

Reserve Price Optimization for First Price Auctions

Zhe Feng^{*a}, Sébastien Lahaie^b, Jon Schneider^b, and Jinchao Ye^b

^aJohn A. Paulson School of Engineering and Applied Sciences, Harvard University

`zhe_feng@g.harvard.edu`

^bGoogle Inc, NYC

`slahaie, jschnei, jinchao@google.com`

June 28, 2020

Abstract

The display advertising industry has recently transitioned from second- to first-price auctions as its primary mechanism for ad allocation and pricing. In light of this, publishers need to re-evaluate and optimize their auction parameters, notably reserve prices. In this paper, we propose a gradient-based algorithm to adaptively update and optimize reserve prices based on estimates of bidders’ responsiveness to experimental shocks in reserves. Our key innovation is to draw on the inherent structure of the revenue objective in order to reduce the variance of gradient estimates and improve convergence rates in both theory and practice. We show that revenue in a first-price auction can be usefully decomposed into a *demand* component and a *bidding* component, and introduce techniques to reduce the variance of each component. We characterize the bias-variance trade-offs of these techniques and validate the performance of our proposed algorithm through experiments on synthetic data and real display ad auctions data from Google ad exchange.

1 Introduction

A reserve price in an auction specifies a minimum acceptable winning bid, below which the item remains with the seller. The reserve price may correspond to some outside offer, or the value of the item to the seller itself, and more generally may be set to maximize expected revenue [25]. In a data-rich environment like online advertising auctions it becomes possible to learn a revenue-optimal reserve price over time, and there is a substantial literature on optimizing reserve prices for *second-price* auctions, which have been commonly used to allocate ad space [27, 23, 24].

In this work we examine the problem of reserve price optimization in *first-price* (i.e., pay-your-bid) auctions, motivated by the fact that all the major ad exchanges have recently transitioned to this auction format as their main ad allocation mechanism [10, 6]. First-price auctions have grown in favor because they are considered more transparent, in the sense that there is no uncertainty in the final price upon winning [5].¹ Unless restrictive assumptions are met, there is in theory no revenue ranking between first- and second-price auctions [19], and there is no guarantee that reserve prices optimized for second-price auctions will continue to be effective in a first-price setting.

^{*}Supported by a NSF award CCF-1841550 and a Google PhD Fellowship. This work was done when the first author was an intern in Google Inc, NYC.

¹The full reasons for the transition are complex, and include the rise of “header bidding” [30]. A header bidding auction is a first-price auction usually triggered by code in a webpage header (hence the name).

From a learning standpoint the shift from second- to first-price auctions introduces several new challenges. In a second-price auction, truthful bidding is a dominant strategy no matter what the reserve. The bidders’ value distributions are therefore readily available, and bids stay static (in principle) as the reserve is varied. In a first-price auction, in contrast, bidders have an incentive to shade their values when placing their bids, and bid-shading strategies can vary by bidder. The gain from setting a reserve price now comes if (and only if) it induces higher bidding, so an understanding of bidder responsiveness becomes crucial to setting effective reserves.

Bid adjustments in response to a reserve price can occur at different timescales. If a bidder observes that it wins too few auctions because of the reserve price, it may increase its bid in the long-term (in a matter of hours up to weeks). Our focus here is on setting reserves prices by taking into account *immediate* bidder responses to reserves. We assume that each bidder has a fixed, unknown bidding function $b(r, v)$ that depends on its private value v and the observed auction reserve r . This agrees with practice in display ad auctions because the reserve r is normally sent out in the ‘bid request’ message to potential bidders [18]. To the extent that the bid function responds to r , first-price reserves can potentially show an immediate positive effect on revenue.

Our Results

We propose a gradient-based approach to adaptively improve and optimize reserve prices, where we perturb current reserves upwards and downwards (e.g., by 10%) on random slices of traffic to obtain gradient estimates.

Our key innovation is to draw on the inherent structure of the revenue objective in order to reduce the variance of gradient estimates and improve convergence rates in both theory (e.g., see Corollary 4.1) and practice. We show that revenue in a first-price auction can be usefully decomposed into two terms: a *demand curve* component which depends only on the bidder’s value distribution; and a *bidding* component whose variance can be reduced based on natural assumptions on bidding functions.

A demand curve is a simpler, more structured object than the original revenue objective (e.g., it is downward-sloping), so the demand component lends itself to parametric modeling to reduce the variance. We offer two variance reduction techniques for the bidding component², referred to as *bid truncation* and *quantile truncation*. Bid truncation can strictly decrease variance with no additional bias assuming the right bidding model, whereas quantile truncation may introduce bias but is less sensitive to assumptions on the bidding model.

We evaluate our approach over synthetic data where bidder values are drawn uniformly, and also over real bid distributions collected from the logs of the Google ad exchange with different bidder response models. Our experimental results confirm that the combination of variance reduction on both objective components leads to the fastest convergence rate. For the demand component, a simple logistic model works well over the synthetic (i.e., uniform) data, but a flexible neural net is needed over the semi-synthetic data. For the bidding component, we find that quantile truncation is much more robust to assumptions on the bidding model.

Related Work

This paper connects with the rich literature on *reserve price optimization for auctions*, e.g., [25, 29]. How to set optimal reserve prices in *second price auctions* based on access to bidders’ historical bid data has been an increasingly popular research direction in Machine Learning community, e.g., [26, 23, 24]. Another related line of work uses no-regret learning in second price auctions with

²Variance reduction of the bidding component relies on the insight that bids far above the reserves are little affected by them (under natural bidding models), so these bids can be filtered out when computing gradient estimates—changes in such bids are likely due to noise rather than any effect of reserves.

partial information feedback to optimize reserve prices, e.g., [7, 9]. All of the works cited so far rely on the fact that the seller can directly learn the valuation distribution from historical bid data, since the second price auction is truthful.

For first-price auctions, we have found little work on setting optimal reserves for asymmetric bidders, since there are no characterizations of equilibrium strategies for this case. Results are only available for limited environments, such as bidders with uniform valuation distributions [19, 22]. Recently, there has been a line of work regarding *revenue optimization against strategic bidders in repeated auctions*, e.g., [3, 17]. In this paper, instead of assuming bidders act strategically, we assume each bidder has a fixed bidding function in response to reserves. This is a common assumption in large market settings and in the dynamic pricing literature [21].

The algorithms developed in this paper are related to the literature on *online convex optimization with bandit feedback* [11, 16, 1, 2]. However, there are two key differences with our work: (1) the revenue function in a first price auction is non-convex, and (2) the seller cannot obtain perfect revenue feedback under perturbed reserves with just a single query (i.e., auction)—the seller needs multiple queries to achieve accurate estimates with high confidence. Our algorithm is also related to *zeroth-order stochastic gradient methods* [14, 4, 13, 20], which we discuss in detail later in Section 3.

2 Preliminaries

We consider a setting where a seller repeatedly sells a single item to a set of m bidders via a first price auction. In such an auction, the seller first sends out a reserve price r to all bidders. Each bidder i then submits a bid b_i . The bidder with the highest bid larger than r wins the item and pays their bid; if no bidder bids above r , the item goes unallocated. Note that the type of reserve price we consider in this work is *anonymous* in the sense that each bidder sees the same reserve price.

Each bidder i has a private valuation $v_i \in [0, 1]$ for the item, where each value v_i is drawn independently (but not necessarily identically) from some unknown distribution.³ In a first-price auction, only the highest bid matters for both allocation and pricing. Thus, to simplify the notation, we write $v = \max_i v_i$ to denote the maximum value and v is drawn i.i.d. from an unknown distribution F across each auction. Our analysis from here on will refer to this ‘representative’ highest bidder. (See Appendix A for a rigorous justification of why we can reduce multiple bidders to a single bidder.)

We write $b(r, v)$ to denote the maximum bid when the reserve price is r and the maximum value is v , and $\mathcal{B}(r)$ to denote the distribution of $b(r, v)$ for a fixed r when v is drawn according to \mathcal{F} . The main goal of the seller considered in this work is to learn the optimal reserve price $r \in [0, 1]$ that maximizes expected revenue:

$$\mathbb{E}_{v \sim \mathcal{F}} [b(r, v) \cdot \mathbb{I}\{b(r, v) \geq r\}]. \quad (1)$$

Note that there is no reason for a bidder to bid a positive value less than the reserve r : such a bid is guaranteed to lose. Therefore, without loss of generality we can assume that if $b(r, v) < r$, then $b(r, v) = 0$. This allows us to write the revenue simply as:

$$\mu(r) = \mathbb{E}_{b \sim \mathcal{B}(r)} [b] = \mathbb{E}_{v \sim \mathcal{F}} [b(r, v)].$$

In this paper, we focus on maximizing the revenue function $\mu(r)$ in the steady state.

³This is without loss of generality, our analysis can easily be applied to any bounded valuation setting.

Response Models

We begin by describing some general properties of bidding functions that hold for any utility-maximizing bidders (see [22] for further discussion).

Definition 2.1. A bidding function $b(r, v)$ satisfies the following properties: 1) $b(r, v) \leq v$ for all v ; 2) $b(r, v) \geq r$ for $v \geq r$; 3) $b(r, v) = 0$ for $v < r$; 4) $b(r, v)$ is non-decreasing in v for all r .

In some of our algorithms, we would like to impose additional constraints on the response model which, while not a consequence of utility-maximizing behavior, are likely to hold in practice. One such constraint is the *diminishing sensitivity in value of bid to reserve*. This says that bidders with a larger value will change their bid less in response to a change in reserves.

Definition 2.2 (Diminishing sensitivity of bid to reserve). If $v_H > v_L \geq r$, then for $\delta > 0$ and $v_L \geq r + \delta$ we have $b(r + \delta, v_H) - b(r, v_H) \leq b(r + \delta, v_L) - b(r, v_L)$.

One natural and concrete example of a response model is a bidder that increases its bid to the reserve as long as the reserve is below its value. We refer to this as the *perfect response model*, formally defined as follows.

Definition 2.3. A perfect response bidding function takes the form:

$$b(r, v) = \begin{cases} b(0, v) & \text{if } b(0, v) \geq r \\ r & \text{if } b(0, v) < r \leq v \\ 0 & \text{if } v < r \end{cases}$$

Note that the perfect response model is based on the *original bid* of the bidder under reserve price 0, namely $b(0, v)$. If $b(0, v)$ is already above the reserve, then this bidder is unaffected by the reserve. Note that the perfect response model satisfies the diminishing sensitivity property.

In practice, bidders are unlikely to exactly follow the perfect response model; for example, bidders will often increase their bid to some amount strictly above the reserve r so as to remain competitive with other bidders. For this reason, we propose a relaxation of the perfect response model which we call the ε -bounded response model: the bid is at most ε greater than what it would have been under the perfect response model if $b(0, v) < r \leq v$ (see also Definition C.6). Note that the ε -bounded response model becomes the perfect response model when $\varepsilon = 0$.

