

# An Empirical Study of Reserve Price Optimisation in Real-Time Bidding

Shuai Yuan, Jun Wang,  
Boweï Chen  
Department of Computer  
Science, University College  
London  
{s.yuan, j.wang,  
bowei.chen}@cs.ucl.ac.uk

Peter Mason  
Advance International Media  
peter@advance.uk.com

Sam Seljan  
AppNexus  
sseljan@appnexus.com

## ABSTRACT

In this paper, we report the first empirical study and live test of the reserve price optimisation problem in the context of Real-Time Bidding (RTB) display advertising from an operational environment. A reserve price is the minimum that the auctioneer would accept from bidders in auctions, and in a second price auction it could potentially uplift the auctioneer’s revenue by charging winners the reserve price instead of the second highest bids. As such it has been used for sponsored search and been well studied in that context. However, comparing with sponsored search and contextual advertising, this problem in the RTB context is less understood yet more critical for publishers because 1) bidders have to submit a bid for each individual impression, which mostly is associated with user data that is subject to change over time. This, coupled with practical constraints such as the budget, campaigns’ life time, etc. makes the theoretical result from optimal auction theory not necessarily applicable and a further empirical study is required to confirm its optimality from the real-world system; 2) in RTB an advertiser is facing nearly unlimited supply and the auction is almost done in “last second”, which encourages spending less on the high cost ad placements. This could imply the loss of bid volume over time if a correct reserve price is not in place. In this paper we empirically examine several commonly adopted algorithms for setting up a reserve price. We report our results of a large scale online experiment in a production platform. The results suggest the our proposed game theory based OneShot algorithm performed the best and the superiority is significant in most cases.

## Categories and Subject Descriptors

H.3.5 [Information Systems]: Information Storage and Retrieval - Online Information Services

## Keywords

Display Advertising; Reserve Price; Revenue Optimisation; Online Advertising; Real-Time Bidding

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
KDD’14, August 24–27, 2014, New York, NY, USA.  
Copyright 2014 ACM 978-1-4503-2956-9/14/08 ...\$15.00.  
<http://dx.doi.org/10.1145/2623330.2623357>.

## 1. INTRODUCTION

In display and mobile advertising, the most significant development in recent years is the growth of Real-Time Bidding (RTB) [34], which allows selling and buying online display advertising in real-time one ad impression at a time. It is projected to reach \$8.49 billion in US alone in 2017, accounting for 29% of all digital display advertising [12].

In RTB, a publisher generates a bid request for each individual impression in real time and then sends a request to advertisers via ad exchanges [23]. In practice, the reserve price may be disclosed or not. Advertisers would submit bids for this impression based on their bidding algorithms. Usually they do so even when the bids are lower than the disclosed reserve to avoid the *time-out penalty* which reduces the bid request volume over time. Those bidding algorithms usually enquire third-party data providers to obtain the user’s interest segments to assist in computing bids. When the time limit expires (usually 100ms) the exchange holds a second price auction [2], and the publisher gets the winner, saves the final price, and displays the ad from the winning advertiser to the user. The whole process is represented in Figure 1.

Along with other fixed-price advertising channels [29, 4], revenue from RTB makes an important source to support various (online) publishers, including but not limited to content providers, social networks, and personal blogs. They provide free services or content and subsidise the maintaining cost by ad revenue. Therefore, uplifting revenue by employing various yield management tools makes one of the key topics on the publisher side. An important tool among them is the *reserve price* optimisation. A reserve price defines the minimum that a publisher would accept from bidders. It reflects the publisher’s private valuation of the inventory: bids will be discarded if they are below the reserve price. In the second price auction, which is commonly used in RTB, the reserve price could potentially uplift the revenue. Figure 2 illustrates how the final price is calculated from bids with a reserve price. Let  $b_1, \dots, b_K$  denote the descending bids and  $\alpha$  the reserve price. Then, the desirable case is  $b_1 \geq \alpha > b_2$  where the publisher gains extra payoff of  $\alpha - b_2$ ; the neutral case is  $b_1 > b_2 \geq \alpha$  where the publisher has no extra gain; and the undesirable case is  $\alpha > b_1$  where the publisher suffers from a loss of  $b_2$ . These cases directly motivate our work.

