Hybrid Dual Censored Joint Learning of Reserve Prices and Bids for Upstream Auctioneers

Piyush Paliwal
LoopMe Ltd.
London, United Kingdom
piyush.paliwal@loopme.com

Lampros C. Stavrogiannis
LoopMe Ltd.
London, United Kingdom
lampros@loopme.com

ABSTRACT

Open Marketplace is an environment where buyers and sellers come together to openly exchange ad inventory. A discerning characteristic of this marketplace is that inventory is traded in a sequence of auctions, known as the supply path. We consider the profit maximization problem from the standpoint of an upstream marketplace (exchange) for opportunities coming from downstream marketplaces. To make a profit, the exchange extends the reserve price originated from the downstream marketplace by a reserve multiplier in its auction for its demand partners and reduces the clearing price of that auction by shading its bid when bidding downstream. We introduce, for the first time, a framework for jointly learning the reserve multiplier and bid shading factor for first-price auctions that generalizes to second-price auctions too, taking into account the censoring of prices on both sides of the market. We also provide an elegant strategy based on the Revenue Equivalence Theorem to deal with the co-existence of both auction types for the same inventory. A/B tests in LoopMe exchange demonstrate the effectiveness of our framework in practice.

CCS CONCEPTS

• **Applied computing** → **Online auctions**.

KEYWORDS

Open Marketplace, Real-Time Bidding, Online Advertising, Upstream Exchange, Sequential Auctions, Censored Joint Learning.

ACM Reference Format:

Piyush Paliwal and Lampros C. Stavrogiannis. 2021. Hybrid Dual Censored Joint Learning of Reserve Prices and Bids for Upstream Auctioneers. In *ADKDD 2021, August 14–18, 2021, Singapore*. ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Online ads that constitute the largest source of revenue for web publishers have predominantly been traded with the use of auctions. The introduction of ad exchanges and real-time bidding (RTB) have made it possible to scale and improve the efficiency of this continuously evolving marketplace. The concept of an ad exchange has been adopted from the financial world whereby brokers are trading assets in stock exchanges on behalf of their clients. Similarly, in online advertising, agencies are using so-called demand-side (DSPs) and supply-side platforms (SSPs), i.e. broker platforms, to buy and sell ad space on behalf of advertisers and publishers respectively, in real time, through ad exchanges in a market that is known as Open Marketplace (OMP). In reality, a chain of such intermediaries (called the supply path) are involved in this process.

With the introduction of sponsored-search ad auctions at the beginning of this millennium, ad space was sold using first-price auctions that were replaced with generalized second-price auctions due to their stability and strategic simplicity. Second-price auctions dominated display advertising and OMP for the same reasons. However, the need for more transparency along with the introduction of header bidding in 2017 where intermediaries compete directly with each other (compared to the previously-used waterfall system) brought first-price auctions back to the forefront of RTB [5]. At the moment, the industry is going through a transition phase where both auction types are taking place and this is likely to persist for the foreseeable future. This is why a number of large platforms enforce both types of auctions based on the type of inventory.

Inspired by the above, in this work, we study the problem of an upstream intermediary (called the *exchange*) that is buying inventory from multiple *downstream* marketplaces (called *SSPs*) on behalf of one or more upstream demand partners (called *upstream DSPs*). The SSPs are using first- or second-price auctions (or both) and the exchange is also running an auction (called the *upstream auction*) among its upstream clients. The main objective of the intermediary is to maximize its profit by setting an appropriate reserve price for its upstream market and a bid price for the downstream market where it is competing with other intermediaries.

We propose a framework for learning the distribution of upstream bids and downstream win prices and to jointly optimize reserve prices and bids, taking into account the censoring patterns of the upstream and downstream auctions. More specifically, the introduction of an upstream reserve price hides the demand below it. Similarly, win prices are not typically disclosed to bidders in the downstream market. Hence, for a second-price auction, the downstream highest opponent bid is only known when winning. This censoring problem is more evident in first-price auctions where the exchange only knows that the highest opponent bid was higher than its own bid in case of a loss, lower in case of a win, but never by how much.

Our contributions are as follows.

- We propose a novel, computationally efficient framework for jointly learning upstream reserve prices and downstream bids for the exchange that takes into account censoring on

AdKDD 2021, August 14–18, 2021, Singapore

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

both sides (upstream and downstream) for first- as well as second-price auctions.

- For *hybrid* inventory, i.e. inventory where both types of auctions are present, we apply an elegant strategy for first-price auctions using information from second-price auctions that is based on the Revenue Equivalence Theorem (RET), a standard result in auction theory.
- We provide empirical results based on A/B tests that have been performed at LoopMe exchange.

