Fuzzy logic based pricing combined with adaptive search for reserve price optimization in online ad auctions

Citation for published version (APA):

Rhuggenaath, J., Akcay, A., Zhang, Y., & Kaymak, U. (2019). Fuzzy logic based pricing combined with adaptive search for reserve price optimization in online ad auctions. In *2019 IEEE International Conference on Fuzzy Systems, FUZZ 2019* [8858975] Institute of Electrical and Electronics Engineers. <https://doi.org/10.1109/FUZZ-IEEE.2019.8858975>

DOI:

[10.1109/FUZZ-IEEE.2019.8858975](https://doi.org/10.1109/FUZZ-IEEE.2019.8858975)

Document status and date:

Published: 01/06/2019

Document Version:

Accepted manuscript including changes made at the peer-review stage

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Fuzzy Logic based Pricing combined with Adaptive Search for Reserve Price Optimization in Online Ad Auctions

Jason Rhuggenaath, Alp Akcay, Yingqian Zhang and Uzay Kaymak
School of Industrial Engineering, Eindhoven University of Technology
Eindhoven, The Netherlands
Email: {j.s.rhuggenaath, a.e.akcay, yqzhang, u.kaymak}@tue.nl

Abstract—In this paper we consider an online publisher that sells advertisement space and propose a method for learning optimal reserve prices in second-price auctions. We study a limited information setting where the values of the bids are not revealed and no historical information about the values of the bids is available. Our proposed method combines an adaptive search procedure with a fuzzy logic pricing step to set reserve prices and is suitable for non-stationary environments. In the fuzzy logic pricing step, we take the gap between the winning bid and second highest bid into account and show that this leads to better decisions for the reserve prices. Experiments using real-life ad auction data show that the proposed method outperforms popular bandit algorithms.

I. INTRODUCTION

One of the main mechanisms that web publishers use in online advertising in order to sell their advertisement space is the real-time bidding (RTB) mechanism [1]. In RTB there are three main platforms: supply side platforms (SSPs), demand side platforms (DSPs) and an ad exchange (ADX) which connects SSPs and DSPs. The SSPs collect inventory of different publishers and thus serve the supply side of the market. Advertisers which are interested in showing online advertisements are connected to DSPs. When a user visits a webpage with an advertisement (ad) slot, the publisher sends a request to the ADX (via an SSP) indicating that an impression can potentially be displayed in this particular ad slot. At the same time, advertisers that are connected to DSPs send bid requests to the ADX indicating that they are willing to bid for this impression. A real-time auction then decides which advertiser is allowed to display its ad and the amount that the advertiser needs to pay. The most popular auction mechanism is the second-price auction, where the winning advertiser pays the second highest bid in the auction.

Publishers can set a reserve price for their inventory in the second-price auction. Due to the reserve price, all bids below the reserve price are disregarded, and as a consequence, there is a possibility that the ad slot is not sold. If the auction does have a winner, the winner pays the maximum of the second highest bid and the reserve price. In this paper we take the perspective of an online publisher that submits his inventory of advertisement space to an SSP and needs to decide on the optimal value of the reserve price. We assume that the publisher has limited information about the winning

bid and second highest bid in the auction. More specifically, the publisher does not observe the actual values of the winning bid and second highest bid. After each sale attempt on the RTB market, the publisher only knows whether the sale was successful and the revenue that is received from that sale. This setting is relevant for publishers that are small and medium size enterprises (SMEs), since the ADX and the connected SSPs typically do not reveal the actual bids placed in the auction but only the result of the auction. Due to the limited feedback, the publisher faces an exploration-exploitation trade-off. He needs to experiment with different reserve prices to figure out which one works best, but at the same time, he does not want to explore too much since he wants to use the best reserve price as much as possible (exploitation). In this paper we present a method called FL-ETC-RAP (fuzzy logic based explore-then-commit algorithm with risk-aware pricing) that addresses the problem of the publisher. Compared to previous studies our method does not assume that the publisher observes the actual values of the winning bid and second highest bid. In FL-ETC-RAP we cycle through four steps where, the first two steps are exploration phases, the third step is a commit phase, and the fourth step is a test phase. In the first step we estimate the value of the second highest bid and in the second step we estimate a reference value of the winning bid that represents a critical boundary: reserve prices above this reference value tend to be too high and reserve prices below the reference value tend to lead to successful sales of impressions. Our method uses an adaptive search procedure to learn this critical boundary. In the third step the publisher uses a fuzzy logic approach to combine the estimates for the winning and second highest bid in order to determine a reserve price. He subsequently commits to using this reserve price for a fixed amount of periods. In the fourth step, a test is done to see whether we should start another commit phase, or, go back and start with another exploration phase.

We summarize the main contributions of this paper as follows:

- We propose a method for learning optimal reserve prices in a limited information setting where the values of the bids are not revealed and no historical information about the values of the bids is available. Our proposed method

is suitable for non-stationary environments.

- We show that taking the gap between the winning bid and second highest bid into account leads to better decisions for the reserve prices (we refer to this as risk-aware pricing).
- Our proposed method combines an adaptive search procedure with a fuzzy logic pricing step to set reserve prices. To the best of our knowledge, the problem of reserve price optimization has not been studied before by using approaches from fuzzy logic.
- Experiments using real-life ad auction data show that the proposed method outperforms popular bandit algorithms.

The remainder of this paper is organized as follows. In Section II we discuss the related literature. Section III provides a formal formulation of the problem. In Section IV we present the our proposed method for setting reserve prices. In Section V we perform experiments and compare our method with baseline strategies in order to assess the quality of our proposed method. Section VI concludes our work and provides some interesting directions for further research.