The optimisation problem has been previously studied in the context of sponsored search [11, 13, 26, 31]. However, the

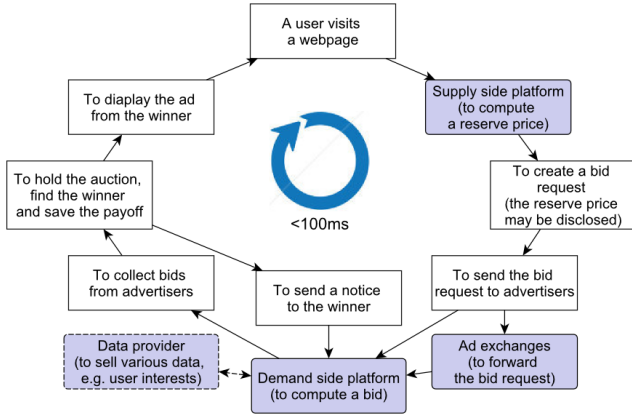


Figure 1: The simplified process of Real-Time Bidding advertising. Major entities are coloured in blue where processes are transparent. One iteration usually finishes in less than 100ms to achieve users’ satisfactory.

problem in the RTB context is different and unique. Firstly, the optimal auction theory requires to know the distribution of the advertisers’ private yet true assessments of the impression before calculating the optimal reserve price [11]. In RTB, it becomes a lot harder to learn the distribution. In RTB an advertiser is required to submit a bid for each individual impression using his own algorithm, which is never disclosed to publishers and could rely heavily on privately-owned user interest segments. Besides, various practical constraints such as the budget, campaign life time, irrationality, divert advertisers from bidding at private values. This difference makes the private value based algorithm inefficient in practice. Thus, it is of great interest to empirically study the subject and examine several commonly adopted algorithms in the real-world as we report in the paper. Secondly, unlike sponsored search, an advertiser does not have the keyword constraint and faces almost unlimited supply of impressions in RTB. Setting up an aggressive reserve price would easily move the advertisers away from those placements and force them to look for something cheaper.

We analyse this possible consequence in the online experiment and reject the significant attrition hypothesis in today’s RTB marketplace, which could be due to the complexity of our scheduling or less sensitivity of common bidding algorithms. The little attrition fact implies good chance of implementing the optimisation in the current eco-system.

In this paper, we represent the first field study of reserve price optimisation in the RTB context. It is based on the analysis of real-world data and online experiments on a production platform. We provide thorough discussions about the commonly adopted algorithms for reserve price optimisation including both the private-value-free and the private-value-based ones [24, 11, 26], and their variations.

## 2. RELATED WORK

Up to now the reserve price problem is mainly studied for sponsored search (SS). For instance, in [26], the authors first studied the reserve price problem in a real-world online advertising system (Yahoo! sponsored search). The authors test optimal auction theory on 450k keywords and the results diverge by number of impressions: for keywords with high search volume the uplift on revenue is good and statistically significant. However, for other keywords the total revenue declines by 2.2% on average. We find this result




Figure 2: The decision process of second price auctions on the publisher side. The desirable case is  $b_1 \geq \alpha > b_2$  where the publisher gains extra payoff of  $\alpha - b_2$ . We ignored *soft floor prices* which make the process a lot more complicated. Interested readers may refer to [34] for further discussion.

consistent with our work, although the optimisation is not helpful when there are few bidders. Additionally, we also perform the bidders’ attrition analysis which is a more significant problem in RTB.

