Lecture Note 11: Budgets in Auction Markets - Dynamics

Christian Kroer*

March 8, 2022

1 Introduction

In the last lecture note we studied auctions with budgets and repeated auctions. However, we ignored one important aspect: time. In this lecture note we consider an auction market setting where a buyer is trying to adaptively pace their bids over time. The goal is to hit the "right" pacing multiplier as before, but each bidder has to learn that multiplier as the market plays out. We'll see how we can approach this problem using ideas from regret minimization.

2 Dynamic Auctions Markets

In this setting we have n buyers who repeatedly participate in second-price auctions. At each time period $t = 1, \dots, T$ a single second-price auction is run. At time t, each bidder samples a valuation v_{it} independently from a cumulative distribution function F_i which is assumed to be absolutely continuous and with bounded density f_i whose support is $[0, \bar{v}_i]$. As usual, we assume that each buyer has some budget B_i that they should satisfy, and we denote by $\rho_i = B_i/T$ the per-period target expenditure; we assume $\rho_i \leq \bar{v}_i$. We may think of each buyer as being characterized by a type $\theta_i = (F_i, \rho_i)$.

At each time period t buyer i observes their valuation v_{it} and then submits a bid b_{it}. We will use $d_{it} = \max_{k \neq i} b_{it}$ to denote the highest bid other than that of i. As before the utility of an buyer is quasi-linear and thus if they win auction t they get utility $v_{it} - d_{it}$. We may write the utility using an indicator variable as $u_{it} = \mathbb{1}\{d_{it} \leq b_{it}\}(v_{it} - d_{it})$, and the expenditure $z_{it} = \mathbb{1}\{d_{i,t} \leq b_{it}\}d_{it}$.

It is assumed that each buyer has no information on the valuation distributions, including their own. Instead, they just know their own target expenditure rate ρ_i and the total number of time periods T. Buyers also do not know how many other buyers are in the market.

At time t, buyer i knows the *history* $(v_{i\tau}, b_{i\tau}, z_{i\tau}, u_{i\tau})_{\tau=1}^{t-1}$ of own values, bids, payments, and utilities. Furthermore, they know their current value v_{it}. Based on this history, they choose a bid b_{it}. We will say that a bidding strategy for buyer i is a sequence of mappings $\beta = \beta_1, \dots$ where β_t maps the current history to a bid (potentially in randomized fashion). The strategy β is budget feasible if the bids b_{it}^{β} generated by β are such that

$$\sum_{t=1}^{T} \mathbb{1}\{d_{it} \leq b_{it}^{\beta}\}d_{it} \leq B_i$$

under any vector of highest competitor bids d_i.

*Department of Industrial Engineering and Operations Research, Columbia University. Email: christian.kroer@columbia.edu.

For a given realization of values $v_i = v_{i1}, \ldots, v_{iT}$ and highest competitor bids d_i we denote the expected value of a strategy β as

$$\pi_i^\beta(v_i, d_i) = \mathbb{E}\left[\sum_{t=1}^{T} \mathbb{1}\{d_{it} \leq b_{it}^\beta\}(v_{it} - d_{it})\right],$$

where the expectation is taken with respect to randomness in β.

We would like to compare our outcome to the *hindsight optimal* strategy. We denote the expected value of that strategy as

$$\begin{aligned} \pi_i^H(v_i, d_i) := \quad \max_{x_i \in \{0,1\}^T} \quad & \sum_{t=1}^{T} x_{it}(v_{it} - d_{it}) \\ s.t. \quad & \sum_{t=1}^{T} x_{it} d_{it} \leq B_i \end{aligned} \tag{1}$$

The hindsight-optimal strategy has a simple structure: we simply choose the optimal subset of items to win while satisfying our budget constraint. In the case where the budget constraint is binding, this is a knapsack problem.