II. RELATED LITERATURE

The problem of maximizing revenues in online advertising has received increasing attention in the machine learning literature over the last decade. In [2] an online learning approach is used to derive a policy for setting optimal reserve prices, but the analysis makes the assumption that the environment is stationary. Other works such as [3], [4], [5], [6] use historical data to directly predict the optimal reserve price or the winning bid (which can indirectly be used to set a reserve price). A drawback of using models based on historical data is that they may not perform well in non-stationary environments. Most of the studies mentioned above do not set reserve prices in an adaptive way that adjusts to changing environments. Some studies such as [5], [7], [8] do study adaptive reserve prices, but they assume that the winning bid and/ or second highest bid are observed. In this paper we do not make this assumption. In [9] a different but related problem is studied, namely optimizing revenues for an SSP using a header-bidding strategy. Finally, we note that [10] studied a problem in a similar setting and proposed a parametric method based on Thompson sampling. In terms of methodology, there are a number of studies in the online advertising domain that use techniques from fuzzy logic. There are studies that focus on bidding behavior estimation in keyword auctions [11], budget allocation in search auctions [12], and bidding price optimization in keyword auctions [13]. To the best of our knowledge, the problem of reserve price optimization has not been studied before by using approaches from fuzzy logic. To summarize, the main differences between this paper and previous works are that: (i) we show how to set adaptive reserve prices in possibly non-stationary environments; (ii) we do not assume that the publisher observes the top two bids in the ad auction, but only observes the revenue of each auction; (iii) we show how fuzzy logic combined with an adaptive

search procedure can be used as an effective strategy in order to set reserve prices in online ad auctions.

III. PROBLEM STATEMENT

We consider a publisher that owns a single advertisement slot and that there is a sequence of impressions (corresponding to this advertisement slot) arriving over time. Time is discretized, and time periods are denoted by $t \in \mathbb{N}$. At the beginning of each time period (that is, upon arrival of an impression) the publisher has to decide on a reserve price $p_t \in [p_l, p_h] = [0, p_h]$. The prices p_l, p_h are the minimum and maximum reserve prices that are acceptable to the publisher. After deciding on a reserve price, the impression is offered for sale on the RTB-market via a Supply Side Platform (SSP). The SSP runs a second-price auction for the impression and the revenue of the publisher depends on the outcome of this auction. Let X_t and Y_t denote the highest and second highest bid respectively in the auction for the impression at time t . Then the revenue (or return) of the publisher at time t is given by $R_t = \mathbb{I}\{p_t \leq X_t\} \cdot \max\{Y_t, p_t\}$. Here $\mathbb{I}\{A\} = 1$ if A is true and $\mathbb{I}\{A\} = 0$ otherwise. The expression for R_t says that if the reserve price p_t is higher than the winning bid ($p_t > X_t$) then the publisher receives zero revenue. If the reserve price does not exceed the winning bid ($p_t \leq X_t$) then the revenue equals the maximum of the second highest bid and the reserve price. Note that, in general, the publisher does not observe the value of X_t and Y_t after a (successful) sale.

Assumption 1: We assume that all bids are non-negative, that is, $X_t > Y_t \geq 0$. If a sale was not successful, that is, if $p_t > X_t$, then the publisher does not observe X_t and Y_t . If a sale was successful, that is, if $p_t \leq X_t$, then the publisher observes $\max\{Y_t, p_t\}$. ■

Assumption 1 formalizes the setting that is relevant for SME publishers, since the ADX and the connected SSPs typically do not reveal the actual bids placed in the auction but only the result of the auction. The objective of the publisher is to maximize the cumulative revenue over the sales horizon of length T . Thus the revenue optimization problem over T time periods or impressions can be expressed as follows:

$$\max_{p_1, \dots, p_T} \mathbb{E} \left\{ \sum_{t=1}^T \mathbb{I}\{p_t \leq X_t\} \cdot \max\{Y_t, p_t\} \right\} \quad (1)$$

Remark 1: In the literature on online advertising and the RTB-market, the reserve price is sometimes also referred to as the *floor price*. In the remainder of this paper we will use the term *top bid* to refer to the winning bid in the online auction and we will use the term *second bid* to refer to the second highest bid. ■

Remark 2: In order to simplify the exposition of our method, we focus on the case where there is a single ad slot. However, in practice, the publisher may want to set a reserve price depending on the characteristics of the user and the ad slot. Our method can also be applied in such a setting by, for example, making segments of users and applying our method for each segment. ■

IV. PROPOSED APPROACH: FL-ETC-RAP

In this section we present a method to sequentially set the reserve price p_t . We will refer to the method developed in this section as FL-ETC-RAP (fuzzy logic based explore-then-commit algorithm with risk-aware pricing). It uses an adaptive search procedure to learn an estimate of the top bid. This estimate is then combined with an estimate of the second bid to set a reserve price.

Our method follows the following main steps:

- 1) Determine a reference value for the second bid.
- 2) Using the reference value for the second bid as a reference point, determine the reference value for the top bid.
- 3) Apply a *risk-aware* reserve price for E periods.
- 4) If the environment is favorable, apply the risk-aware reserve price for another E periods. Otherwise, go to step 1.

Our method cycles through 4 steps. The first two steps are called exploration phases, the third step is called a commit phase. In the fourth step, a test is done to see whether we should start another commit phase, or, go back and start with another exploration phase. We now elaborate on each of these phases.