To the best of our knowledge, this is the first work for an intermediary that optimizes profit for first-price auctions and, also, implements a RET-based strategy for hybrid inventory. The rest of this work is organized as follows. In Section 2, we discuss related work. We define the problem in Section 3. Section 4 details our methodology. In Section 5, we present our experimental results on real traffic. Finally, Section 6 concludes.

2 RELATED WORK

The works that are closest in spirit to ours are those of [1, 6] where the authors describe a similar setting of profit maximization for intermediaries. In [1], the authors propose a new methodology for setting reserve prices for upstream bidders in the presence of a downstream auction. The work in [6] is motivated by the rise of header bidding and first-price auctions and casts the bidding as a contextual multi-armed bandit problem where they apply Thompson Sampling. Despite the intermediary setting, each of the aforementioned works focuses on one side of the optimization (upstream reserve prices and first-price bidding strategies). The focus of our work is on the joint optimisation of both supply and demand instead. [1] also studies second-price upstream auctions under revenue sharing with the SSP whereas we provide a general framework for both first- and second-price auctions.

From an auction theoretic perspective, our work is related to the recent literature on auctions with intermediaries. This problem was first described for a single item in [4] where the authors provide results on optimal reserve prices for both the intermediaries and the downstream auctioneer under second-price auctions. [8] extends that work by considering different auction mechanisms for the intermediaries as well as non-captive upstream bidders. An extensive summary of this stream of work can be found in [2].

3 PROBLEM DEFINITION

The flow of an ad-request can be described as follows. First, the publisher sends a request to an SSP which then sends the request to several bidders (downstream DSPs and/or exchanges). Each exchange in turn runs its upstream auction for the same opportunity by asking its upstream DSPs to submit a bid and then submits a bid at the SSP. The winner of the downstream auction (downstream DSP or exchange) pays the downstream clearing price to the SSP and the winner's ad is delivered on the publisher's ad slot (Figure 1).

We model the aforementioned process from the standpoint of a single exchange where opportunities arrive from multiple SSPs. The exchange acts as a bidding agent for the SSPs and as a selling auctioneer for its upstream DSPs. Each SSP may hold first- or second-price sealed-bid auctions and the exchange holds its own upstream auction. For simplicity, we assume that the upstream auction type is the same as that of the downstream SSP. However, the exchange has the flexibility to determine the type of auction to run based on other considerations (KPIs or business related).

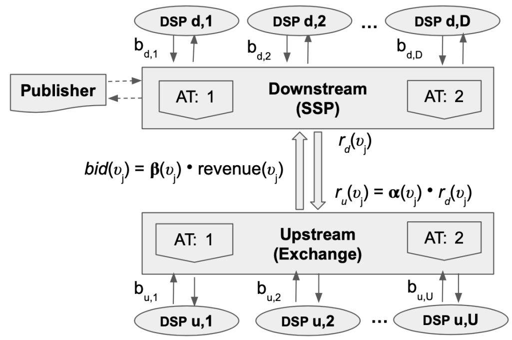

Figure 1: Open Marketplace Setting.

The exchange makes a profit from the difference of its upstream (revenue) and downstream (cost) clearing prices. It achieves this by extending the SSP's reserve price in its upstream auction and reducing the upstream clearing price when bidding downstream.

Each request $j \in \{1, \dots J\}$ is represented by a feature vector v_j. This vector encodes information about the user (such as user agent, location) and context (like URL, app category), a reserve price, $r_d(v_j) > 0$, i.e. the minimum price below which no bids are accepted by the SSP, and the auction type, AT:

$$AT(v_j) = \begin{cases} 1, & \textit{first-price auction} \\ 2, & \textit{second-price auction} \end{cases} \tag{1}$$

Let $\mathcal{W}$ represent the set of all winning bids and $\mathcal{L}$ the set of all losing bids for our modeled exchange. The exchange receives the downstream ad-request j and extends the SSP's downstream reserve price, $r_d(v_j)$, with the following upstream reserve price:

$$r_u(v_j) = \alpha(v_j) \cdot r_d(v_j) \tag{2}$$

where $\alpha(v_j) \geq 1.0$ is the *reserve multiplier*, i.e. we consider a linear function for the optimization. For simplicity, it will be denoted as α_j for the remainder of this paper. $\alpha_j = \alpha_{min} = 1$ constitutes a lower bound since any upstream reserve price below $r_d(v_j)$ can risk low, non-acceptable bids at the SSP (the submitted bids cannot be lower than $r_d(v_j)$). v_j is then sent out across U participating upstream DSPs at the exchange, each of which in turn submits a bid, resulting in the set $\{b_{u,1}, b_{u,2} \dots, b_{u,U}\}$. Without loss of generality, we assume that bids are ordered as $b_{u,1} > b_{u,2} > \dots > b_{u,U}$. So, when $b_{u,1} \geq r_u(v_j)$, the highest bidder bidding $b_{u,1}$ wins the upstream auction and agrees to pay the upstream clearing price, *contingent on* the exchange also winning at the downstream auction. This price constitutes the revenue for the exchange in case of winning downstream too and can be expressed as:

$$revenue(v_j) = \begin{cases} b_{u,1}, & AT(v_j) = 1 \\ \max(b_{u,2}, r_u(v_j)), & AT(v_j) = 2 \end{cases} \tag{3}$$

Next, the exchange submits a bid, $bid(v_j)$, to the SSP, assumed to be a linear function of its private valuation (i.e. revenue).

$$bid(v_j) = \beta(v_j) \cdot revenue(v_j) \tag{4}$$

where $\beta(v_j) \in [r_d(v_j)/revenue(v_j), 1]$ is the *bid shading factor*. For simplicity, it will be denoted as β_j for the remainder of this paper. $\beta_j = \beta_{min} = r_d(v_j)/revenue(v_j)$ constitutes a lower bound since the exchange cannot bid below $r_d(v_j)$ to be considered in the auction. $\beta_j = \beta_{max} = 1$ is an upper bound since the exchange cannot bid more than its valuation.

We assume D downstream DSPs which are connected to the SSP, in addition to the exchange. Given the ordered set of their bids, $b_{d,1} > b_{d,2} > \ldots > b_{d,D}$, if $bid(v_j) < b_{d,1}$, the exchange loses the downstream auction. Otherwise, the exchange wins and pays the downstream auction's clearing price, incurring a cost:

$$cost(v_j) = \begin{cases} bid(v_j), & AT(v_j) = 1 \\ \max(b_{d,1}, r_d(v_j)), & AT(v_j) = 2 \end{cases} \tag{5}$$

From the perspective of the exchange, winning the auction is determined by first receiving at least one bid not less than $r_u(v_j)$, i.e. the upstream auction clears, and subsequently winning at the downstream auction. Not receiving an acceptable bid at its upstream auction automatically disqualifies the exchange from participating downstream. When both of the these conditions are true ($j \in \mathcal{W}$), its profit can be expressed as:

$$profit(v_j) = revenue(v_j) - cost(v_j) \tag{6}$$

For lost bids ($j \in \mathcal{L}$), the profit equals 0. The exchange aims to maximize profit by setting an optimal upstream reserve price and downstream bid. This is achieved by jointly learning the two key ingredients: reserve multiplier α_j and bid shading factor β_j. For the remainder of the paper profit will be referred as $profit(v_j; \alpha_j, \beta_j)$ to denote explicitly that decision variables are α_j and β_j.

In the following section, we provide the details of our framework that takes into account the time constraints due to the real-time nature of bidding as well as the censoring of prices both in the downstream and upstream markets.

4 PROPOSED DUAL CENSORED JOINT LEARNING FRAMEWORK

One of the challenges for the exchange is that the reserve price and bid value must be jointly learned; profit is contingent on winning downstream and this latter cost must be considered when setting an optimal reserve price. Furthermore, one of the requirements of RTB is the need to respond in less than 100 milliseconds. If not, the bidder is not considered in the SSP's auction, translating to immediate profit reduction. SSPs are also selectively calling out bidders based on these timeouts, meaning that such delays can have a significant impact on long-term profitability. This is challenging enough for DSPs which need to calculate a bid. However, for an exchange, responding additionally requires running an upstream auction first. Calculating these values asynchronously, i.e. optimizing the bid value only after the winning upstream bid is known, translates to an additional layer of computations that risks further timeouts. For this reason, the proposed framework comprises determining α_j and β_j *synchronously*, before the upstream auction takes place.

We develop our strategy inspired by the exploration-exploitation trade-off where an agent simultaneously attempts to acquire new knowledge (*exploration*) and optimize its decisions given what it has learned (*exploitation*). The problem is challenging due to the simultaneous learning of the two bid distributions from the downstream and upstream auctions that is necessary for optimizing the two key ingredients of the agent, i.e. the upstream reserve price and the downstream bid. Data censoring makes the problem more difficult. There are three types of censoring in the survival analysis literature, namely left, right and interval [3]. Left censoring or right censoring occurs when observations are censored below or above a censoring point, respectively. Interval censoring occurs when observations are censored between lower and upper censoring points.