Ideally we would like to choose a strategy such that π_i^β approaches π_i^H. However, this turns out not to be possible. We will use the idea of asymptotic γ-competitiveness to see this. Formally, β is asymptotic γ-competitive if

$$\limsup_{\substack{T\to\infty, \\ B_i = \rho_i T}} \sup_{\substack{v_i \in [0, \bar{v}_i]^T, \\ d_i \in \mathbb{R}_+^T}} \frac{1}{T}\left(\pi_i^H(v_i, d_i) - \gamma \pi_i^\beta(v_i, d_i)\right) \leq 0$$

Intuitively, the condition says that asymptotically, β should achieve at least $1/\gamma$ of the hindsight-optimal expected value.

For any $\gamma < \bar{v}_i/\rho_i$, asymptotic γ-competitiveness turns out to be impossible to achieve. Thus, if our target expenditure ρ_i is much smaller than our maximum possible valuation, we cannot expect to do anywhere near as well as the hindsight-optimal strategy.

The general proof is quite involved, but the high-level idea is not too complicated. Here we show the construction for $\bar{v}_i = 1, \rho_i = 1/2$, and thus the claim is that $\gamma < \bar{v}_i/\rho_i = 2$ is unachievable. The impossibility is via a worst-case instance. In this instance, the highest other bid comes from one of the two following sequences:

$$\begin{aligned} d^1 &= (d_{high}, \ldots, d_{high}, \bar{v}_i, \ \ldots\ldots \ , \bar{v}_i) \\ d^2 &= (d_{high}, \ldots, d_{high}, d_{low}, \ldots, d_{low}), \end{aligned}$$

for $\bar{v}_i \geq d_{high} > d_{low} > 0$. The general idea behind this construction is that in the sequence d^1, buyer i must buy many of the expensive items in order to maximize their utility, since they receive zero utility for winning items with price $\bar{v}_i$. However, in the sequence d^2, buyer i must save money so that they can buy the cheaper items priced at d_{low}.

For the case we consider here, there are $T/2$ of each type of highest other bid (assume T is even for convenience). Now, we may set $d_{high} = 2\rho_i - \epsilon$ and $d_{low} = 2\rho_i - k\epsilon$, where ϵ and k are constants that can be tuned. For sufficiently small ϵ, i can only afford to buy $T/2$ items total, no matter the combination of items. Furthermore, buying an item at price d_{low} yields k times as much utility as buying an item at d_{high}.

Now, in order to achieve at least half of the optimal utility under d^1, buyer i must purchase at least $T/4$ of the items priced at d_{high}. Since they don't know whether d^1 or d^2 occurred until after deciding whether to buy at least $T/4$ of the d_{high} items, this must also occur under d^2. But then buyer i can at most afford to buy $T/4$ of the items priced at d_{low} when they find themselves in the d^2 case. Now for any $\gamma < 2$, we can pick k and ϵ such that achieving $\gamma \pi_i^H$ requires buying at least $T/4 + 1$ of the d_{low} items.

It follows that we cannot hope to design an online algorithm that competes with $\gamma \pi_i^H$ for $\gamma < \bar{v}_i / \rho_i$. However, it turns out that a subgradient descent algorithm can achieve exactly $\gamma = \bar{v}_i / \rho_i$

3 Adaptive Pacing Strategy

The idea is to construct a pacing multiplier $\alpha_i = \frac{1}{1+\mu}$ by running a subgradient descent scheme on the value for μ that allows i to smoothly spend their budget across the T time periods.

The algorithm takes as input a stepsize $\epsilon_i > 0$ and some initial value $\mu_1 \in [0, \bar{\mu}_i]$ (where $\bar{\mu}_i$ is some upper bound on how large μ needs to be). We use $P_{[0,\bar{\mu}_i]}$ to denote projection onto the interval $[0, \bar{\mu}_i]$. The algorithm, which we call APS, proceeds as follows

