
Revenue-Incentive Tradeoffs in Dynamic Reserve Pricing

Yuan Deng¹ Sébastien Lahaie¹ Vahab Mirrokni¹ Song Zuo¹

Abstract

Online advertisements are primarily sold via repeated auctions with reserve prices. In this paper, we study how to set reserves to boost revenue based on the historical bids of strategic buyers, while controlling the impact of such a policy on the incentive compatibility of the repeated auctions. Adopting an incentive compatibility metric which quantifies the incentives to shade bids, we propose a novel class of dynamic reserve pricing policies and provide analytical tradeoffs between their revenue performance and bid-shading incentives. The policies are inspired by the exponential mechanism from the literature on differential privacy, but our study uncovers mechanisms with significantly better revenue-incentive tradeoffs than the exponential mechanism in practice. We further empirically evaluate the tradeoffs on synthetic data as well as real ad auction data from a major ad exchange to verify and support our theoretical findings.

1. Introduction

Online advertising is the practice of placing text, banner, or video ads on search engines, publisher websites, or social media for the purpose of raising brand awareness and generating sales to site visitors. This form of advertising represents a major source of revenue for both online publishers and Internet companies; the revenue of this market exceeded \$124.6 billion in 2019, witnessing a 15.9% increase from 2018 (PricewaterhouseCoopers, 2020). A large fraction of online ads are sold via repeated auctions in which the advertisers bid in real time.

Reserve pricing is a key tool used to boost revenue in repeated auctions (Ostrovsky & Schwarz, 2011; Paes Leme et al., 2016). However, according to classical auction theory, setting the reserve in an effective manner requires knowledge of the advertisers' valuation distributions (Myerson,

1981). A natural idea to circumvent this difficulty in practice is to learn a reserve price from the advertisers' historical bids, a practice known as dynamic reserve pricing. Although this approach is intuitive, one must be cautious with this kind of pricing policy since it complicates the long-term incentives of the advertisers. In particular, an advertiser may have an incentive to shade her bids in order to lower future reserve prices, even if each auction in isolation is truthful (e.g., second-price).

In the past decade, there has been a growing body of literature on dynamic reserve pricing with strategic advertisers under different settings. Two typical assumptions made in the literature are that the advertisers are impatient, so that they discount their utilities in the future (Amin et al., 2014; Drutsa, 2018; Golrezaei et al., 2019; Deng et al., 2019; 2020a), and that the market is large enough that each advertiser contributes a negligible fraction of the total revenue (Liu et al., 2018; Epasto et al., 2018). It has been shown that a constant fraction of revenue loss against the revenue-optimal benchmark is inevitable if these two assumptions do not hold (Amin et al., 2013). Nonetheless, in practice, many ad slots see competition from only a few large advertisers and it is likely that these advertisers care about their long-run utilities as much as their instantaneous payoffs in each auction (Nedelec et al., 2019).

In this paper, we consider an environment in which the advertisers are patient and the market might be small with only a few advertisers. Our objective is to understand *the tradeoff between bid-shading incentives and revenue performance* for different reserve pricing policies. Intuitively, in a single-buyer second-price auction environment, a reserve pricing policy that does not learn from the advertisers' historical bids and always sets a zero reserve would result in zero revenue, and induce no bid-shading incentive; on the other hand, a reserve pricing policy setting the Myersonian (i.e. revenue-optimal) reserve (Myerson, 1981) with respect to the advertisers' historical bid distribution maximizes revenue (provided that the advertisers bid truthfully in the past), but it induces a significant bid-shading incentive.

We consider a multi-stage model in which the auctioneer learns a reserve from the bids in the last stage and applies it in the current stage (see Section 2.1). Incentive compatibility is classically defined as a binary notion: An auction is either

¹Google Research, New York City, NY, USA. Correspondence to: Yuan Deng <dengyuan@google.com>.

incentive compatible or not. To achieve a more nuanced comparison between different reserve pricing policies, we adopt an incentive compatibility metric which can be interpreted as a quantitative index of the advertiser’s incentives to shade its bids (see Section 2.2 and 2.3). The advantage of this metric is that it is simple to compute via black-box simulations, and more importantly, it decomposes into static and dynamic components so that the incentives specifically induced by dynamic reserve prices can be evaluated.

Our Contributions. We first provide analytical revenue-incentive tradeoff for *linear* reserve pricing policies. A policy is linear if the reserve is scaled by a factor β when all the advertisers’ bids are scaled by β . Linear policies include the policy of setting the Myersonian reserve, and more generally any policy of setting the reserve as a fixed quantile of the bid distribution. We show that for such policies, there is precisely a linear tradeoff between revenue and bid-shading incentives: One must sacrifice one unit of the incentive metric in order to gain one unit of revenue. In other words, we obtain the intuitive result that the larger the revenue, the stronger the bid-shading incentive on the part of the advertiser. Our result further implies that although the Myersonian reserve maximizes revenue, it also minimizes the incentive metric among all linear policies in a single-buyer environment.

Motivated by our results for linear policies, we propose a novel class of *non-linear* policies based on the Box-Cox transformations familiar from statistics (Box & Cox, 1964). These policies randomize over a set of quantiles to select the reserve, where the weight ascribed to a quantile is proportional to some function of its revenue performance. Therefore, in these policies, bid shading affects not just the quantiles but also their relative weights. We provide a closed-form characterization of the revenue-incentive tradeoffs for our class of non-linear policies, demonstrating that they can result in significantly better revenue-incentive tradeoffs than linear policies.

Our analysis reveals the possibility of inducing bid-raising incentives via a non-linear policy’s mixture weights, which can counteract the bid-shading incentives. However, there is no dominant non-linear policy and designing a non-linear policy to induce bid-raising incentives is an empirical question that must take into account the actual valuation distributions at hand. Our class of non-linear policies is inspired by the exponential mechanism from the literature on differential privacy, which is a fundamental technique used to mitigate incentives to deviate from truthfulness when the market is large (McSherry & Talwar, 2007; Epasto et al., 2018). Our results shed light on how to adapt and generalize the exponential mechanism to achieve a significantly better revenue-incentive tradeoffs in a small market.

We empirically evaluate the revenue-incentive tradeoffs of our class of nonlinear policies on both synthetic data and real ad auction data from a major ad exchange. The experiments verify and support our theoretical findings, and uncover nonlinear policies with significantly better revenue-incentive tradeoffs than both linear policies and the exponential mechanism for realistic bid distributions.

Related Work. The recent line of study on revenue-optimal dynamic reserve pricing with strategic advertisers was initiated by Amin et al. (2013) and Medina & Mohri (2014); see den Boer (2015) for a survey. They design no-regret policies in a non-contextual environment with a single impatient advertiser, and the regret guarantee is later improved by Drutsa (2017; 2018; 2020). Amin et al. (2014) develop a no-regret policy in a contextual setting and Golrezaei et al. (2019) enrich the model by incorporating market noise. Deng et al. (2019; 2020a) generalize these results by designing a robust dynamic mechanism which competes favourably against the optimal dynamic mechanism. All these results concern impatient advertisers and Amin et al. (2013) show that no learning algorithm can achieve sublinear revenue loss with a patient advertiser. For a large market, Liu et al. (2018) design no-regret policies using techniques from differential privacy, while McSherry & Talwar (2007) and Epasto et al. (2018) provide general frameworks for designing mechanisms that are approximately incentive compatible. For patient buyers and small markets, Nedelec et al. (2019) investigate the advertiser’s bidding strategies against different dynamic reserve pricing policies.

