Dynamic Mechanism Design in the Field [*]

Vahab Mirrokni[1], Renato Paes Leme[1], Rita Ren[1], and Song Zuo[†2]

[1]Google, New York City, NY [‡]
[2]Tsinghua University, Beijing, China [§]

February 21, 2018

Abstract

Dynamic mechanisms are a powerful technique in designing revenue-maximizing repeated auctions. Despite their strength, these types of mechanisms have not been widely adopted in practice for several reasons, e.g., for their complexity, and for their sensitivity to the accuracy of predicting buyers' value distributions. In this paper, we aim to address these shortcomings and develop simple dynamic mechanisms that can be implemented efficiently, and provide theoretical guidelines for decreasing the sensitivity of dynamic mechanisms on prediction accuracy of buyers' value distributions. We prove that the dynamic mechanism we propose is provably dynamic incentive compatible, and introduce a notion of buyers' regret in dynamic mechanisms, and show that our mechanism achieves bounded regret while improving revenue and social welfare compared to a static reserve pricing policy. Finally, we confirm our theoretical analysis via an extensive empirical study of our dynamic auction on real data sets from online adverting. For example, we show our dynamic mechanisms can provide a +17% revenue lift with relative regret less than 0.2%.

1 Introduction

The majority of online advertising platforms run a sequence of repeated auctions to sell their inventory of page-views. While much of the existing literature for such ad auctions discuss static one-shot auctions, using dynamic auctions optimized across different time periods could potentially bring significant gains both in terms of revenue and social welfare. The power of dynamic mechanisms has been observed by a number of recent papers [Bergemann and Välimäki, 2002; Parkes and Singh, 2004; Cavallo, 2008; Athey and Segal, 2013; Kakade *et al.*, 2013; Pai and Vohra, 2013; Pavan *et al.*, 2014; Devanur *et al.*, 2015; Chawla *et al.*, 2016]. We refer to the survey by Bergemann and Said [2011] for a comprehensive treatment on the subject.

Even though dynamic mechanisms can be much more effective in maximizing revenue and social welfare, they have not been widely adopted in practice. The main issues therein are partially the shortcomings of dynamic mechanisms: One major issue is the complexity induced by the exponentially growing design space, making it difficult to solve or even to describe such mechanisms. Another issue is that the incentives under dynamic environments rely on the consistency between the seller and buyers on the impact of current actions on future outcomes, while such consistency requires a higher order common knowledge assumption that generally does not hold in practice. For example, the accuracy of the seller's prediction of the buyer's valuations plays a key role in the level of dynamic incentive compatibility ensured by the mechanism; the less accurate the seller's belief about buyer's valuation, the more incentive constraints are violated. This becomes more challenging in high-dimensional settings where learning buyers' value distributions is very hard, and the estimates may be noisy.

Recently, a series of work has made significant progress on the first issue of complexity described above. For example, Ashlagi *et al.* [2016] and Mirrokni *et al.* [2016c] show that an ϵ-approximate optimal dynamic mechanism

---

can be efficiently computed via a dynamic program. However, such an approach relies on the forecast of the sequence of future items and then solves large linear programs recursively via backward induction. Mirrokni *et al.* [2016a] introduce the so-called *bank account mechanisms* that have simple structure and can achieve the optimal revenue with ex-ante individual rationality in single buyer cases. Later on, the authors extend the bank account framework to guarantee ex-post individual rationality and show a simple randomized construction that 5-approximates the optimal revenue for general multiple buyer cases [Mirrokni *et al.*, 2016b]. Balseiro *et al.* [2017] consider a similar single buyer setting with i.i.d. values over the items across different periods and design simple and deterministic mechanisms that approach the optimal revenue in the limit.

Despite all these improvements made on the complexity issue, the existing approaches above for multiple buyer cases still require randomized allocation rules[1] that are undesirable in real applications. More importantly, none of them has discussed the unavailability of the accurate prior knowledge on the distributions of buyer values in practice, yet all of them adopt the dynamic incentive notion that is sensitive to the prediction errors on the buyer values.

Impacts of Learning on Incentives   A central problem in the intersection of game theory and machine learning is how to learn optimal mechanism from data (see for example the papers by Morgenstern and Roughgarden [2015, 2016]). For static mechanism design (the case studied in previous papers), errors in estimating distributions can hurt revenue but don't hurt incentive constraints. In dynamic mechanism design, the estimation error also harms incentive compatibility. To the best of our knowledge, we are the first paper to formally define and estimate a measure of the robustness of the mechanism with respect to learning errors (buyer regret). This is particularly important since dynamic mechanisms offer a promise of quite significant improvements in both revenue and allocation efficiency (both in theory and in data).

Our Contribution   Firstly, we propose a family of special bank account mechanisms, coined *Dynamic Second Price Auctions (DSP)*, for general multiple buyer settings, which addresses the practical issues raised in the previous paragraph. The auctions have a very common auction format (second price auctions with reserves) except that the reserves are adjusted dynamically depending on the history of the auction.[2] Our main goal in this paper is *not* to design the revenue optimal auction but instead to offer a pragmatic design of dynamic mechanisms that can be implemented in practice and is easy to switch from a second price auction with reserves, which is a default auction run by most advertising exchanges.

We assume a baseline static second price auction with personalized reserves. We are agnostic to the nature of those reserves, they can be either the revenue optimal reserves or chosen with a different goal, e.g., to maximize a combination of revenue and welfare. Our main theorem shows how to process existing reserves and produce an auction that is dynamic incentive compatible (Theorem 3.1) and strictly improves over the baseline auction both in terms of revenue and social welfare (Proposition 3.2). Moreover, we confirm on data that this improvement is quite significant (Section 5).

The idea of the DSP is as follows: in static second price auctions, reserves improve revenue at the expense of efficiency. DSP first improves the efficiency by introducing dynamic discounts on the reserves based on how much utility the buyers accumulated in past periods. DSP then also introduces an extra payment (termed "bank account spent") that gives the buyer the right to participate in the auction with lower reserves. Such payments will compensate for the loss in revenue by lowering reserves. The discount policy and the accompanying fees should be carefully calculated not to harm incentive compatibility and individual rationality.

Secondly, we introduce the notion of *buyer regret* (or simply regret) to quantify how much the dynamic incentive compatibility is violated (Section 4). Intuitively, the regret of a buyer means how much the buyer can increase his/her utility by rejecting all the dynamic discounts. We import the *bank account limits* studied by Mirrokni *et al.* [2016a] to enable a trade-off between revenue lifts and buyer regrets: the regrets can be bounded by the product of bank account limits and prediction errors (Theorem 4.2). Such insights provide theoretical guidelines to reduce the regret.

Finally, we conduct empirical studies to justify our insights on real data from online ad auctions (Section 5). To the best of our knowledge, this is the first work that examines the power of simple dynamic auctions on real data sets. In the first part, we simulate DSPs with different baseline reserves and dynamic discounts and show how the revenue lifts, regrets, and buyer utilities change as bank account limits increase. In particular, the revenue lift could be as much as +150% comparing to the second price auctions with monopoly reserves [Hartline and Roughgarden, 2009]

---

[1]The dynamic mechanisms by Balseiro *et al.* [2017] and Mirrokni *et al.* [2016a] are deterministic but both restricted to single buyer cases.

[2]Although the idea is ported from Mirrokni *et al.* [2016a,b], the mechanism is carefully redesigned to meet the practical requirements (Section 3).

2

while incurring relative regret 13% (of the corresponding buyer utility), or +17% revenue lift with relative regret less than 0.2%. In the second part, we compare the statistics for the same auctions but with empirical distributions of different levels of accuracy (including simulations on synthetic data generated from known distributions, so we have access to the ground truth distributions in this case) and observe that the regrets decrease significantly while keeping the revenue lifts almost unchanged. Combining the observations from both parts, we conclude that the guidelines from Theorem 4.2 do help in reducing the regrets.

