AdKDD
August 15, 2021

Virtual Conference
KDD2021
August 14th – 18th

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

Piyush Paliwal, Lampros C. Stavrogiannis

LoopMe Ltd.

AdKDD 2021

Outline

- Introduction
- Related Work
- Problem Formulation
- Proposed Methodology
- Experiments
- Conclusions

Introduction

Supply Path

- Downstream SSPs hold first- or second-price auctions (or both) for their demand partners (downstream DSPs + exchanges) -> *Downstream auction*.
- Exchange resells the opportunity for its demand partners (upstream DSPs) -> *Upstream auction*.
- Exchange submits a bid to compete downstream.

Goal

- Maximize the exchange’s profit by learning the upstream reserve price and downstream *bid*.

Challenges

- **Synchronous/joint** learning of *reserve* and *bid*; profit is contingent on winning downstream.
- **Dual censoring** on upstream and downstream markets.
- Presence of both auctions; **hybrid**.

Related Work

- Lisbona, KDD 16 BrightRoll Exchange with second-price upstream auction under revenue sharing with the SSP. Focus on reserve prices.

- Jauvion, KDD 18 First-price bidding strategies. Focus on bid optimization.

Our work is the joint optimisation of both supply and demand.

Problem Formulation

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

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

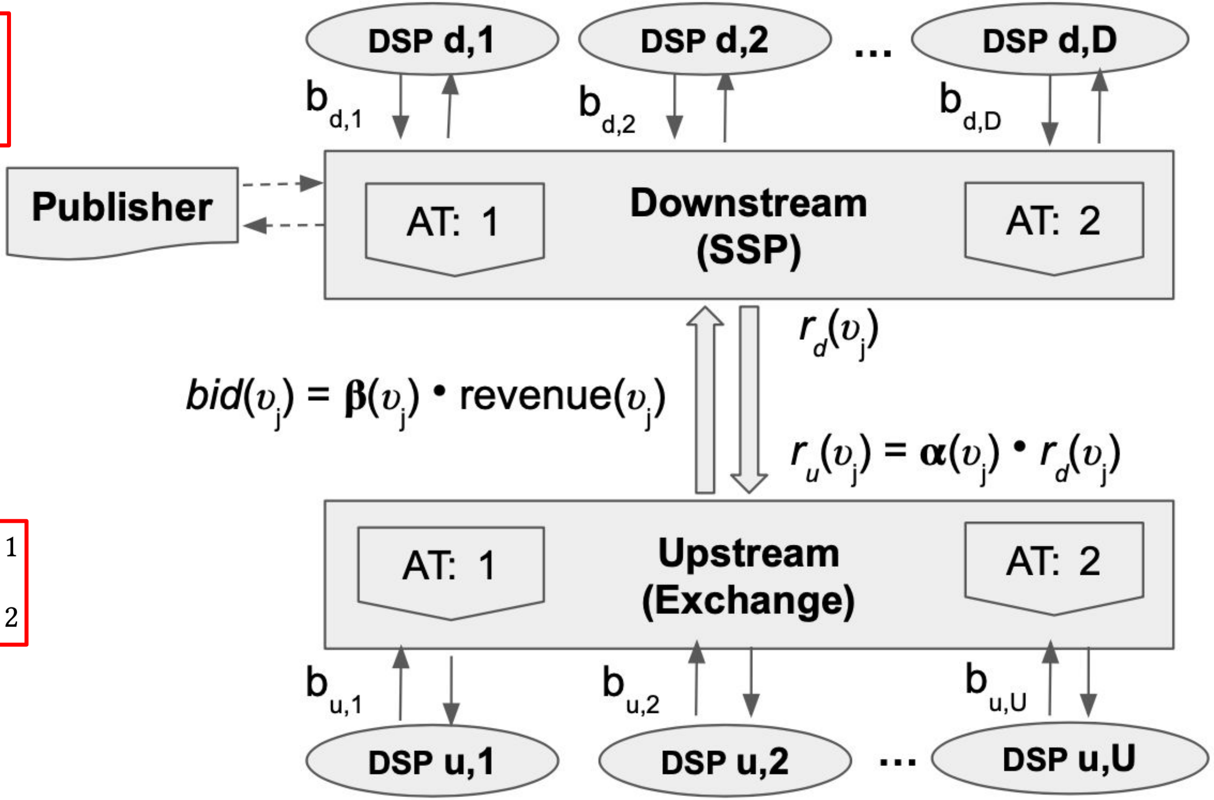

Open Marketplace Setting

Problem Formulation

- Reserve multiplier $\boldsymbol{\alpha(v_j)}$ **>=1.0** and bid shading factor $\boldsymbol{\beta(v_j) \in [r_d(v_j)/\mathrm{revenue}(v_j), 1]}$ are decision variables.