3 Gradient Descent Framework

The first-price auction setting introduces several challenges for setting reserve prices. First, the seller cannot observe true bidder values because truthful bidding is not a dominant strategy in a first-price auction. Second, how the bidders will react to different reserves is unknown to the seller—the only information that the seller receives is bids drawn from distribution $\mathcal{B}(r)$ when the seller sets a reserve price r .

One natural idea, and the approach we take in this paper, is to optimize the reserve price via gradient descent. Gradient descent is only guaranteed to converge to the optimal reserve when our objective is convex (or at least, unimodal), which is not necessarily true for an arbitrary revenue function. However, gradient descent has a number of practical advantages for reserve price optimization, including:

1. Gradient descent allows us to incorporate prior information we may have about the location of a good reserve price (possibly significantly reducing the overall search cost).

Algorithm 1 Zeroth-order stochastic projected gradient framework for reserve optimization.

Input: Initial reserve $r_1 \in (0, 1)$, total number of iterations T (a variable to be fixed later).

Output: Reserve prices r_2, r_3, \dots, r_{T+1} .

for $t = 1, 2, \dots, T$ **do**

Set a reserve price of $r_t^+ = (1 + \beta_t)r_t$ in n_t auctions.

Set a reserve price of $r_t^- = (1 - \beta_t)r_t$ in n_t auctions.

Construct an estimate \hat{G}_t of the gradient of revenue at r_t , based on the feedback of experiments.

Update reserve: $r_{t+1} = \Pi(r_t + \alpha_t \hat{G}_t)$, where

$$\Pi(x) = \arg \min_{z \in (0,1)} |z - x|.$$

end for

2. The adaptivity of gradient descent allows us to quickly converge to a local optimum and follow this optimum if it changes over time, significantly saving on search cost (over global methods such as grid search).
3. In practice, many revenue curves have a unique local optimum (see Section 5), so gradient descent is likely to converge to the optimal reserve.

More specifically, since the seller has no direct access to the gradients (i.e, first-order information) of $\mu(r)$, we consider approaches that fit in the framework of *zeroth-order stochastic optimization*. Our framework, summarized in Algorithm 1, proceeds in rounds. In round t where the current reserve is r_t , the seller selects a perturbation size β_t and randomly sets the reserve price to either $(1 + \beta_t)r_t$ or $(1 - \beta_t)r_t$ on separate slices of experiment traffic, until it has received n_t samples from both $\mathcal{B}((1 + \beta_t)r_t)$ and $\mathcal{B}((1 - \beta_t)r_t)$. The seller then uses these $2n_t$ samples to estimate the gradient \hat{G}_t of the revenue curve $\mu(r)$ at r_t and updates the reserve price based on this gradient estimate using learning rate (step size) α_t .

We assume that we have access to a fixed total number of samples $N = \sum_{t=1}^T n_t$ (the number of iterations T is a variable that will be fixed later). There is then a trade-off between n_t (i.e, the number of samples per iteration) and T (the number of iterations available to optimize the reserve price).

Zeroth-order stochastic gradient descent is a well-studied problem [14, 4, 13, 20]. In this paper, we focus on taking advantage of the structure of $b(r, v)$ to construct good *discrete gradient* estimates \hat{G}_t , as this aspect is specific to the problem of reserve price optimization. Specifically, we tackle the following problem which we term the *discrete gradient problem*:

- **Input:** n samples X_1^+, \dots, X_n^+ drawn i.i.d from $\mathcal{B}(r^+)$ and n samples X_1^-, \dots, X_n^- drawn i.i.d from $\mathcal{B}(r^-)$, for known $r^+ > r^-$.
- **Output:** An estimator \hat{G} for the discrete derivative $(\mu(r^+) - \mu(r^-))/(r^+ - r^-)$. This estimator has bias $\text{Bias}(\hat{G})$ and variance $\text{Var}(\hat{G})$, where $\text{Bias}(\hat{G}) = \left| \mathbb{E}[\hat{G}] - \frac{\mu(r^+) - \mu(r^-)}{r^+ - r^-} \right|$.

Solutions to the discrete gradient problem with small bias and variance directly translate into faster convergence rates for our gradient descent. We provide a detailed convergence result in Theorem C.2 in Appendix C.1. We summarize this result informally as follows.

Theorem 3.1 (Informal Restatement of Theorem C.2). *If for all t , $\text{Bias}(\hat{G}_t) \leq B$ and $\text{Var}(\hat{G}_t) \leq V$*

then for optimal choices of α_t and n_t (and fixing $\beta_t = \delta/2r_t$), Algorithm 1 satisfies

$$\min_{t \in [T]} |\mathcal{P}_C^t|^2 = \tilde{O} \left(T^{-1/2} + \delta^2 + B^2 + V + (T/N)^2 \right).$$

Here \mathcal{P}_C^t can be thought of as the true gradient at round t (see Definition C.1 in Appendix).

Intuitively, we want to design an estimator and choose our parameters α_t, β_t, n_t , so as to trade off between δ, B , and V . In the following sections, we show how to do this for a variety of bidder response models.

Naive Gradient Estimation

The simplest method for estimating the discrete gradient is to take the difference between the average revenue from bids from $\mathcal{B}(r^+)$ and the average revenue from bids from $\mathcal{B}(r^-)$. More formally, we compute discrete gradient as,

$$\hat{G} = \frac{\sum_{i=1}^n X_i^+ - \sum_{i=1}^n X_i^-}{n(r^+ - r^-)}. \quad (2)$$

We show that \hat{G} has the following properties.

Theorem 3.2. *Assume that $r^+ - r^- = \delta$, then $\text{Bias}(\hat{G}) = 0$, $\text{Var}(\hat{G}) \leq \frac{1}{2\delta^2 n}$.*

This leads to the following convergence rate via Theorem 3.1.

Corollary 3.1. *Using this estimator \hat{G} , and setting $T = N^{1/2}$ and $\delta = \Theta(N^{-1/8})$, Algorithm 1 achieves convergence, $\min_{t \in [T]} |\mathcal{P}_C^t|^2 \leq \tilde{O}(N^{-1/4})$.*

Although there are no matching lower bounds, this is the best known asymptotic convergence rate for zeroth-order optimization over a non-convex objective [14, 4]. The naive gradient estimation approach has the advantage that it works regardless of response model, is simple to compute (it uses only revenue information and not individual bids), and leads to an *unbiased* estimator for the discrete derivative. The disadvantage is that the variance of this estimator can be large (especially as we take δ small). In the following section, we show how to address this by taking into account the inherent structure of the revenue objective based on an underlying bidder response model.

4 Variance Reduced Gradient Estimation

In this section, we first introduce another representation of the revenue formula by decomposing it into a *demand* component and a *bidding* component. We then propose techniques to reduce the variance of the discrete gradient of each component.

4.1 Revenue Decomposition

We can decompose the revenue $\mu(r)$ in the following way.

Theorem 4.1. *We have that*

$$\mu(r) = \mathbb{E}_{v \sim \mathcal{F}}[\max(b(r, v) - r, 0)] + r \Pr_{v \sim \mathcal{F}}[v \geq r]. \quad (3)$$

Define $E(r) = \mathbb{E}_{v \sim \mathcal{F}}[\max(b(r, v) - r, 0)]$ and $D(r) = \Pr_{v \sim \mathcal{F}}[v \geq r]$, so that $\mu(r) = E(r) + rD(r)$. These two terms capture two different aspects of bidder behavior which contribute to revenue. The function $D(r)$ amounts to a “demand curve” which gives the proportion of values that clear the reserve r , and therefore the proportion of auctions that are bid on at r . If the auction were just a simple posted-price auction (i.e., the winner is charged the quoted price r), then the *demand* component $rD(r)$ would be the associated revenue. However, in a first-price auction the winning bidder pays its bid, not the reserve. Therefore the *bidding* component $E(r)$ captures the excess contribution from bids greater than the reserve.

To construct a good estimator \hat{G} for the discrete gradient of $\mu(r)$, it suffices to construct good estimators \hat{G}_E and \hat{G}_D for the discrete gradients of $E(r)$ and $rD(r)$ respectively, and then output $\hat{G} = \hat{G}_E + \hat{G}_D$. Note that $\text{Bias}(\hat{G}) \leq \text{Bias}(\hat{G}_E) + \text{Bias}(\hat{G}_D)$ and $\text{Var}(\hat{G}) \leq 2(\text{Var}(\hat{G}_D) + \text{Var}(\hat{G}_E))$, so it suffices to bound the bias and variance of each component separately.

4.2 Estimating the Demand Component Gradient

We begin by discussing how to estimate the gradient \hat{G}_D of the demand component of revenue. One method of doing so is by estimating $D(r)$ with a parametric function $f_\theta(r)$, and using this approximation to estimate the gradient \hat{G}_D . (See Appendix B for additional justification for why this is likely to be possible and helpful). Suppose that we have access to additional historical data \mathcal{S} with which we can fit our parametric class to $D(r)$; let $\hat{\theta}$ be the resulting learned parameter. This learned demand function gives rise to the following estimator \hat{G}_D :

$$\hat{G}_D = \frac{r^+ f_{\hat{\theta}}(r^+) - r^- f_{\hat{\theta}}(r^-)}{r^+ - r^-} \quad (4)$$

Note that this decreases overall variance, the variance of \hat{G}_D is 0 because the randomness of \hat{G}_D only comes from historical samples \mathcal{S} , which are independent of the samples obtained in the current round, at the cost of a possible increase in bias (due to inaccuracy in estimating $D(r)$).

4.3 Estimating the Bidding Component Gradient

In this section we propose a variance reduction method to achieve a better estimator for \hat{G}_E for a variety of bidder models.

Variance reduction via bid truncation. We first consider the special case of the perfect response (and more generally, the ε -bounded response) bidding model. In the perfect response model, if you were going to bid $b > r^+$ when the reserve was r^+ , you will bid the same bid b when the reserve is r^- . This means that large bids (bids larger than r^+) do not contribute in expectation to $\mu(r^+) - \mu(r^-)$, but they do add noise to our gradient estimation. By filtering these out, we can reduce the variance of our estimator while keeping our estimator unbiased.