2 Preliminaries

We study a setting where a seller repeatedly interacts with n buyers selling one item per period over T periods. The value of each agent $i \in [n]$ for the item in period $t \in [T]$ is $v_t^i \in \mathcal{V}$. If $x_t^i \in [0, 1]$ represents the probability that agent i is allocated the item in period t, his utility is $v_t^i \cdot x_t^i$.

Throughout this paper, (i) we use superscripts as the indices of agents, bold fonts for vectors of all agents (i.e., $\mathbf{a} = (a^1, \ldots, a^n)$), and superscript $-i$ for the vector except the i-th element (i.e., $a^{-i} = (a^1, \ldots, a^{i-1}, a^{i+1}, \ldots, a^n)$); (ii) we use subscripts as the indices of periods and $a_{1..t}$ to denote the sequence $a_1, \ldots, a_t$.

The values v_t^i are assumed to be drawn from independent[3] distributions F_t^i. The distributions F_t^i are assumed to be common knowledge but the realizations of the random variables are initially unknown for both the agents and the designer. At each period, the following events happen:

1. each agent i learns his value $v_t^i \sim F_t^i$;

2. each agent i reports value $\hat{v}_t^i$ to the designer;

3. the designer implements an outcome $\boldsymbol{x}_t \in [0, 1]^n$ and charges the agents $\boldsymbol{p}_t$;

4. each agent accrues utility $u_t^i = v_t^i \cdot x_t^i - p_t^i$.

A dynamic mechanism can then be described in terms of a pair of maps for each period, which associate the history of reports $\hat{\boldsymbol{v}}_{1..t} = (\hat{\boldsymbol{v}}_1, \hat{\boldsymbol{v}}_2, \ldots, \hat{\boldsymbol{v}}_t)$ to an outcome $\boldsymbol{x}_t$ and payment $\boldsymbol{p}_t$, i.e., $\boldsymbol{x}_t : \mathcal{V}^{nt} \to [0, 1]^n$ and $\boldsymbol{p}_t : \mathcal{V}^{nt} \to \mathbb{R}^n$. Therefore we can define:

$$u_t^i(v_t^i; \hat{\boldsymbol{v}}_{1..t}) = v_t^i \cdot x_t^i(\hat{\boldsymbol{v}}_{1..t}) - p_t^i(\hat{\boldsymbol{v}}_{1..t})$$

If $\boldsymbol{x}_t$ and $\boldsymbol{p}_t$ are only functions of the valuations $\boldsymbol{v}_t$ reported in time t, the mechanism is static.

2.1 Dynamic Incentive Compatibility

A mechanism is incentive compatible if it provides incentives for agents to reveal their true types in each iteration. Such conditions for dynamic mechanisms can be easily defined by backward induction: in the last period, regardless of the history so far and other agents' reports, it should be incentive compatible for each agent to report his true value. This corresponds to the usual notion of incentive compatibility in (static) mechanism design:

$$v_T^i = \arg\max_{\hat{v}_T^i} u_T^i(v_T^i; \hat{\boldsymbol{v}}_{1..T})$$

for all $i, \hat{\boldsymbol{v}}_{1..T-1}, \hat{v}_T^{-i}, v_T^i$. To simplify notations, from now on we will omit the "for-all" quantification and assume all expressions are quantified as "for-all" in its free variables. For the next-to-last-period, it should be incentive compatible for the agent given that he will report truthfully in the following period:

$$v_{T-1}^i = \arg\max_{\hat{v}_{T-1}^i} \left\{ u_{T-1}^i(v_{T-1}^i; \hat{\boldsymbol{v}}_{1..T-1}) + \mathbb{E}_{v_T^i} \left[u_T^i(v_T^i; \langle \hat{\boldsymbol{v}}_{1..T-1}, v_T^i, \hat{v}_T^{-i} \rangle) \right] \right\}.$$

Proceeding by backward induction for all periods, we require that:

$$v_t^i = \arg\max_{\hat{v}_t^i} u_t^i(v_t^i; \hat{\boldsymbol{v}}_{1..t}) + U_t^i(\hat{v}_{1..t}^i | \hat{v}_{1..T}^{-i}) \tag{DIC}$$

---

[3]In practice, it is still fair to assume that the values of the agents (buyers) are independent conditional on each particular inventory and cookie. If not independent, the weak version of truthfulness still holds, i.e., truthful if all other buyers bid truthfully.

where the second term is the *continuation utility*, i.e., the expected utility obtained from the subsequent periods of the mechanism assuming the agent reports truthfully:

$$U_t^i(\hat{v}_{1..t}^i | \hat{v}_{1..T}^{-i}) := \mathop{\mathbb{E}}_{v_{t+1..T}^i} \left[\sum_{\tau=t+1}^{T} u_\tau^i(v_\tau^i; \langle \hat{\boldsymbol{v}}_{1..t}, v_{t+1..\tau}^i, \hat{v}_{t+1..\tau}^{-i} \rangle) \right]$$

A well-known fact in dynamic mechanism design is that DIC implies that agent i's expected overall utility $U_0^i(\hat{v}_{1..T}^{-i})$ is maximized by reporting truthfully in each period.

2.2 Ex-Post Individual Rationality

Another desirable constraint is ex-post individual rationality which says that an agent should derive non-negative utility from the mechanism for every realization of the values:

$$\sum_{t=1}^{T} u_t^i(v_t^i; \boldsymbol{v}_{1..t}) \geq 0 \tag{eP-IR}$$

We will focus on the problem of maximizing revenue subject to DIC, eP-IR, and feasibility constraints:

$$\max \quad \text{Rev} = \mathbb{E} \left[\sum_{t=1}^{T} \sum_{i=1}^{n} p_t^i(\boldsymbol{v}_{1..t}) \right]$$

$$\text{s.t.} \quad \text{(DIC), (eP-IR), and feasibility:} \sum_{i=1}^{n} x_t^i(\boldsymbol{v}_{1..t}) \leq 1$$

2.3 Bank Account Mechanisms

The space of mechanisms satisfying DIC and eP-IR is very broad and unstructured. We restrict our attention in this section to a subclass of dynamic mechanisms introduced by Mirrokni *et al.* [2016a] called *bank account mechanisms*. The mechanisms are simple, dynamic incentive compatible by design and have the notable feature that for any distribution, there is a revenue-optimal bank account mechanisms. More precisely, for any dynamic mechanism satisfying DIC and eP-IR there is a bank account mechanism with at least the same welfare and revenue.

Bank account mechanisms keep a state for each buyer, which is a scalar called *balance*. Each period depends on the previous periods through the vector of buyer balances. Another main feature is that in this framework, the designer needs to specify single-period auctions that are single-period incentive compatible together with a valid balance update policy. That is, once a valid balance update policy is in place, all the designer needs to worry about are single-period incentive compatibility constraints.

A bank account mechanism B is defined in terms of the following functions for each period:

- A static single-period auction parametrized by a balance vector $\boldsymbol{b} \in \mathbb{R}_+^n$, $x_t^B(\boldsymbol{v}_t, \boldsymbol{b})$, $p_t^B(\boldsymbol{v}_t, \boldsymbol{b})$, that is (single-period) *incentive-compatible* for each $\boldsymbol{b}$, i.e.:

$$v_t^i \cdot x_t^{B,i}(\boldsymbol{v}_t, \boldsymbol{b}) - p_t^{B,i}(\boldsymbol{v}_t, \boldsymbol{b}) \geq v_t^i \cdot x_t^{B,i}(\langle \hat{v}_t^i, v_t^{-i} \rangle, \boldsymbol{b}) - p_t^{B,i}(\langle \hat{v}_t^i, v_t^{-i} \rangle, \boldsymbol{b}). \tag{IC}$$

  Note that we do not require the mechanism to be (single-period) individually rational. We also require the utility of the agent to be *balance independent in expectation*, i.e., that:

$$\mathop{\mathbb{E}}_{v_t^i \sim F_t^i} \left[v_t^i \cdot x_t^{B,i}(\boldsymbol{v}_t, \boldsymbol{b}) - p_t^{B,i}(\boldsymbol{v}_t, \boldsymbol{b}) \right] \text{ is a non-negative constant not depending on } \boldsymbol{b}. \tag{BI}$$

