Auction Prediction With Incomplete Information

AlephD

May 2014

1 Presentation

We study ad-auctions, auctions on which ad spaces are sold by publishers to advertisers in milliseconds. We focus on the simulation of the gain of publishers, should they have set a different floor price for a specific set of auctions.

1.1 Second price ad-auctions

We denote:

- by $\mathcal{B} = \{b_1, \dots, b_B\}$ the set of advertisers (buyers or bidders);
- and by $\mathcal{A} = \{a_1, \dots, a_A\}$ the set of successive auctions, happening at time

$$\{t_{a_1}, \dots, t_{a_A}\}$$

The advertisers (also called bidders) are sent *bid requests* for every auction a. We assume they take part in all of them by simultaneously announcing a price. Note that a price of 0, can be considered equivalent to a decision not to participate to a specific auction. The bidder with the highest price for auction a[1] can display his ad and pays the maximum of the second price and the reserve price for the auction (provided his price is above the reserve price).

Let us consider an auction a happening at time t_a, let

$$p_a = (p_{a,b_1}, p_{a,b_2}, \dots, p_{a,b_B})$$

be the bids of the B bidders in the auction a. We denote:

- by $\rho_a : \mathcal{B} \to \{1, \dots, B\}$ the function which gives the rank of an advertiser in the auction a (1 will be the rank of the highest price);
- by $\beta_a : \{1, \dots, B\} \to \mathcal{B}$ the function which gives the advertiser corresponding to a rank in the auction a. We have: $\beta_a = \rho_a^{-1}$.

We denote:

$$(p_{a,1}, p_{a,2}, \dots, p_{a,B}) = \left(p_{a,\beta_a(1)}, p_{a,\beta_a(2)}, \dots, p_{a,\beta_a(B)}\right)$$

[1] The case of equality is extremely rare. We assume it does not happen

1.2 The publisher

The publisher: s (seller) can act on an auction a by setting a reserve-price r_a before it happens. The payoff of s at auction a is a function of $p_{a,1}$, $p_{a,2}$ and r_a:

$$\pi_a(p_{a,1}, p_{a,2}, r_a) = \begin{cases} p_{a,2} & \text{if } p_{a,2} \geq r_a \\ r_a & \text{if } p_{a,1} \geq r_a > p_{a,2} \\ d_a & \text{if } r_a > p_{a,1} \end{cases}$$

Where d_a is the value received by the publisher for the display of a *default ad*. It can be assumed to be 0.

After the auction, he observes the result of a imperfectly. Let's assume he observes the n-uplet:

$$o_a = \begin{cases} (a, u_a, r_a, p_{a,2}, i_a) & \text{if } p_{a,2} \geq r_a \\ (a, u_a, r_a, r_a, i_a) & \text{if } p_{a,1} \geq r_a > p_{a,2} \\ (a, u_a, r_a, 0, i_a) & \text{if } r_a > p_{a,1} \end{cases}$$

where u_a is a *user id* and i_a is an *information set*. An information set is a set of auctions to which *AlephD* will respond with the same reserve price.

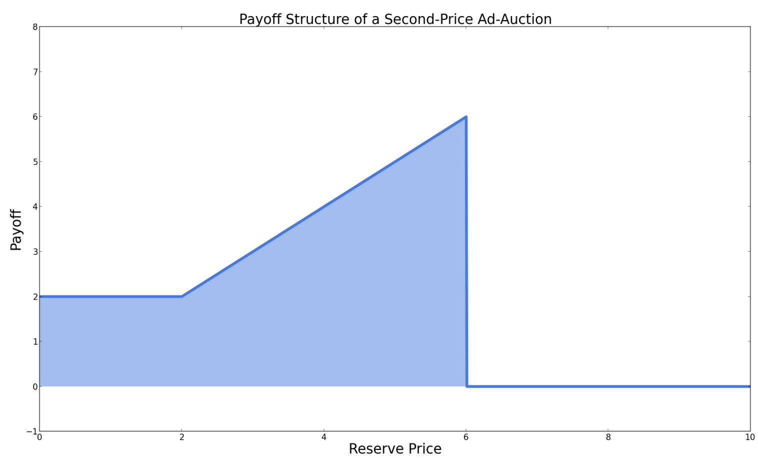

Figure 1: Payoff Structure of a Second-Price Ad-Auction

2 The objective

The publisher will receive successive auction requests and will have to respond with a reserve-price to each of them. Each request is characterized by a time $t = 1, \ldots, n$ and an information set $i_t \in \mathcal{I}$ where $\mathcal{I}$ is finite. The publisher has to design a reserve-price strategy that maximizes his revenue.

2.1 The information available

Before the auction takes place, the information available to the publisher are the following:

- The user id, that characterizes the user that is loading the page;
- The tag id, that characterized the placement where the impression will appear;
- Geographical data, such as the city, the country, the longitude and latitude when available;
- The user agent, a string that contains information about the OS, the browser and the configurations of the user;
- The time, unix epoch from which we can extract the hour and the day;
- Historical data, such as previous closing prices for this user and this tag, previous tags visited by this user.

After the auction takes place, additional information is gathered:

- The closing price, the price at which the impression has been sold;
- The creative id, that characterizes the ad that won the auction, it includes the advertiser, the campaign and the format of the ad.

2.2 The problem

The challenge in this setup is that the first-bid is not observed after the auction takes place. Therefore, the payoff as a function of the reserve-price is not entirely defined for each auction. However, we have a partial information about the payoff with the closing price p_a.

$$\begin{cases} \text{if } p_a = 0 \ \& \ r_a = 0, & p_{a,1} = 0, \text{ the payoff is 0 for all } r \\ \text{if } p_a = 0 \ \& \ r_a > 0, & p_{a,1} < r_a, \text{ the payoff is 0 for } r \geq r_a \\ \text{if } p_a > 0 \ \& \ p_a = r_a, & p_{a,2} \leq r_a \leq p_{a,1}, \text{ the payoff is known for } r \leq r_a \\ \text{if } p_a > 0 \ \& \ p_a > r_a, & r_a < p_{a,2}, \text{ the payoff is known for } r \leq r_a \end{cases}$$

Thus, to gain some information about the payoff for a given information set, we must explore the space of reserve-price.

Ideally, the reserve-price has to be as close as possible to the first-bid while remaining smaller. Due to the sharp asymmetry of the payoff function, it is preferable not to try to be too close to the first-bid given the high cost of exceeding it. In some cases, the first-bid distribution given a set of information can be approximated by a log-normal distribution.

We are faced with the exploration/exploitation dilemma where we need to find a balance between setting many different reserve-prices to span the space of payoffs and maximizing the revenue using the best reserve-price given the information available.

At each time t, the publisher is given an information set i_t that it uses to determine the reserve-price r_t, that goes in input of the payoff function $\pi(p_{t,1}, p_{t,2}, r_t)$ together with the first price $p_{t,1}$ and the second price $p_{t,2}$. The set of possible values for the reserve-price r_t is $\mathbb{R}^+$. Hence, the publisher must learn the best mapping $g : \mathcal{I} \to \mathbb{R}^+$ that will maximize his revenue on the whole period. We can define the regret as:

$$R_n = \max_{g:\mathcal{I}\to\mathbb{R}^+} \sum_{t=1}^{n} \pi(p_{t,1}, p_{t,2}, g(i_t)) - \sum_{t=1}^{n} \pi(p_{t,1}, p_{t,2}, r_t)$$

Therefore, we must find a strategy that minimizes the upper bound of the regret.