For the exchange as a selling auctioneer, the upstream DSPs' bid prices are left-censored as they are only observed above the reserve price r_u. For the exchange as a bidding agent, the win price is right-censored for the lost bids in second-price auctions; the exchange only knows the win price is above its own bid value. In first-price auctions, when the auction is won, the clearing price equals the exchange's bid value. The minimum price to win for the won bids is below the bid value, i.e. below the left-censoring point, that is, winning would have been possible even below that point. When the auction is lost, similar to second-price auctions, the minimum price to win for the lost bids is above the bid value, i.e. above the right-censoring point; winning would only be possible above that.

We use the term **dual censoring exploration** to refer to finding the censoring patterns of the two distributions that considers both types of censoring on the upstream and downstream sides. As we will describe in Section 5, a fixed percentage of the data is randomly chosen for this exploration. The process can be described as follows.

Initially, to deal with cold start, we use random exploration of $\alpha_j \sim \mathcal{U}(\alpha_{min}, \alpha_{min} \cdot (1 + \epsilon_\alpha))$ where $\epsilon_\alpha > 0$ is a parameter that determines the exploration bounds that can be tuned accordingly. Similarly, we use random exploration for $\beta_j \sim \mathcal{U}(\beta_{min}, \beta_{max})$. This exploration data is used to derive optimal values, α^* and β^*, resulting in optimal upstream reserve price r_u^* and downstream bid bid^*, respectively, which will be used in the main process below.

- To find a better upstream reserve price, the exchange, as a selling auctioneer, uses a random exploration of $\alpha_j \sim \mathcal{U}(\alpha^* \cdot (1 - \epsilon_\alpha), \alpha^* \cdot (1 + \epsilon_\alpha))$ where α^* is the currently exploited reserve multiplier. Exploring in $[\alpha^* \cdot (1 - \epsilon_\alpha), \alpha^*]$ is necessary to learn the distribution of upstream bids below r_u^*, i.e. the left-censoring point. Exploring in $[\alpha^*, \alpha^* \cdot (1 + \epsilon_\alpha)]$ is also necessary since, despite the fact that bids are observed above r_u^*, the exchange must learn the upstream DSPs' behaviour; their bidding strategy often incorporates the incoming reserve price. Increasing the latter enables the exchange to learn more about the bidders' private valuation in order to maximize its revenue from the upstream auction. Any values below α_{min} are set to α_{min}. This process can be applied for both auction types.
- Independent to the aforementioned process, to find a better downstream bid for first-price auctions, the exchange, as a bidding agent, uses a random exploration of $\beta_j \sim \mathcal{U}(\beta^* \cdot (1 - \epsilon_\beta), \beta^* \cdot (1 + \epsilon_\beta))$ where β^* is the the currently exploited bid shading factor, and $\epsilon_\beta > 0$ is a tunable parameter that determines the exploration bounds. Exploring in $[\beta^* \cdot (1 - \epsilon_\beta), \beta^*]$ is necessary to learn the distribution of the minimum bid to win below bid*, i.e., the left-censoring point. Exploring in

Algorithm 1 DCJL: Dual Censored Joint Learning of upstream reserve (via α^*) and downstream bid (via β^*) in first-price auctions

Input: $D_E(v_j; \alpha_j, \beta_j)$: a set of Exploration data for training
Output: $(\alpha^*, \beta^*) = argmax_{\alpha,\beta} \sum_{v \in D_E} profit(v; \alpha, \beta)$

α_j and β_j are bucketized in NB_α and NB_β number of bins, respectively, which result into $D_E(v_j; \alpha_j, \beta_j, bin_\alpha, bin_\beta)$
SET $max\_effective\_profit = 0$; $\alpha^* = 1$; $\beta^* = 1$
for $i_\alpha = 1$ to NB_α **do**
  for $i_\beta = 1$ to NB_β **do**
    SET $D_G = D_E(v_j; \alpha_j, \beta_j, bin_\alpha = i_\alpha, bin_\beta = i_\beta)$
    SET $e\_profit = \sum_{v \in D_G} profit(v; \alpha, \beta) \, / \, |D_G|$
    if $e\_profit > max\_effective\_profit$ **then**
      SET $max\_effective\_profit = e\_profit$
      SET $\alpha^* = \bar{\alpha}, \forall \alpha \in \{\alpha | v \in D_G\}$
      SET $\beta^* = \bar{\beta}, \forall \beta \in \{\beta | v \in D_G\}$
    end if
  end for
end for
return (α^*, β^*)

$[\beta^*, \beta^* * (1 + \epsilon_\beta)]$ helps to learn the distribution of the minimum bid to win above bid^*, i.e., the right-censoring point. Any values below β_{min} and above β_{max} are set to β_{min} and β_{max}, respectively. Since truth-telling is a dominant strategy in second-price auctions, the exchange should always bid its private valuation, $revenue(v_j)$, to maximize its profit, i.e. $\beta_j = 1$. Hence, no exploration is performed in these auctions.