In [11] the optimal auction problem is characterised for sponsored search. The authors demonstrate the calculation of optimal reserve price for multi-item auctions and conclude that it is independent of the number of bidders. However, in [31] the authors conclude that both the number of bidders and the number of ad links affect the optimal reserve price. Some assumptions are made in the first paper: advertisers know their value per click; they have the same Click-Through Rate (CTR) at a given position; they share the common knowledge of position CTRs; they all maximise the expected profit. These assumptions might be realistic when search engines bids for advertisers, but however are far from achievable in RTB. Also in the SS context, in [13] the advertiser specific reserve prices are discussed within the generalised second-price auction (GSP) and the Vickrey-Clarke-Groves (VCG) auction [10] settings. As proved in the paper the bidder specific reserve prices could result in losing truthfulness in the VCG pricing; but could lead to envy-free equilibrium in the GSP pricing.

We plan our work in a RTB environment. In [23] the abstract model of the ad exchange is proposed and several problems are discussed, including auction mechanism, call out optimisation [8], publisher revenue optimisation [15], arbitrage bidding and risk analysis [7] etc. As mentioned by the authors, the reserve price problem taking an essential role in yield optimisation yet remains unsolved. In this paper, we study the problem at placement-level (a group of impressions) and leave the impression-level dynamic reserve prices to future work.

There is also rich literature from bidders’ perspective on how to bid when a randomised reserve price in ad network is present [16, 14]. The optimal auction design is studied in [14] with two layers: central seller to intermediaries and intermediaries to bidders. The authors conclude that revenue-maximizing intermediaries will use an auction with a randomized reserve price chosen from an interval in equilibrium, and the optimal reserve price decreases with the number of buyers. In our paper, although no intermediary is employed, we also observe and discuss the relationship of the optimal reserve price and number of bidders.

Due to the similarity in mechanisms, RTB also can borrow a lot from established research on auction theories [19, 28, 22], esp. electronic commerce [3, 32]. In [18] the detection and reaction towards reserve prices cheating are first discussed for online auctions. The authors of [22] discuss the situation that an item could be resold if previous auctions fail due to high reserve price. An equilibrium reserve price was proven to approach the one in an optimal static auction. Meanwhile, there are papers on scoring inventories [27] considering unique characteristics of RTB auctions

(e.g., number of bidders, bid time, bidders' identity). However, few of them take the reserve price into account.

### 3. THE RESERVE PRICE PROBLEM AND ALGORITHMS

Suppose a publisher tries to maximise the ad revenue for a single placement (i.e., single-item). Impressions from this placement are sold using the second price auction in RTB. The publisher could set a reserve price (also known as the hard floor price)  $\alpha$  before an auction is conducted.

To simplify the discussion, we assume there is only one impression at each step  $t$  and the publisher is optimising the revenue over horizon  $T$ . We also assume that at each step  $t$  there are  $K \geq 2$  bidders participating in the auction. If there is only one bidder it is equivalent to having the second highest bid equals to the minimal bid (usually \$0.01). Note bidders could be different each time and could change quite a lot—this introduces great noise into the system and challenges the estimation of the distribution of the private evaluation values from the advertisers, where most auction theories assume there is a fix distribution of bids.

For each auction, we denote the final bids of the placement as  $b_1(t), b_2(t), \dots, b_K(t)$ . Without loss of generality we assume  $b_1(t) \geq b_2(t) \geq \dots \geq b_K(t)$ . Therefore, without a reserve price ( $\alpha = 0$ ) the payoff could be denoted as  $r(t) = b_2(t)$ . Now suppose the publisher sets a non-zero reserve price at each step, denoted by  $\alpha(t)$ . The payoff function becomes:

$$r'(t) = \begin{cases} \alpha(t), & b_1(t) \geq \alpha(t) > b_2(t) \\ b_2(t), & b_2(t) \geq \alpha(t) \\ 0, & \alpha(t) > b_1(t) \end{cases} \quad (1)$$

The overall income is  $R(T) = \sum_t^T r'(t)$ . We assume it is zero payoff when the reserve price is too high. In practice, publishers usually redirect these impressions to managed campaigns or other ad marketplaces for re-selling. This reduces the risk of over-optimisation.

