Dynamic Second Price Auctions with Low Regret*

Vahab Mirrokni
Google Research
mirrokni@google.com

Renato Paes Leme
Google Research
renatoppl@google.com

Rita Ren
Google
rren@google.com

Song Zuo
Tsinghua University
songzuo.z@gmail.com

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 [5, 18, 6, 2, 11, 17, 19, 8, 7]. We refer to the survey by Bergemann and Said [4] 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

*A full version could be found at https://ssrn.com/abstract=2956713.

31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA.

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. [1] and Mirrokni et al. [14] 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. [12] introduce the so-called *bank account mechanisms* that have simple structure and can achieve the optimal revenue under interim individually rational constraints. Later on, the authors extend the bank account framework to guarantee ex-post individual rationality and show a simple construction that 5-approximates the optimal revenue for general multiple buyer cases [13]. Balseiro et al. [3] consider a similar setting with a single buyer who has 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 still require randomized allocation rules[2] 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 [16, 15]). 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)* 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. 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 could be either the revenue optimal reserve prices or be chosen with a different goal, e.g., to maximize a combination of revenue and welfare. Our main theorem shows how to process existing reserve prices and produce an auction that is dynamic incentive compatible and strictly improves over the baseline auction both in terms of revenue and social welfare (Proposition 1.1). Moreover, we confirm on data that this improvement is quite significant.

The idea of the dynamic second price auction is as follows: reserve prices improve revenue at the expense of efficiency. Our auction modifies the static second price auction by introducing dynamic discounts on the reserves to the buyers, where the discounts are based on how much utility buyers accumulated in past periods. The auction also introduces an extra payment (which we call bank account spent) which 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 reserve lowering policy and their 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. 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. [12] to enable a trade-off between revenue lifts and buyer regrets:

[2] The dynamic mechanisms by [3] are deterministic but restricted to the setting of single buyer with i.i.d. values.

the regrets can be bounded by the product of bank account limits and prediction errors (Theorem 1.2). 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. To the best of our knowledge, this is the first paper 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 [10] while incurring relative regret 13% (comparing with 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 1.2 do help in reducing the regrets.

1.1 Hightlighting Some Theoretical Results

We highlight some of the theoretical results here and refer the readers to the full version for more details.

The DSP auction can always generate strictly higher revenue and social welfare than the second price auction with baseline reserves, except for some extreme cases.

Proposition 1.1 (Revenue & social welfare guarantees)**.** *For any given baseline reserves $\overline{\mathbf{r}}_t$, we can always design a DSP that strictly outperforms the second price auction with reserves being $\overline{\mathbf{r}}_t$ 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, i.e.,*
$$\forall i,\ \Pr[v_t^i > \overline{r}_t^i] = 0 \textit{ or } \Pr[v_t^i < \overline{r}_t^i] = 0;$$
- *the probability density near $\overline{r}_t^i$ is zero, i,e,*
$$\forall i,\ \delta > 0,\ \Pr[\overline{r}_t^i - \delta \le v_t^i < \overline{r}_t^i] = 0.$$

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 1.2 (Regret bounds)**.**

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

References

[1] 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.

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

[3] Santiago Balseiro, Vahab Mirrokni, and Renato Paes Leme. Dynamic mechanisms with martingale utilities. 2016.

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

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

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

[7] 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.

[8] 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.

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

[10] 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.

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

[12] 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.

[13] Vahab Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. Oblivious dynamic mechanism design. 2016.

[14] 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.

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

[16] 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.

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

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

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