The exploration process sacrifices some traffic to observe data that would otherwise be censored which in turn helps the process learn optimal α^* and β^*. Given the collected exploration data, D_E, these values are then chosen such as to maximize profit as described in Algorithm 1 (DCJL). More specifically, we utilize binning to group the continuous values into a smaller number of bins. The explored α and β values are binned in NB_α and NB_β bins, respectively. For each joint combination of bins (a total of $NB_\alpha \times NB_\beta$ combinations), the corresponding profit is calculated to find its maximum value. One heuristic is to set $NB_\alpha = NB_\beta = \lfloor \sqrt{|D_E|/data_c} \rfloor$ where $data_c$ is a minimum traffic threshold for each combination.

DCJL bucketizes the data in $NB_\alpha \times NB_\beta$ chunks. For second-price auctions $\beta = 1$ and hence the optimization becomes $\alpha^* = argmax_\alpha \sum_{v \in D_E} profit(v; \alpha, \beta = 1)$. So, DCJL can be simplified to bucketizing the data in NB_α chunks. In this case, $NB_\alpha = \lfloor |D_E|/data_c \rfloor$.

In a **hybrid auction setup**, where the auction type can vary by request for the same inventory, we can separately optimize first- and second-price traffic using the aforementioned approach. However, another way to exploit the co-existence of both auction types is to infer the average winning price, $\bar{w_2}$, from second-price traffic to utilize it for first-price traffic by adjusting the β_{min} (was $r_d(v_j)/revenue(v_j)$ before) and β_{max} (was 1) as:

$$\beta_{min} = \max(\bar{w_2} * (1 - \epsilon), r_d(v_j))/revenue(v_j) \tag{7}$$

$$\beta_{max} = \min(1, \bar{w_2} * (1 + \epsilon)/revenue(v_j)) \tag{8}$$

where $\epsilon \in (0, 1]$ is a tunable parameter to deal with the estimation error of $\bar{w_2}$. This method is inspired by the Revenue Equivalence Theorem (RET) [7] whereby the expected clearing price is the same in both auction types for the same inventory under general conditions. Making use of this inference, we avoid unnecessarily exploring in non-interesting areas which reduces the cost of the exploration. We refer to this method as DCJL_RET since it still uses the same DCJL except that D_E is now collected with the adjusted exploration process that uses Eq. (7) and (8). The method is applicable for hybrid first-price traffic.

5 EXPERIMENTS

5.1 Dataset and Experimental Settings

We verify our results using A/B tests. More specifically, we ran our experiment at LoopMe exchange during the last 3 weeks of March 2021. The experiment comprised 50 placements from 10 SSPs. 20 placements had first-price (*FP* placements) and 20 second-price (*SP* placements) auctions. Each of the remaining 10 placements' traffic involved both first- and second-price auctions. We will call these placements *hybrid* to differentiate from the former *non-hybrid* ones.

Since this is the first work solving this problem and also due to the synchronous setting that is not present in previous works, we compare our methods against a baseline strategy used by account managers with fixed α and β values tuned based on business objectives. For this strategy, $\alpha_{base} = 1.25$ and $\beta_{base} = 0.8$ for first-price auctions, i.e. the downstream reserve price and upstream clearing price are multiplied and divided respectively by the same 1.25. For second-price auctions, $\alpha_{base} = 1.25$ and $\beta_{base} = 1$.

For all placements, we randomly split the traffic into three groups: 10% is allocated to the *Baseline* group that implements the baseline strategy, 10% is allocated to the *Exploration* group that performs the dual censoring exploration, and the remaining 80% is allocated to the *Exploitation* group that optimizes reserve prices and bids. We call the combination of exploration and exploitation groups the *Test* group, and compare our results between test and baseline.

The main evaluation metric is $ProfitLift = 100 \cdot (\frac{e\_profit\_1}{e\_profit\_0} - 1)$ where e_profit_0 and e_profit_1 are the profit per request for the baseline and test groups, respectively. This makes the two groups comparable by standardizing the number of requests. We are not reporting the actual profit values for confidentiality reasons.