Comparing with the original payoff function, the case  $r'(t) = \alpha(t)$  provides an extra gain, whereas the case  $r'(t) = 0$  incurs a loss. We plan our work at the placement level; in other words, we set up the reserve price periodically. From a control theory's point of view, throughout the planning horizon the publisher wants to explore (learn) the optimal  $\alpha$  of the placement; the same time the publisher also exploits (predict) the known best  $\alpha$  to get as much payoff as possible and to compensate the loss.

#### 3.1 Optimal Auction Theory

Regardless of the existence of reserve price, bidders are encouraged to bid their private values in the second price auctions [24, 21]. Note that this dominant strategy does not hold in modern sponsored search where quality scores are generally used [10] in ad ranking. Without quality scores, the strategy of bidding at the private value forms part of the Nash equilibrium of the system, meaning as time elapses advertisers have no incentive to change their bids, given that all other factors remain the same. In this non-cooperative game [25], the winner could, but would not, lower his bid to let other competitors win because losing the auction is not beneficial in either short-term or long-term (lowering the bid while still winning has no effect since the winner always pays the second highest bid).

Suppose the publisher knows the bidders' private value distribution. The optimal auction theory mathematically defines the optimal reserve price [19, 20, 31, 24]. There is also a large scale experiment in sponsored search to test this theory [26]. We briefly introduce the theory to make the paper self-contained. Again suppose there are  $K$  bidders and they are risk-neutral and symmetric, i.e. having identical value distributions. Each bidder  $k \in K$  has private information on the value of an impression, drawn from distribution  $F_k(x)$ , where  $F_k(x)$  denotes the probability that the advertiser's private evaluation value is less than or equal to a certain number  $x$ . Usually it is assumed Log-normal [26] or Uniform distribution [24]. Assuming private values are independently distributed, the distribution over value vector is

$$F(\cdot) = F_1(\cdot) \times \dots \times F_K(\cdot),$$

and then the optimal reserve price is given as (see [25] for details):

$$\alpha = \frac{1 - F(\alpha)}{F'(\alpha)} + v_P, \quad (2)$$

where  $F'(\alpha)$  is the density function, the first order derivative of  $F(\alpha)$  and  $v_P$  is the publisher's private value. In practice,  $v_P$  could be obtained from a guaranteed contract with a flat CPM, or from another ad network where the average revenue is known.

In the experiment we implement this theory as OPTAUC and follow the Log-normal distribution assumption of bidders' private values. We also adopt the symmetric assumption, i.e., there is only one distribution for all bidders. Under these assumptions the optimality of the auction under GSP is proved in [11]. The estimation of Log-normal's mean and standard deviation was obtained using the training dataset (impression-level logs from 14 Dec 2012 to 18 Jan 2013).

##### 3.1.1 Drawbacks in RTB Practice

In practice, there are drawbacks of the optimal auction theory mostly due to the difficulty of learning bidders' private values, e.g.,  $F(x)$ . Firstly, a bidder could have a complex private value distribution for impressions. In RTB an advertiser computes a bid for each individual impression based on the contextual [6] and behavioural [33] data. The data is fed into their private valuation models which are never disclosed to publishers or other advertisers. This is especially true in RTB comparing with SS where search engines run bidding algorithms for advertisers and host auctions as a publisher at the same time. Also, in SS the auctions are based on keywords, so the population of the bidders are relatively more stable, whereas in RTB, the auctions are in the impression level and the advertisers are more flexible in terms of choosing the impressions to bid.

We tested the Uniform distribution at placement level and Log-normal distribution at both placement and impression level. Although these distributions are widely adopted in research literature [24, 26], only a small portion of tests returned positive results as shown in Figures 7 and 8.

Secondly, it is *assumed* that advertisers bid at their private values in the second price auction [19, 26]. However, in practice, an advertiser may not know clearly his private value of an impression. Instead, he wants to achieve the best possible performance. Also in different stages (learning, prospecting, retargeting, etc.) of an advertising campaign, the bidding strategy could change. This makes the bidding

activity vary greatly across the limited flight time of a campaign. See a plot from our experiment in Figure 5.

Thirdly, there are other practical constraints including accessibility of auction details, noise introduced by the frequent change of auction winners, c.f. Figure 6. We believe these drawbacks lead to undesirable performance of an optimal auction theory based algorithm in our real-world experiment.