There has been a growing amount of interest in developing metrics that quantify incentive compatibility. The common approach is to evaluate the advertiser’s regret, which is her utility difference between best-responding and truthful reporting (Parkes et al., 2001; Duetting et al., 2015; Balcan et al., 2019). However, computing a regret-based metric is often a complex optimization task of solving for the best response, and moreover, it is unclear how to generalize regret to repeated auctions. In this work, we instead adopt an incentive compatibility metric recently introduced in Deng et al. (2020b), which generalizes the idea of the marginal incentive to deviate (Erdil & Klemperer, 2010). The metric has been adopted to develop variants of generalized second-price auction using deep learning (Zhang et al., 2020).

2. Preliminaries

We consider an environment with multiple buyers participating in repeated auctions with a single seller (e.g., an ad exchange). The repeated auctions sell a sequence of *queries*, each granting the winning buyer a chance to display ads on the website. These queries arrive in an online manner and they must be sold once they arrive. For convenience, we

assume that there is exactly one query per round, so the t -th query arrives at round t .

For ease of presentation, we take the perspective of one fixed buyer. We assume that the buyer's private valuation $v_t \in V = \mathbb{R}_{\geq 0}$ for round t is drawn independently from a continuous distribution F_t over V with density function f_t . Upon the arrival of the t -th query, the buyer first learns her private valuation v_t drawn from F_t and then submits her bid $b_t \in V$ to the seller. After receiving the bids from all buyers, the seller decides how to allocate and charge for the ad slot according to an allocation function x_t and a payment function p_t . As we take the perspective of one fixed buyer, x_t and p_t subsume the mechanism (in particular, the reserve price applied) and other buyers' bidding behavior. In line with the literature, we assume that the buyer's utility is quasi-linear such that her utility under bid b_t at round t with private valuation v_t is $u_t(b_t; v_t) = v_t \cdot x_t(b_t) - p_t(b_t)$. For convenience, let $\hat{u}_t(v_t) = v_t \cdot x_t(v_t) - p_t(v_t)$ be the buyer's utility under truthful bidding.

We consider a widely adopted auction format: *the second-price auction with lazy personalized reserve* (Ostrovsky & Schwarz, 2011; Paes Leme et al., 2016), in which the buyer wins the ad slot if she bids higher than the highest bid from others b_t^- and the reserve r_t for round t . If a buyer wins, she pays $\max(b_t^-, r_t)$; otherwise, she pays nothing. We further assume that at each round, the highest bid among other buyers is drawn independently from a continuous distribution G_t over V with density function g_t . In practice, many other auction formats are used in online advertising (e.g., first-price auction, generalized second-price auction). In this work, we are interested in the incentives across auctions, rather than within each individual auction, so we consider second-price auctions (which are truthful) to focus on the dynamic incentives.

2.1. A Multi-stage Model for Reserve Pricing

In a setting where the seller has no prior knowledge about the buyers' valuations, a canonical way to apply reserve pricing is to learn a reserve to apply in the future from the buyers' historical bids (Kanoria & Nazerzadeh, 2014; den Boer, 2015). To analyze the tradeoff between incentives and revenue under different reserve pricing policies, we propose a multi-stage model where each stage consists of sufficiently many queries (rounds), and the reserve is set to be the same within each stage. In particular,

- The reserve is set to be 0 for all queries in stage 1;
- For each stage $t > 1$, the seller applies a reserve pricing policy to compute a reserve r_t from the buyer's bids from stage $t - 1$; The seller sets the reserve to r_t for queries in stage t .

One can hardly hope for a good revenue performance guarantee for dynamic reserve pricing if the environments are quite different across both stages. For the purpose of theoretical analysis, we further assume that $F_t = F$ and $G_t = G$ for all t , so that they are identical across queries. We will examine how the revenue-incentive tradeoffs generalize across stages in our experiments for realistic bid distributions that arise practice.

2.2. Incentive Compatibility Metric

An auction $\langle x, p \rangle$ is incentive compatible (IC) if reporting truthfully is always an optimal strategy for the buyer regardless of her private valuation and others' bids. Formally, for all $v \in V$, $v \in \arg \max_b v \cdot x(b) - p(b)$. Myerson's lemma (Myerson, 1981) pins down the relationship between the allocation and payment rule for any truthful auction: (1) The allocation rule x is non-decreasing; (2) The payment rule satisfies $p(v) = v \cdot x(v) - \int_0^v x(z) dz$. In particular, the latter condition implies that for each valuation, the derivative of the buyer's utility equals the allocation probability, i.e., $\frac{d\hat{u}(v)}{dv} = x(v)$ for all v .

Inspired by Myerson's lemma, Deng et al. (2020b) proposes a metric to quantify incentive compatibility.

Theorem 2.1 (Individual Stage-IC Metric (Deng et al., 2020b)). *The stage-IC metric $\text{SIC}^{\langle x, p \rangle}$ for an agent with valuation distribution F in an auction $\langle x, p \rangle$ is defined as*

$$\lim_{\alpha \rightarrow 0} \frac{\mathbb{E} [\hat{u}((1 + \alpha)v)] - \mathbb{E} [\hat{u}((1 - \alpha)v)]}{2\alpha \cdot \mathbb{E} [v \cdot x(v)]}.$$

The expectation here is taken over the valuation distribution F . Intuitively, the numerator of SIC is the difference between the agent's expected utility when her valuations are perturbed multiplicatively up versus down by α , while the denominator is 2α times the expected welfare. This metric can be viewed as a test of the Myerson's condition $\frac{d\hat{u}(v)}{dv} = x(v)$ for all v , except that only the mean value of the test is reported.

Deng et al. (2020b) demonstrates that the metric takes the form of an index between 0 and 1 for reasonable auctions: it is non-negative if the buyer's utility under truthful bidding is non-decreasing in her true valuation, and it is at most 1 if overbidding is a weakly dominated strategy. Additionally, the metric is 1 for incentive-compatible auctions while it is 0 for first-price auctions. More importantly, the metric captures *the marginal benefit of uniform bidding*, in which the buyer always bids $b = \beta v$ for a constant factor β approaching to 1. In particular, when the metric is below 1, the buyer has incentive to shade her bids. The smaller the metric, the stronger the bid-shading incentive. On the other hand, when the metric is above 1, the buyer has the incentive to raise her bids. The larger the metric, the stronger

the bid-raising incentive. This observation is important in practice, as uniform bidding has been shown to be a typical (and sometimes optimal) strategy for buyers to participate in ad auctions in complex environments (Aggarwal et al., 2019; Balseiro & Gur, 2019; Deng et al., 2021a).

2.3. Dynamic Incentive Compatibility Metric

Intuitively, a mechanism is dynamic-IC if for any stage t , reporting truthfully at stage t is an optimal strategy for the buyer regardless of her private valuation and others' bids, assuming the buyer always bids truthfully in the future (Mirrokni et al., 2020; Deng et al., 2020b; 2021b).¹

The IC metric can be generalized to multi-stage auctions by taking the future stages into account to verify the definition of dynamic-IC. In our multi-stage model, when measuring dynamic-IC for stage t , only stage $t + 1$ is relevant since the bids at stage t only affect the reserve price at stage $t + 1$; and the reserve prices for all future stages are the same provided that the buyer always reports truthfully from stage $t + 1$ onwards, under the dynamic-IC definition.