A. Exploration phase I: Reference value for the second bid

In the first exploration phase, the goal is to determine a *reference value* for the second bid. The reference value represents a value that is supposed to be close to the second bid. In order to determine the reference value for the second bid we can exploit the structure of the second-price auction. More specifically, there is a risk-free way to actually observe the second bid. It can be accomplished by setting a reserve price equal to zero: $p_t = 0$. If $p_t = 0$, then (by Assumption 1) the observed revenue will be equal to the second bid since $R_t = \max\{Y_t, p_t\}$. In order to estimate the second bid, we apply a reserve price of zero for M periods and use the average of the observed revenues as our reference value. That is, we set $p_k = 0$ for $k = t, t+1, \dots, t+M-1$ and use $p^Y = \sum_{k=t}^{t+M-1} R_k / M$ as our reference value. Here M is a parameter of the algorithm that is chosen by the publisher.

B. Exploration phase II: Reference value for the top bid

In the second exploration phase, the goal is to determine a good reference value for the top bid. Determining a good reference value for the top bid is much harder since (by Assumption 1) the top bid is (in general) never observed. In Algorithm 1 we propose an estimator that performs well in our numerical experiments. The main idea behind the estimator in Algorithm 1 is as follows. The estimator takes as input a sample of K auction outcomes $\{(\mathbb{I}\{p_k \leq X_k\}, p_k)\}_{k=1}^K$ and returns a reference value p^X that represents a *critical boundary*. The interpretation of the critical boundary is that it is the highest reserve price such that reserve prices above p^X tend to lead to unsuccessful sales but reserve prices below p^X are successful.

Note that in order to use Algorithm 1 the values for p_k need to be chosen and the choices for these values will determine the quality and usefulness of the reference value that is obtained. In order to guide the process of choosing appropriate values for p_k , we use an adaptive search procedure.

Algorithm 1 Compute-Ref-X

Require: A sample of K reserve prices and outcomes $\{(\mathbb{I}\{p_k \leq X_k\}, p_k)\}_{k=1}^K$.

- 1: Set $A^1 = \{p_k | \mathbb{I}\{p_k \leq X_k\} = 1, 1 \leq k \leq K\}$.
- 2: Set $A^2 = \{p_k | \mathbb{I}\{p_k \leq X_k\} = 0, 1 \leq k \leq K\}$.
- 3: **if** $A^1 = \emptyset$ **then**
- 4: Set $p^X = \min\{p | p \in A^2\}$.
- 5: Set $I_f = -1$.
- 6: **else if** $A^2 = \emptyset$ **then**
- 7: Set $p^X = \max\{p | p \in A^1\}$.
- 8: Set $I_f = 1$.
- 9: **else**
- 10: Set $p^X = \min\{\max\{p | p \in A^1\}, \min\{p | p \in A^2\}\}$.
- 11: Set $I_f = -1$.
- 12: **end if**
- 13: **return** p^X and I_f .

C. Exploration phase II: Learning the reference value for the top bid using adaptive search

We now discuss the adaptive search (AS) procedure used to learn the reference value for the top bid. The length of this exploration phase is $S \cdot m$, where S and m are parameters of the algorithm that are specified by the publisher. The adaptive search explores reserve prices in a search neighborhood that is determined by the reference value for the second bid and a *search radius*. More specifically, if p_t^Y is the reference value available at time t , then the search neighborhood is: $[p_t^Y, p_t^Y + W]$. Here $W > 0$ is the search radius.

The full procedure for the AS step is described in Algorithm 2.

Algorithm 2 AS-Ref-X

Require: $S, m, t_s, W, p^Y, \kappa$

- 1: Set $t = t_s$.
- 2: **for** $k = 1$ **to** m **do**
- 3: **for** $i = 1$ **to** S **do**
- 4: Set $t = t + 1$.
- 5: **if** $k = 1$ **then**
- 6: Draw L from uniform distribution on $[p^Y, p^Y + W]$.
- 7: **else**
- 8: Set $L = G \cdot (1 + I_f \cdot \kappa)$.
- 9: **end if**
- 10: Set reserve price $p_t = L$.
- 11: Observe outcome of auction: $\mathbb{I}\{p_t \leq X_t\}$.
- 12: Run Compute-Ref-X (Algorithm 1) with input $\{(\mathbb{I}\{p_k \leq X_k\}, p_k)\}_{k=t_s}^t$ and p^X and I_f as output.
- 13: Set $G = p^X$.
- 14: **end for**
- 15: **end for**
- 16: **return** G

The main idea is to first choose S initial reserve prices uniformly at random from $[p^Y, p^Y + W]$ and then use the initial reserve prices in order to get an initial estimate for p^X .

After we have an initial estimate for p^X , we proceed to refine it by choosing the next $(m-1)S$ reserve prices in the direction that is most promising for learning the critical boundary (Line 8 in Algorithm 2). The motivation for the dynamics in Line 8 comes from the definition of the critical boundary as specified in Algorithm 1. If the sample used to determine p^X contains an unsuccessful sale, then we know that the next best estimate for the reference value p^X will be at most p^X , and therefore we need to explore reserve prices lower than p^X . The opposite reasoning holds for the case where there is no unsuccessful sale while determining p^X .

D. Commit phase: Risk-aware pricing using fuzzy logic

In the *commit phase* the publisher uses the reference values p_t^X and p_t^Y for the top bid and second bid in order to determine a reserve price p^* . The publisher then *commits* to using p^* for E periods, that is, $p_{t+1} = \dots = p_{t+E} = p^*$. Here E is a parameter of the algorithm that is chosen by the publisher. There are in general many ways to select p^* . In this paper we present a simple scheme that performs well in our numerical experiments. The intuition behind the scheme is as follows: (i) if the gap between Y_t and X_t is believed to be large, then try to set a reserve price above the second bid (but not too high); (ii) if the gap between Y_t and X_t is believed to be small, then choose a reserve price close to Y_t .