Since we only apply this filtering when estimating the bidding component $E(r)$ but not the demand component $rD(r)$, we must be careful when implementing this. Note that a large bid $b > r^+$ contributes $b - r^+$ to $E(r^+)$ and $b - r^-$ to $E(r^-)$, and therefore $r^+ - r^-$ to $E(r^+) - E(r^-)$. We can therefore construct an unbiased estimator for $E(r^+) - E(r^-)$ by computing the contribution of unfiltered bids ($b < r^+$) from both $\mathcal{B}(r^+)$ or $\mathcal{B}(r^-)$ and then adding $r^+ - r^-$ for each filtered bid in $\mathcal{B}(r^-)$ (or equivalently, each filtered bid in $\mathcal{B}(r^+)$; under perfect response, the fraction of filtered bids is equal in both models in expectation). Note that every bid from $\mathcal{B}(r^+)$ is either filtered or has excess 0, so we can write this gradient \hat{G}_E entirely in terms of bids from $\mathcal{B}(r^-)$. Formally, we define truncated bid Y_i^- as

$$Y_i^- = \begin{cases} \max(X_i^- - r^-, 0) & \text{if } X_i^- \leq r^+ \\ (r^+ - r^-) & \text{otherwise} \end{cases}$$

Our estimate for the gradient of $E(r)$ is then given by

$$\hat{G}_E = -\frac{\sum_{i=1}^n Y_i^-}{n(r^+ - r^-)}. \quad (5)$$

Since any bid in an ε -bounded model only differs from one in the perfect response model by at most ε , we can apply this same estimator to an ε -bounded response model. The following theorem characterizes the bias and variance of the estimator for the ε -bounded response model.

Theorem 4.2. *Assume that $r^+ - r^- = \delta$, then the estimator \hat{G}_E in Eq. (5) for ε -bounded response model, satisfies: $\text{Bias}(\hat{G}_E) = \frac{2\varepsilon}{\delta}$, $\text{Var}(\hat{G}_E) \leq \frac{1}{4n}$.*

Note that the bias of estimator \hat{G}_E is 0 for the perfect response model. The complete proof is given in Appendix C.5. Combining the above results for \hat{G}_E and \hat{G}_D , we have the following improved convergence result for the ε -bounded response model.

Corollary 4.1. *Suppose $\text{Bias}(\hat{G}_D) \leq \varepsilon_D/\delta$. Using the estimator \hat{G}_E proposed in Eq. (5) for the ε -bounded response model, setting $T = N^{2/3}$ and $\delta = \Theta(\sqrt{\varepsilon + \varepsilon_D})$, Algorithm 1 achieves convergence, $\min_{t \in [T]} |\mathcal{P}_C^t|^2 \leq \tilde{O}(\varepsilon + \varepsilon_D + N^{-1/3})$.*

For perfect response bidding models, the above convergence rate is strictly faster than the convergence rate of naive estimator in Corollary 3.1 (state-of-the-art convergence rate for zeroth-order stochastic gradient descent), but with additional bias coming from demand estimation. However, we show this bias has practically negligible effect on the revenue in our experiments.

Variance reduction via quantile truncation. In Eq. (5), we reduced the variance of \hat{G}_E by truncating all bids at the fixed threshold of $t = r^+$. In general, this does not quite work: for bidder response models that are far from perfect response, this truncation can introduce a very large bias. Here we demonstrate one technique for constructing good estimators \hat{G}_E as long as the bidding function $b(r, v)$ possesses diminishing sensitivity in value to reserve.

Instead of truncating in *bid space*, we will instead want to truncate in *value space* to reduce the variance. Even though we cannot directly truncate by values, since $b(r, v)$ is monotonically increasing in v , quantiles of bids (e.g., of $\mathcal{B}(r^+)$ and $\mathcal{B}(r^-)$) directly correspond to quantiles of values (of \mathcal{F}). Instead of setting a threshold t directly on the value, it is therefore equivalent to truncate at a *fixed quantile of the bid distribution*.

To achieve this, we first sort X_i^+ and X_i^- in ascending order. Then we compute \hat{G}_E as

$$\hat{G}_E = \frac{\sum_{i=1}^{qn} \max(X_i^+ - r^+, 0) - \sum_{i=1}^{qn} \max(X_i^- - r^-, 0)}{n(r^+ - r^-)} - (1 - q), \quad (6)$$

where $q \in [0, 1]$ is the *quantile threshold* used to truncate bids. The following theorem characterizes the bias and variance of the above \hat{G}_E ,

Theorem 4.3. *Let $r^+ - r^- = \delta$, $t = \mathcal{F}^{-1}(q)$, and $\tilde{t} = \mathcal{F}^{-1}(q + n^{-2/3})$. Then the estimator \hat{G}_E in Eq. (6) satisfies, $\text{Bias}(\hat{G}_E) \leq \frac{(1-q)(b(r^+, t) - b(r^-, t))}{\delta} + O(n^{-2/3})$, $\text{Var}(\hat{G}_E) \leq \frac{2\tilde{t}^2}{n\delta^2} + O(n^{-5/3}\delta^{-2})$.*

Unlike with bid truncation, with quantile truncation we have a clear bias-variance tradeoff as we change q : larger values of q decrease the bias (both by decreasing $(1-q)$ and $b(r^+, t) - b(r^-, t)$, which is decreasing due to diminishing sensitivity) but lead to larger variance. Since one can estimate this bound on the bias (by approximating $b(r^+, t) - b(r^-, t)$ via $Y_{qn}^+ - Y_{qn}^-$), it is possible to choose q to optimize this bias-variance tradeoff as one sees fit (for example, to minimize $B^2 + V$ in Theorem 3.1). We show a convergence rate result for this *quantile truncation* approach in Corollary C.1 in Appendix C.7.

5 Experiments

We evaluate the performance of our algorithms on synthetic and semi-synthetic data sets. Due to space limitations, we present the complete experimental results in Appendix D.

5.1 Data Generation

The data generation process consists of two parts: a base bid distribution specifying the distribution of bids when no reserve is set, and a response model describing how a bidder with bid b would update its bid in response to a reserve of r .

Response models. We assume that in the absence of a reserve bidders bid a constant fraction γ of their value v (i.e., $b = \gamma v$), which we refer to as *linear shading*. We consider linear shading combined with perfect response and with ϵ -bounded response, which we implement by adding a uniform $[0, \epsilon]$ random variable to the bid. We also examine equilibrium bidding for n i.i.d. bidders with uniformly distributed valuation [19]: $b = \frac{r^n + (n-1)v^n}{nv^{n-1}}$.

Synthetic data. In our synthetic data sets, the (base) bid distribution is the uniform $[0, 1]$ distribution. We apply the perfect response model, ϵ -bounded response model and equilibrium bidding model. In the simulations, we apply a constant shading factor of 0.4 for the perfect response model and ϵ -bounded response model. For equilibrium bidding, we assume that each auction contains $n = 2$ bidders.


Semi-synthetic data. For our semi-synthetic data sets, we separately collected the empirical distributions of winning bids over one day for 20 large publishers on a major display ad exchange. Each distribution was filtered for outliers and normalized to the interval $[0, 1]$. For this semi-synthetic data we only test the perfect-response model and ϵ -bounded response model, since there is no closed-form solution for the equilibrium bidding strategy. We use 0.3 as the constant shading factor for semi-synthetic data.

5.2 Methodology

Gradient descent algorithms. We examine five different gradient descent algorithms: (I) *Naive GD*: naive gradient descent using the gradient estimator in Eq. (2); (II) *Naive GD with bid truncation*: gradient descent using the gradient estimator in Eq. (5) for the bidding component, and a naive estimate⁴ of the demand component; (III) *Naive GD with quantile truncation*: gradient descent using the gradient estimator in Eq. (6) for the bidding component, and naive estimate of the demand component; (IV) *Demand modeling with bid truncation*: Same as the second variant, but with a parametric model of the demand curve to estimate demand component of gradient; (V) *Demand modeling with quantile truncation*: Same as the third variant, but with a parametric model of the demand curve to estimate demand component of gradient. The parameters used in these algorithms are specified in Appendix D.

Demand curve estimation. To reduce variance following the ideas of Section 4, we need a model \hat{G}_D for the demand component of the discrete gradient. Instead of estimating \hat{G}_D from historical data, we adaptively learn the demand curve during the training process. Concretely, at each round t , we observe new (reserve, demand) pairs from $2n_t$ samples and retrain our demand curve using all the samples observed up to the current round. We use this trained demand curve to compute \hat{G}_D based on (4). For the synthetic data, a simple logistic regression can effectively learn the demand curve. However, the semi-synthetic data required a more flexible model so for

⁴We can form a naive unbiased estimator $\hat{G}_D = \frac{\hat{D}(r^+) - \hat{D}(r^-)}{r^+ - r^-}$, where $\hat{D}(r^+) = \frac{1}{n} \sum_i \mathbb{I}\{x_i^+ \geq r^+\}$ and similarly for $\hat{D}(r^-)$.



(a) Synthetic data with perfect re- (b) Synthetic data with equilibrium (c) Semi-synthetic data with per-
 sponse. response. fect response.

Figure 1: Revenue as a function of round t for (a) synthetic data with perfect response, (b) synthetic data with equilibrium response, and (c) semi-synthetic data with perfect response.




Figure 2: Revenue as a function of round t for synthetic data with equilibrium response.

this case we model demand using a fully connected neural network with 1 hidden layer, 15 hidden nodes and ReLU activations.

5.3 Evaluation

Effectiveness of gradient descent. First, we confirm that gradient descent can effectively find optimal reserves in our models. For each semi-synthetic model, we construct the revenue curve as a function of reserve with assumed response models. We find that 19 out of the 20 revenue curves have a clear single local maximum (the remaining curve has 2). In all cases (synthetic and semi-synthetic models), the revenue learned by the naive gradient descent algorithm is at least 95% of the revenue at the optimal reserve, which indicates that gradient descent can efficiently find the optimal reserve in these cases despite the lack of convexity.

Effectiveness of variance reduction methods. We first evaluate the performance of the quantile-based variance reduction method. We run the algorithm variants (I), (III) and (V) under synthetic data and semi-synthetic data with multiple bidder response models. Figures (1a) and (1c) show the revenue achieved by the three algorithms over time under the perfect response model. We find that quantile-based variance reduction leads to a more stable training process which converges faster than naive gradient descent. Figure (1b) evaluates the performance of the three algorithm variants under synthetic data and an equilibrium response model, with similar conclusions. Overall, quantile-based variance reduction outperforms naive gradient descent. Moreover, with the addition of demand curve estimation, algorithm variant (V) achieves better revenue and converges to an optimal reserve faster than the other two algorithms, in agreement with our theoretical guarantees.

We next consider variance reduction using bid truncation, which is used in algorithm variants (II) and (IV). Bid truncation is tailored to perfect response and performs the best overall for this response model, in accordance with the theoretical guarantees, but quantile truncation is competitive and often performs as well over the semi-synthetic data (see Appendix D for a detailed comparison). Under the equilibrium response model, bid truncation can in fact hinder the training

process and lead to a substantially suboptimal reserve price (see Figure 2). In summary, quantile-based variance reduction coupled with a good demand-curve estimation is the method of choice to achieve good reserve prices under a range of different bid distributions and bidder response models.