We therefore focus, without loss of generality, on a multi-stage model with two stages. Note that with the reserve learned from stage 1 fixed, stage 2 would be IC, and therefore, the buyer has no incentive to misreport in stage 2. However, the buyer may still have incentive to misreport in stage 1 to lower the reserve in stage 2.

To define the metric for dynamic-IC in stage 1, let $x_1(b)$ and $p_1(b)$ be the expected allocation and payment when the buyer bids b in stage 1. Moreover, let $x_2(\beta, b)$ and $p_2(\beta, b)$ represent the expected allocation and payment when she submits bids uniformly scaled by a factor β from her private valuations in stage 1, i.e., βv , and bids b in stage 2. In addition, let $\hat{u}_2(\beta, v) = v \cdot x_2(\beta, v) + p_2(\beta, v)$. Therefore, the buyer's expected cumulative utility under truthful bidding is $\mathbb{E}[\hat{u}_1(v)] + \mathbb{E}[\hat{u}_2(1, v)]$. We are now ready to define the dynamic-IC metric adapted from (Deng et al., 2020b).

Definition 2.2 (Dynamic-IC Metric). *The dynamic-IC metric DIC for a buyer with valuation distribution F in a dynamic auction in stage 1 is defined as*

$$\text{DIC}^{(x_1, p_1)} + \lim_{\alpha \rightarrow 0} \frac{\mathbb{E}[\hat{u}_2(1 + \alpha, v)] - \mathbb{E}[\hat{u}_2(1 - \alpha, v)]}{2\alpha \cdot \mathbb{E}[v \cdot x_1(v)]}.$$

The dynamic-IC metric (DIC) additionally takes the utility in stage 2 into account: the additional term shares the same denominator as the stage-IC metric, while the numerator computes the utility difference in stage 2 under truthful bidding when the buyer's valuations are perturbed multiplicatively up versus down by α in stage 1.

¹By a backward induction argument, this is equivalent to assuming the buyer bids in a way to optimize her utility in the future.

Since our auction format is IC for stage 1 in isolation², $\text{SIC}^{(x_1, p_1)}$ is always 1 (Deng et al., 2020b). Therefore, we can rewrite the metric as follows:

$$\text{DIC} = 1 + \lim_{\alpha \rightarrow 0} \frac{\mathbb{E}[\hat{u}_2(1 + \alpha, v)] - \mathbb{E}[\hat{u}_2(1 - \alpha, v)]}{2\alpha \cdot \mathbb{E}[v \cdot x_1(v)]}.$$

The second term is usually negative under dynamic reserve pricing since dynamic reserve pricing induces the buyer to shade her bids.

2.4. Reserve Pricing Policy

In general, a reserve pricing policy maps a distribution of the buyer's bids to a distribution of reserves. Given a distribution F over V with finite and non-zero density function f , there is a bijection between values in V and quantiles $\kappa \in [0, 1]$ such that $\kappa = F(v)$. For convenience, denote the quantile function that maps a quantile to a value by q such that $q(\kappa) = F^{-1}(\kappa)$. As a result, it is without loss of generality to represent a reserve pricing policy $M : \Delta(V) \rightarrow \Delta([0, 1])$ as a function that maps a distribution of bids to a distribution of quantiles, where $\Delta(V)$ is the set of distributions over V .

Denote by M_F the distribution of quantiles returned by M under truthful bidding when the buyer's valuation distribution is F . Notice that the dynamic-IC metric only depends on the bid distributions generated by the uniform bidding strategy with factor $(1 - \alpha)$ and $(1 + \alpha)$. We denote by M_F^β the distribution of quantiles returned by the policy when the buyer adopts the uniform bidding strategy with factor β , and the reserve price will be set to $\beta \cdot q(\kappa)$ with a quantile κ selected by the policy. A policy is called *linear* if it is independent of β : $M_F^\beta = M_F$ for all β . In a linear policy, the distribution over quantiles remains unchanged, while the reserve price is scaled accordingly by β .

3. Linear Policy

We begin by analyzing linear policies. Observe that given a quantile distribution M_F , the revenue contribution from the buyer under truthful bidding is $\text{REV} = \text{REV}_{\text{reserve}} + \text{REV}_{\text{compt}}$, where $\text{REV}_{\text{reserve}}$ is the contribution from reserve (when the highest bid from others is below the reserve) and $\text{REV}_{\text{compt}}$ is the contribution from competition (when the highest bid from others is above the reserve):

$$\begin{aligned} \text{REV}_{\text{reserve}} &= \mathbb{E}_{\kappa \sim M_F} \left[(1 - \kappa)q(\kappa) \cdot G(q(\kappa)) \right], \\ \text{REV}_{\text{compt}} &= \mathbb{E}_{\kappa \sim M_F} \left[\int_0^\infty \int_{q(\kappa)}^v z \cdot g(z) dz f(v) dv \right]. \end{aligned}$$

²More precisely, the auction for stage 1 is IC if the buyer were to participate only in stage 1 but not stage 2.

Here, $(1 - \kappa) \cdot G(q(\kappa))$ is the probability that the buyer's valuation is above $q(\kappa)$ and the highest bid from others is below $q(\kappa)$. In other words, it is the probability that the reserve $q(\kappa)$ is effective under truthful bidding.

Let $\hat{u}_2(\beta, v; \kappa)$ be the buyer's utility in stage 2 when (i) in stage 1, she adopts the uniform bidding strategy with factor β and (ii) in stage 2, she truthfully bids her private value v while the reserve is the κ -th quantile from stage 1, i.e., $\beta \cdot q(\kappa)$. Note that $\hat{u}_2(\beta, v; \kappa) = 0$ if $v \leq \beta \cdot q(\kappa)$; and

$$\hat{u}_2(\beta, v; \kappa) = (v - \beta \cdot q(\kappa)) \cdot G(\beta \cdot q(\kappa)) + \int_{\beta \cdot q(\kappa)}^v (v - z) \cdot g(z) dz,$$

if $v > \beta \cdot q(\kappa)$, where the first term is the utility when the highest bid from others is below $\beta \cdot q(\kappa)$ and the second term is the utility when the highest bid from others is above $\beta \cdot q(\kappa)$. We further define

$$\text{util}(\kappa, \beta) = \mathbb{E}_{v \sim F} [\hat{u}_2(\beta, v; \kappa)]$$

as the buyer's expected utility. We now proceed to provide analytical revenue-incentive tradeoffs for linear policies. For convenience, let $W = \mathbb{E}[v \cdot x_1(v)]$ for the rest of the paper.