- A balance update policy $b_t^B(\boldsymbol{v}_t, \boldsymbol{b})$ which maps the previous balances and the reports to the current balances, satisfying the following *balance update* conditions:

$$0 \leq b_t^{B,i}(\boldsymbol{v}_t, \boldsymbol{b}) \leq b^i + v_t^i \cdot x_t^{B,i}(\boldsymbol{v}_t, \boldsymbol{b}) - p_t^{B,i}(\boldsymbol{v}_t, \boldsymbol{b}). \tag{BU}$$

Given the balance update functions, we can define $\boldsymbol{b}_t : \mathcal{V}^t \to \mathbb{R}^n_+$ recursively as:

$$\boldsymbol{b}_0 = \boldsymbol{0} \quad \text{and} \quad \boldsymbol{b}_1(\boldsymbol{v}_1) = \boldsymbol{b}_1^B(\boldsymbol{v}_1, \boldsymbol{0}) \quad \text{and} \quad \boldsymbol{b}_t(\boldsymbol{v}_{1..t}) = \boldsymbol{b}_t^B(\boldsymbol{v}_t, \boldsymbol{b}_{t-1}^B(\boldsymbol{v}_{1..t-1})).$$

which allows us to define a dynamic mechanism in the standard sense as:

$$\boldsymbol{x}_t(\boldsymbol{v}_{1..t}) = \boldsymbol{x}_t^B(\boldsymbol{v}_t, \boldsymbol{b}_{t-1}), \quad \boldsymbol{p}_t(\boldsymbol{v}_{1..t}) = \boldsymbol{p}_t^B(\boldsymbol{v}_t, \boldsymbol{b}_{t-1}).$$

In what follows we will abuse notations by dropping the superscript B and refer to $\boldsymbol{x}_t(\boldsymbol{v}_{1..t})$ and $\boldsymbol{x}_t(\boldsymbol{v}_t, \boldsymbol{b}_{t-1})$ interchangeably. One important theorem from previous studies is that any bank account mechanism satisfies the stronger notions of DIC and eP-IR.

Theorem 2.1 (Mirrokni *et al.* [2016b]). *Any bank account mechanism satisfying IC, BI, and BU is dynamic incentive compatible* (DIC) *and ex-post individually rational* (eP-IR).

2.4 Limits on Bank Accounts

The balance of each buyer is a non-negative variable b_i that changes over time. Mirrokni *et al.* [2016a] show that the balance can be used as a measure of inter-period dependency. One extreme are mechanisms that always keep balance at zero for all buyers. Since bank account mechanisms can only depend on previous periods via balance, if balances are constant equal to zero the mechanism must run independent static auctions in each period.

Based on this observation, the authors propose the design of bank account mechanisms with a limit as a way to limit inter-period dependencies. Formally, the limit L^i on the bank account of buyer i means that a bank account mechanism satisfies an additional bank account limit constraint,

$$\forall t, \ 0 \leq b_t^i \leq L^i. \tag{BL}$$

Later in Section 4, we show how bank account limits could help on establishing a trade-off between the DIC guarantee and the revenue lifts under the presence of prediction errors on buyer values.

3 Dynamic Second Price Auctions

Mirrokni *et al.* [2016b] propose simple auctions following the bank account framework that are provably near-optimal, but still rely on randomized allocations for the revenue guarantee. The auction formats in each period are also non-standard. In contrast, most auctions used in practice are variations over second price auctions with reserves.

Our focus on this paper is to design a family of practical dynamic mechanisms that could be easily implemented in practice using the current infrastructure for running second price auctions with reserves. Our auction will be still based on the bank account framework and the twist with respect to usual second price auctions is that the reserves will be dynamic and adjusted based on reports from past periods.

Before we describe our dynamic auction, we recall the two implementations of second price auctions with personalized reserve [Dhangwatnotai *et al.*, 2015]. Given reported values v_i for each agent and reserves r_i the following auctions are single-period incentive compatible:

- *Eager reserves:* The winner is the highest bidder such that $v_i \geq r_i$ and his payment is the maximum between his reserve and the bid of the second highest buyer above reserves. If no buyer is above his reserve, there is no winner.

- *Lazy reserves:* We select the buyer with the highest bid. If his bid is above reserve he is the winner and his payment is the maximum of his reserve and the second highest bid (regardless of whether the second highest is below or above reserve). If the selected buyer is below his reserve, there is no winner.

They differ in the ordering of filtering and ranking. In the eager we first filter buyer below reserve and then rank them. In the lazy we first rank buyers based on their bid and then we filter the top buyer based on his reserve. The auction we propose is a hybrid where we have for each buyer a vector of eager reserves r_t^E and a vector of lazy reserves r_t^L with $r_t^{i,E} \leq r_t^{i,L}$. Our auction will have two filtering phases, one before the winner selection and one after:

- Step 1: run a second price auction with reserves r_t^E, tie-breaking lexicographically. Let agent i be the winner and use w_t^i to denote the second highest bid.

- Step 2: for the winner i, run an additional posted price auction at $r_t^{i,L} \geq r_t^{i,E}$. If $v_t^i \geq r_t^{i,L}$, agent i wins the item and pay $\max\{r_t^{i,L}, w_t^i\}$.

We will make the auction dynamic by adjusting the lazy reserves $r_t^{i,L}$ dynamically over time as a function of the balance b_{t-1} accumulated in previous steps. Besides the payment defined above, each agent will be charged an additional amount of money (paid by his/her bank account), which we call *spent*, denoted by s_t^i. The intuition is that we will charge an upfront fee for the buyer in each step in exchange for a more favorable lazy reserve.

3.1 Auction Description

Our goal in this paper is to design practical dynamic mechanism that are competitive against simple static second price auctions with reserves. We are agnostic to how reserve prices are computed and we assume instead there is a baseline auction applying reserve price $\bar{r}_t^i$ in each iteration. Those could be the revenue-optimal reserves or could be reserves that are designed, for example, to optimize revenue subject to a constraint on the welfare loss. Our goal will be to process those reserves and come up with a dynamic auction that improves the original one in terms of both welfare and revenue.

Our auction will take as input for each buyer i and step t a baseline reserve price $\bar{r}_t^i$, distributions F_t^i and bank account limts L^i. Based on $\bar{r}_t^i$ and F_t^i we will have a pre-processing step that generates reserve lower bounds $\underline{r}_t^i \leq \bar{r}_t^i$. In Section 3.2 we will discuss which properties are required from the pre-processing step. For now, we just assume the DSP auction has access to a range of dynamic reserves $[\underline{r}_t^i, \bar{r}_t^i]$ for each i and t.

Dynamic Second Price Auction (DSP)
Input: F_t^i and $[\underline{r}_t^i, \bar{r}_t^i]$ for each i, t and L^i for each i.
Initialize balance of all buyers to zero $b_0^i = 0$ each timestep t:

1. Agents report their value v_t^i and let j be the agent with the highest bid above $\underline{r}_t^j$.

2. Let w_t^i be the highest bid from other agents $i' \neq i$ with bids $v_t^{i'}$ above $\underline{r}_t^{i'}$ or zero in case there is no other agent i' bids $v_t^{i'}$ above $\underline{r}_t^{i'}$. (Or equivalently, w_t^i is the winning bid from step 1, if agent i is removed from this auction.)