For each placement, we receive new data and update α^* and β^* every hour, to deal with non-stationarity, for the exploitation group using historical exploration data of the last 48 hours. Recall the exploration process where the explored β, if greater than β_{max} (which is 1), was set to β_{max}. Due to this, DCJL will not return β^* above it. DCJL may however return β^* below β_{min} (because β_{min} varies on request level), which, similar to the exploration process, is set to β_{min} for the applicable requests. What's more, for the duration of the experiment we use $\epsilon_\alpha = \epsilon_\beta = 0.5$, $\epsilon = 0.1$ and $data_c = 20000$, as we have found these values to perform well on historical data. For each hybrid placement, the learning for second- and first-price traffic uses the exploration data of the second- and first-price portions only, respectively. In every run, we additionally compute $\bar{w_2}$ from the entire second-price data of the last 48 hours.

For hybrid placements, we also want to compare the performance of DCJL against DCJL_RET on the same set of placements for first-price auctions. To mitigate interference, given that both methods are sharing historical exploration data for the optimization, we follow a different experiment design. More specifically, every 4

Table 1: A|B Test Results for Non-Hybrid First- and Second-Price Placements by Scenarios.

Placement Scenario, %	AT	Group	Requests (in Billions)	Successes / Requests (%)	Wins / Successes (%)	Avg_r_d ($eCPM)	Profit Lift (%)	$\bar{\alpha^*}$	$\bar{\beta^*}$
$\bar{\alpha^*} > \alpha_{base}$ & $\bar{\beta^*} > \beta_{base}$, (21.2%)	1	1	3.456	21.06	18.73	0.63	10.14	1.41	0.84
		0	0.384	27.90	11.52				
$\bar{\alpha^*} < \alpha_{base}$ & $\bar{\beta^*} < \beta_{base}$, (11.3%)	1	1	1.845	35.07	8.83	0.52	31.11	1.23	0.64
		0	0.205	30.69	13.51				
$\bar{\alpha^*} > \alpha_{base}$ & $\bar{\beta^*} < \beta_{base}$, (39.9%)	1	1	6.516	19.96	15.12	0.48	37.81	1.63	0.67
		0	0.724	34.59	13.48				
$\bar{\alpha^*} < \alpha_{base}$ & $\bar{\beta^*} > \beta_{base}$, (27.6%)	1	1	4.5	39.88	13.08	1.05	18.19	1.19	0.82
		0	0.5	31.93	11.72				
$\bar{\alpha^*} > \alpha_{base}$, (60.1%)	2	1	10.35	19.75	21.31	0.59	30.92	1.57	1.0
		0	1.15	33.30	12.10				
$\bar{\alpha^*} < \alpha_{base}$, (39.9%)	2	1	6.876	39.62	9.94	0.66	24.53	1.21	1.0
		0	0.764	26.19	11.72				

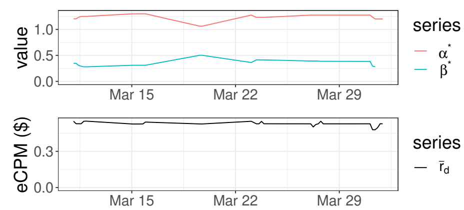

Figure 2: Hourly updated values for an FP placement.

days, we alternate between DCJL and DCJL_RET but use only the most recent 48 hours to measure lift. This makes sure that the right exploration data is used for each method, discarding the exploration data that had been exposed to the previous strategy

5.2 Evaluation

5.2.1 ***Non-Hybrid Placements.*** Table 1 groups non-hybrid placements based on their auction type, AT, and how the average α^* and β^* values during exploitation within the test (Group=1) group compare against the baseline (Group=0), and details the corresponding results. *Requests* denotes the number of ad-requests from the SSPs. *Successes*/*Requests* (called the success rate) is the ratio of Successes to Requests, where Successes is defined as the number of requests that end up with a winner in the upstream auction. $Wins/Successes$ is the ratio of total Wins to Successes where Wins are the requests for which the exchange wins at the SSP. SSPs may vary their reserve prices, r_d, over time both at an intra- and inter-placement level; their average is reported as *Avg*_r_d. $\bar{\alpha^*}$ and $\bar{\beta^*}$ are the averages of all hourly updated values of α^* and β^*, respectively, which were used in the exploitation group. Figure 2 illustrates α^* and β^* values for an example FP placement, as these were updated at every hourly run, along with the average received r_d. Note that placements differ in terms of publishers, bidders, number of requests or other targeting attributes. Hence, the effects of the optimization can be very different for each placement. For this reason, we report the overall profit lift and not the average across placements.

In all cases where $\bar{\alpha^*} > \alpha_{base}$, the success rate of the test group is lower than that of the baseline group. This is expected as the higher the reserve price, the lower the probability that the upstream auction clears (i.e. there is at least one bid above the reserve price). Despite the lower success rate, a higher reserve price enforces the bidders to bid above it and hence results in higher upstream clearing price per success. Conversely, $\bar{\alpha^*} < \alpha_{base}$ increases the success rate but lowers the upstream clearing price per success.