Theorem 3.1. *For any linear policy, the dynamic-IC metric of each buyer with valuation distribution F and competing bid distribution G satisfies*

$$\text{DIC} = 1 - \text{REV}_{\text{reserve}}/W.$$

Proof. We focus on $\lim_{\alpha \rightarrow 0} \frac{\mathbb{E}[\hat{u}_2(1+\alpha, v)] - \mathbb{E}[\hat{u}_2(1-\alpha, v)]}{2\alpha}$ in DIC since other terms are constant. For a linear policy, this can be written as

$$\begin{aligned} & \lim_{\alpha \rightarrow 0} \frac{\mathbb{E}[\hat{u}_2(1 + \alpha, v)] - \mathbb{E}[\hat{u}_2(1 - \alpha, v)]}{2\alpha} \\ &= \lim_{\alpha \rightarrow 0} \frac{\mathbb{E}_{\kappa \sim M_F^{1+\alpha}} [\text{util}(\kappa, 1 + \alpha)] - \mathbb{E}_{\kappa \sim M_F^{1-\alpha}} [\text{util}(\kappa, 1 - \alpha)]}{2\alpha} \\ &= \lim_{\alpha \rightarrow 0} \frac{\mathbb{E}_{\kappa \sim M_F} [\text{util}(\kappa, 1 + \alpha)] - \mathbb{E}_{\kappa \sim M_F} [\text{util}(\kappa, 1 - \alpha)]}{2\alpha} \\ &= \left. \frac{\partial \mathbb{E}_{\kappa \sim M_F} [\text{util}(\kappa, \beta)]}{\partial \beta} \right|_{\beta=1}, \end{aligned}$$

where the second equation follows that the policy is linear. All that remains to show is that

$$\left. \frac{\partial \text{util}(\kappa, \beta)}{\partial \beta} \right|_{\beta=1} = -(1 - \kappa) \cdot q(\kappa) \cdot G(q(\kappa))$$

for any κ (details deferred to Appendix A.1). \square

Theorem 3.1 reveals a linear tradeoff between the dynamic-IC metric and the revenue contribution from the reserve: The larger $\text{REV}_{\text{reserve}}$ is, the stronger bid-shading incentive

the buyer has in stage 1. Theorem 3.1 also shows that the metric is fully determined by $\text{REV}_{\text{reserve}}$, and as a result, for any two linear policies M_F and M'_F that achieves the same $\text{REV}_{\text{reserve}}$, their dynamic-IC metrics must be the same. Moreover, as $\text{REV} = \text{REV}_{\text{reserve}}$ when there is only one buyer, we immediately have the following corollary for the single-buyer environment:

Corollary 3.2. *For any linear policy under the single-buyer two-stage model, the dynamic-IC metric of a buyer with valuation distribution F satisfies $\text{DIC} = 1 - \text{REV}/W$.*

Note that the policy that applies the Myersonian reserve (Myerson, 1981) estimated from the bid distribution is also linear. Hence, the Myersonian reserve maximizes revenue but also minimizes the dynamic-IC metric (i.e., creates the strongest bid shading incentives) among all linear policies in the single-buyer environment.

4. Non-linear Policy Based on the Box-Cox Transformation

In the previous section, we analyzed the revenue-incentive tradeoff for linear reserve pricing policies. We now turn to reserve pricing policies that are not linear. A canonical non-linear policy is inspired by the exponential mechanism (McSherry & Talwar, 2007; Epasto et al., 2018), which is widely used to establish differential privacy and leads to approximate-IC mechanisms when the market is large. Formally, there is a weight function $w : [0, 1] \rightarrow \mathbb{R}_{\geq 0}$ that assigns each quantile κ a weight $w(\kappa)$, and the probability of selecting each quantile κ is proportional to $\exp(\theta \cdot w(\kappa))$, where $\theta \in \mathbb{R}_{\geq 0}$ is an adjustable parameter. For example, $w(\kappa)$ can be the revenue in stage 1 when the reserve is $q(\kappa)$. Another canonical mechanism is the power mechanism (Epasto et al., 2020), in which the probability of selecting each quantile κ is proportional to $w(\kappa)^\theta$.

We consider a broad class of non-linear policies that contains both the exponential mechanism and the power mechanism. It is inspired by the power transformation proposed by Box & Cox (1964), known as *the Box-Cox transformation*. The single-parameter Box-Cox transformation function is:

$$\tau_\lambda(y) = \begin{cases} (y^\lambda - 1)/\lambda, & \lambda \in \mathbb{R} \setminus \{0\}; \\ \ln y, & \lambda = 0. \end{cases}$$

Definition 4.1 (Non-linear Policy based on the Box-Cox Transformation). *A non-linear policy M is specified by a tuple $\langle \lambda, \theta, w \rangle$, where $\lambda \in \mathbb{R}$ and $\theta \in \mathbb{R}_{\geq 0}$. The probability of selecting each quantile κ is proportional to $\exp(\theta \cdot \tau_\lambda(w(\kappa)))$.*

Observe that when $\lambda = 0$, the policy becomes the power mechanism, and when $\lambda = 1$, it becomes the exponential mechanism. Recall that the dynamic-IC metric concerns the

bid distribution generated by the uniform bidding strategy for factor $(1 - \alpha)$ and $(1 + \alpha)$. For convenience, denote by $w(\kappa, \beta)$ the weight function when the buyer adopts the uniform bidding strategy with factor β in stage 1. We assume that $w(\kappa, \beta)$ is continuous and differentiable with respect to β , but we do not require it to be continuous with respect to κ . This enables different weight function formats for different κ . The following theorem characterizes the dynamic-IC metric for our class of non-linear policies.

Theorem 4.2. *For any non-linear policy $M = \langle \lambda, \theta, w \rangle$, the dynamic-IC metric of a buyer with valuation distribution F and competing bid distribution G satisfies*

$$\text{DIC} = 1 - (\text{REV}_{\text{reserve}} - \theta \cdot \Psi) / W,$$

where

$$\Psi = \text{cov}_{\kappa \sim M_F} \left[\text{util}(\kappa, 1), \frac{\partial w(\kappa, \beta)}{\partial \beta} \Big|_{\beta=1} \cdot w(\kappa, 1)^{\lambda-1} \right].$$

Proof. We focus on $\lim_{\alpha \rightarrow 0} \frac{\mathbb{E}[\hat{u}_2(1+\alpha, v)] - \mathbb{E}[\hat{u}_2(1-\alpha, v)]}{2\alpha}$ in DIC since other terms are constant. For a non-linear policy, this can be written as

$$\begin{aligned} & \lim_{\alpha \rightarrow 0} \frac{\mathbb{E}[\hat{u}_2(1+\alpha, v)] - \mathbb{E}[\hat{u}_2(1-\alpha, v)]}{2\alpha} \\ &= \lim_{\alpha \rightarrow 0} \frac{\mathbb{E}_{\kappa \sim M_F^{1+\alpha}} [\text{util}(\kappa, 1+\alpha)] - \mathbb{E}_{\kappa \sim M_F^{1-\alpha}} [\text{util}(\kappa, 1-\alpha)]}{2\alpha} \\ &= \frac{\partial \mathbb{E}_{\kappa \sim M_F^\beta} [\text{util}(\kappa, \beta)]}{\partial \beta} \Big|_{\beta=1} \\ &= \frac{\partial \mathbb{E}_{\kappa \sim M_F} [\text{util}(\kappa, \beta)]}{\partial \beta} \Big|_{\beta=1} + \frac{\partial \mathbb{E}_{\kappa \sim M_F^\beta} [\text{util}(\kappa, 1)]}{\partial \beta} \Big|_{\beta=1} \\ &= -\text{REV}_{\text{reserve}} + \frac{\partial \mathbb{E}_{\kappa \sim M_F^\beta} [\text{util}(\kappa, 1)]}{\partial \beta} \Big|_{\beta=1}, \end{aligned}$$

where the last equation would follow the proof of Theorem 3.1. The equation $\frac{\partial \mathbb{E}_{\kappa \sim M_F^\beta} [\text{util}(\kappa, 1)]}{\partial \beta} \Big|_{\beta=1} = \theta \cdot \Psi$ then concludes the proof (details deferred to Appendix A.2). \square

Theorem 4.2 shows that compared with linear policies (Theorem 3.1), the dynamic-IC metric for non-linear policies has an additional covariance term of $\theta \cdot \Psi$ under the quantile distribution M_F . Note that Theorem 4.2 generalizes Theorem 3.1 since a linear mechanism can be implemented with a weight function independent of β so that $\frac{\partial w(\kappa, \beta)}{\partial \beta} = 0$ for any κ and β .