3. We use the balance b_{t-1}^i to compute a dynamic lazy reserve $r_t^i \in [\underline{r}_t^i, \bar{r}_t^i]$. Let

$$\bar{s}_t^i = \int_{\max(\underline{r}_t^i, w_t^i)}^{\max(\bar{r}_t^i, w_t^i)} \left(1 - F_t^i(v)\right) dv$$

   If $b_{t-1}^i \geq \bar{s}_t^i$ then we set: $r_t^i = \underline{r}_t^i$ and $s_t^i = \bar{s}_t^i$.
   Otherwise, we set $s_t^i = b_{t-1}^i$ and find r_t^i such that:

$$b_{t-1}^i = \int_{\max(r_t^i, w_t^i)}^{\max(\bar{r}_t^i, w_t^i)} \left(1 - F_t^i(v)\right) dv.$$

4. If $v_t^j \geq r_t^j$, we allocate the item to j and set $x_t^j = 1$. For other $i \neq j$, $x_t^i = 0$.

5. Charge buyer i the amount $p_t^i = x_t^i \max(w_t^i, r_t^i) + s_t^i$.

6. Update the bank account to

$$b_t^i = \min\left(b_{t-1}^i + v_t^i x_t^i - p_t^i, L^i\right).$$

Theorem 3.1. *The DSP defined above satisfies IC, BI, and BU. Hence it is DIC and eP-IR. In addition, the bank limit constraint* (BL) *is also satisfied.*

Proof of Theorem 3.1. According to the definition of DSP, the 2-step auction at each period is incentive compatible, and the additional payment s_t^i is clearly independent of v_t^i. Hence constraint (IC) is satisfied.

By substituting p_t^i into the construction of b_t^i, we get

$$b_{t-1}^i - s_t^i \leq b_t^i \leq b_{t-1}^i + v_t^i \cdot x_t^i(\boldsymbol{v}_t, \boldsymbol{b}_t) - p_t^i(\boldsymbol{v}_t, \boldsymbol{b}_t).$$

Since $s_t^i \leq b_{t-1}^i$, constraint (BU) is satisfied.

We then verify the constraint (BI). Consider the expected utility of agent i from the 2-step auction.

$$\mathbb{E}_{v_t^i}[\hat{u}_t^i] = \mathbb{E}_{v_t^i}\left[[v_t^i - \max\{r_t^i, w_t^i\}]^+\right] = \int_{\max\{r_t^i, w_t^i\}}^{\infty} \left(v - \max\{r_t^i, w_t^i\}\right) dF_t^i(v) = \int_{\max\{r_t^i, w_t^i\}}^{\infty} \left(1 - F_t^i(v)\right) dv, \quad (1)$$

where $[x]^+ := \max\{0, x\}$ and the last equality is from integration by parts. Then by the construction of s_t^i,

$$\mathbb{E}_{v_t^i \sim F_t^i}\left[v_t^i \cdot x_t^{B,i}(\boldsymbol{v}_t, \boldsymbol{b}) - p_t^{B,i}(\boldsymbol{v}_t, \boldsymbol{b})\right] = \mathbb{E}_{v_t^i}[\hat{u}_t^i] - s_t^i = \int_{\max\{\overline{r}_t^i, w_t^i\}}^{\infty} \left(1 - F_t^i(v)\right) dv, \quad (2)$$

which is a non-negative constant not depending on $\boldsymbol{b}$. Thus constraint (BI) is satisfied.

Finally, the bank account limit constraint (BL) is directly implied by the construction of b_t^i. $\qquad\square$

3.2 Revenue and Social Welfare Guarantees

We show that the DSP auction can always generate strictly higher revenue and social welfare than the second price auction with baseline reserves $\overline{\boldsymbol{r}}_{1..T}$, as long as (i) the dynamic lazy reserves $\boldsymbol{r}_t$ and $\underline{\boldsymbol{r}}_t$ are properly designed, (ii) the prior distributions $F_{1..T}^i$ have positive probability densities near the baseline reserves,[4] (iii) and the baseline reserves are neither obviously too high nor too low.

The DSP auction is particularly effective with respect to static second price auctions when bank account balances $\boldsymbol{b}_t$ are sufficiently large such that the dynamic reserves could be set as $\boldsymbol{r}_t = \underline{\boldsymbol{r}}_t$ without breaking the BU constraint. In this case, the allocation in each period is exactly the same as the second price auction with reserves being $\underline{\boldsymbol{r}}_t$, hence higher social welfare than the baseline with larger reserves $\overline{\boldsymbol{r}}_t \geq \underline{\boldsymbol{r}}_t$. The improvement is strict if the allocation probability gets increased by the reduction of reserves, which can be easily achieved via properly chosen $\underline{\boldsymbol{r}}_t$ as long as the baseline reserves are not obviously too low, i.e., the social welfare has not been maximized by the baseline.

Meanwhile, the revenue strictly increases because the expected buyer utilities decrease. For each buyer i, according to constraint BI, his/her expected utility is a constant with respect to the balance $\boldsymbol{b}_{t-1}$, in particular, the same as the expected utility when $b_{t-1}^i = 0$. In this case, the auction is equivalent to the second price auction with reserves $\overline{r}_t^i$ for buyer i and $\underline{r}_t^{-i}$ for others from his/her perspective, because the spent $s_t^i = 0$ ($s_t^i \leq b_{t-1}^i$) and hence the dynamic reserve r_t^i is set equal to $\overline{r}_t^i$ to satisfy constraint BI. Therefore, comparing with the baseline, the expected utility of buyer i is hurt by the reduction of the reserves of other buyers, independent of the spent s_t^i.[5]

If the balances $\boldsymbol{b}_t$ are large, the algorithm will choose lazy reserves $r_t^i = \underline{r}_t^i$ and the social welfare will be an improvement over the social welfare of the baseline auction. If the balances $\boldsymbol{b}_t$ are too low, however, the lazy reserves r_t^i could be closer to $\overline{r}_t^i$ and cause the social welfare to be even worse. Let sW be the welfare of the static auction with reserves $\overline{r}_t^i$ and dW($\boldsymbol{b}_t$) the welfare of the dynamic auciton with balances $\boldsymbol{b}_t$. Then we define:

$$\text{lift} = \max_{\boldsymbol{b}} \text{dW}(\boldsymbol{b}) - \text{sW} \qquad \text{loss} = [\text{sW} - \min_{\boldsymbol{b}} \text{dW}(\boldsymbol{b})]^+$$

Now, if we can guarantee that the balance is high most of the time, then the welfare lift will be closer to lift than $-$loss.

The balance b_t^i is a random walk in the interval $[0, L^i]$, where:

$$\mathbb{E}[b_{t+1}^i - b_t^i] = \mathbb{E}\left[x_i^t(v_t^i - \max\{\overline{r}_t^i, w_t^i\})\right]^+ > 0,$$

---

[4]In fact, the assumption here could be even weaker, as long as the probability that the buyer value v_t^i falls in $[\underline{r}_t^i, \overline{r}_t^i]$ is positive, i.e., $F_t^i(\underline{r}_t^i) < F_t^i(\overline{r}_t^i)$.

[5]The readers might wonder if it is possible to design a DSP that generates higher buyer utility at the same time. As we will show in the Section 5 by simulation, it can be done by using slightly lower baseline reserves instead of the given ones.

and $b^i_{t+1} - b^i_t \geq -s^i_t \geq -(\bar{r}^i_t - \underline{r}^i_t) \geq -\delta$ for $\delta = \max_i\{\bar{r}^i_t - \underline{r}^i_t\}$. By standard arguments from stochastic processes, the balance is far from L^i with probability exponentially decreasing in $\mathbb{E}[b^i_{t+1} - b^i_t]/\delta$, therefore:

$$\mathbb{E}[dW(\boldsymbol{b}) - sW] = \text{lift} + \exp(-1/\delta)(\text{lift} - \text{loss})$$

which is positive for sufficiently small δ. In other words, the social welfare is guaranteed to be strictly increased in a DSP, when $\underline{r}^i_t$ is sufficiently close to $\bar{r}^i_t$.

To summarize, we have the following proposition.