References

- [1] Alekh Agarwal, Ofer Dekel, and Lin Xiao. Optimal algorithms for online convex optimization with multi-point bandit feedback. In *Proceedings of the 23rd Annual Conference on Learning Theory (COLT)*, June 2010.
- [2] Alekh Agarwal, Dean P Foster, Daniel J Hsu, Sham M Kakade, and Alexander Rakhlin. Stochastic convex optimization with bandit feedback. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 1035–1043, 2011.
- [3] Kareem Amin, Afshin Rostamizadeh, and Umar Syed. Learning prices for repeated auctions with strategic buyers. In *Advances in Neural Information Processing Systems 26*, 2013.
- [4] Krishnakumar Balasubramanian and Saeed Ghadimi. Zeroth-order nonconvex stochastic optimization: Handling constraints, high-dimensionality and saddle-points, 2018.
- [5] Ross Benes. How SSPs use deceptive price floors to squeeze ad buyers. <https://digiday.com/marketing/ssps-use-deceptive-price-floors-squeeze-ad-buyers/>, September 2017. Accessed: 2020-01-29.
- [6] Jason Bigler. Rolling out first price auctions to google ad manager partners digi-day. <https://www.blog.google/products/admanager/rolling-out-first-price-auctions-google-ad-manager-partners>, September 2019. Accessed: 2020-01-27.
- [7] Avrim Blum, Vijay Kumar, Atri Rudra, and Felix Wu. Online learning in online auctions. *Theor. Comput. Sci.*, 324:137–146, 2003.
- [8] Stéphane Boucheron, Maud Thomas, et al. Concentration inequalities for order statistics. *Electronic Communications in Probability*, 17, 2012.
- [9] N. Cesa-Bianchi, C. Gentile, and Y. Mansour. Regret minimization for reserve prices in second-price auctions. *IEEE Transactions on Information Theory*, 61(1):549–564, Jan 2015.
- [10] Yuyu Chen. Programmatic advertising is preparing for the first-price auction era. <https://digiday.com/marketing/programmatic-advertising-readying-first-price-auction-era>, October 2017. Accessed: 2020-01-29.
- [11] Abraham D. Flaxman, Adam Tauman Kalai, Adam Tauman Kalai, and H. Brendan McMahan. Online convex optimization in the bandit setting: Gradient descent without a gradient. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '05*, pages 385–394, Philadelphia, PA, USA, 2005.
- [12] James E Gentle. *Computational statistics*, volume 308. Springer, 2009.
- [13] Saeed Ghadimi. Conditional gradient type methods for composite nonlinear and stochastic optimization. *Mathematical Programming*, 173(1):431–464, Jan 2019.

- [14] Saeed. Ghadimi and Guanghui. Lan. Stochastic first- and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- [15] Saeed Ghadimi, Guanghui Lan, and Hongchao Zhang. Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Mathematical Programming*, 155: 267–305, 2013.
- [16] Elad Hazan and Kfir Y. Levy. Bandit convex optimization: Towards tight bounds. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’14, pages 784–792, 2014.
- [17] Zhiyi Huang, Jinyan Liu, and Xiangning Wang. Learning optimal reserve price against non-myopic bidders. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS18, 2018.
- [18] IAB. OpenRTB API Specification, Version 2.5. www.iab.com/wp-content/uploads/2016/03/OpenRTB-API-Specification-Version-2-5-FINAL.pdf, December 2016. Accessed: 2020-01-26.
- [19] Vijay Krishna. *Auction theory*. Academic press, 2009.
- [20] S. Liu, X. Li, P. Chen, J. Haupt, and L. Amini. Zeroth-order stochastic projected gradient descent for nonconvex optimization. In *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 1179–1183, Nov 2018.
- [21] Jieming Mao, Renato Leme, and Jon Schneider. Contextual pricing for lipschitz buyers. In *Advances in Neural Information Processing Systems 31*, pages 5643–5651, 2018.
- [22] Steven A. Matthews. A Technical Primer on Auction Theory I: Independent Private Values. Discussion Papers 1096, Northwestern University, Center for Mathematical Studies in Economics and Management Science, May 1995.
- [23] Mehryar Mohri and Andrés Muñoz Medina. Learning algorithms for second-price auctions with reserve. *J. Mach. Learn. Res.*, 17(1):2632–2656, January 2016.
- [24] Andres Munoz and Sergei Vassilvitskii. Revenue optimization with approximate bid predictions. In *Advances in Neural Information Processing Systems 30*, pages 1858–1866, 2017.
- [25] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [26] Michael Ostrovsky and Michael Schwarz. Reserve prices in internet advertising auctions: A field experiment. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, EC ’11, pages 59–60, 2011.
- [27] Renato Paes Leme, Martin Pal, and Sergei Vassilvitskii. A field guide to personalized reserve prices. In *Proceedings of the 25th international conference on world wide web*, pages 1093–1102, 2016.
- [28] Sashank J. Reddi, Suvrit Sra, Barnabás Póczos, and Alexander J. Smola. Proximal stochastic methods for nonsmooth nonconvex finite-sum optimization. In *NIPS*, 2016.
- [29] John G Riley, William F Samuelson, et al. Optimal auctions. *American Economic Review*, 71(3):381–392, 1981.

- [30] Mark Weiss. Digiday research: Header bidding and first-price auctions boost publisher revenues. <https://digiday.com/media/digiday-research-header-bidding-and-first-price-auctions-boost-publisher-revenues>, January 2019. Accessed: 2020-01-29.

Reserve Price Optimization for First Price Auctions

Appendix

A From multiple bidders to a single bidder

Our auction can contain multiple bidders, each with their own value distribution F_i and bid function $b_i(r, v_i)$. But when setting reserve prices, we only care about the maximum bid; more specifically, the distribution of maximum bid at each reserve. Thus it is useful to abstract away the set of bidders in the auction as a single “mega-bidder” whose value is the maximum of all the bidders’ values and who always bids the maximum of all the bids.

Theorem A.1. *Let \mathcal{F} be the distribution of $\max(v_1, v_2, \dots, v_m)$ (where each v_i is independently drawn from F_i) and let $\mathcal{B}(r)$ be the distribution of $\max(b_1(r, v_1), b_2(r, v_2), \dots, b_m(r, v_m))$. Then there exists a bid function $b(r, v)$ such that the distribution of $b(r, v)$ when $v \sim \mathcal{F}$ is equal to the distribution $\mathcal{B}(r)$.*

Proof. If $B_r(b)$ is the CDF of $\mathcal{B}(r)$ and $F(v)$ is the cdf of \mathcal{F} , let $b(r, v) = B_r^{-1}(F(v))$. This guarantees that if $v \sim \mathcal{F}$, then $b(r, v) \sim \mathcal{B}(r)$. \square

Note that this reduction also preserves the properties of Definition 2.1. For example, if $b_i(r, v_i) \leq v_i$ for every bidder i , then the induced $b(r, v)$ also satisfies $b(r, v) \leq v$.

B Demand Function Estimation

As in Section 3, it is possible to form a naive unbiased estimate of the demand component via the estimator $\hat{D}(r) = \frac{1}{n} \sum_{i=1}^n \mathbb{I}(X_i \geq r)$. The variance of the resulting unbiased estimator \hat{G}_D is then bounded by (see Theorem C.5), $\text{Var}(\hat{G}_D) \leq \frac{(r^+)^2}{2\delta^2 n}$.

Note that for small r , the variance guarantee here is significantly better than the variance guarantee in Theorem 3.2. Thus, in instances where the optimal reserve is small (and hence we mostly test small r^+), combining this naive estimator with better estimators for \hat{G}_E (like the ones we explore in the next section) can already lead to better convergence rates overall.

To obtain even better estimators, we can leverage the following two facts about the demand function. First, the demand function only depends on the value distribution \mathcal{F} of the bidders, and not their specific bidding behavior. Since we expect value to be relatively stable in comparison to bidding behavior, this means that we can reasonably use data from previous rounds to learn the demand function and inform calculation of \hat{G}_D (whereas the naive gradient update only uses data from the current round). Second, we expect the demand function $D(r)$ to be simpler and more nicely structured than the full revenue function $\mu(r)$ —for example, $D(r)$ is weakly decreasing in r —and therefore more amenable to parametric modeling.

C Omitted Proofs

C.1 Formal convergence rate

To show a convergence result for a non-convex problem with constraints, a measure called *gradient mapping* is widely used in the literature e.g. [15, 28, 20]. We define the gradient mapping used in this paper as follows,

Definition C.1 (Gradient Mapping). Let $f(x)$ be a differentiable function defined on $[0, 1]$, $\mathcal{C} \in (0, 1)$ is a convex space, and $\Pi_{\mathcal{C}}$ is the projection operator defined as

$$\Pi_{\mathcal{C}}(x) = \arg \min_{z \in \mathcal{C}} |z - x|, \forall x \in \mathbb{R}. \quad (7)$$

The gradient mapping $\mathcal{P}_{\mathcal{C}}$ is then defined as

$$\mathcal{P}_{\mathcal{C}}(r, \hat{g}, \alpha) := \frac{1}{\alpha} (\Pi_{\mathcal{C}}(r + \alpha \hat{g}) - r)$$

where \hat{g} is a gradient estimate of $\mu(r)$ (can be biased), $r \in \mathcal{C}$ is the reserve and α is the learning rate.

The gradient mapping $\mathcal{P}_{\mathcal{C}}$ can be interpreted as the *projected gradient*, which offers a feasible update from the previous reserve r . Indeed, if the projection operator is the identity function, the gradient mapping just returns the gradient.

Theorem C.2. Suppose $\mu(r)$ is L -smooth. Let $\hat{G}_t = \frac{1}{n_t} \sum_{i=1}^{n_t} \hat{G}_{i,t}$ be the gradient estimator at time t , where $\forall i \in [n_t], |\hat{G}_{i,t}| \leq C$ almost surely. Fix $\alpha_t = \Theta(\frac{1}{\sqrt{T}})$, $\beta_t = \delta/2r_t$ and $n_t = N/T$. Assume for all t that $\text{Bias}(\hat{G}_t) \leq B$ and $\text{Var}(\hat{G}_t) \leq V$. Then with probability at least $1 - \delta$, Algorithm 1 satisfies that

$$\begin{aligned} & \min_{t \in [T]} |\mathcal{P}_{\mathcal{C}}(r_t, \mu'(r_t), \alpha_t)|^2 \\ & \leq \mathcal{O} \left(\sqrt{\frac{1}{T}} + \left(\delta + B + \frac{CT \ln(2T/\delta)}{N} + \sqrt{V \ln(2T/\delta)} \right)^2 \right). \end{aligned}$$

where $\mathcal{P}_{\mathcal{C}}$ is the gradient mapping.

To prove Theorem C.2, we first show some useful inequalities summarized in Lemma C.3 and Lemma C.4.

Lemma C.3 (Bernstein Inequality). Let $G = \frac{1}{n} \sum_{i=1}^n G_i$ be the random variable of estimation of revenue's gradient (G_i can be correlated), where $|G_i| \leq C$ almost surely, and $V = \text{Var}[G]$. Then, we have

$$\left| \frac{1}{n} \sum_{i=1}^n G_i - \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n G_i \right] \right| \leq \frac{2C \ln(2T/\delta)}{3n} + \sqrt{2V \ln(2T/\delta)}$$

holds with probability at least $1 - \frac{\delta}{T}$.