In order to quantify the gap between Y_t and X_t we look at relative gap $\gamma = |p_t^X - p_t^Y|/p_t^Y$ and absolute gap $\eta = |p_t^X - p_t^Y|$. Depending on the values of γ and η the publisher sets a reserve price according to $p^* = \omega p_t^X + (1 - \omega)p_t^Y$ for some $\omega \in (0, 1)$. By controlling the parameter ω the publisher can decide the degree to which he wants to exploit the fact that the gap between Y_t and X_t is large. In order to determine the value of ω we use an approach based on fuzzy logic [14], [15]. We assume that γ and η take on the linguistic values “small”, “medium” and “large” and define corresponding fuzzy sets. Similarly, we assume that ω takes on the linguistic values “small”, “medium” and “large”, and we define rules that relate the linguistic values of γ and η to those of ω . The rules of the fuzzy system are as follows:

- 1) IF γ is *small* THEN ω is *small*.
- 2) IF γ is *medium* THEN ω is *medium*.
- 3) IF γ is *large* THEN ω is *large*.
- 4) IF γ is *small* AND η is *medium* THEN ω is *medium*.
- 5) IF γ is *small* AND η is *large* THEN ω is *medium*.
- 6) IF γ is *large* AND η is *small* THEN ω is *small*.
- 7) IF γ is *medium* AND η is *small* THEN ω is *small*.

Together these rules model a situation where the gap between X_t and Y_t can be “small”, “medium” and “large”, and for larger gaps p^* is closer to p^X . The first three rules are baseline rules and the last four rules are exceptions to the baseline. The baseline rules state that if the relative gap increases then more weight should be put on p_t^X . The last four rules state that the publisher can deviate from the baseline, if the absolute gaps are large enough. The main motivation for using a fuzzy logic approach is that it allows us to (i) incorporate domain/expert knowledge and (ii) risk-preferences of the publisher in a

natural way. From the structure of the second-price auction, it follows that, reserve price optimization is only lucrative if the gap between the second bid and top bid is sufficiently large. However, is often hard to specify what constitutes a “large” or “small” gap using only crisp boundaries. Intuitively, we expect that there is a gradual transition from a “small” gap to a “large” gap. The fuzzy rule-base allows for a natural way to incorporate this knowledge about the structure of the problem. Another aspect is that publishers may have different risk-preferences: some publishers may be more aggressive and try to set the reserve price as close as possible to their estimate of the top bid, while others may be more conservative and have a preference for having a higher fill-rate (percentage of ad slots sold).

For the (i) linguistic value “small” we use a decreasing Z-shaped membership function (denoted by $\mu_{ZL}(x; a, b)$); (ii) linguistic value “medium” we use a trapezoidal membership function (denoted by $\mu_T(x; a, b, c, d, e)$); (iii) linguistic value “large” we use an increasing Z-shaped membership function (denoted by $\mu_{ZR}(x; a, b)$). For a decreasing Z-shaped membership function, the membership degree is 1 if $x \leq a$, linearly decreasing for $a < x < b$ and equal to zero for $x \geq b$. For an increasing Z-shaped membership function, the membership degree is 0 if $x \leq a$, linearly increasing for $a < x < b$ and equal to 1 for $x \geq b$. For a trapezoidal membership function the membership degree is zero if $x \leq a$, linearly increasing from zero to e for $a < x < b$, constant at e for $b \leq x \leq c$, linearly decreasing from e to zero for $c < x < d$, zero for $x \geq d$. The exact values are discussed in Section V.

We refer to this process as *risk-aware pricing* since the publisher explicitly takes information about both Y_t and X_t into account while setting reserve prices in order to reduce the risk of setting a reserve price that is too high. Together the rules described above model the risk-preferences of the publisher that is trying to set the reserve price. In principle, the publisher can choose to set reserve prices solely based on p_t^X , so why would he be interested in risk-aware pricing? The main motivation for risk-aware pricing is related to estimation error. Recall that the reference values are approximations that are based on realizations from the distribution of Y_t and X_t . The reference value p_t^X is an approximation for the critical boundary for the top bid, but this value still suffers from errors. In particular, it depends on (i) the sample of reserve prices used to determine its value, and (ii) the distribution of X_t . Assuming a fixed distribution for X_t , a different sample of reserve prices will lead to a different value of p_t^X . Furthermore, as the distribution of X_t might change over time, the value of p_t^X might not be representative for future time periods.

E. Test phase: Monitoring the environment

After using p^* for E periods, that is, after using $p_{t+1} = \dots = p_{t+E} = p^*$ the publisher can use the observed revenues in order to determine whether he wants to start another commit phase. Suppose that $\{(\mathbb{I}\{p_k \leq X_k\}, R_k)\}_{k=t+1}^{t+E}$ is observed during the commit phase. The publisher can conduct a test to see whether the environment has changed substantially and

use this test to determine whether it is worthwhile to start another commit phase, or, to go back and start another exploration phase. The full procedure for the test is described in Algorithm 3. The idea is to use a fraction $0 < f < 1$ to divide

Algorithm 3 Test-Commit-Phase

Require: A sample of size E of revenues, auction outcomes $(\{\mathbb{I}\{p_k \leq X_k\}, R_k\})_{k=1}^E$, the commit reserve price p^* , f , θ_S , d_R , d_S , θ_L , θ_H .

Calculate sample statistics.