For FP placements, looking at the first scenario ($\bar{\alpha^*} > \alpha_{base}$ & $\bar{\beta^*} > \beta_{base}$) that comprises 21.2% of all FP placements, the wins-to-successes rate in the test (18.73%) is higher than the baseline (11.52%). This happens since higher reserve prices lead to fewer successes but higher upstream clearing prices per success, followed by higher $\bar{\beta^*}$ values that increase the chance of winning at the SSP given the upstream auction has cleared. The opposite happens when $\bar{\alpha^*} < \alpha_{base}$ & $\bar{\beta^*} < \beta_{base}$; lower reserve prices lead to more successes but lower upstream clearing prices per success which, followed by lower $\bar{\beta^*}$ values, reduce the wins-to-successes rate.

When $\bar{\alpha^*} > \alpha_{base}$ & $\bar{\beta^*} < \beta_{base}$, or $\bar{\alpha^*} < \alpha_{base}$ & $\bar{\beta^*} > \beta_{base}$, it is not only the relative differences but also the actual values of α and β that dictate the ranking of the wins-to-successes rate for the two groups. For both scenarios, in our experiment, the wins-to-successes rate was higher for the test group compared to the baseline. In the former case, smaller $\bar{\beta^*}$ values reduce the winning probability at the SSP but the set reserve prices were high enough (yielding higher upstream clearing prices per success) to compensate and bids were competitive enough to win at the SSP. That is, the impact of higher $\bar{\alpha^*}$ values dominated the contribution of smaller $\bar{\beta^*}$ values. However, the results could have been in the opposite direction had the contribution of $\bar{\beta^*}$ dominated instead. When $\bar{\alpha^*} < \alpha_{base}$ and $\bar{\beta^*} > \beta_{base}$, lowering the reserve price resulted in a lower upstream clearing price per success but $\bar{\beta^*}$ was high enough to lead to competitive bids at the SSP. As before, it could have been possible for the test group to achieve a lower wins-to-successes rate than the baseline had the contribution of $\bar{\beta^*}$ been dominated by that of the lower reserve price instead.

For SP placements, there are only two cases, $\bar{\alpha^*} > \alpha_{base}$ and $\bar{\alpha^*} < \alpha_{base}$. The former scenario comprises 60.1% of all SP placements. Here, the success rate in the test group is lower than that of the baseline but the wins-to-successes rate is higher; a higher reserve price results in a higher upstream clearing price per success which, when passed to the SSP, increases the probability of winning downstream given the upstream has cleared. Conversely, when $\bar{\alpha^*} < \alpha_{base}$, the success rate in the test group is higher than that of the baseline group but the wins-to-successes rate is lower.

Table 2: A|B Test Results for Hybrid Placements.

Type of Placement	Group	Requests (in Billions)	Successes / Requests (%)	Wins / Successes (%)	Avg_r_d ($eCPM)	Profit Lift (%)	$\bar{\alpha^*}$	$\bar{\beta^*}$
Hybrid SP	1	5.633	24.77	15.36	0.63	24.19	1.42	1.0
	0	0.625	32.83	10.58				
Hybrid FP (DCJL)	1	1.265	19.14	15.87	0.70	23.26	1.31	0.74
	0	0.140	25.68	13.09				
Hybrid FP (DCJL_RET)	1	1.252	17.46	16.29	0.69	27.37	1.35	0.75
	0	0.139	24.13	12.47				

Table 3: Aggregated A|B Test Results.

Type of Placement	Nb Placements	Profit Lift (%)
FP	20	25.15
SP	20	28.27
Hybrid	10	24.16

Note that maximizing the number of requests that convert into wins does not necessarily maximize the exchange's profit since there exists an inherent trade-off between win rate (Wins/Requests) and profit per win; a smaller number of wins but with higher profit per win might be better than a higher number of wins but lower profit per win depending on the characteristics of the placements. In our experiment, profit lift was the highest when $\bar{\alpha^*} > \alpha_{base}$ & $\bar{\beta^*} < \beta_{base}$ and $\bar{\alpha^*} > \alpha_{base}$ for FP and SP placements, respectively.