Intuitively, under linear policies, a buyer who adopts the uniform bidding strategy can only change the scale of the reserve while the distribution of quantiles remains the same, which gives the buyer the incentive to shade her bids to lower the reserve. Such an incentive is captured by the

partial derivative

$$\frac{\partial \mathbb{E}_{\kappa \sim M_F} [\text{util}(\kappa, \beta)]}{\partial \beta} \Big|_{\beta=1} = -\text{REV}_{\text{reserve}},$$

where β appears in the utility of the second stage. In contrast, under non-linear policies, a buyer can further change the distribution of the quantiles, and such an effect is captured by the partial derivative

$$\frac{\partial \mathbb{E}_{\kappa \sim M_F^\beta} [\text{util}(\kappa, 1)]}{\partial \beta} \Big|_{\beta=1} = \theta \cdot \Psi,$$

where β appears in the distribution M_F^β of the quantiles.

A properly designed non-linear policy, resulting in a positive covariance term Ψ , will create an incentive for the buyer to raise her bids, offsetting the bid-shading incentive captured by $\text{REV}_{\text{reserve}}$. Note that $\text{util}(\kappa, 1)$ is non-increasing as κ increases, and therefore, the covariance term is likely to be positive if $\frac{\partial w(\kappa, \beta)}{\partial \beta} \Big|_{\beta=1} \cdot w(\kappa, 1)^{\lambda-1}$ is also non-increasing for most quantiles.

In particular, when $w(\kappa, \beta)$ is linear in β , i.e., $w(\kappa, \beta) = \beta \cdot w(\kappa, 1)$, we immediately have:

Corollary 4.3. *For any non-linear policy $M = \langle \lambda, \theta, w \rangle$ with $w(\kappa, \beta)$ linear in β , the dynamic-IC metric of a buyer with valuation distribution F and competing bid distribution G satisfies $\text{DIC} = 1 - (\text{REV}_{\text{reserve}} - \theta \cdot \Psi) / W$ where*

$$\Psi = \text{cov}_{\kappa \sim M_F} [\text{util}(\kappa, 1), w(\kappa, 1)^\lambda].$$

Particularly, for a power mechanism in which $\lambda = 0$, the dynamic-IC metric satisfies $\text{DIC} = 1 - \text{REV}_{\text{reserve}} / W$.

Corollary 4.3 implies that the revenue-incentive tradeoff in a power mechanism with a weight function linear in β is exactly the same as in a linear policy. In the single-buyer environment, all the revenue comes from reserve pricing, i.e., $\text{REV} = \text{REV}_{\text{reserve}}$, and is hence linear in β . Similarly, the welfare performance is also linear in β . Therefore, if one uses a convex combination between the revenue and the welfare performance as the weight function, the tradeoff given in Corollary 4.3 still applies.

5. Experiments

In this section we report on results from simulations of the non-linear policies. When simulating a policy, we use REV (i.e., the buyer's expected payment) as the weight function. By varying the parameter θ (in Definition 4.1), we obtain different combinations of $(\text{REV}_{\text{reserve}} / W, \text{DIC})$ which yields a tradeoff curve that we plot. Note that the first coordinate is (normalized) revenue under truthful bidding, but when $\text{DIC} < 1$ one would expect buyers to shade their bids. The plots are best interpreted as showing dominance relationships: for two mechanisms that yield the same revenue




Figure 1: Revenue curve




Figure 2: Log-normal




Figure 3: Shifted log-normal

Figure: Synthetic data experiment: revenue-incentive tradeoff *without* competing bids under non-linear policies with different λ ; the dashed line corresponds to the line of $\text{DIC} = 1 - \text{REV}_{\text{reserve}}/W$.

under truthful bidding, the one with higher DIC should in practice be closer to this predicted revenue.

We first present synthetic-data experiments to verify the theory, and next report on experiments over real bid data from a major ad exchange to understand which particular policies work well over realistic bid distributions in practice.

5.1. Synthetic Data

We first provide results using synthetic data. We consider a standard log-normal distribution F_1 with density function $f_1(x) = \frac{1}{x\sqrt{2\pi}} \cdot \exp(-\frac{\ln x}{2})$ and a shifted log-normal distribution F_2 with density function $f_2(x) = f_1(x-1)$ for $x > 1$ and $f_2(x) = 0$ for $x \in (0, 1]$. We use a log-normal distribution because it is commonly used to model bid and value distributions in practice (Lahaie & Pennock, 2007; Ostrovsky & Schwarz, 2011; Thompson & Leyton-Brown, 2013; Golrezaei et al., 2019).

We consider a single-buyer environment without competing bids. We compute the 100-quantiles, excluding the 100-th quantile (i.e., the maximum), and their revenues as the weights. Figure 1 shows the revenue curve for both distributions. Observe that the revenue for the log-normal distribution is increasing as the quantile increases for most of the support, while the revenue for the shifted log-normal distribution is decreasing over most of the support. According to the covariance term Ψ in Corollary 4.3, this suggests that a non-linear policy with negative λ should have a better tradeoff than one with positive λ for the log-normal distribution, but have a worse tradeoff for the shifted log-normal distribution. Figures 2 and 3 show the tradeoffs and confirm these theoretical findings. In addition, we see that for the log-normal distribution, the smaller λ , the better the tradeoff; and the pattern for the shifted log-normal distribution is

the opposite. Moreover, the tradeoff curve for $\lambda = 0$ coincides with the diagonal line of $\text{DIC} = 1 - \text{REV}_{\text{reserve}}/W$ in agreement with the theory. Note that the curves do not start from 0 revenue (and 1 dynamic-IC metric) since all the non-linear policies return a uniform distribution over quantiles when $\theta = 0$, resulting in a positive revenue. Moreover, the rightmost points in Figures 2 and 3 in fact correspond to the Myerson’s auctions.

We defer the experiments with competing bids and with λ that has a large absolute value to Appendix B. The trends for both distributions are similar to the setting without competing bids, but the curve for $\lambda = 0$ deviates from the diagonal line since the revenue is no longer linear in the bid-shading factor β . Moreover, for a λ with a large absolute value, the dynamic-IC metric could be above 1, resulting in an incentive for the buyer to raise her bids beyond its true value. Therefore, one must be cautious with the choice of λ to balance the bid-raising and bid-shading incentives.