Proof. By Bernstein's inequality, we have

$$\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n G_i - \frac{1}{n} \sum_{i=1}^n \mathbb{E}[G_i] \right| \leq z \right) \geq 1 - 2 \exp \left(\frac{-z^2 n}{2(V + Cz/3n)} \right)$$

Let $2 \exp \left(\frac{-z^2}{2(V + Cz/3n)} \right) = \frac{\delta}{T}$ and solving for z , we have

$$z = \frac{\frac{2C \ln(2T/\delta)}{3n} + \sqrt{\left(\frac{2C \ln(2T/\delta)}{3n} \right)^2 + 8V \ln(2T/\delta)}}{2} \leq \frac{2C \ln(2T/\delta)}{3n} + \sqrt{2V \ln(2T/\delta)} := z'$$

where the last inequality is based on the fact $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$. Therefore, we have

$$\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n G_i - \frac{1}{n} \sum_{i=1}^n \mathbb{E}[G_i] \right| \leq z' \right) \geq \mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n G_i - \frac{1}{nt} \sum_{i=1}^n \mathbb{E}[G_i] \right| \leq z \right) \geq 1 - \frac{\delta}{T}.$$

□

Lemma C.4. For any $t \in [T]$, we have

$$\begin{aligned} \hat{G}_t \cdot \mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t) &\geq |\mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t)|^2 \\ |\mathcal{P}_{\mathcal{C}}(r_t, \mu'(r_t), \alpha_t) - \mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t)| &\leq |\mu'(r_t) - \hat{G}_t| \end{aligned}$$

Proof. Since \mathcal{C} is a convex space, then for any $x \in \mathcal{C}, z \in \mathbb{R}$, we have $(x - \Pi_{\mathcal{C}}(z)) \cdot (z - \Pi_{\mathcal{C}}(z)) \leq 0$. Let $z = r_t + \alpha_t \hat{G}_t$ and $x = r_t$, we have

$$(r_t - \Pi_{\mathcal{C}}(r_t + \alpha_t \hat{G}_t)) \cdot (r_t + \alpha_t \hat{G}_t - \Pi_{\mathcal{C}}(r_t + \alpha_t \hat{G}_t)) \leq 0,$$

which implies

$$\alpha_t \hat{G}_t (\Pi_{\mathcal{C}}(r_t + \alpha_t \hat{G}_t) - r_t) \geq |r_t - \Pi_{\mathcal{C}}(r_t + \alpha_t \hat{G}_t)|^2,$$

Thus, we have $\hat{G}_t \cdot \mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t) \geq |\mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t)|^2$. Again, since \mathcal{C} is a convex space, $\forall x, z \in \mathbb{R}, |\Pi_{\mathcal{C}}(x) - \Pi_{\mathcal{C}}(z)| \leq |x - z|$. Then we can prove the second inequality,

$$|\mathcal{P}_{\mathcal{C}}(r_t, \mu'(r_t), \alpha_t) - \mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t)| \leq \frac{1}{\alpha_t} |\Pi_{\mathcal{C}}(r_t + \alpha_t \mu'(r_t)) - \Pi_{\mathcal{C}}(r_t + \alpha_t \hat{G}_t)| \leq |\mu'(r_t) - \hat{G}_t|$$

□

Proof of Theorem C.2. Denote $r_t^+ = (1 + \beta_t)r_t$ and $r_t^- = (1 - \beta_t)r_t$. First we bound the bias of \hat{G}_t compared with $\mu'(r_t)$. $|\mu'(r_t) - \hat{G}_t|$ can be decomposed as follows,

$$\left| \mu'(r_t) - \hat{G}_t \right| \leq \left| \mu'(r_t) - \frac{\mu(r_t^+) - \mu(r_t^-)}{r_t^+ - r_t^-} \right| + \left| \mathbb{E}[\hat{G}_t] - \frac{\mu(r_t^+) - \mu(r_t^-)}{r_t^+ - r_t^-} \right| + \left| \mathbb{E}[\hat{G}_t] - \hat{G}_t \right| \quad (8)$$

Then we bound the three terms above separately. For the first term, we have

$$\mu'(r_t) - \frac{\mu(r_t^+) - \mu(r_t^-)}{r_t^+ - r_t^-} = \frac{\mu'(r_t)\beta_t r_t - \mu(r_t^+) + \mu'(r_t)\beta_t r_t + \mu(r_t^-)}{r_t^+ - r_t^-}$$

By the smoothness of $\mu(r)$,

$$\begin{aligned} |\mu(r_t^+) - \mu(r_t) - \mu'(r_t) \cdot \beta_t r_t| &\leq \frac{L}{2} \cdot \beta_t^2 r_t^2 \\ |\mu(r_t^-) - \mu(r_t) + \mu'(r_t) \cdot \beta_t r_t| &\leq \frac{L}{2} \cdot \beta_t^2 r_t^2, \end{aligned}$$

Thus, we get

$$\left| \mu'(r_t) - \frac{\mu(r_t^+) - \mu(r_t^-)}{r_t^+ - r_t^-} \right| \leq \frac{1}{2\beta_t r_t} \cdot L\beta_t^2 r_t^2 = \frac{L\beta_t r_t}{2}$$

By assumption, the second term is bounded by B . Combining Lemma C.3 in Appendix, for any fixed $t \in [T]$,

$$\left| \mu'(r_t) - \hat{G}_t \right| \leq \frac{L\beta_t r_t}{2} + B + \frac{2C \ln(2T/\delta)}{3n_t} + \sqrt{V \ln(2T/\delta)} = B_{\delta,t}, \quad (9)$$

holds with probability at least $1 - \frac{\delta}{T}$.

The L -smoothness of revenue function $\mu(r)$ implies the following inequalities, for any $t = 1, \dots, T$,

$$\begin{aligned} & \mu(r_{t+1}) \\ \geq & \mu(r_t) + \mu'(r_t) \cdot (r_{t+1} - r_t) - \frac{L}{2}(r_{t+1} - r_t)^2 \\ \geq & \mu(r_t) + \alpha_t \mu'(r_t) \cdot \mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t) - \frac{L\alpha_t^2}{2} |\hat{G}_t|^2 \\ = & \mu(r_t) + \alpha_t \hat{G}_t \cdot \mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t) + \alpha_t (\mu'(r_t) - \hat{G}_t) \cdot \mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t) - \frac{L\alpha_t^2}{2} |\hat{G}_t|^2 \\ \geq & \mu(r_t) + \alpha_t |\mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t)|^2 + \alpha_t (\mu'(r_t) - \hat{G}_t) \cdot \mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t) - \frac{L\alpha_t^2}{2} |\hat{G}_t|^2 \end{aligned}$$

The first inequality is because $|r_{t+1} - r_t| \leq \alpha_t \hat{G}_t$ and the last inequality is based on first inequality in Lemma C.4. Rearranging the above inequalities, we have,

$$\begin{aligned} \alpha_t |\mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t)|^2 & \leq \mu(r_{t+1}) - \mu(r_t) + \alpha_t (\hat{G}_t - \mu'(r_t)) \cdot \mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t) + \frac{L\alpha_t^2}{2} |\hat{G}_t|^2 \\ & \leq \mu(r_{t+1}) - \mu(r_t) + \frac{\alpha_t |\hat{G}_t - \mu'(r_t)|^2}{2} + \frac{\alpha_t |\mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t)|^2}{2} + \frac{L\alpha_t^2}{2} |\hat{G}_t|^2 \end{aligned} \quad (10)$$

The last inequality is based on Cauchy-Schwartz inequality and the second statement in Lemma C.4. Rearranging Equation (10), we have

$$\alpha_t |\mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t)|^2 \leq 2(\mu(r_{t+1}) - \mu(r_t)) + \alpha_t |\hat{G}_t - \mu'(r_t)|^2 + L\alpha_t^2 |\hat{G}_t|^2 \quad (11)$$

Using Cauchy-Schwartz inequality and Lemma C.4, we can bound $|\mathcal{P}_{\mathcal{C}}(r_t, \mu'(r_t), \alpha_t)|$ in the following way,

$$\begin{aligned} |\mathcal{P}_{\mathcal{C}}(r_t, \mu'(r_t), \alpha_t)|^2 & \leq 2|\mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t)|^2 + 2|\mathcal{P}_{\mathcal{C}}(r_t, \mu'(r_t), \alpha_t) - \mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t)|^2 \\ & \leq 2|\mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t)|^2 + 2|\mu'(r_t) - \hat{G}_t|^2 \end{aligned}$$

Then we can write down the following bound,

$$\begin{aligned} \alpha_t |\mathcal{P}_{\mathcal{C}}(r_t, \mu'(r_t), \alpha_t)|^2 & \leq 2\alpha_t |\mathcal{P}_{\mathcal{C}}(r_t, \hat{G}_t, \alpha_t)|^2 + 2\alpha_t |\mu'(r_t) - \hat{G}_t|^2 \\ & \leq 4(\mu(r_{t+1}) - \mu(r_t)) + 4\alpha_t |\hat{G}_t - \mu'(r_t)|^2 + 2L\alpha_t^2 |\hat{G}_t|^2 \\ & \leq 4(\mu(r_{t+1}) - \mu(r_t)) + 4\alpha_t |\hat{G}_t - \mu'(r_t)|^2 + 4L\alpha_t^2 (|\hat{G}_t - \mu'(r_t)|^2 + |\mu'(r_t)|^2) \end{aligned} \quad (12)$$

Let $r^* \in \mathcal{C}$ be the point with the minimum absolute gradient, i.e., $r^* = \arg \min_{r \in \mathcal{C}} |\mu'(r)|$. Notice $|\mu'(r_t)|^2 \leq 2|\mu'(r^*)|^2 + 2|\mu'(r^*) - \mu'(r_t)|^2 \leq 2L + 2|\mu'(r^*)|^2 := L^*$. Summing up the inequality (12) from $t = 1$ to T , based on inequality (9) with probability at least $1 - \delta$, we have

$$\sum_{t=1}^T \alpha_t |\mathcal{P}_{\mathcal{C}}(r_t, \mu'(r_t), \alpha_t)|^2 \leq 4(\mu(r_{T+1}) - \mu(r_1)) + 4 \sum_{t=1}^T \alpha_t |B_{\delta,t}|^2 + 4L \sum_{t=1}^T \alpha_t^2 (|B_{\delta,t}|^2 + L^*)$$