Proposition 3.2 (Revenue & social welfare guarantees). *For any given baseline reserves $\bar{\boldsymbol{r}}_t$, we can always design a DSP that strictly outperforms the second price auction with reserves being $\bar{\boldsymbol{r}}_t$ both in terms of revenue and social welfare, except for the following extreme cases:*

- *the given baseline reserves are obviously too high or too low,*

$$\forall i, \ \Pr[v^i_t > \bar{r}^i_t] = 0 \ \text{or} \ \Pr[v^i_t < \bar{r}^i_t] = 0;$$

- *the probability density near $\bar{r}^i_t$ is zero, i,e,*

$$\forall i, \ \exists \delta > 0, \ \Pr[\bar{r}^i_t - \delta \leq v^i_t < \bar{r}^i_t] = 0.$$

4 Regret and Bank Account Limits

In the theoretical analysis we made in previous sections, some assumptions are not necessarily satisfied in practice. In particular, the ground truth prior distributions F^i_t are unavailable in general. Instead, empirical estimations $\tilde{F}^i_t$ are commonly used in revenue maximization. In this case, the DIC constraint is no longer guaranteed since the spent s^i_t, which is part of the payment, is potentially miscalculated due to the inconsistency between F^i_t and $\tilde{F}^i_t$.

4.1 Regret

In order to empirically justify the incentive guarantees of dynamic mechanisms, we introduce the notion of buyer *regret*, which measures the utility gain of a buyer by disabling his/her bank account (i.e., deposits $d^i_t \equiv 0$). For fixed actions of other buyers, the sequence of items, and the valuations of buyer i, the regret of buyer i refers to his/her utility change from the dynamic mechanism to a copy of this mechanism with *the dynamic components disabled for him/her only*. By disabling the dynamic components for a certain buyer, we mean his/her bank account is disabled, i.e., the balance is fixed to be zero. Hence his/her reserve price of the t-th item will be fixed at $r^i_t = \bar{r}^i_t$, independent of the past periods.

One practical feature of this mechanism is that by disabling deposits for one buyer, i.e. $d^i_t = 0$, the buyer will face a static auction, in the sense that his/her bids in one iteration don't affect his allocation probability or pricing in subsequent rounds. This provides a smooth transition where buyers are allow to opt-in to participate in the dynamic mechanism or otherwise keep participating in a static auction. If we have perfect access to the distributions of buyers, it should be indifferent for the buyer's utility to opt-in to the dynamic component. However, if the priors are not perfectly calculated, it is possible for buyers to regret the decision to participate in a dynamic auction. In this section we formally define and bound the notion of buyer regret in a dynamic mechanism.

Definition 4.1 (Regret). *For any fixed sequence of items and valuations, the regret of buyer i from a DSP M is defined as the difference of his/her utilities from M (denoted by $\tilde{U}_{TL}^i(\boldsymbol{v}_{1..T})$) and another DSP M' (denoted by $\tilde{U}_{TL}'^i(\boldsymbol{v}_{1..T})$) that is the same as M except that the dynamic lazy reserves for buyer i are fixed, i.e., $r'^i_t = \bar{r}^i_t$. Formally,*

$$R^i(\boldsymbol{v}_{1..T}) = \tilde{U}_{TL}'^i(\boldsymbol{v}_{1..T}) - \tilde{U}_{TL}^i(\boldsymbol{v}_{1..T}).$$

4.2 Regret, Limits, and Prediction errors

As we mentioned previously, the bank account limits can upper bound the expected regrets under the presence of prediction errors.

In particular, the following theorem quantifies an upper bound on the absolute value of the expected regret by the prediction errors and bank account limits.

Theorem 4.2 (Regret bounds).

$$\left|\mathbb{E}[R^i]\right| \le T \cdot limit^i \cdot prediction\text{-}error^i \le TL^i \max_{t,\,r' \in [\underline{r}^i_t, \overline{r}^i_t]} \left| 1 - \frac{\int_{r'}^{\overline{r}^i_t}(1 - F^i_t(v))\mathrm{d}v}{\int_{r'}^{\overline{r}^i_t}(1 - \tilde{F}^i_t(v))\mathrm{d}v} \right|$$

Let us start with the utility that buyer i obtains from period t of mechanism M using estimation $\tilde{F}^i_t$:

$$\tilde{u}^i_t(v^i_t; \boldsymbol{v}_t, \tilde{\boldsymbol{b}}_t) = v^i_t \cdot \tilde{x}^i_t(\boldsymbol{v}_t, \tilde{\boldsymbol{b}}_t) - \tilde{p}^i_t(\boldsymbol{v}_t, \tilde{\boldsymbol{b}}_t) = [v^i_t - \max\{\tilde{r}^i_t, w^i_t\}]^+ - \tilde{s}^i_t(\tilde{\boldsymbol{b}}_t), \tag{3}$$

where the˜sign indicates the terms that might be influenced by the estimation $\tilde{F}^i_t$:

- $\tilde{\boldsymbol{b}}_t$ is the balance vector achieved by M at the beginning of period t,
- $\tilde{r}^i_t$ is the dynamic reserve price of buyer i in period t,
- and $\tilde{s}^i_t(\tilde{\boldsymbol{b}}_t)$ is the spent of buyer i in period t.

Note that the highest bid from other agents, w^i_t, is independent of both the balance vector $\tilde{\boldsymbol{b}}_t$ and the estimation $\tilde{F}^i_t$.

The utility from mechanism M' can be given as $\tilde{u}^i_t(v^i_t; \boldsymbol{v}_t, \tilde{\boldsymbol{b}}'_t)$, where the only difference with the utility from M is that the bank account of buyer i is disabled, meaning that (i) the dynamic reserve is fixed at $\overline{r}^i_t$ and (ii) the spent is always 0, i.e.,

$$\tilde{u}^i_t(v^i_t; \boldsymbol{v}_t, \tilde{\boldsymbol{b}}'_t) = [v^i_t - \max\{\overline{r}^i_t, w^i_t\}]^+. \tag{4}$$

The following lemma captures the key insight to prove Theorem 4.2:

Lemma 4.3. *The expected regret of buyer i from each single period t only comes from the computation error of the spent caused by the imperfect estimation $\tilde{F}^i_t$, i.e.,*

$$\mathbb{E}_{v^i_t \sim F^i_t}\left[\tilde{u}^i_t(v^i_t; \boldsymbol{v}_t, \tilde{\boldsymbol{b}}'_t) - \tilde{u}^i_t(v^i_t; \boldsymbol{v}_t, \tilde{\boldsymbol{b}}_t)\right] = \tilde{s}^i_t(\tilde{\boldsymbol{b}}_t) - s^i_t(\tilde{\boldsymbol{b}}_t).$$

Proof of Lemma 4.3. Taking expectation of (3) and (4) over $v^i_t \sim F^i_t$, we get the expected utilities[6]

$$\mathbb{E}_{v^i_t \sim F^i_t}\left[\tilde{u}^i_t(v^i_t; \boldsymbol{v}_t, \tilde{\boldsymbol{b}}_t)\right] = \int_{\max\{\tilde{r}^i_t, w^i_t\}}^{\infty}\left(1 - F^i_t(v)\right)\mathrm{d}v - \int_{\max\{\tilde{r}^i_t, w^i_t\}}^{\max\{\overline{r}^i_t, w^i_t\}}\left(1 - \tilde{F}^i_t(v)\right)\mathrm{d}v,$$

$$\mathbb{E}_{v^i_t \sim F^i_t}\left[\tilde{u}^i_t(v^i_t; \boldsymbol{v}_t, \tilde{\boldsymbol{b}}'_t)\right] = \int_{\max\{\overline{r}^i_t, w^i_t\}}^{\infty}\left(1 - F^i_t(v)\right)\mathrm{d}v.$$

Hence their difference is

$$\mathbb{E}_{v^i_t \sim F^i_t}\left[\tilde{u}^i_t(v^i_t; \boldsymbol{v}_t, \tilde{\boldsymbol{b}}'_t) - \tilde{u}^i_t(v^i_t; \boldsymbol{v}_t, \tilde{\boldsymbol{b}}_t)\right] = \int_{\max\{\tilde{r}^i_t, w^i_t\}}^{\max\{\overline{r}^i_t, w^i_t\}}\left(1 - \tilde{F}^i_t(v)\right) - \left(1 - F^i_t(v)\right)\mathrm{d}v = \tilde{s}^i_t(\tilde{\boldsymbol{b}}_t) - s^i_t(\tilde{\boldsymbol{b}}_t).$$