5.2. Ad Auction Data

In this section we evaluate the revenue-incentive tradeoffs within our family of mechanisms over real data from the auction logs of a major display ad exchange. Our goal here is not to evaluate the exact auction mechanism and pricing policies used by the exchange, but rather to validate our theory over realistic bid distributions.³ The purpose of this

³In particular, the exchange from which the data is drawn runs a first-price auction, but we use the bids to simulate second-price auctions with reserves, consistent with our model so far. One could first try to back out values from bids, but we view this as an orthogonal problem which would not substantially alter the experimental conclusions. To empirically evaluate the tradeoff, what matters is the shape of the distribution; in ad auctions, bid and value distributions are likely to have a similar shape since one is just a rescaling of the other under the common strategy of uniform




Figure 4: Tradeoff on training data.




Figure 5: Tradeoff on testing data.

Figure: Real data experiment: revenue-incentive tradeoff *without* competing bids under non-linear policies with different λ ; the dashed line corresponds to the line of $DIC = 1 - \text{REV}_{\text{reserve}}/W$.

evaluation is twofold. In Section 5.1, we saw that depending on the bid distributions, higher λ could lead to better or worse tradeoffs. Therefore, we are interested in uncovering the pattern that holds under realistic bid distributions. Next, we are interested in confirming that the tradeoffs can be measured to a useful level of accuracy in practice (i.e., the confidence intervals are narrow enough to distinguish mechanisms using different values of λ).

Our dataset consists of auction records of bids placed on queries from a single publisher over two consecutive days. We focus on a single bidder and only retain the auctions that involve this bidder, yielding over 100K auction records in total. The first day is used as the training set to compute the bid quantiles and their associated revenue performance as reserves under a second-price auction. Specifically, we compute the (5, 10, 25, 50, 75, 90, 95)-th bid quantiles and their counterfactual revenue. The non-linear policies from Section 4 are then implemented based on this information, and the revenue-incentive tradeoffs are evaluated on both the training (same-day) data and the testing (next-day) data. The policies take mixtures of the bid quantiles along with the option of no reserve price. Note that our model has assumed that bids are i.i.d. across days; this does not hold exactly in practice, so our empirical results will clarify the extent to which the relationships characterized in Theorem 4.2 and Corollary 4.3 extend to future time periods.

Results assuming just a single bidder (i.e., discarding all competing bidders in the simulation) are presented in Figure 4 and Figure 5. (In all plots, we present 95% confidence intervals.) We observe that lower λ achieve better revenue-incentive tradeoffs, in the sense that at every revenue level, lower λ dominate higher λ in terms of the DIC metric; in particular, the exponential mechanism ($\lambda = 1$) exhibits the bidding which motivates our analysis.

worst tradeoffs. Following the intuition from Corollary 4.3, this is because revenue increases with the reserve (and is negatively correlated with utility) for a large portion of the support of the empirical bid distribution. By using negative values like $\lambda = -1$, the resulting bid-raising incentives can counteract the bid-shading incentives, as the DIC metric initially stays quite flat even as the revenue performance increases, as seen in Figure 4. This figure also confirms that the power mechanism ($\lambda = 0$) follows a specific linear tradeoff curve, validating the theory. In Figure 5 we evaluate the tradeoffs on the testing data; the trends are extremely similar, but the power mechanism no longer tracks the predicted line because the distributions are not exactly identical day-to-day.

Plots of the tradeoff curves when there are competing bids are deferred to Appendix B. Again, we find that lower λ yield better tradeoffs. The trends are similar to the single-bidder case, but because the revenue-weighting is now non-linear the power mechanism deviates from the diagonal line. Comparing the range of the dynamic-IC metric to the single-bidder case, we find that the dynamic-IC metric is higher, because the competing bids mute the bid-shading incentives of the dynamic reserve prices.

6. Conclusions

This work provides analytical characterizations of the revenue-incentive tradeoffs of dynamic reserve pricing policies, and introduces a novel class of policies that randomize over bid quantiles as reserves.

For linear policies, we proved that bid-shading incentives are directly proportional to the revenue extracted by a reserve. For non-linear policies based on the Box-Cox transformation, we showed that an additional term is added which

depends on the covariance between quantile weight and quantile utility. As a result, by choosing an appropriate weighting scheme, it is possible to counteract bid-shading incentives with bid-raising incentives. The mechanism exhibiting the best tradeoff depends on the bidder's valuation distribution, and in particular, on the extent to which revenue increases with quantile over the distribution's support. For realistic distributions like the lognormal, and on an empirical distribution of bids from a major ad exchange, the exponential mechanism in fact accentuates incentives to shade bids and significantly better tradeoffs can be achieved with log-concave weighting schemes.

In future work, it would be interesting to consider more general reserve policies, for example, a combination between quantile-based reserve and constant reserves, such as, the 20-th quantile of the bid distribution plus a constant. Other reserve pricing strategies, such as anonymous reserves, are also interesting to consider. Different strategies may come with different revenue-incentive tradeoffs, and it is an intriguing question to characterize the landscape of revenue-incentive tradeoffs.

References

- Aggarwal, G., Badanidiyuru, A., and Mehta, A. Autobidding with constraints. In *International Conference on Web and Internet Economics*, pp. 17–30. Springer, 2019.
- Amin, K., Rostamizadeh, A., and Syed, U. Learning prices for repeated auctions with strategic buyers. In *Advances in Neural Information Processing Systems*, pp. 1169–1177, 2013.
- Amin, K., Rostamizadeh, A., and Syed, U. Repeated contextual auctions with strategic buyers. In *Advances in Neural Information Processing Systems*, pp. 622–630, 2014.
- Balcan, M.-F., Sandholm, T., and Vitercik, E. Estimating approximate incentive compatibility. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19, pp. 867–867, New York, NY, USA, 2019. ACM.
- Balseiro, S. R. and Gur, Y. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, 65(9):3952–3968, 2019.
- Box, G. E. and Cox, D. R. An analysis of transformations. *Journal of the Royal Statistical Society: Series B (Methodological)*, 26(2):211–243, 1964.
- den Boer, A. V. Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in operations research and management science*, 20(1): 1–18, 2015.
- Deng, Y., Lahaie, S., and Mirrokni, V. A robust non-clairvoyant dynamic mechanism for contextual auctions. In *Advances in Neural Information Processing Systems*, pp. 8654–8664, 2019.
- Deng, Y., Lahaie, S., and Mirrokni, V. Robust pricing in dynamic mechanism design. In *International Conference on Machine Learning*, pp. 2494–2503. PMLR, 2020a.
- Deng, Y., Lahaie, S., Mirrokni, V., and Zuo, S. A data-driven metric of incentive compatibility. In *Proceedings of The Web Conference 2020*, pp. 1796–1806, 2020b.
- Deng, Y., Mao, J., Mirrokni, V., and Zuo, S. Towards efficient auctions in an auto-bidding world. In *Proceedings of the Web Conference 2021*, pp. 3965–3973, 2021a.
- Deng, Y., Mirrokni, V., and Zuo, S. Non-clairvoyant dynamic mechanism design with budget constraints and beyond. In *Proceedings of the 2021 ACM Conference on Economics and Computation*, EC '21. ACM, 2021b.
- Drutsa, A. Horizon-independent optimal pricing in repeated auctions with truthful and strategic buyers. In *Proceedings of the 26th International Conference on World Wide Web*, pp. 33–42. International World Wide Web Conferences Steering Committee, 2017.
- Drutsa, A. Weakly consistent optimal pricing algorithms in repeated posted-price auctions with strategic buyer. In *International Conference on Machine Learning*, pp. 1318–1327, 2018.
- Drutsa, A. Reserve pricing in repeated second-price auctions with strategic bidders. In *International Conference on Machine Learning*, 2020.
- Duetting, P., Fischer, F., Jirapinyo, P., Lai, J. K., Lubin, B., and Parkes, D. C. Payment rules through discriminant-based classifiers. *ACM Transactions on Economics and Computation (TEAC)*, 3(1):5, 2015.
- Epasto, A., Mahdian, M., Mirrokni, V., and Zuo, S. Incentive-aware learning for large markets. In *Proceedings of the 2018 World Wide Web Conference*, pp. 1369–1378, 2018.
- Epasto, A., Mahdian, M., Mirrokni, V., and Zampetakis, E. Optimal approximation-smoothness tradeoffs for softmax functions. In *Advances in Neural Information Processing Systems*, 2020.
- Erdil, A. and Klempner, P. A new payment rule for core-selecting package auctions. *Journal of the European Economic Association*, 8(2-3):537–547, 2010.
- Golrezaei, N., Javanmard, A., and Mirrokni, V. Dynamic incentive-aware learning: Robust pricing in contextual auctions. In *Advances in Neural Information Processing Systems*, pp. 9756–9766, 2019.