- 1: Set $n^* = \lfloor f \cdot E \rfloor$.
- 2: Set $\mu_G = \sum_{k=1}^E R_k / E$. Set $\mu_B = \sum_{k=1}^{n^*} R_k / n^*$.
- 3: Set $\mu_A = \sum_{k=n^*+1}^E R_k / (E - n^*)$.
- 4: Set $\rho_G = \sum_{k=1}^E \mathbb{I}\{p_k \leq X_k\} / E$.
- 5: Set $\rho_B = \sum_{k=1}^{n^*} \mathbb{I}\{p_k \leq X_k\} / n^*$. Set $\rho_A = 1 - \rho_B$.
- 6: **if** $\rho_G \geq \theta_S$ **then**
- 7: Set $V = \text{true}$.
- 8: Set $G = |\mu_A - \mu_B| / \min\{\mu_A, \mu_B\}$.
- 9: **if** $G \leq d_R$ **then**
- 10: **if** $\mu_G > p^*$ **then**
- 11: Set $V = \text{false}$.
- 12: **end if**
- 13: **else**
- 14: **if** $\mu_A > p^*$ **then**
- 15: Set $V = \text{false}$.
- 16: **end if**
- 17: **end if**
- 18: **else**
- 19: **if** $|\rho_A - \rho_B| \leq d_S$ **then**
- 20: Set $V = \text{true}$.
- 21: **if** $\rho_G \leq \theta_L$ **then**
- 22: Set $V = \text{false}$.
- 23: **end if**
- 24: **else**
- 25: Set $V = \text{false}$.
- 26: **if** $\rho_A \geq \theta_H$ **then**
- 27: Set $V = \text{true}$.
- 28: **end if**
- 29: **end if**
- 30: **end if**
- 31: **return** V

the sample of observed results $(\{\mathbb{I}\{p_k \leq X_k\}, R_k\})_{k=t+1}^{t+E}$ into two groups, where the first group uses the first $n^* = \lfloor f \cdot E \rfloor$ observations and the other group the remaining $E - n^*$ observations. Next, two main cases are distinguished: case (i) where the commit price was not close to the critical boundary and where most sales were successful; case (ii) where most of the sales are not successful and the commit price is most likely too high. In case (i) we start another commit phase, unless the observed revenues exceed the commit price, because this is an indication that the second bid will be higher than the commit price (in particular, it is expected to be higher than the previous reference value p_t^Y) in the future and that this commit price will be too low. In case (ii) we check how frequent the unsuccessful sales are. If the success rate in the two groups are similar and unsuccessful sales are very frequent, then we start another exploration phase. If the success rate in the two groups differ but in the last (most recent) group of observations the success rate meets a minimum requirement, then we start another commit phase. Finally, if there are N_{test} or more

commit phases in a row, then we start an exploration phase with probability P_{test} (see Lines 24-38 in Algorithm 4).

F. Full algorithm

The details of the full procedure of FL-ETC-RAP are described in Algorithm 4.

Algorithm 4 Pseudocode for FL-ETC-RAP

Require: $M, E, S, m, W, \kappa, f, \theta_S, d_R, d_S, \theta_L, \theta_H, N_{test}, P_{test}$.

- 1: Set $t = 0$. Set $N = 0$.
- 2: **Exploration Phase I.**
- 3: Set $A = \emptyset$.
- 4: **for** $j = 1$ **to** M **do**
- 5: Set $t = t + 1$.
- 6: Apply reserve price $p_t = 0$.
- 7: Observe outcome of auction: $(\{\mathbb{I}\{p_t \leq X_t\}, R_t\})$.
- 8: Set $A = A \cup \{R_t\}$.
- 9: **end for**
- 10: Set $p^Y = \sum_{R \in A} R / M$.
- 11: **Exploration Phase II.**
- 12: Set $t_s = t$.
- 13: Run AS-Ref-X (Algorithm 2) with inputs $S, m, t_s, W, p^Y, \kappa$ and G as output.
- 14: Set $p^X = G$.
- 15: **Risk-aware pricing.**
- 16: Determine p^* using fuzzy inference system from Section IV-D.
- 17: Set $t = t + S \cdot m$.
- 18: Set $p_t = p^*$.
- 19: Set $B = \emptyset$.
- 20: **for** $j = 1$ **to** E **do**
- 21: Set $t = t + 1$.
- 22: Set reserve price p_t .
- 23: Observe outcome of auction: $(\{\mathbb{I}\{p_t \leq X_t\}, R_t\})$.
- 24: Set $B = B \cup \{(\mathbb{I}\{p_t \leq X_t\}, R_t)\}$.
- 25: **end for**
- 26: **Test for change in environment.**
- 27: Run Test-Commit-Phase (Algorithm 3) with inputs $B, p^*, t_s, f, \theta_S, d_R, d_S, \theta_L, \theta_H$ and Q as output.
- 28: **if** $Q = \text{true}$ **then**
- 29: Set $N = N + 1$.
- 30: **if** $N < N_{test}$ **then**
- 31: Go to Line 16.
- 32: **else**
- 33: Go to Line 16.
- 34: **end if**
- 35: **end if**
- 36: **else**
- 37: Go to Line 2.
- 38: **end if**

V. NUMERICAL EXPERIMENTS

In this section we conduct experiments to evaluate the effectiveness of our proposed approach. The experiments are implemented in Python 2.7 and run on Intel(R) Core(TM) CPU i5-6300U @ 2.40GHz with 8GB RAM under Windows 7 environment.

A. Dataset Description

In order to evaluate our method we use real-life data from ad auction markets. We use the publicly available iPinYou dataset [16], which contains information from the perspective of nine advertisers on a DSP. It contains information about bids placed by advertisers on a DSP for impressions during a week. For each bid there is information about the ad slot (height, visibility, etc.), time of day, the ad exchange, and the result of the bid. It is important to note that the dataset only contains information about the top bid and the second bid if the advertiser actually wins the auction. As a consequence the bid records represent a biased sample from the distribution of the top bid and second bid. However, the dataset contains information for several advertisers and could still be used to get a general understanding of the dynamics on the ad auction market. We use the iPinYou dataset to construct synthetic data for the top bid and second bid in order to test our proposed approach. Note that the values of the bids are not revealed to any of the algorithms: they are only needed in order to test which methods yield the highest returns.