$\square$

Proof of Theorem 4.2. We use Lemma 4.3 to bound the expected total regret:

$$\left|\mathbb{E}[R^i]\right| \le \sum_{t=1}^{T}\left|\mathbb{E}_{v^i_t \sim F^i_t}\left[\tilde{u}^i_t(v^i_t; \boldsymbol{v}_t, \tilde{\boldsymbol{b}}'_t) - \tilde{u}^i_t(v^i_t; \boldsymbol{v}_t, \tilde{\boldsymbol{b}}_t)\right]\right| \le \sum_{t=1}^{T}\left|\tilde{s}^i_t(\tilde{\boldsymbol{b}}_t) - s^i_t(\tilde{\boldsymbol{b}}_t)\right| \le \sum_{t=1}^{T}\tilde{s}^i_t(\tilde{\boldsymbol{b}}_t) \cdot \max_{t',b}\left|1 - \frac{s^i_{t'}(\tilde{\boldsymbol{b}})}{\tilde{s}^i_{t'}(\tilde{\boldsymbol{b}})}\right|.$$

Note that the spent $\tilde{s}^i_t(\tilde{\boldsymbol{b}}_t)$ is no more than the balance $\tilde{b}^i_t$ and the balance is further bounded by the limit L^i, hence

$$\sum_{t=1}^{T}\tilde{s}^i_t(\tilde{\boldsymbol{b}}_t) \le TL^i.$$

---

[6]Intermediate steps are omitted, which are similar to (1) and (2) in the proof of Theorem 3.1.

For the prediction error term, if $w_t^i \geq \bar{r}_t^i$, either spent will be zero, hence no prediction error; otherwise, the integrations are both from some dynamic reserve $r' \in [\underline{r}_t^i, \bar{r}_t^i]$ to the baseline reserve $\bar{r}_t^i$. Therefore,

$$\max_{t, \boldsymbol{b}} \left| 1 - \frac{s_t^i(\tilde{\boldsymbol{b}})}{\tilde{s}_t^i(\tilde{\boldsymbol{b}})} \right| = \max_{t, \tilde{r}_t^i} \left| 1 - \frac{\int_{\max\{\tilde{r}_t^i, w_t^i\}}^{\max\{\bar{r}_t^i, w_t^i\}} 1 - F_t^i(v)\mathrm{d}v}{\int_{\max\{\tilde{r}_t^i, w_t^i\}}^{\max\{\bar{r}_t^i, w_t^i\}} 1 - \tilde{F}_t^i(v)\mathrm{d}v} \right| = \max_{t, r' \in [\underline{r}_t^i, \bar{r}_t^i]} \left| 1 - \frac{\int_{r'}^{\bar{r}_t^i} 1 - F_t^i(v)\mathrm{d}v}{\int_{r'}^{\bar{r}_t^i} 1 - \tilde{F}_t^i(v)\mathrm{d}v} \right|.$$

$\square$

Theorem 4.2 then suggests three possible ways for the mechanism designer to reduce the buyer regret:

- to reduce the bank account limit L;

- to make a better prediction $\tilde{F}_t^i$;

- to tune the dynamic region of the lazy reserves $[\underline{r}_t^i, \bar{r}_t^i]$ such that the expected prediction error is small in there.

In the next section, we compare these three methods empirically.

5 Empirical Study

Finally, we use anonymized real bid data from a major ad exchange to evaluate the simple dynamic auction we proposed and establish the trade-offs between the regret and the revenue-lift via the limits on bank accounts.

Data Sets   The data set is a collection of auction records, where each record corresponds to a real time auction for an impression and consists of a seller (publisher) id and a set of bid records. Each bid record corresponds to a buyer id and the value of the bid submitted by that buyer to the auction. Our data set will consist of auctions with bids from 5 large buyers over the period of 2 days. We will partition the data set into a *training set* consisting of data for the first day and a *testing set* consisting of data for the second day.

Learning the Valuation Distributions   Before running the simulation, we need to do some preprocessing of the data set to learn the valuation distributions of the buyers. In order to justify our insights about the impact of prediction errors on regrets from Section 4, we learn several distribution estimators $\tilde{F}$ with different levels of accuracies. In Section 5.2, we conduct same auctions with these estimators and compare the simulation results.

Empirical Distributions:

- `NAIVE`: one empirical distribution for the bids of each buyer;

- `PUBLISHER`: one empirical distribution for the bids of each buyer-publisher pair.

We learn distributions based on real data from the first day and evaluate using real data form the second day. This provides an evaluation for our methods when we have a noisy estimation of distribution.

Parameterized Distributions:

- `LOGNORMAL`: for each buyer, we fit a lognormal distribution $\ln \mathcal{N}(\mu, \sigma^2)$, by maximum likelihood estimation.

For parametric distributions, we fit a log normal using real ad data from the first day and sample from those distributions to generate synthetic data for the second day. We do so to evaluate our methods in the case we have a perfect estimation of the distribution. Hence the only source of regret must be the random noises from generated data and the numerical errors of the simulation.

baseline $\bar{r}$	$\alpha = \underline{r}/\bar{r}$	limit L	revenue lift	social welfare lift	buyer utility loss	relative regret
MINOPT r^{M}	0%	20	+150%	+41.9%	−12.0%	12.7%
CAPPED r^{C}	0%	20	+91.5%	+19.7%	−6.24%	6.25%
ROBUST r^{R}	0%	20	+62.8%	+11.6%	−4.02%	3.74%
MINOPT r^{M}	25%	0.2	+17.2%	+5.04%	−1.01%	0.19%
CAPPED r^{C}	25%	0.2	+19.20%	+4.07%	−1.38%	1.02%
ROBUST r^{R}	25%	0.2	+18.78%	+3.35%	−1.34%	0.93%

Table 1: The *revenue lifts*, *social welfare lifts*, *buyer utility losses*, and *relative regrets* for selected DSPs with $\tilde{F} = \texttt{NAIVE}$.

Baseline Reserve Prices and $\underline{r}$ We use the following three reserve prices as the baseline $\bar{r}$. All of them are estimated based on the $\tilde{F} = \texttt{NAIVE}$ distribution, since the main purpose of having different valuation estimators is to compare the effects of distribution estimation errors on buyer regrets instead of on the baseline reserve selection.

- MINOPT: the minimum reserve that maximizes the expected monopoly revenue,

$$r^{\mathrm{M}} = \min\{\arg\max_r r(1 - \tilde{F}(r))\}.$$

- CAPPED: the same with MINOPT but *the selling probability must be no less than* 15%,

$$r^{\mathrm{C}} = \min\{\arg\max_{r:\tilde{F}(r)\leq 0.85} r(1 - \tilde{F}(r))\}.$$

- ROBUST: the minimum reserve that achieves at least 80% of the revenue of r^{M},

$$r^{\mathrm{R}} = \min\{r : r(1 - \tilde{F}(r)) \geq 0.8 r^{\mathrm{M}}(1 - \tilde{F}(r^{\mathrm{M}}))\}.$$

The attempt of having CAPPED and ROBUST is to reduce the regrets under the guidance of Theorem 4.2: to avoid $\bar{r}^i$ falling into the region where $1 - \tilde{F}(v)$ is too small.

The matching $\underline{r}$ for those baselines, namely, $\underline{r}^{\mathrm{M}}, \underline{r}^{\mathrm{C}}, \underline{r}^{\mathrm{R}}$, are defined to be the corresponding baselines rescaled by a constant α, where α takes a value from $\{0\%, 10\%, 25\%, 50\%\}$ in simulations, $\underline{r}^{\mathrm{M}} = \alpha r^{\mathrm{M}}, \underline{r}^{\mathrm{C}} = \alpha r^{\mathrm{C}}, \underline{r}^{\mathrm{R}} = \alpha r^{\mathrm{R}}$.