### 3.2 A Simplified Dynamic Game

To address the above-mentioned issue, in this paper, we take an alternative view and propose a simple auction game between the publisher and auction winners and identify the dominant strategies. We simplify the game by dropping the *repeated* nature of auctions. Thus the publisher only considers the current auction and do not learn the private values from historical knowledge. In fact, the result of this simplification follows the instinct and is easy to implement; it also performs the best in most of cases in our online large scale real-world experiment.

First we give the extensive form representation of this dynamic game:

- Player: the winner of auctions (advertisers)  $w$  and the publisher  $p$ .
- The information set  $I$  before acting is the same for the winner and the publisher. It has two decision nodes:  $I_1$ , the winning bid  $b$  is equal to or higher than the current reserve price  $\alpha$ ;  $I_2$ , the winning bid is lower than the reserve price.
- The action set of the winner  $A_w$ :  
 $a_{w1}$ , to increase  $b$  to higher than  $\alpha$ ;  
 $a_{w2}$ , to increase  $b$  to lower than  $\alpha$ ;  
 $a_{w3}$ , to decrease or hold  $b$  to higher than  $\alpha$ ;  
 $a_{w4}$ , to decrease or hold  $b$  to lower than  $\alpha$ .
- The action set of the publisher  $A_p$ :  
 $a_{p1}$ , to increase or hold  $\alpha$  to higher than  $b$ ;  
 $a_{p2}$ , to increase or hold  $\alpha$  to lower than  $b$ ;  
 $a_{p3}$ , to decrease  $\alpha$  to higher than  $b$ ;  
 $a_{p4}$ , to decrease  $\alpha$  to lower than  $b$ .
- The sequence of move: first the publisher, then the winner.

The game tree representation, as well as the payoff function, is given in Figure 3. Note for some nodes the payoff of the winner consists two numbers, e.g. 30/−10, when the winner choose to increase the bid, i.e., the action  $a_{w1}$ . The positive value stands for the case where it is still profitable to increase the bid, while the negative value stands for the possible loss if increasing the bid, since the advertiser has reached the maximum affordable price. In the latter case, an advertiser would choose other actions like  $a_{w2}$  or  $a_{w4}$ . These values have been carefully selected to reflect the positions of bidder. For example,  $I_1 \rightarrow I_1|a_{p2} \rightarrow I_1|a_{w1}$  would give less payoff to the advertiser than  $I_1 \rightarrow I_2|a_{p1} \rightarrow I_1|a_{w1}$  because he has been increasing bids to win auctions.

When deducing dominant strategies we assume both cases happen with an equal chance since both publishers and advertisers do not use historical knowledge but only the last state. Therefore, payoff of the publisher in these cases is discounted by 0.5 since a rational advertiser would not choose a negative payoff.

#### 3.2.1 The Publisher's Dominant Strategy

Based on the formulation above we analyse the case where the publisher and the winner play the game for only one

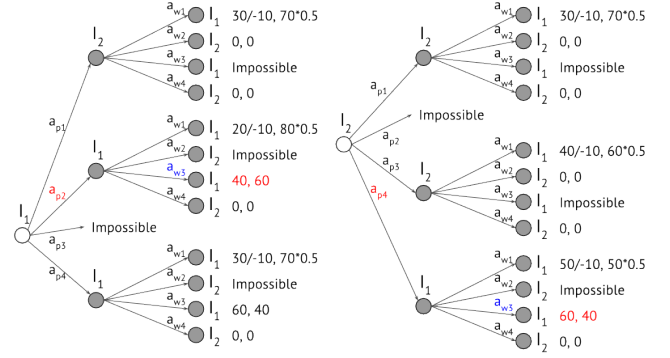


Figure 3: The game between the winner and the publisher in the reserve price problem. At the leaf nodes we give the result information set as well as the payoffs of (winner, publisher). Note for the action  $a_{w1}$  the payoff of the winner could be negative if he has already been bidding the maximal affordable price. We assume these cases happen at a chance of 50% due to no utilisation of historical data. Thus, the payoff of the publisher is discounted by 0.5 in these cases.