Setting $\alpha_t = \Theta\left(\frac{1}{\sqrt{T}}\right)$, $\beta_t = \delta/2r_t$ and $n_t = N/T$, By Cauchy-Schwartz inequality and the fact that $\sum_{t=1}^T \alpha_t = \Theta(\sqrt{T})$ and $\sum_{t=1}^T \alpha_t^2 = \Theta(1)$. Therefore, we get

$$\begin{aligned} \min_{t=1, \dots, T} |\mathcal{P}_{\mathcal{C}}(r_t, \mu'(r_t), \alpha_t)|^2 &\leq \frac{1}{\sum_{t=1}^T \alpha_t} \sum_{t=1}^T \alpha_t |\mathcal{P}_{\mathcal{C}}(r_t, \mu'(r_t), \alpha_t)|^2 \\ &\leq \mathcal{O}\left(\sqrt{\frac{1}{T}} + \left(\delta + B + \frac{CT \ln(2T/\delta)}{N} + \sqrt{V \ln(2T/\delta)}\right)^2\right) \end{aligned}$$

□

C.2 Proof of Theorem 3.2

Proof. Since X_i^+ and X_i^- are independent random samples from \mathcal{B}_{r^+} and \mathcal{B}_{r^-} respectively, $\mathbb{E}[\hat{G}] = \frac{\mu(r^+) - \mu(r^-)}{r^+ - r^-}$. For the variance, since X_i^+ is bounded by $[0, 1]$, then the variance of each X_i^+ and X_i^- is at most $1/4$, which implies,

$$\text{Var}(\hat{G}) \leq \frac{2}{4\delta^2 n} = \frac{1}{2\delta^2 n}$$

□

C.3 Variance bounds for \hat{G}_D

In this subsection we bound the variance of the unbiased demand estimator

$$\hat{G}_D = \frac{r^+ \sum_{i=1}^n \mathbb{I}(X_i^+ \geq r^+) - r^- \sum_{i=1}^n \mathbb{I}(X_i^- \geq r^-)}{n(r^+ - r^-)}.$$

Since $\mathbb{E}\left[\frac{1}{n} \sum_{i=1}^n \mathbb{I}(X_i^+ \geq r^+)\right] = D(r^+)$ and $\mathbb{E}\left[\frac{1}{n} \sum_{i=1}^n \mathbb{I}(X_i^- \geq r^-)\right] = D(r^-)$, it follows immediately that

$$\mathbb{E}[\hat{G}_D] = \frac{r^+ D(r^+) - r^- D(r^-)}{r^+ - r^-}$$

and therefore that $\text{Bias}(\hat{G}_D) = 0$. In the following theorem, we bound $\text{Var}(\hat{G}_D)$.

Theorem C.5. *Let $\delta = r^+ - r^-$. We have that*

$$\text{Var}(\hat{G}_D) \leq \frac{(r^+)^2}{2n\delta^2}.$$

Proof. Since the X_i^+ and X_i^- are independent random variables, their variances are additive so

$$\begin{aligned} \text{Var}(\hat{G}_D) &= \frac{1}{n^2 \delta^2} (n(r^+)^2 \text{Var}(\mathbb{I}(X_i^+ \geq r^+)) + n(r^-)^2 \text{Var}(\mathbb{I}(X_i^- \geq r^-))) \\ &\leq \frac{(r^+)^2}{2n\delta^2}, \end{aligned}$$

where in the last line we have used the fact that the variance of a Bernoulli random variable is bounded above by $1/4$. □

C.4 Proof of Theorem 4.1

Proof. Note that

$$\begin{aligned}
\mu(r) &= \mathbb{E}_{v \sim \mathcal{F}}[b(r, v)] \\
&= \mathbb{E}_{v \sim \mathcal{F}}[\max(b(r, v), r) - r\mathbb{I}\{b(r, v) = 0\}] \\
&= \mathbb{E}_{v \sim \mathcal{F}}[\max(b(r, v), r)] - r\mathbb{E}_{v \sim \mathcal{F}}[\mathbb{I}\{b(r, v) = 0\}] \\
&= \mathbb{E}_{v \sim \mathcal{F}}[\max(b(r, v), r)] - r \Pr_{v \sim \mathcal{F}}[v < r] \\
&= \mathbb{E}_{v \sim \mathcal{F}}[\max(b(r, v), r)] - r(1 - \Pr_{v \sim \mathcal{F}}[v \geq r]) \\
&= \mathbb{E}_{v \sim \mathcal{F}}[\max(b(r, v), r) - r] + r \Pr_{v \sim \mathcal{F}}[v \geq r]
\end{aligned}$$

□

C.5 Proof of Theorem 4.2

In the following proof we slightly abuse notation, by denoting \mathcal{B}_r to be the CDF of the bid distribution w.r.t reserve price r . We also formally state the definition of ε -bounded response model here.

Definition C.6 (ε -bounded response model). *A ε -bounded response bidding function takes the form,*

$$b(r, v) = \begin{cases} b(0, v) & \text{if } b(0, v) \geq r \\ r + z & \text{if } b(0, v) \leq r \leq v \\ 0 & \text{if } v < r \end{cases}$$

where $z \in [0, \varepsilon]$, and z can be a random variable.

Proof. To bound the variance, note that each Y_i^- is constrained to the interval $[0, r^+ - r^-]$. Since this interval has length at most δ , the variance of each Y_i^- is at most $(\delta^2/4)$, then

$$\text{Var}(\hat{G}_E) \leq \frac{n(\delta^2/4)}{n^2\delta^2} \leq \frac{1}{4n}.$$

Then we focus on bounding the bias,

$$\begin{aligned}
&\mathbb{E}_v[\max(b(r^+, v) - r^+, 0) - \max(b(r^-, v) - r^-, 0)] \\
&= \int_0^1 \max(b, r^+) d\mathcal{B}_{r^+} - \int_0^1 \max(b, r^-) d\mathcal{B}_{r^-} - (r^+ - r^-) \\
&= \int_0^{r^+ + \varepsilon} \max(b, r^+) d\mathcal{B}_{r^+} - \int_0^{r^+ + \varepsilon} \max(b, r^-) d\mathcal{B}_{r^-} + \int_{r^+ + \varepsilon}^1 \max(b, r^+) d\mathcal{B}_{r^+} - \int_{r^+ + \varepsilon}^1 \max(b, r^-) d\mathcal{B}_{r^-} \\
&\quad - (r^+ - r^-) \\
&= \int_0^{r^+} \max(b, r^+) d\mathcal{B}_{r^+} - \int_0^{r^+} \max(b, r^-) d\mathcal{B}_{r^-} + \int_{r^+}^{r^+ + \varepsilon} \max(b, r^+) d\mathcal{B}_{r^+} - \int_{r^+}^{r^+ + \varepsilon} \max(b, r^-) d\mathcal{B}_{r^-} \\
&\quad - (r^+ - r^-) \\
&= r^+ \mathcal{B}_{r^+}(r^+) - \int_0^{r^+} \max(b, r^-) d\mathcal{B}_{r^-} - (r^+ - r^-) + \int_{r^+}^{r^+ + \varepsilon} b d\mathcal{B}_{r^+} - \int_{r^+}^{r^+ + \varepsilon} b d\mathcal{B}_{r^-}
\end{aligned}$$

Here the third equality holds because the fact that if $b > r^+ + \varepsilon$, $\mathcal{B}_{r^+}(b) = \mathcal{B}_{r^-}(b)$ by property of the ε -bounded response. The fourth equality is based on $\max(b(r^+, v), r^+) = r^+$ when $v \leq r^+$ and $\mathcal{F}(r^+) = \mathcal{B}_{r^+}(r^+)$. Then we consider $\mathbb{E}[Y_i^-]$, where

$$\begin{aligned}\mathbb{E}[Y_i^-] &= \mathbb{E}[\max(X_i^- - r^-, 0)\mathbb{I}\{X_i^- \leq r^+\}] + \mathbb{E}[(r^+ - r^-)\mathbb{I}\{X_i^- > r^+\}] \\ &= \int_0^{r^+} \max(b, r^-)d\mathcal{B}_{r^-} + r^+(1 - \mathcal{B}_{r^-}(r^+)) - r^-\end{aligned}$$

Before bounding the bias of \hat{G}_E , we state some useful equations based on integral by part.

$$\int_{r^+}^{r^++\varepsilon} bd\mathcal{B}_{r^+} = (r^+ + \varepsilon)\mathcal{B}_{r^+}(r^+ + \varepsilon) - r^+\mathcal{B}_{r^+}(r^+) - \int_{r^+}^{r^++\varepsilon} \mathcal{B}_{r^+}(b)db \quad (13)$$

$$\int_{r^+}^{r^++\varepsilon} bd\mathcal{B}_{r^-} = (r^+ + \varepsilon)\mathcal{B}_{r^-}(r^+ + \varepsilon) - r^+\mathcal{B}_{r^-}(r^+) - \int_{r^+}^{r^++\varepsilon} \mathcal{B}_{r^-}(b)db \quad (14)$$

Based on definition of ε -bounded response, $\mathcal{B}_{r^-}(r^+ + \varepsilon) = \mathcal{B}_{r^+}(r^+ + \varepsilon)$. Then we have

$$\begin{aligned}& \left| \mathbb{E}[\hat{G}_E] - \frac{E(r^+) - E(r^-)}{r^+ - r^-} \right| \\ &= \left| -\frac{\mathbb{E}[Y_i^-]}{r^+ - r^-} - \frac{E(r^+) - E(r^-)}{r^+ - r^-} \right| \\ &= \frac{1}{r^+ - r^-} \left| r^+\mathcal{B}_{r^+}(r^+) - r^+\mathcal{B}_{r^-}(r^+) + \int_{r^+}^{r^++\varepsilon} bd\mathcal{B}_{r^+} - \int_{r^+}^{r^++\varepsilon} bd\mathcal{B}_{r^-} \right| \\ &= \frac{1}{r^+ - r^-} \left| \int_{r^+}^{r^++\varepsilon} \mathcal{B}_{r^+}(b)db - \int_{r^+}^{r^++\varepsilon} \mathcal{B}_{r^-}(b)db \right| \\ & \quad (\text{Based on Equations (13) and (14) as well as } \mathcal{B}_{r^-}(r^+ + \varepsilon) = \mathcal{B}_{r^+}(r^+ + \varepsilon)) \\ &\leq \frac{2\varepsilon}{r^+ - r^-}\end{aligned}$$

where the inequality is because $\forall b \in [r^+, r^+ + \varepsilon], \mathcal{B}_{r^+}(b), \mathcal{B}_{r^-}(b) \leq 1$. □

C.6 Proof of Theorem 4.3

We start with the following helpful auxiliary lemmas.