Parameterized DSPs To summarize, each DSP we use in the simulations has four parameters: (i) distribution estimators $\tilde{F} \in \{\texttt{NAIVE}, \texttt{PUBLISHER}, \texttt{LOGNORMAL}\}$; (ii) baseline reserves $\bar{r} \in \{r^{\mathrm{M}}, r^{\mathrm{C}}, r^{\mathrm{R}}\}$; (iii) eager reserves $\underline{r} = \alpha\bar{r}$ where $\alpha \in \{0\%, 10\%, 25\%, 50\%\}$; (iv) finally the bank account limits $L \in \{0, 0.1, 0.2, 0.5, 1, 2, 5, 10, 20\}$.[7]

In particular, for part I (Section 5.1), we run simulations for the DSPs with $\tilde{F} = \texttt{NAIVE}$; for part II (Section 5.2), we run simulations for all the DSPs (for those with $\tilde{F} = \texttt{LOGNORMAL}$, we run simulations on synthetic bids sampled from $\tilde{F}$).

5.1 Part I: Evaluating the Performance

For each DSP M, we evaluate its *revenue lift*, *welfare lift*, and *buyer utility losses* comparing with the second price auction with reserves being the baseline reserves of M, as well as the *relative buyer regrets* (buyer regret / buyer utility). All these statistics are plotted in planes (see Figure 1, Figure 2, Figure 3, and Figure 4 grouped by each statistic and each baseline reserve) with x-axis being the bank account limit (in log-scale, $L = 0$ excluded) and points with the same α are connected as different curves. Therefore, these curves show the trends of the statistics as the bank account limit increases.

---

[7]The limits are shown in relative numbers for privacy reasons. The actual limit applied in simulation is $L \cdot (\bar{r}^i + \mathbb{E}_{v \sim \tilde{F}^i}[v])$.
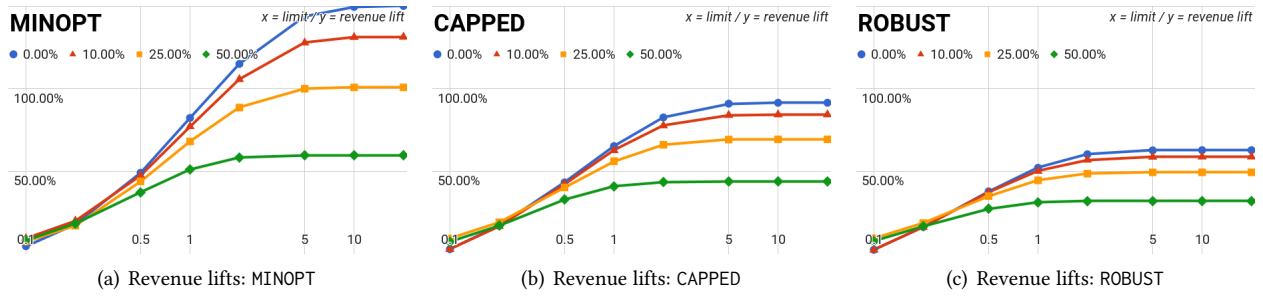
Figure 1: Revenue (y-axis) of DSPs with various baseline reserves changing with bank account limits (x-axis).
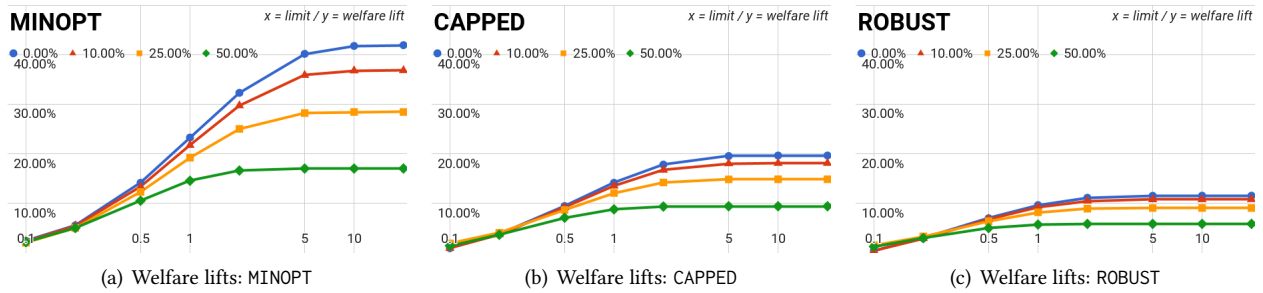

Figure 2: Social welfare (y-axis) of DSPs with various baseline reserves changing with bank account limits (x-axis).
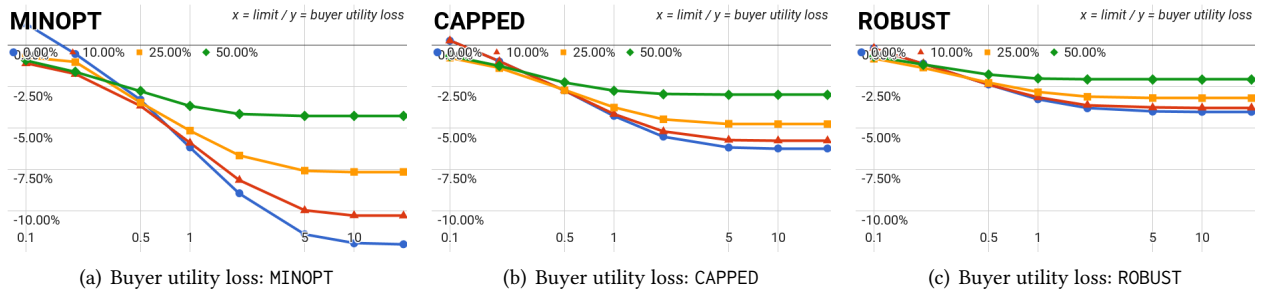

Figure 3: Buyer utility losses (y-axis) of DSPs with various baseline reserves changing with bank account limits (x-axis).
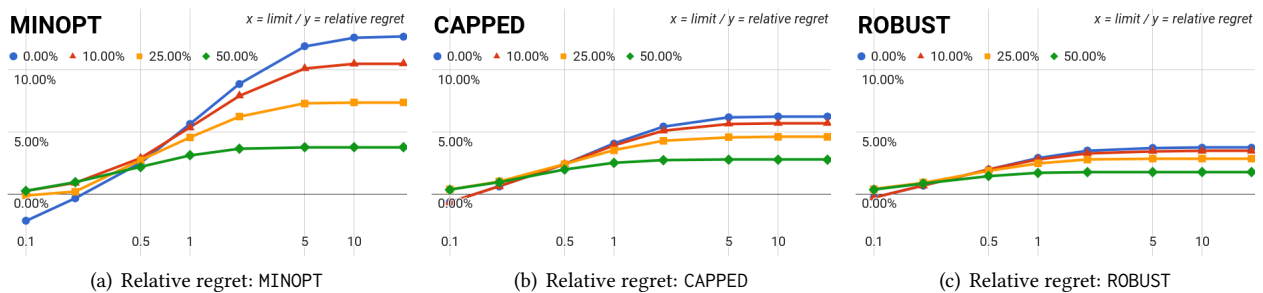

Figure 4: Relative regrets (y-axis) of DSPs with various baseline reserves changing with bank account limits (x-axis).