round. We claim the dominant strategy for the publisher is:

$$s_p^*(I) = \begin{cases} a_{p2}, & \text{if } I = I_1 \\ a_{p4}, & \text{if } I = I_2 \end{cases} \quad (3)$$

which gives the expected payoff:

$$R(s_p^*|I) = \begin{cases} 60 & \text{if } I = I_1 \\ 40 & \text{if } I = I_2 \end{cases} \quad (4)$$

We eliminate the proof here since it could be easily acquired by following nodes iteratively in the game tree, c.f. Figure 3. This strategy will serve as the foundation of the ONESHOT algorithm we propose in the following section. On the other hand, the dominant strategy for the winner is

$$s_w^*(I) = \begin{cases} a_{w3}, & \text{if } I = I_1 \\ a_{w1}, & \text{if } I = I_2 \end{cases} \quad (5)$$

which indicates that the bid price should be gradually reduced but increased again when lost the auction.

Based on these analysis we have the Nash equilibrium of the game:  $s^* = (s_p^*, s_w^*)$ . This indicates the publisher should keep the reserve price below the winning bid, while trying to approach the winning bid gradually. If the reserve price is too high, it should be reduced drastically to make sure it will be below the winning bid. Following these actions, the final state of the system comes to the second price auction without a reserve price, i.e., the reserve price is always equal to the second highest bid. In this state, the publisher does not need to lower  $\alpha$  any more. The winner will not lower the bid otherwise he would lose the auction which gives a zero payoff. However, this state has never been observed in our experiment due to the frequent change of bidders, variation of bids, possibly different strategies adopted by advertisers (e.g. always bid at the private value).

#### 3.2.2 Introducing the Randomness

One may argue that since the winner of auction is constantly changing this is no longer a game between two players, but between a player and a group. The frequency of changing is indeed high. See Figure 6 from the experiment

Table 1: The payoff matrix of the publisher if the new winner bid randomly.  $I_1$  and  $I_2$  denote the starting information set.  $a_{p*}$  stands for publishers’ actions. The publisher’s dominant strategy remains the same and is illustrated in bold font.

|          | $I_1$                      | $I_2$                      |
|----------|----------------------------|----------------------------|
| $a_{p1}$ | 1/3 * 70                   | 1/3 * 70                   |
| $a_{p2}$ | <b>1/3 * 80 + 1/3 * 60</b> | -                          |
| $a_{p3}$ | -                          | 1/3 * 60                   |
| $a_{p4}$ | 1/3 * 70 + 1/3 * 40        | <b>1/3 * 50 + 1/3 * 40</b> |

section. But, since the knowledge of auctions is not shared among the bidder group, each player of this group only possesses an imperfect information set. This could result in the randomness of winner’s actions: the new winner does not have the outcome of the last auction, so has to bid randomly, which is usually based on private valuation of the impression or simply instinct.

In this case we claim that the dominant strategy of the publisher is still the one defined in Equation 3. We remove the negative payoffs of the winner from the game tree in Figure 3 and allow him to bid randomly regardless of what action has been chosen by the publisher. We remove the negative payoffs because the new winner should not bid above his private valuation of the impression in the first time step. The single-step payoff matrix of the publisher is shown in Table 1.

By playing this auction game with advertisers’ random actions the publisher has the payoff:

$$R'(s_p^*|I) = \begin{cases} \frac{1}{3} * 80 + \frac{1}{3} * 60, & \text{if } I = I_1 \\ \frac{1}{3} * 50 + \frac{1}{3} * 40, & \text{if } I = I_2 \end{cases} \quad (6)$$

We consider this case describes the real-world marketplace more closely. In fact, the convergence to the second price auction state hardly happens due to two reasons. Firstly, the winner of auctions keeps changing constantly as illustrated in Figure 6, especially when there are many bidders in auctions. Secondly, the detection of the reserve price could be difficult (when it is not disclosed) and costly. It becomes even harder if certain randomisation is introduced to the final reserve price setting as we did in the experiment. The winner could suffer from distinguishing a reserve price from his competitors’ bids. From the bidder’s perspective, the dominant strategy of competing with other bidders (bidding the private value) and with the publisher (keep lowering the bid before losing) is clearly different.