Lemma C.7. *Let Y_1, Y_2, \dots, Y_n be n iid uniform random variables. Let $Y^{(k)}$ be the k th largest Y_i . Then with probability at least $1 - n^{-2/3}$,*

$$\left| Y^{(k)} - \frac{k}{n+1} \right| \leq n^{-2/3}.$$

Proof. From the theory of order statistics [12], we know that $Y^{(k)} \sim \text{Beta}(k, n+1-k)$. It is known that $\mathbb{E}[Y^{(k)}] = \frac{k}{n+1}$ and that $\text{Var}(Y^{(k)}) \leq 1/(8n)$. The statement immediately follows from Chebyshev's inequality. □

Lemma C.8. Let $f : [0, 1] \rightarrow [0, 1]$ be an increasing function and let Y_1, Y_2, \dots, Y_n be n iid uniform random variables. Let $X_i = f(Y_i)$, and let S_k be the r.v. equal to the sum of the k smallest X_i . Then

$$\left| \frac{1}{n} \mathbb{E}[S_{qn}] - \int_0^q f(x) dx \right| \leq 3n^{-2/3}$$

Proof. Let Z_1, Z_2, \dots, Z_{qn} be (a random permutation) of the qn smallest Y_i (so $S_k = \sum f(Z_i)$). Note that conditioned on $Z_{qn+1} = r$, the Z_i are independently distributed according to $U([0, r])$. In particular, we have that

$$\frac{1}{n} \mathbb{E}[S_{qn} | Z_{qn+1} = r] = \int_0^r f(x) dx.$$

From Lemma C.7, we know that with probability at least $1 - n^{-2/3}$, $Z_{qn+1} \in [q - 2n^{-2/3}, q + 2n^{-2/3}]$. Since $f(x) \in [0, 1]$, it follows that $\int_0^r f(x) dx$ is 1-Lipshitz and therefore (conditioned on $Z_{qn+1} \in [q - 2n^{-2/3}, q + 2n^{-2/3}]$),

$$\left| \int_0^q f(x) dx - \int_0^{Z_{qn+1}} f(x) dx \right| \leq 2n^{-2/3}.$$

On the other hand, in the ($n^{-2/3}$ probability) case where $Z_{qn+1} \notin [q - 2n^{-2/3}, q + 2n^{-2/3}]$, $\frac{1}{n} \mathbb{E}[S_{qn}]$ is still bounded in $[0, 1]$. The theorem statement immediately follows. \square

Lemma C.9. Let X_i be an iid collection of n rvs. Let $X^{(k)}$ be the k th smallest of the X_i (so $X^{(1)} \leq X^{(2)} \leq \dots \leq X^{(n)}$). Then, if $S_k = \sum_{i=1}^k X^{(i)}$, we have that

$$\text{Var}(S_k) \leq n \mathbb{E}[(X^{(k)})^2].$$

Proof. The Efron-Stein inequalities (see Theorem 2 of [8]) state that for any collection of n random variables X_i and any measurable functions $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and $f_i : \mathbb{R}^{n-1} \rightarrow \mathbb{R}$ we have that

$$\text{Var}(f(X)) \leq \sum_{i=1}^n \mathbb{E}[(f(X) - f_i(X_{-i}))^2],$$

where X_{-i} is the $(n-1)$ -tuple of rvs $(X_1, X_2, \dots, X_{i-1}, X_{i+1}, \dots, X_n)$.

Let $f(X)$ equal the sum of the k smallest entries in X , and let $f_i(X')$ equal the sum of the $k-1$ smallest entries in X' . Note that for this choice of f and f_i , $f(X) = S_k$, and $0 \leq f(X) - f_i(X_{-i}) \leq X^{(k)}$ (since the $k-1$ smallest entries in X_{-i} are a subset of the k smallest entries in X). It follows that $\text{Var}(S_k) \leq n \mathbb{E}[(X^{(k)})^2]$, as desired. \square

We can now proceed to prove Theorem 4.3.

Proof of Theorem 4.3. We begin by bounding the variance of our estimator \hat{G}_E . Let us begin by focusing on $\text{Var}(\sum_{i=1}^{qn} Y_i^-)$. Since these Y_i are sorted, Lemma C.9 implies that $\text{Var}(\sum_{i=1}^{qn} Y_i^-) \leq n \mathbb{E}[(Y_{qn}^-)^2]$. Since $\tilde{t} = \mathcal{F}^{-1}(q + n^{-2/3})$, by Lemma C.7, with probability at least $1 - n^{-2/3}$, $Y_{qn}^- \leq \tilde{t}$, and therefore $\text{Var}(\sum_{i=1}^{qn} Y_i^-) \leq n \tilde{t}^2 + n^{1/3}$. Similarly, $\text{Var}(\sum_{i=1}^{qn} Y_i^+) \leq n \tilde{t}^2 + n^{1/3}$. Since the sets of rvs Y_i^- and Y_i^+ are independent, we have that

$$\text{Var}(\hat{G}_E) \leq \frac{2\tilde{t}^2}{n\delta^2} + O(n^{-5/3}\delta^{-2}).$$

We now proceed to bound the bias of \hat{G}_E . First, note that by Lemma C.8, we have that

$$\left| \frac{1}{n} \mathbb{E} \left[\sum_{i=1}^{qn} Y_{qn}^- \right] - \int_0^q \max(b(r^-, \mathcal{F}^{-1}(x)) - r^-, 0) dx \right| \leq 3n^{-2/3},$$

and therefore

$$\left| \frac{1}{n} \mathbb{E} \left[\sum_{i=1}^{qn} Y_{qn}^- \right] - \int_0^t \max(b(r^-, v) - r^-, 0) d\mathcal{F}(v) \right| \leq 3n^{-2/3}.$$

Likewise

$$\left| \frac{1}{n} \mathbb{E} \left[\sum_{i=1}^{qn} Y_{qn}^+ \right] - \int_0^t \max(b(r^+, v) - r^+, 0) d\mathcal{F}(v) \right| \leq 3n^{-2/3}.$$

It follows that

$$\left| \hat{G}_E - \left(\frac{1}{\delta} \int_0^t (\max(b(r^+, v) - r^+, 0) - \max(b(r^-, v) - r^-, 0)) d\mathcal{F}(v) - (1 - q) \right) \right| \leq 6n^{-2/3}. \quad (15)$$

On the other hand, note that

$$\begin{aligned} E(r^+) - E(r^-) &= \int_0^1 (\max(b(r^+, v) - r^+, 0) - \max(b(r^-, v) - r^-, 0)) d\mathcal{F}(v) \\ &= \int_0^t (\max(b(r^+, v) - r^+, 0) - \max(b(r^-, v) - r^-, 0)) d\mathcal{F}(v) \\ &\quad + \int_t^1 (\max(b(r^+, v) - r^+, 0) - \max(b(r^-, v) - r^-, 0)) d\mathcal{F}(v). \end{aligned}$$

Now, note that for $v \geq r^+$, $\max(b(r^+, v) - r^+, 0) - \max(b(r^-, v) - r^-, 0) = b(r^+, v) - b(r^-, v) - (r^+ - r^-)$. Since $\Pr[v \geq t] = 1 - q$, it follows that

$$\begin{aligned} &\int_t^1 (\max(b(r^+, v) - r^+, 0) - \max(b(r^-, v) - r^-, 0)) d\mathcal{F}(v) \\ &= \int_t^1 (b(r^+, v) - b(r^-, v) - (r^+ - r^-)) d\mathcal{F}(v) \\ &= \int_t^1 (b(r^+, v) - b(r^-, v)) d\mathcal{F}(v) - (1 - q)(r^+ - r^-) \\ &\in [-(1 - q)\delta, -(1 - q)\delta + (b(r^+, t) - b(r^-, t))(1 - q)]. \end{aligned}$$

Here the last line follows since $b(r^+, v) - b(r^-, v)$ is decreasing in v (due to diminishing sensitivity to reserve) but always non-negative. Combining this with equation 15, we have that:

$$\left| \hat{G}_E - \frac{1}{\delta} (E(r^+) - E(r^-)) \right| \leq \frac{(b(r^+, t) - b(r^-, t))(1 - q)}{\delta} + 6n^{-2/3},$$

as desired. □

C.7 Convergence Rate of Quantile Truncation

Corollary C.1. *Suppose $\text{Bias}(\hat{G}_D) \leq \varepsilon_D/\delta$. Using the estimator \hat{G}_E proposed in Eq. (5) for the response model with diminishing sensitivity property, for any fixed quantile q , setting $T = N^{2/3}$ and $\delta = \Theta(\sqrt{\varepsilon_D + 1 - q})$, Algorithm 1 achieves convergence,*

$$\min_{t \in [T]} |\mathcal{P}_C^t|^2 \leq \tilde{O} \left(\varepsilon_D + 1 - q + \left(1 + \frac{\mathcal{F}^{-1}(q + N^{-2/9})}{\varepsilon_D + 1 - q} \right) \cdot N^{-1/3} \right)$$

D Additional Experiments

In this section, we show the parameters used in the algorithms and some additional experiments. For additional experiments, we compare the two truncation methods for perfect response models, and then we show the complete results of 20 semi-synthetic data sets with perfect response. Finally, we test for other different response models in synthetic data and one of the semi-synthetic data sets (the first data set). In addition to the figures about the revenue curve learned by our algorithms, we also report the average revenue of the first several rounds learned by our algorithms in Tables (see Table 1, 2, and 3).

Set up. For all the algorithms, we set the learning rate to 0.05, the minimum reserve price to 0.1, the maximum reserve price to 5.0, and the perturbation size to $\beta_t = 0.1$ at each round. For quantile truncation, we use the 80% quantile as the threshold to discard bids. We run the algorithms for 200 rounds with access to 100 samples at each round (50 for r^- and 50 for r^+), which forms 1 trial. We repeat 50 trials for each algorithm and report the mean (solid line) and 95% confidence interval (translucent and colored area) of the revenue achieved during training for each algorithms shown in the figures. To obtain revenue curves learned by the algorithms over time, the revenue at each point is estimated through 10,000 bids randomly drawn from the bid distribution and bidder response model.

Comparison of bid truncation and quantile truncation for perfect response models.

We show the average revenue achieved by algorithms with bid truncation and quantile truncation for synthetic data with perfect response in Table 1 (first 50th rounds) and semi-synthetic data with perfect response in Table 2 (first 20th rounds) and Table 3 (first 50th rounds). We show that the bid truncation (Algorithm (IV)) performs the best in most of cases, but quantile truncation (Algorithm (V)) is also competitive and performs well for perfect response in both synthetic data and semi-synthetic data sets. Bid truncation doesn't significantly outperform than quantile truncation significantly, since the algorithms only take small number of the samples at each round and the quantile truncation is more stable in this setting.

Performance of all 20 semi-synthetic data sets with perfect response. We evaluate the performance of five algorithms for 20 semi-synthetic data sets with perfect response. Here we still repeat 50 trials for each algorithm and report the mean and 95% confidence interval of the revenue. The results are summarized in Table 2 and Table 3, where each table records the average revenue over the first 20 rounds and 50 rounds, respectively. The revenue is normalized by the optimal revenue of each data set (empirically evaluated by grid search). We find in semi-synthetic data, the variance reduction methods (bid truncation and quantile truncation) improve the revenue achieved by the algorithms. Interestingly, we find Naive gradient descent with bid truncation also works well in several semi-synthetic data sets, this is because our demand modeling approach relies on a good estimator \hat{G}_D and sometimes, the simple neural network cannot learn the demand curve very accurately in the beginning.

