) > 0$, where the inequality holds because by Lemma 7.4 and $\lambda > \frac{1}{\gamma + 1}$.

8.6 Proof of Lemma 7.6

By Eq. (16), λ_i^* , $i \in [n]$, solves the following equation

$$\begin{aligned} (\lambda_1^*, \dots, \lambda_n^*) = \arg \min_{\lambda_i \in [0, \frac{1}{1 + \gamma_i}]} & \mathbb{E}\left[\sum_{i=1}^n q_i^*(\mathbf{v})\bar{\phi}_i(v_i, \lambda_i, \gamma_i)\right] \\ & - \sum_{i=1}^n u_i^*(v)(1 - \lambda_i(\gamma_i + 1)), \end{aligned}$$

Since, the above optimization is convex, λ_i^* satisfies the KKT condition. That is,

$$\begin{aligned} \lambda_i^* \frac{\partial \left(\mathbb{E}\left[\sum_{i=1}^n q_i^*(\mathbf{v})\bar{\phi}_i(v_i, \lambda_i^*, \gamma_i)\right] - \sum_{i=1}^n u_i^*(v)(1 - \lambda_i^*(\gamma_i + 1)) \right)}{\partial \lambda_i^*} & = 0 \\ \Rightarrow \lambda_i^* \left(\mathbb{E}\left[(-1 + \gamma_i)\bar{\phi}_i(v_i) + v_i q_i^*(\mathbf{v})\right] + u_i^*(v)(\gamma_i + 1) \right) & = 0. \end{aligned}$$