- Kanoria, Y. and Nazerzadeh, H. Dynamic reserve prices for repeated auctions: Learning from bids. In *Web and Internet Economics: 10th International Conference, WINE 2014, Beijing, China, December 14-17, 2014, Proceedings*, volume 8877, pp. 232. Springer, 2014.
- Lahaie, S. and Pennock, D. M. Revenue analysis of a family of ranking rules for keyword auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*, pp. 50–56, 2007.
- Liu, J., Huang, Z., and Wang, X. Learning optimal reserve price against non-myopic bidders. In *Advances in Neural Information Processing Systems*, pp. 2038–2048, 2018.
- McSherry, F. and Talwar, K. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pp. 94–103. IEEE, 2007.
- Medina, A. M. and Mohri, M. Learning theory and algorithms for revenue optimization in second price auctions with reserve. In *Proceedings of the 31st International Conference on Machine Learning*, pp. 262–270, 2014.
- Mirroknj, V., Paes Leme, R., Tang, P., and Zuo, S. Non-clairvoyant dynamic mechanism design. *Econometrica*, 88(5):1939–1963, 2020.
- Myerson, R. B. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- Nedelec, T., El Karoui, N., and Perchet, V. Learning to bid in revenue-maximizing auctions. In *International Conference on Machine Learning*, pp. 4781–4789, 2019.
- Ostrovsky, M. and Schwarz, M. Reserve prices in internet advertising auctions: A field experiment. In *Proceedings of the 12th ACM conference on Electronic commerce*, pp. 59–60, 2011.
- Paes Leme, R., Pal, M., and Vassilvitskii, S. A field guide to personalized reserve prices. In *Proceedings of the 25th international conference on world wide web*, pp. 1093–1102, 2016.
- Parkes, D. C., Kalagnanam, J., and Eso, M. Achieving budget-balance with vickrey-based payment schemes in exchanges. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence, IJCAI'01*, pp. 1161–1168, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc.
- PricewaterhouseCoopers. IAB Internet advertising revenue report: Full year 2019 results & Q1 2020 revenues. https://www.iab.com/wp-content/uploads/2020/05/FY19-IAB-Internet-Ad-Revenue-Report_Final.pdf, 2020. Accessed: 2021-01-20.
- Thompson, D. R. and Leyton-Brown, K. Revenue optimization in the generalized second-price auction. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pp. 837–852, 2013.
- Zhang, Z., Liu, X., Zheng, Z., Zhang, C., Xu, M., Pan, J., Yu, C., Wu, F., Xu, J., and Gai, K. Optimizing multiple performance metrics with deep gsp auctions for e-commerce advertising. *arXiv preprint arXiv:2012.02930*, 2020.

Appendix

A. Omitted Proofs

A.1. Proof of Theorem 3.1

Proof of Theorem 3.1 continued. Recall that

$$\text{util}(\kappa, \beta) = \int_{\beta \cdot q(\kappa)}^{\infty} \left((v - \beta \cdot q(\kappa)) \cdot G(\beta \cdot q(\kappa)) + \int_{\beta \cdot q(\kappa)}^v (v - z) \cdot g(z) dz \right) f(v) dv.$$

As a result, taking the partial derivative of $\text{util}(\kappa, \beta)$ with respect to β , we have

$$\begin{aligned} \frac{\partial \text{util}(\kappa, \beta)}{\partial \beta} &= -q(\kappa) \cdot \left[\left((v - \beta \cdot q(\kappa)) \cdot G(\beta \cdot q(\kappa)) + \int_{\beta \cdot q(\kappa)}^v (v - z) \cdot g(z) dz \right) f(v) \right] \Big|_{v=\beta \cdot q(\kappa)} \\ &\quad + \int_{\beta \cdot q(\kappa)}^{\infty} \frac{\partial \left((v - \beta \cdot q(\kappa)) \cdot G(\beta \cdot q(\kappa)) + \int_{\beta \cdot q(\kappa)}^v (v - z) \cdot g(z) dz \right)}{\partial \beta} \cdot f(v) dv. \end{aligned}$$

Observe that the first term is simply 0, while for the second term, we have

$$\begin{aligned} &\frac{\partial \left((v - \beta \cdot q(\kappa)) \cdot G(\beta \cdot q(\kappa)) + \int_{\beta \cdot q(\kappa)}^v (v - z) \cdot g(z) dz \right)}{\partial \beta} \\ &= -q(\kappa) \cdot G(\beta \cdot q(\kappa)) + (v - \beta \cdot q(\kappa)) \cdot q(\kappa) \cdot g(\beta \cdot q(\kappa)) - q(\kappa) \cdot (v - \beta \cdot q(\kappa)) \cdot g(\beta \cdot q(\kappa)) \\ &= -q(\kappa) \cdot G(\beta \cdot q(\kappa)). \end{aligned}$$

Thus, we have that

$$\frac{\partial \text{util}(\kappa, \beta)}{\partial \beta} \Big|_{\beta=1} = \int_{q(\kappa)}^{\infty} -q(\kappa) \cdot G(q(\kappa)) \cdot f(v) dv = -(1 - \kappa) \cdot q(\kappa) \cdot G(q(\kappa)),$$

which concludes the proof. \square

A.2. Proof of Theorem 4.2

Proof of Theorem 4.2 Continued. For convenience, let $h(\kappa, \beta) = \exp(\theta \cdot \tau_{\lambda}(\mathbf{w}(\kappa, \beta)))$ and $Z = \int_0^1 h(\kappa, \beta) d\kappa$. As a result, we have that

$$\mathbb{E}_{\kappa \sim M_F^{\beta}} [\text{util}(\kappa, 1)] = \frac{\int_0^1 \text{util}(\kappa, 1) \cdot \exp(\theta \cdot \tau_{\lambda}(\mathbf{w}(\kappa, \beta))) d\kappa}{\int_0^1 \exp(\theta \cdot \tau_{\lambda}(\mathbf{w}(\kappa, \beta))) d\kappa} = \frac{\int_0^1 \text{util}(\kappa, 1) \cdot h(\kappa, \beta) d\kappa}{Z}.$$