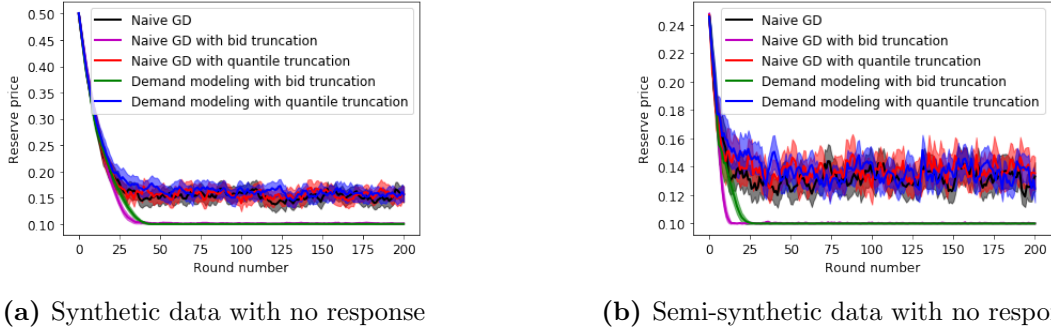


Figure 3: Plots of reserve price as a function of round t for (a) synthetic data with no response, and (b) one semi-synthetic data with no response.


No response model. In no response model, bidders don't change their bids based on reserve prices, i.e. $\forall r > 0, r \geq 0, \mathcal{B}_r(r') = \mathcal{B}_0(r')$. This can be regarded as a perfect response model with linear shading factor 1. In this case, we expect the algorithms converge to the lower bound of the reserve prices 0.1. For no response models, we set perturbation be 0.3 to speed up convergence and we still use logistic regression to learn demand curve. Figure 3 shows that all the algorithms almost converge to minimum reserve price 0.1 and the bid truncation method works the best for no response model.

Mixture of no response and perfect response models. We also consider another non-perfect response model, which is a mixture of perfect response and no response models. We assume the bidder will not respond to reserve price with probability 0.1 and use perfect response with probability 0.9. Figure 4 shows the revenue curve learned by the algorithms for synthetic data and one semi-synthetic data. We find quantile-based variance reduction speed up the training and converges to optimal reserve faster than naive gradient descent. Since this is a mixture of perfect response and no response model, the revenue achieved by bid truncation methods is worse than the quantile-based approach. Through this experiments, we find quantile truncation is not sensitive with different response models, whereas, the bid truncation method is very sensitive to a slightly non-perfect response model.


ε -bounded response. In the experiments for ε -bounded response model, we set $\varepsilon = 0.05$ and the bias term $z \sim \text{Unif}[0, \varepsilon]$ in ε -bounded response model (see definition C.6). We visualize the revenue learned by algorithms for synthetic data and semi-synthetic data in Figure 5. The figures show that for ε -bounded response, the quantile truncation still works better than naive gradient descent. Since we have demonstrated that bid truncation is sensitive to non-perfect response model, the performance of the bid truncation is worse than quantile truncation in ε -bounded response.

Response model	Algorithm (I) <i>rev</i>	Algorithm (II) <i>rev</i>	Algorithm (III) <i>rev</i>	Algorithm (IV) <i>rev</i>	Algorithm (V) <i>rev</i>
Perfect response	86.7± 0.8%	86.5± 0.7%	87.1± 0.9%	88.9± 0.3%	89.4± 0.6%
Equilibrium response	95.5± 0.5%	90.1± 0.2%	96.3± 0.4%	93.5± 0.6%	96.4± 0.4%
ε -bounded response	86.4± 0.8%	83.0± 0.6%	86.6± 1.0%	85.1± 0.8%	87.1± 0.8%
Mixture response	90.5± 0.3%	90.3± 0.3%	90.8± 0.4%	90.8± 0.3%	91.0± 0.4%

Table 1: Average revenue of the first 50th rounds of five algorithms for synthetic data with different response models. The revenue is normalized by the optimal revenue and we repeat 50 trials for each algorithm to report the 95% confidence interval.




(a) Synthetic data with mixture response




(b) Semi-synthetic data with mixture response

Figure 4: Plots of reserve price and revenue as a function of round t for (a) synthetic data and (b) one semi-synthetic data, with mixture response, where the bidder uses perfect response model with probability 0.9 and doesn't respond to the reserve, otherwise.



(a) Synthetic data with ε -bounded response



(b) Semi-synthetic data with ε -bounded response

Figure 5: Plots of revenue as a function of round t for (a) synthetic data with ε -bounded response, and (b) one semi-synthetic data with ε -bounded response, where $\varepsilon = 0.05$ and the bias term $z \sim \text{Unif}[0, \varepsilon]$.

Semi-synthetic data sets	Algorithm (I) <i>rev</i>	Algorithm (II) <i>rev</i>	Algorithm (III) <i>rev</i>	Algorithm (IV) <i>rev</i>	Algorithm (V) <i>rev</i>
(1)	73.9± 0.4%	73.8± 0.2%	74.1± 0.4%	75.3± 0.1%	75.3± 0.3%
(2)	91.8± 0.8%	92.9± 0.3%	92.9± 0.6%	94.5± 0.1%	94.3± 0.6%
(3)	94.1± 0.2%	94.7± 0.2%	93.5± 0.3%	94.4± 0.1%	93.3± 0.3%
(4)	82.9± 0.4%	83.4± 0.0%	82.9± 0.3%	83.5± 0.0%	83.2± 0.2%
(5)	93.0± 0.6%	94.8± 0.0%	93.4± 0.4%	94.4± 0.0%	93.9± 0.4%
(6)	93.8± 0.7%	91.8± 0.5%	94.0± 0.6%	95.6± 0.1%	95.9± 0.4%
(7)	86.4± 0.4%	85.1± 0.4%	87.4± 0.5%	87.9± 0.1%	88.6± 0.5%
(8)	95.0± 0.7%	96.0± 0.2%	95.3± 0.6%	96.7± 0.1%	95.6± 0.3%
(9)	83.6± 1.0%	84.0± 0.5%	84.9± 0.9%	89.1± 0.1%	88.6± 1.1%
(10)	94.3± 0.8%	94.8± 0.3%	94.9± 0.8%	96.4± 0.0%	95.9± 0.5%
(11)	52.0± 1.1%	53.1± 1.0%	52.2± 1.6%	52.1± 1.0%	51.2± 1.6%
(12)	86.6± 0.7%	86.9± 0.3%	87.9± 0.7%	89.8± 0.1%	89.6± 0.5%
(13)	89.7± 1.5%	89.8± 0.8%	89.1± 1.5%	92.1± 0.1%	91.2± 1.0%
(14)	93.3± 0.6%	93.0± 0.4%	93.1± 0.7%	95.7± 0.1%	95.2± 0.7%
(15)	96.2± 0.7%	95.0± 0.5%	96.3± 0.7%	97.9± 0.1%	97.8± 0.3%
(16)	88.4± 0.8%	87.8± 0.4%	88.6± 0.8%	92.0± 0.1%	92.1± 0.7%
(17)	94.6± 0.6%	94.9± 0.1%	93.9± 0.7%	92.5± 0.1%	91.6± 0.6%
(18)	91.0± 0.6%	90.9± 0.3%	90.8± 0.6%	93.0± 0.1%	92.7± 0.4%
(19)	83.4± 0.6%	83.4± 0.4%	84.7± 0.6%	85.9± 0.1%	85.8± 0.6%
(20)	94.7± 0.9%	97.3± 0.1%	95.6± 0.7%	93.9± 0.1%	92.2± 0.4%

Table 2: Average revenue of first 20th rounds of five algorithms for semi-synthetic data sets with perfect response. The revenue is normalized by the optimal revenue of each data set. We repeat 50 trials to get the 95% confidence interval.

Semi-synthetic data sets	Algorithm (I) <i>rev</i>	Algorithm (II) <i>rev</i>	Algorithm (III) <i>rev</i>	Algorithm (IV) <i>rev</i>	Algorithm (V) <i>rev</i>
(1)	75.1± 0.6%	74.5± 0.3%	75.7± 0.6%	76.8± 0.2%	77.5± 0.7%
(2)	93.4± 0.5%	94.1± 0.2%	94.4± 0.4%	94.8± 0.1%	95.4± 0.4%
(3)	93.9± 0.2%	95.0± 0.1%	93.3± 0.2%	94.9± 0.1%	93.4± 0.3%
(4)	83.3± 0.3%	83.7± 0.0%	83.3± 0.3%	83.9± 0.0%	83.7± 0.1%
(5)	95.7± 0.3%	96.8± 0.0%	95.9± 0.2%	96.4± 0.0%	96.3± 0.2%
(6)	95.2± 0.4%	93.7± 0.5%	95.4± 0.4%	96.6± 0.0%	96.6± 0.2%
(7)	88.8± 0.5%	87.1± 0.5%	90.0± 0.6%	90.0± 0.1%	91.1± 0.4%
(8)	97.3± 0.3%	98.1± 0.1%	97.0± 0.4%	98.7± 0.1%	97.4± 0.2%
(9)	91.5± 0.8%	91.9± 0.6%	92.8± 0.8%	95.6± 0.1%	95.3± 0.7%
(10)	96.4± 0.5%	97.0± 0.2%	96.9± 0.4%	97.7± 0.0%	97.4± 0.2%
(11)	74.4± 0.7%	75.1± 0.6%	73.3± 2.3%	75.6± 0.7%	73.5± 2.6%
(12)	89.2± 0.6%	89.3± 0.3%	90.5± 0.6%	91.6± 0.1%	91.9± 0.4%
(13)	95.6± 0.7%	96.0± 0.3%	94.8± 0.9%	96.0± 0.1%	94.8± 0.3%
(14)	94.3± 0.5%	94.2± 0.3%	94.6± 0.5%	96.1± 0.1%	96.1± 0.5%
(15)	97.4± 0.5%	97.0± 0.3%	97.4± 0.3%	99.1± 0.0%	98.3± 0.1%
(16)	91.8± 0.7%	90.4± 0.5%	92.2± 0.8%	93.9± 0.2%	94.7± 0.5%
(17)	96.9± 0.2%	97.3± 0.1%	96.3± 0.3%	94.1± 0.0%	93.7± 0.2%
(18)	92.1± 0.4%	92.0± 0.2%	92.5± 0.5%	93.5± 0.1%	93.6± 0.3%
(19)	84.7± 0.6%	84.5± 0.5%	86.1± 0.6%	87.2± 0.1%	87.4± 0.5%
(20)	97.0± 0.5%	99.3± 0.1%	96.9± 0.6%	97.1± 0.1%	95.1± 0.2%

Table 3: Average revenue of first 50th rounds of five algorithms for semi-synthetic data sets with perfect response. The revenue is normalized by the optimal revenue of each data set. We repeat 50 trials to get the 95% confidence interval.