- For won bids, the profit $\boldsymbol{\mathrm{profit}(v_j; \alpha_j, \beta_j)} = \boldsymbol{\mathrm{revenue}(v_j) - \mathrm{cost}(v_j)}$.

- For lost bids, the profit equals 0.

- Revenue, cost and profit are *contingent* on winning downstream.

Profit is maximized by jointly learning the two ingredients: $\boldsymbol{\alpha_j}$ and $\boldsymbol{\beta_j}$.

Proposed Methodology

Dual Censoring

- ❖ Upstream DSPs’ bids are left-censored.
- ❖ Downstream winning feedback is right-censored (second-price) and left- plus right-censored (first-price).

To deal with dual censoring, exchange performs the two following explorations, independently:

$$\alpha_j \sim U(\alpha^* \cdot (1 - \varepsilon_\alpha), \alpha^* \cdot (1 + \varepsilon_\alpha))$$

$$\beta_j \sim U(\beta^* \cdot (1 - \varepsilon_\beta), \beta^* \cdot (1 + \varepsilon_\beta))$$

Resulting exploration data is used to find the optimal α and β values that maximize the profit.

Proposed Methodology

Heuristic:

$$NB_\alpha = NB_\beta = \lfloor \sqrt{|D_E|/data_c} \rfloor$$

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; α_j, β_j, bin_α, bin_β)
SET max_effective_profit = 0; α* = 1; β* = 1
for i_α = 1 to NB_α do
    for i_β = 1 to NB_β do
        SET D_G = D_E(v_j; α_j, β_j, bin_α = i_α, bin_β = i_β)
        SET e_profit = Σ_{v∈D_G} profit(v; α, β) / |D_G|
        if e_profit > max_effective_profit then
            SET max_effective_profit = e_profit
            SET α* = ᾱ, ∀α ∈ {α|v ∈ D_G}
            SET β* = β̄, ∀β ∈ {β|v ∈ D_G}
        end if
    end for
end for
return (α*, β*)
```

Proposed Methodology

- Following the dominant strategy in second-price, the optimization becomes:

$$argmax_{\alpha} \sum_{v \in D_E} profit(v; \alpha, \beta = 1)$$

- Hybrid inferential learning (DCJL_RET) utilizing Revenue Equivalence Theorem.

$$\boxed{\begin{aligned} \beta_{min} &= \max(\bar{w_2} * (1 - \epsilon), r_d(v_j))/revenue(v_j) \\ \beta_{max} &= \min(1, \bar{w_2} * (1 + \epsilon)/revenue(v_j)) \end{aligned}}$$

This reduces the cost of the exploration process.

loopMe

Experiments: Settings

- Last 3 weeks of March 2021 at LoopMe exchange.

- 20 first-price (FP placements), 20 second-price (SP placements) and 10 hybrid placements.

- 10% Baseline, 10% Exploration and 80% Exploitation traffic. Test group combines exploration and exploitation

- Measured profit lift between test (group 1) and baseline (group 0):

$$(e\_profit\_1 / e\_profit\_0 - 1) * 100$$

Experiments: Settings

- Baseline strategy:

  $\alpha_{base} = 1.25$ and $\beta_{base} = 0.8$ in first-price.
  $\alpha_{base} = 1.25$ and $\beta_{base} = 1$ in second-price.

- Parameters:

  $\varepsilon_{\alpha} = \varepsilon_{\beta} = 0.5$, $\varepsilon = 0.1$ and $data_c = 20000$

- Hybrid setting:

  To compare DCJL against DCJL_RET on the same set of FP placements, we alternate every 4 days between DCJL and DCJL_RET but use only the most recent 48 hours to measure lift.

Experiments: Results

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				

loopme

Experiments: Results

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

loopme

Experiments: Results

- FP placements are harder to optimize in comparison to SP placements.

    - ★ SP optimization requires just α to be learned; a lower exploration cost.
    - ★ Joint learning of α and β in FP increases the exploration cost.
    - ★ Heterogeneous nature of placements may also justify the differences.

- For hybrid placements, inferential learning using RET has played a significant role increasing the profit.

    - ★ In addition to overall improvement using RET (Table 2), exploration_profit_lift, 100*(e_profit_I/e_profit_NI - 1) is also measured, which equals 36.37%. This is measured only using the exploration sets of DCJL_RET and DCJL for hybrid FP.

Conclusions

- Introduced an efficient framework for jointly learning upstream reserve prices and downstream bids for first- and second-price auctions under dual censoring.
- Proposed an elegant strategy based on the RET to deal with hybrid inventory.
- A/B tested methods at LoopMe exchange.

Future work:

- Could include non-linear functions of the bids and reserve prices.

Thank you!

Piyush Paliwal

piyush.paliwal@loopme.com

piyushpaliwal90@gmail.com

Lampros C. Stavrogiannis

lampros@loopme.com

lstavr@gmail.com

Please reach out with questions/feedback