Notice that the partial derivative of $h(\kappa, \beta)$ with respect to β is

$$\frac{\partial h(\kappa, \beta)}{\partial \beta} = \theta \cdot (\mathbf{w}(\kappa, \beta))^{\lambda-1} \cdot \frac{\partial \mathbf{w}(\kappa, \beta)}{\partial \beta} \cdot h(\kappa, \beta).$$

Therefore, with a few steps of algebraic manipulations, we have

$$\begin{aligned}
 & \frac{\partial \mathbb{E}_{\kappa \sim M_F^\beta} [\text{util}(\kappa, 1)]}{\partial \beta} \\
 &= \theta \cdot \frac{\int_0^1 \text{util}(\kappa, 1) \cdot (w(\kappa, \beta))^{\lambda-1} \cdot \frac{\partial w(\kappa, \beta)}{\partial \beta} \cdot h(\kappa, \beta) d\kappa}{Z} \\
 & \quad - \theta \cdot \frac{\left(\int_0^1 \text{util}(\kappa, 1) \cdot h(\kappa, \beta) d\kappa \right) \left(\int_0^1 (w(\kappa, \beta))^{\lambda-1} \cdot \frac{\partial w(\kappa, \beta)}{\partial \beta} \cdot h(\kappa, \beta) d\kappa \right)}{Z^2} \\
 &= \theta \cdot \mathbb{E}_{\kappa \sim M_F^\beta} \left[\text{util}(\kappa, 1) \cdot \frac{\partial w(\kappa, \beta)}{\partial \beta} \cdot (w(\kappa, \beta))^{\lambda-1} \right] \\
 & \quad - \theta \cdot \mathbb{E}_{\kappa \sim M_F^\beta} [\text{util}(\kappa, 1)] \cdot \mathbb{E}_{\kappa \sim M_F^\beta} \left[\frac{\partial w(\kappa, \beta)}{\partial \beta} \cdot (w(\kappa, \beta))^{\lambda-1} \right] \\
 &= \theta \cdot \text{cov}_{\kappa \sim M_F^\beta} \left[\text{util}(\kappa, 1), \frac{\partial w(\kappa, \beta)}{\partial \beta} \cdot (w(\kappa, \beta))^{\lambda-1} \right].
 \end{aligned}$$

We can conclude the proof by plugging in $\beta = 1$. □

B. Omitted Experiment Results

B.1. Tradeoffs with competing bids




Figure 6: Log-normal distribution




Figure 7: Shifted log-normal distribution

Figure: Synthetic data experiment: revenue-incentive tradeoff *with* competing bids under non-linear policies with different λ ; the dashed line corresponds to the line of $\text{DIC} = 1 - \text{REV}_{\text{reserve}}/W$.

We report here on the tradeoff curves when there are competing bids (i.e., a fixed competing bid distribution). First we consider the synthetic-data simulations. For the log-normal distribution, we set the competing bid to be 0.1 and for the shifted log-normal distribution, we set the competing bid to be 1.1. We use the buyer's contribution to the revenue (i.e., her expected payment) as the weight function, which is no longer linear in the bid-shading factor β given the existence of the competing bid. Figure 6 and 7 show the tradeoffs: the trends for both distributions are similar to the setting without competing bids, but the curve for $\lambda = 0$ deviates from the diagonal line.

Tradeoff curves based on real empirical bid distributions from an ad exchange are given in Figure 8 and Figure 9. The patterns are similar to the single-bidder case: smaller (negative) values of λ yield better tradeoffs, and their confidence intervals for the dynamic-IC metric are also narrower. Because of the presence of competing bids, the curve for the power mechanism no longer tracks the diagonal line, as observed in the synthetic-data experiments. Distributions on day 1 and day 2 are not identical, but similar enough that the relative rankings of the mechanisms remain identical.

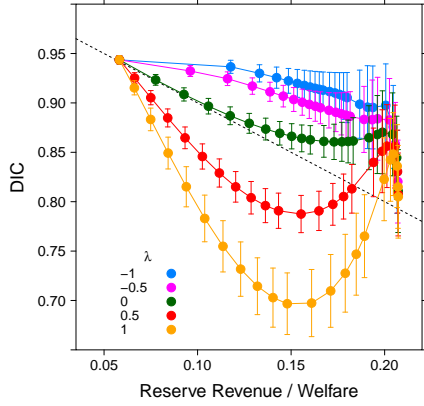


Figure 8: Tradeoff on training data.

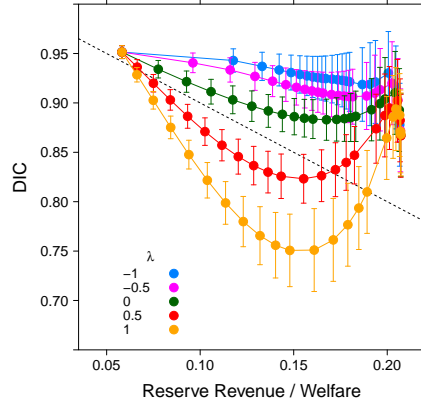


Figure 9: Tradeoff on testing data.

Figure: Real data experiment: revenue-incentive tradeoff *with* competing bids under non-linear policies with different λ ; the dashed line corresponds to the line of $DIC = 1 - REV_{\text{reserve}}/W$.

B.2. Pushing λ further




Figure 10: Log-normal distribution




Figure 11: Shifted log-normal distribution

Figure: Synthetic data experiment: revenue-incentive tradeoff *without* competing bids under non-linear policies with large λ ; the dashed line corresponds to the line of $DIC = 1 - REV_{\text{reserve}}/W$.

One may wonder whether we can push λ further to get an even better tradeoff. Figure 10 and Figure 11 show the experiment for λ with large absolute values without competing bids. For the shifted log-normal distribution, Figure 11 shows that for $\lambda = 7.5$ and $\lambda = 9$, the dynamic-IC metric could be above 1, indicating an incentive of bid raising for the buyer. Therefore, one must be cautious with the choice of λ to balance the bid-raising and bid-shading incentives. As for the log-normal distribution, Figure 10 shows that it is not true that the IC metric shares the same trend of monotonicity for all revenue performance, even under the same distribution. As a result, it suggests that, one needs to take both the revenue performance and the underlying distribution into account when choosing λ .

We then conduct the experiments with competing bids for λ with large absolute values. Again, for the log-normal distribution, we set the competing bid to be 0.1 and for the shifted log-normal distribution, we set the competing bid to be 1.1. Figure 12 and 13 show the tradeoffs: the trends for both distributions are similar to the setting without competing bids, but the curve for $\lambda = 0$ deviates from the diagonal line.




Figure 12: Log-normal distribution

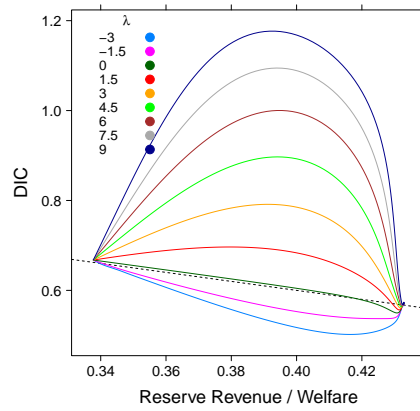


Figure 13: Shifted log-normal distribution

Figure: Synthetic data experiment: revenue-incentive tradeoff *with* competing bids under non-linear policies with large λ ; the dashed line corresponds to the line of $DIC = 1 - REV_{reserve}/W$.