- Initialize the remaining budget at $\tilde{B}_{i1} = B_i$
- For every time period $t = 1, \ldots, T$:
  1. Observe v_{it}, construct a paced bid $b_{it} = \min(\frac{v_{it}}{1+\mu_t}, \tilde{B}_{it})$
  2. Observe spend z_{it}, and update the pacing multiplier:
  $$\mu_{t+1} = P_{[0,\bar{\mu}_i]}(\mu_t - \epsilon_i(\rho_i - z_{it}))$$
  3. Update remaining budget $\tilde{B}_{i,t+1} = \tilde{B}_{it} - z_{it}$

This algorithm is motivated by Lagrangian duality. Consider the following Lagrangian relaxation of the hindsight-optimal optimization problem (1):

$$\max_{x \in \{0,1\}^T} \sum_{t=1}^{T} \left[x_{it}(v_{it} - (1-\mu) d_{it}) + \mu \rho_i \right].$$

The optimal solution for the relaxed problem is easy to characterize: we set $x_{it} = 1$ for all t such that $v_{it} \geq (1-\mu) d_{it}$. Importantly, this is achieved by the bid $b_{it} = \frac{v_{it}}{1+\mu}$ that we use in APS.

The Lagrangian dual is the minimization problem

$$\inf_{\mu \geq 0} \sum_{t=1}^{T} \left[(v_{it} - (1-\mu) d_{it})^+ + \mu \rho_i \right], \tag{2}$$

where $(\cdot)^+$ denotes thresholding at 0. This dual problem upper bounds π_i^H (but we do not necessarily have strong duality since we did not even start out with a convex primal program). The minimizer of the dual problem yields the strongest possible upper bound on ϕ_i^H, however, solving this requires us to know the entire sequences v_i, d_i. APS approximates this optimal μ by taking a subgradient step on the t'th term of the dual:

$$\partial_\mu \left[(v_{it} - (1-\mu) d_{it})^+ + \mu \rho_i \right] \ni \rho_i - d_{it} \mathbb{1}\{b_{it} \geq d_{it}\} = \rho_i - z_{it}.$$

Thus APS is taking subgradient steps based on the subdifferential of the t'th term of the Lagrangian dual of the hindsight-optimal optimization problem.

The APS algorithm achieves exactly the lower bound we derived earlier, and is thus asymptotically optimal:

Theorem 1. *APS with stepsize* $\epsilon_i = O(T^{-1/2})$ *is* $\frac{\bar{v}_i}{\rho_i}$*-asymptotic competitive, and converges at a rate of* $O(T^{-1/2})$.

This result holds under adversarial conditions: for example, the sequence of highest other bids may be as d^1, d^2 in the lower bound. However, in practice we do not necessarily expect the world to be quite this adversarial. In a large-scale ad market, we would typically expect the sequences v_i, d_i to be more stochastic in nature. In a fully stochastic setting with independence, APS turns out to achieve π_i^H asymptotically:

Theorem 2. *Suppose* $(v_{it,d_{it}})$ *are sampled independently from stationary, absolutely continuous CDFs with differentiable and bounded densities. Then the expected payoff from APS with stepsize* $\epsilon_i = O(T^{-1/2})$ *approaches* π_i^H *asymptotically at a rate of* $T^{-1/2}$.

Theorem 2 shows that if the environment is well-behaved then we can expect much better performance from APS.

4 Historical Notes

The material presented here was developed by Balseiro and Gur [1]. Beyond auction markets, the idea of using paced bids based on the Lagrange multiplier μ has been studied in the revenue management literature, see e.g. Talluri and Van Ryzin [3], where it is shown that this scheme is asymptotically optimal as T tends to infinity. There is also recent work on the adaptive bidding problem using multi-armed bandits [2].

References

[1] Santiago R Balseiro and Yonatan Gur. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, 65(9):3952–3968, 2019.

[2] Arthur Flajolet and Patrick Jaillet. Real-time bidding with side information. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 5168–5178. Curran Associates Inc., 2017.

[3] Kalyan Talluri and Garrett Van Ryzin. An analysis of bid-price controls for network revenue management. *Management science*, 44(11-part-1):1577–1593, 1998.