a) Construction of second bid: For a specific advertiser, we take the values of the second bid for the first 310000 impressions (sorted chronologically). We then divide these 310000 impressions into blocks with length 500. Within each block we sample with replacement 500 values of second bid from the 500 impressions in the block. After we are done with the sampling, we take rolling mean with window length of 25 observations of the resulting time series. Finally, we take the last 300000 values of the resulting time series as the values for the second bid. The main reason for taking a rolling mean is that the advertisers in this dataset are bidding on ad slots from different publishers (with different properties etc.), whereas we are interested in a single publisher that is selling a specific ad slot. By taking a rolling mean we are effectively extracting the general trend in the second bids.

b) Construction of top bid: In order to construct the top bid, we take the time series of the second bid and divide the time series into blocks with length B . Within each block we determine the maximum of the values of the second bid (denote the maximum by MB). The value of the top bid within a block is then equal to $MB \cdot (1 + u) \cdot (1 + v)$ where $v \sim \mathcal{U}(0.0, z)$. B , z and u are discrete uniformly distributed with $B \in \{50, 100\}$, $z \in \{0, 0.2, 0.4\}$ and $u \in \{0, 0.1, 0.2\}$. The draws of u are identically independently distributed (i.i.d) between blocks and draws of v are i.i.d within blocks. This construction models a situation where the gap between the top bid and second bid varies over time and is independent of the level of the second bid.

We use data from 4 advertisers and we repeat the above procedure 5 times for each advertiser in order to generate 5 time series for the top bid and second bid.

B. Benchmark Strategies

The MAB framework is a popular framework for decision making under exploration-exploitation trade-offs. In order to judge the quality of our proposed method, we compare its

performance with two MAB algorithms: (i) the UCB algorithm [17] and (ii) the EXP3 algorithm [18]. These are popular bandit algorithms that are simple to implement and have satisfactory performance in a broad range of applications. In the case of i.i.d and bounded rewards for each arm, UCB achieves an order-optimal upperbound on cumulative regret. In the adversarial setting with bounded rewards, EXP3 achieves a worst-case order-optimal upperbound on cumulative regret.

C. Settings and Performance Metrics

In addition to FL-ETC-RAP, we also consider ETC-TB. ETC-TB does not use a risk-aware price, instead it uses the reference value for the top bid as the commit price. In order to measure the performance of the methods, we consider four performance metrics. The first performance metric is the cumulative average return, which is defined as $\sum_{t=1}^T \hat{R}_t / T$, where $\hat{R}_t$ is the observed return in period t . This is our main metric to determine the profitability of a strategy. The second metric is the success rate, which is defined as $\sum_{t=1}^T \mathbb{I}\{p_t \leq X_t\} / T$. This measures how often reserve prices are set too high. The third metric is the revenue rate, which is defined as $\sum_{t=1}^T \mathbb{I}\{p_t \leq X_t\} \hat{R}_t / \sum_{t=1}^T X_t$. This measures the rate at which the top bid is extracted. The fourth metric is the revenue rate given success, which is defined as $\sum_{t=1}^T \mathbb{I}\{p_t \leq X_t\} \hat{R}_t / \sum_{t=1}^T \mathbb{I}\{p_t \leq X_t\} X_t$. This measures the rate at which revenue is extracted given that a sale is successful. We average the three performance metrics over the 5 samples constructed for each advertiser.

We use the following settings for FL-ETC-RAP: $W = 50$, $E = 25$, $M = 15$, $S = 5$, $m = 3$, $\kappa = 0.05$, $f = 0.5$, $\theta_S = 0.8$, $d_R = 0.1$, $d_S = 0.1$, $\theta_L = 0.45$, $\theta_H = 0.8$, $N_{test} = 3$, $P_{test} = 0.05$. With these choices $M = Sm = 15$ and so both exploration phases have the same duration. We use the following membership functions: $\mu_{ZL}^\gamma(x; 0.1, 0.2)$, $\mu_{ZR}^\gamma(x; 0.15, 0.5)$, $\mu_T^\gamma(x; 0.05, 0.1, 0.2, 0.3, 1)$, $\mu_{ZL}^\eta(x; 15, 20)$, $\mu_{ZR}^\eta(x; 40, 50)$, $\mu_T^\eta(x; 15, 20, 40, 45, 1)$, $\mu_{ZL}^\omega(x; 0.05, 0.1)$, $\mu_{ZR}^\omega(x; 0.6, 0.7)$, $\mu_T^\omega(x; 0.4, 0.5, 0.6, 0.7, 1)$. The universe of discourse for γ and ω is set to $[0, 1]$, and for η it is set to $[0, W]$. In order to calculate the firing strength of each rule we used the minimum operator as the t-norm. The output fuzzy sets are unioned using the maximum t-conorm and we used the centroid defuzzification method in order to determine the final output. We use $p_l = 0$, $p_h = 250$ since the resulting time series of the top bid is at most 250. In the UCB and EXP3 algorithms each arm represents a reserve price and we use $N = 100$ arms which are equally spaced in the interval $[p_l, p_h]$.

D. Results: FL-ETC-RAP versus bandits

Table I presents results for FL-ETC-RAP, ETC-TB, UCB and EXP3. The results show that FL-ETC-RAP generally outperforms the other methods. An interesting observation is that UCB outperforms EXP3, even though EXP3 makes no assumption on the sequence of returns. A similar finding was also reported in [9]. As UCB outperforms EXP3, the rest of this subsection focuses on the differences between UCB

(a) Reserve prices selected by FL-ETC-RAP and UCB for advertiser 2259.

(b) Results for advertisers 1458, 3427, 2259 and 2261.

Fig. 1. Cumulative average returns and selected reserve prices.

TABLE I
PERFORMANCE OF FL-ETC-RAP, ETC-TB, UCB AND EXP3 ON SYNTHETIC DATASETS.