(a) NAIVE: revenue-regrets   (b) PUBLISHER: revenue-regrets   (c) LOGNORMAL: revenue-regrets

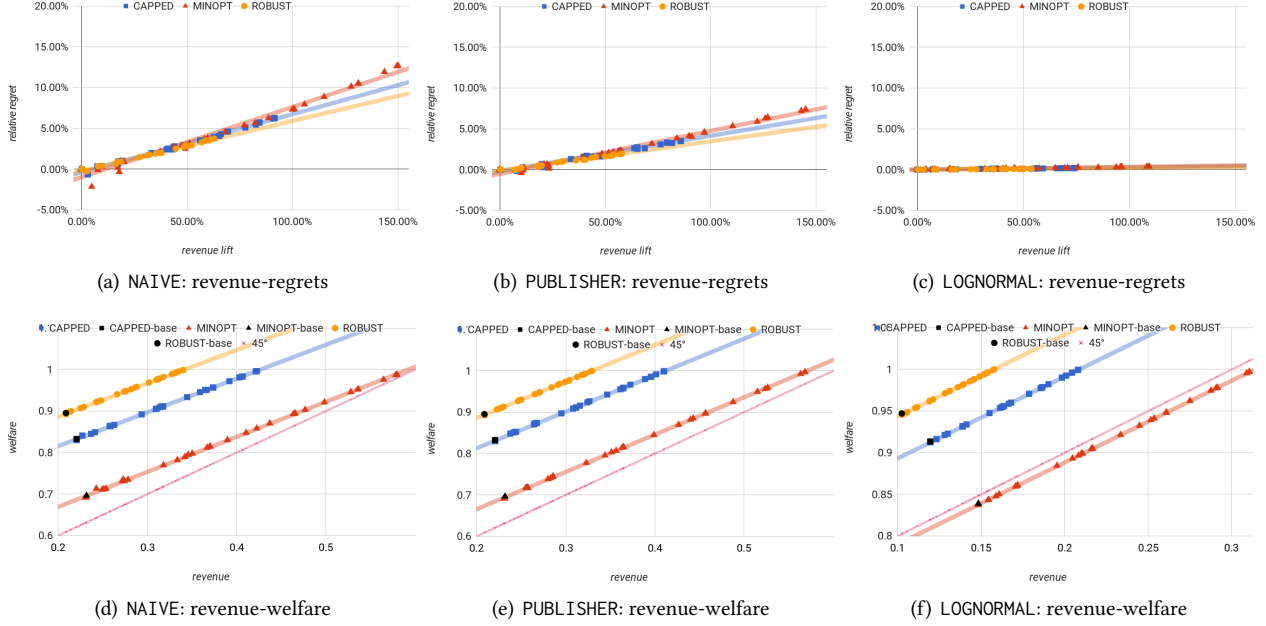(d) NAIVE: revenue-welfare   (e) PUBLISHER: revenue-welfare   (f) LOGNORMAL: revenue-welfare

Figure 5: (a)(b)(c): Trade-off between revenue lifts and relative regrets. (d)(e)(f): Improvements on revenue and social welfare. Absolute numbers are rescaled for privacy reasons.

Interpretation In all these figures, it can be easily observed that the DSPs tend to be farther from the static baseline as the bank account limit L becomes larger. In particular, when L is larger than certain thresholds (e.g., 5 for MINOPT and 1 for ROBUST) the changes become flat. These DSPs are already close to the one with $L = \infty$.

Moreover, as the bank account limit increases, both the revenue and social welfare increase while the buyer utility decreases, consistent with our insights in Section 3.2. In particular, as we expect in Proposition 3.2, for all different $\bar{r}$, both the improvements on revenue and social welfare are strictly positive and significant for sufficiently large L (> 0.2). Concretely, (see Table 1) these improvements can be roughly as large as +150% for revenue and +41.9% for welfare when $\bar{r} = $ MINOPT (+91.5% and +19.7% for CAPPED, +62.8% and +11.6% for ROBUST). Surprisingly, significant improvements (+17.2% revenue and +5.04% welfare) can be achieved even with little buyer utility loss -1.01% and relative regret 0.19%.

In addition, as L becomes small, the regrets get reduced as suggested by the guidelines from Theorem 4.2.

5.2 Part II: Understanding the Regrets

For each estimator $\tilde{F}$, we draw two graphs, each of them containing all the DSPs using $\tilde{F}$ as their estimator. In Figure 5, the first row of graphs shows the trade-offs between revenue lifts (x-axis) and relative buyer regrets (y-axis); the second row of graphs shows the improvements on revenue (x-axis) and social welfare (y-axis).

Interpretation In each of the subgraphs of Figure 5, almost linear trade-offs between revenue lifts and relative regrets can be observed, where the slopes (for estimators NAIVE and PUBLISHER) are roughly 0.03 to 0.08. In particular, the better trade-offs (flatter slopes) can be achieved by either choosing less aggressive baselines (ROBUST $\prec$ CAPPED $\prec$ MINOPT) or using more accurate estimators (LOGNORMAL $\succ$ PUBLISHER $\succ$ NAIVE), in the same spirit of Theorem 4.2.

Similarly, as shown in Figure 5, the revenues and social welfares from different DSPs are quasi-linearly related. The more accurate estimator being used, the closer the slopes of the trends are to 1 (cf. the 45° line), which means the less revenue comes from buyer utility losses.

Finally, attentive readers may notice that the numbers for estimator LOGNORMAL are quite different from the corresponding ones for the other two estimators, like the revenue lifts or the revenue versus social welfare in Figure 5. Such "unexpected" differences in fact come from the differences between the revenue and social welfare of the baseline auctions (second price auctions with reserves) from different simulations, caused by the stochastic differences between

the real bids from the test set and the synthetic bids generated by the lognormal distribution learned from the training set. As we mentioned previously, the training set and test set consist of bids from the first and the second day. The phenomenon we just observed reflects one fundamental difficulty in minimizing the prediction errors, that is, the distribution of buyer values is probably evolving over time. Therefore, the methodology of reducing the regret in the presence of prediction errors is necessary in applying dynamic mechanisms, because the prediction errors might be inevitable.

References

Itai Ashlagi, Constantinos Daskalakis, and Nima Haghpanah. Sequential mechanisms with ex-post participation guarantees. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 213–214. ACM, 2016.

Susan Athey and Ilya Segal. An efficient dynamic mechanism. *Econometrica*, 81(6):2463–2485, 2013.

Santiago Balseiro, Vahab Mirrokni, and Renato Paes Leme. Dynamic mechanisms with martingale utilities. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, page 165, 2017.

Dirk Bergemann and Maher Said. Dynamic auctions. *Wiley Encyclopedia of Operations Research and Management Science*, 2011.

Dirk Bergemann and Juuso Välimäki. Information acquisition and efficient mechanism design. *Econometrica*, 70(3):1007–1033, 2002.

Ruggiero Cavallo. Efficiency and redistribution in dynamic mechanism design. In *Proceedings of the 9th ACM conference on Electronic commerce*, pages 220–229. ACM, 2008.

Shuchi Chawla, Nikhil R Devanur, Anna R Karlin, and Balasubranianian Sivan. Simple pricing schemes for consumers with evolving values. In *SODA 2016*, pages 1476–1490. SIAM, 2016.

Nikhil R Devanur, Yuval Peres, and Balasubramanian Sivan. Perfect bayesian equilibria in repeated sales. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 983–1002. SIAM, 2015.

Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333, 2015.

Jason D Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 225–234. ACM, 2009.

Sham M Kakade, Ilan Lobel, and Hamid Nazerzadeh. Optimal dynamic mechanism design and the virtual-pivot mechanism. *Operations Research*, 61(4):837–854, 2013.

Vahab Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. Dynamic auctions with bank accounts. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.

Vahab Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. Non-clairvoyant dynamic mechanism design. *SSRN*, 2016.

Vahab Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. Optimal dynamic mechanisms with ex-post ir via bank accounts. *arXiv preprint arXiv:1605.08840*, 2016.

Jamie H Morgenstern and Tim Roughgarden. On the pseudo-dimension of nearly optimal auctions. In *Advances in Neural Information Processing Systems*, pages 136–144, 2015.

Jamie Morgenstern and Tim Roughgarden. Learning simple auctions. In *Conference on Learning Theory (COLT)*, 2016.

Mallesh M Pai and Rakesh Vohra. Optimal dynamic auctions and simple index rules. *Mathematics of Operations Research*, 38(4):682–697, 2013.

David C Parkes and Satinder Singh. An mdp-based approach to online mechanism design. 2004.

Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: A myersonian approach. *Econometrica*, 82(2):601–653, 2014.