5.2.2 ***Hybrid Placements.*** Table 2 displays the performance of hybrid placements. The measured lifts are between the test and baseline groups, similar to non-hybrid placements reported in Table 1. In our data, 69% of the hybrid traffic operated using a second-price auction denoted by $Hybrid\ SP$ where we observed a profit lift of 24.19%. The remaining 31% hybrid traffic operated using a first-price auction denoted by $Hybrid\ FP$ (although not all 31% is reported due to unmeasured days; see Section 5.1). Using the inference led to an increased lift (27.37%) for Hybrid FP in comparison to the non-inference setting that gives a lift of 23.26%. In addition, we measure how the inference reduces the exploration cost by comparing the two strategies. For this, we use the inferential exploration data which is a subset of the test group of Hybrid FP (DCJL_RET) and non-inferential exploration data which is a subset of the test group of Hybrid FP (DCJL). The profit per request of these two exploration sets, denoted by e_profit_I and e_profit_NI, respectively, compute $exploration\_profit\_lift$ as 100*(e_profit_I/e_profit_NI - 1), which equals 36.37%.

5.2.3 ***Overall Results.*** Table 3 displays the overall impact on the placements studied. $Nb\ Placements$ denotes the number of placements within a particular type of placement. FP and SP placements achieved a lift of 25.15% and 28.27%, respectively. These are the overall lifts of all the FP and SP placements, respectively, reported in Table 1. The hybrid placements achieved a lift of 24.16%, which is the overall lift of all the hybrid placements.

In general, FP placements are a bit harder to optimize in comparison to SP placements. While SP optimization requires just one parameter, α, to be learned (with the benefit of a lower exploration cost), joint learning of α and β in FP placements increases the exploration cost. As noted before, placements are heterogeneous that may also justify the differences in FP vs SP results. Lastly, for the hybrid placements, as detailed in Section 5.2.2, results suggest that the adjusted exploration process for the first-price traffic using the inference setup can play a significant role increasing profit for the exchange when dealing with hybrid inventories.

6 CONCLUSIONS AND FUTURE WORK

In this work, we have introduced a novel and practical framework for jointly optimizing reserve prices and bids for intermediaries in ad exchanges for both first- and second-price auctions that takes into account the censoring problem for both upstream and downstream bids. To the best of our knowledge, this is the first work that addresses profit optimization for intermediaries in first-price auctions. What's more, we have proposed an elegant strategy based on the Revenue Equivalence Theorem to deal with inventory where both auction types are present. We have finally validated the effectiveness of the proposed strategies with the use of A/B tests in a real exchange (LoopMe). Our framework can be extended in a number of ways. More specifically, different functions can be fitted to learn the intermediary's expected profit and other, non-linear functions of the bids and reserve prices can be used.

REFERENCES

[1] Miguel Angel Alcobendas Lisbona, Sheide Chammas, and Kuang-chih Lee. 2016. Optimal Reserve Prices in Upstream Auctions: Empirical Application on Online Video Advertising. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (San Francisco, California, USA) *(KDD '16)*. Association for Computing Machinery, New York, NY, USA, 1395–1404. https://doi.org/10.1145/2939672.2939877

[2] Santiago R. Balseiro, Ozan Candogan, and Huseyin Gurkan. 2019. *Intermediation in Online Advertising*. Springer International Publishing, Cham, 505–528. https://doi.org/10.1007/978-3-030-01863-4_21

[3] R.C. Elandt-Johnson and N.L. Johnson. 1999. *Survival Models and Data Analysis*. John Wiley and Sons.

[4] Jon Feldman, Vahab Mirrokni, S. Muthukrishnan, and Mallesh M. Pai. 2010. Auctions with Intermediaries: Extended Abstract. In *Proceedings of the 11th ACM Conference on Electronic Commerce* (Cambridge, Massachusetts, USA) *(EC '10)*. Association for Computing Machinery, New York, NY, USA, 23–32. https://doi.org/10.1145/1807342.1807346

[5] Djordje Gligorijevic, Tian Zhou, Bharatbhushan Shetty, Brendan Kitts, Shengjun Pan, Junwei Pan, and Aaron Flores. 2020. Bid Shading in The Brave New World of First-Price Auctions. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management* (Virtual Event, Ireland) *(CIKM '20)*. Association for Computing Machinery, New York, NY, USA, 2453–2460. https://doi.org/10.1145/3340531.3412689

[6] Grégoire Jauvion, Nicolas Grislain, Pascal Dkengne Sielenou, Aurélien Garivier, and Sébastien Gerchinovitz. 2018. Optimization of a SSP's Header Bidding Strategy Using Thompson Sampling. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (London, United Kingdom) *(KDD '18)*. Association for Computing Machinery, New York, NY, USA, 425–432. https://doi.org/10.1145/3219819.3219917

[7] Vijay Krishna. 2010. *Auction Theory*. Academic Press.

[8] Lampros C. Stavrogiannis. 2014. *Competition between demand-side intermediaries in ad exchanges*. Ph.D. Dissertation. University of Southampton. https://eprints.soton.ac.uk/372683/