	FL-ETC-RAP			ETC-TB			UCB			EXP3		
advertiser 1458	min	max	mean	min	max	mean	min	max	mean	min	max	mean
cumulative average return	73.63	74.39	74.04	67.90	68.49	68.28	69.07	72.01	70.90	63.50	63.75	63.63
success rate	0.85	0.85	0.85	0.78	0.78	0.78	0.91	1.00	0.97	0.86	0.86	0.86
revenue rate	0.69	0.70	0.70	0.64	0.64	0.64	0.65	0.68	0.67	0.60	0.60	0.60
revenue rate success	0.80	0.80	0.80	0.79	0.79	0.79	0.65	0.72	0.68	0.68	0.68	0.68
advertiser 3427	min	max	mean	min	max	mean	min	max	mean	min	max	mean
cumulative average return	84.28	85.25	84.64	78.40	79.42	78.79	68.99	77.50	74.31	72.71	73.07	72.96
success rate	0.86	0.86	0.86	0.79	0.80	0.79	0.66	0.93	0.80	0.86	0.87	0.87
revenue rate	0.71	0.71	0.71	0.66	0.66	0.66	0.58	0.65	0.62	0.61	0.61	0.61
revenue rate success	0.80	0.80	0.80	0.80	0.80	0.80	0.68	0.77	0.73	0.69	0.69	0.69
advertiser 2259	min	max	mean	min	max	mean	min	max	mean	min	max	mean
cumulative average return	100.50	101.29	100.80	95.09	96.01	95.38	95.43	97.23	96.28	88.05	88.59	88.33
success rate	0.88	0.89	0.88	0.83	0.84	0.83	0.95	0.98	0.97	0.88	0.89	0.89
revenue rate	0.72	0.72	0.72	0.68	0.68	0.68	0.68	0.69	0.69	0.63	0.63	0.63
revenue rate success	0.79	0.80	0.79	0.79	0.79	0.79	0.68	0.72	0.70	0.69	0.70	0.69
advertiser 2261	min	max	mean	min	max	mean	min	max	mean	min	max	mean
cumulative average return	99.17	99.56	99.35	93.72	94.47	94.03	92.17	96.08	94.70	86.95	87.25	87.13
success rate	0.88	0.88	0.88	0.83	0.83	0.83	0.94	0.99	0.97	0.89	0.90	0.90
revenue rate	0.71	0.71	0.71	0.67	0.67	0.67	0.66	0.69	0.68	0.62	0.62	0.62
revenue rate success	0.79	0.79	0.79	0.78	0.78	0.78	0.67	0.71	0.69	0.68	0.68	0.68

and FL-ETC-RAP. The performance gap differs depending on the specific advertiser, but in general, FL-ETC-RAP has a cumulative average return that is about 4.0 to 10.0 units higher than UCB. This associated difference in revenue is about 3% to 9%. The main explanation for the superior performance of FL-ETC-RAP compared to UCB is that (i) FL-ETC-RAP is able to better track changes in the top bid and (ii) FL-ETC-RAP is better in selecting reserve prices that are closer to the top bid when the gap between the top bid and second bid is large. In Fig. 1a a rolling average with window size 250 of the selected reserve prices by FL-ETC-RAP and UCB relative to the top bid for a specific advertiser are shown (for a specific sample). From Fig. 1a we see that UCB tends to be more conservative and selects reserve prices that are often low, which results in a higher success rate but a lower revenue rate (see also Table I). Fig. 1a shows that UCB generally not very responsive to changes in the top bid. On the other hand, the

reserve prices chosen by FL-ETC-RAP are generally higher than those chosen by UCB and they tend to do a better job at tracking the changes in the top bid. Finally, in Fig. 1b we can see that FL-ETC-RAP tends to outperform UCB quite early within the sales horizon.

E. Results: Impact of risk-aware pricing

If we compare FL-ETC-RAP with ETC-TB, then we see that FL-ETC-RAP significantly outperforms ETC-TB. The performance gap differs depending on the specific advertiser, but in general, FL-ETC-RAP has a cumulative average return that is at least 5.0 higher than ETC-TB. ETC-TB tends to select reserve prices that are higher than those selected by FL-ETC-RAP, and this results in a lower success rate (see Table I). On the other hand, the revenue rate given success is similar, which indicates that, for successful sales, FL-ETC-RAP and ETC-TB tend to select similar reserve prices. The results show that estimation errors relating to the top bid can have a big

impact on performance. The main reason for this is that, in this application, exceeding the top bid by an amount $\Delta > 0$ is more costly than selecting a reserve price Δ below the top bid. Taking the gap between the top and second bid into account leads to better performance.

Remark 3: We have presented results based on basic settings that performed well across all the advertisers in the iPinYou dataset. The results for the other 5 advertisers in the iPinYou dataset are qualitatively similar to those reported in Table I and are omitted due to space limitations. Results with $S \in \{3, 8\}$, $W = 75$, $M = 25$ are also similar. Performance can be improved by tuning the parameters for each advertiser separately. Also, the risk-aware pricing component can be refined to improve performance. In our experiments we did not conduct extensive parameter tuning. One approach to conduct parameter tuning is to use a percentage of the impressions (say the first 10%) to tune parameters and then evaluate the performance on the remaining impressions. ■

VI. DISCUSSION AND CONCLUSION

We proposed a method for learning optimal reserve prices in second-price auctions. We studied a limited information setting where the values of the bids are not revealed and no historical information about the values of the bids is available. Our proposed method combines an adaptive search procedure with a fuzzy logic pricing step to set reserve prices.

Our results indicate that it is important to properly deal with non-stationarity in the distribution of the bids. Another key insight is that incorporating knowledge about the structure of the problem (e.g. in the form of risk-aware pricing) can lead to improved performance. Fuzzy logic provides us with useful tools that can be used in order to incorporate knowledge about the structure of the problem in a principled way.

Our method can be improved in various ways. One drawback of our approach is that it takes a number of control parameters as input. Future work can be directed towards optimizing the parameters of the risk-aware pricing component and to make it adaptive and self-regulatory. Another direction is to incorporate the estimation uncertainty of the top bid and second bid into the fuzzy logic system that determines the reserve price. One approach that we are investigating is to use an interval type-2 fuzzy logic system with adaptive membership functions. Yet another direction for future work is to use swarm intelligence techniques and evolutionary algorithms to tune the parameters of the membership functions (see for example [19]).

REFERENCES

- [1] J. Wang, W. Zhang, and S. Yuan, "Display advertising with real-time bidding (RTB) and behavioural targeting," *Foundations and Trends® in Information Retrieval*, vol. 11, no. 4-5, pp. 297–435, 2017. [Online]. Available: <http://dx.doi.org/10.1561/15000000049>
- [2] N. Cesa-Bianchi, C. Gentile, and Y. Mansour, "Regret minimization for reserve prices in second-price auctions," *IEEE Transactions on Information Theory*, vol. 61, no. 1, pp. 549–564, Jan 2015.
- [3] Z. Xie, K.-C. Lee, and L. Wang, "Optimal reserve price for online ads trading based on inventory identification," in *Proceedings of the ADKDD'17*, ser. ADKDD'17. New York, NY, USA: ACM, 2017, pp. 6:1–6:7. [Online]. Available: <http://doi.acm.org/10.1145/3124749.3124760>
- [4] M. R. Rudolph, J. G. Ellis, and D. M. Blei, "Objective variables for probabilistic revenue maximization in second-price auctions with reserve," in *Proceedings of the 25th International Conference on World Wide Web*, ser. WWW '16. Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee, 2016, pp. 1113–1122. [Online]. Available: <https://doi.org/10.1145/2872427.2883051>
- [5] S. Yuan, J. Wang, B. Chen, P. Mason, and S. Seljan, "An empirical study of reserve price optimisation in real-time bidding," in *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '14. New York, NY, USA: ACM, 2014, pp. 1897–1906. [Online]. Available: <http://doi.acm.org/10.1145/2623330.2623357>
- [6] M. Mohri and A. M. n. Medina, "Learning algorithms for second-price auctions with reserve," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 2632–2656, Jan. 2016. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2946645.3007027>
- [7] D. Austin, S. Seljan, J. Monello, and S. Tzeng, "Reserve price optimization at scale," in *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, Oct 2016, pp. 528–536.
- [8] P. Chahuaara, N. Grislain, G. Jauvion, and J.-M. Renders, "Real-time optimization of web publisher RTB revenues," in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '17. New York, NY, USA: ACM, 2017, pp. 1743–1751.
- [9] G. Jauvion, N. Grislain, P. Dkengne Sielenou, A. Garivier, and S. Gerchinovitz, "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*, ser. KDD '18. New York, NY, USA: ACM, 2018, pp. 425–432. [Online]. Available: <http://doi.acm.org/10.1145/3219819.3219917>
- [10] J. Rhuggenaath, A. Akcay, Y. Zhang, and U. Kaymak, "Optimizing reserve prices for publishers in online ad auctions," in *IEEE Conference on Computational Intelligence for Financial Engineering and Economics (CIFER)*, 2019, to appear.
- [11] M. Kumari and K. K. Bharadwaj, "Fuzzy logic based effective bidding range computation and bidder's behavior estimation in keyword auctions," in *2010 IEEE 2nd International Advance Computing Conference (IACC)*, Feb 2010, pp. 299–303.
- [12] R. Qin, Y. Yuan, and J. Li, "A two-stage fuzzy budget allocation model in search auctions," in *Proceedings of 2014 IEEE International Conference on Service Operations and Logistics, and Informatics*, Oct 2014, pp. 144–148.
- [13] Q. Madera, O. Castillo, M. Garcia-Valdez, and A. Mancilla, "Bidding strategies based on type-1 and interval type-2 fuzzy inference systems for google adwords advertising campaigns," in *2016 IEEE 8th International Conference on Intelligent Systems (IS)*, Sep. 2016, pp. 133–138.
- [14] G. J. Klir and B. Yuan, *Fuzzy sets and fuzzy logic: theory and applications*. Upper Saddle River, NJ: Prentice Hall, 1995.
- [15] J. Mendel, H. Hagsras, W.-W. Tan, W. W. Melek, and H. Ying, *Introduction To Type-2 Fuzzy Logic Control: Theory and Applications*, 1st ed. Wiley-IEEE Press, 2014.
- [16] W. Zhang, S. Yuan, J. Wang, and X. Shen, "Real-time bidding benchmarking with iPinYou dataset," *arXiv preprint arXiv:1407.7073*, 2014.
- [17] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine Learning*, vol. 47, no. 2, pp. 235–256, May 2002. [Online]. Available: <https://doi.org/10.1023/A:1013689704352>
- [18] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. Schapire, "The nonstochastic multiarmed bandit problem," *SIAM Journal on Computing*, vol. 32, no. 1, pp. 48–77, 2002. [Online]. Available: <https://doi.org/10.1137/S0097539701398375>
- [19] M. L. Lagunes, O. Castillo, F. Valdez, J. Soria, and P. Melin, "Parameter optimization for membership functions of type-2 fuzzy controllers for autonomous mobile robots using the firefly algorithm," in *Fuzzy Information Processing*, G. A. Barreto and R. Coelho, Eds. Cham: Springer International Publishing, 2018, pp. 569–579.