### 3.2.3 The OneShot Algorithm

We design an algorithm based on the dominant strategy analysis above. For the publisher, if the winning bid is higher than the reserve price, slowly increase the reserve price; otherwise, decrease drastically. We make the actual implementation slightly different by introducing parameters to control the magnitude of the change under different situations. We describe the algorithm as follows.

$$\begin{cases} \alpha(t+1) = (1 - \epsilon^t \lambda_h) \alpha(t) & \text{if } \alpha(t) > b_1(t) \\ \alpha(t+1) = (1 + \epsilon^t \lambda_e) \alpha(t) & \text{if } b_1(t) \geq \alpha(t) \geq b_2(t) \\ \alpha(t+1) = (1 + \epsilon^t \lambda_l) \alpha(t) & \text{if } b_2(t) > \alpha(t) \end{cases}$$

where  $\epsilon \in (0, 1]$  and  $\lambda_h, \lambda_e, \lambda_l \in [0, 1]$ .  $\epsilon$  is a decay factor w.r.t. time, allowing the reserve price to converge if needed.

$\lambda_h$  controls the cooling speed when the reserve price is too high;  $\lambda_e$  controls the continued exploration when the reserve price is successfully;  $\lambda_l$  controls the heating up speed when the reserve price is too low.

Note the values of these parameters depend on targeting combinations (placements, date and time, geography, etc.). They need to be tuned to achieve the best result. For example, in experiments we found an effective setting of these parameters as  $\epsilon = 1.0$ ,  $\lambda_h = 0.3$ ,  $\lambda_e = 0.01$ , and  $\lambda_l = 0.02$  for placement 834119 and hour-of-day=8. We obtain these parameters using the training dataset (impression-level logs from 14 Dec 2012 to 18 Jan 2013). This algorithm is denoted as ONESHOT in the experiments.

## 3.3 Other Private Value Based Algorithms

Although not optimal in theory, there are other algorithms which make use of private value distributions in practice. Here we propose two simple methods, which are based on Bayes’ rules. An advantage of such algorithms is they could be easily tuned to be more aggressive or less according to the greediness (or risk preference) of the publisher. A disadvantage is that it does not take the future into account, so only the current payoff is maximized.

### 3.3.1 Bivariate Log-normal Distribution

To make the algorithm easier to understand we keep the Log-normal distribution assumption of private values. First we capture the first and second highest bids using a bivariate Log-normal distribution:

$$\begin{bmatrix} B_1 \\ B_2 \end{bmatrix} \sim \text{lognorm}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \quad (7)$$

where  $\boldsymbol{\mu}$  and  $\boldsymbol{\Sigma}$  stand for mean and variance respectively. For simplicity we consider  $\boldsymbol{\Sigma}$  known and use a bivariate Gaussian distribution as the conjugate prior for  $\boldsymbol{\mu}$ :

$$\boldsymbol{\mu}(t) \sim \mathcal{N}(\boldsymbol{\theta}(t), \boldsymbol{\Delta}(t))$$

The priori could be learned using historical data. In our experiment, impression-level logs from 14 Dec 2012 to 18 Jan 2013 were used. During evaluation, each time the publisher observes the highest and 2nd highest bids in an auction, then updates the belief using Bayesian inference [30]:

$$\begin{bmatrix} B_1 \\ B_2 \end{bmatrix} | \mathbf{b}(t) \sim \text{lognorm}(\boldsymbol{\theta}(t+1), \boldsymbol{\Delta}(t+1) + \boldsymbol{\Sigma})$$

where

$$\begin{aligned} \boldsymbol{\theta}(t+1) &= (\boldsymbol{\Delta}^{-1}(t) + \boldsymbol{\Sigma}^{-1})^{-1} (\boldsymbol{\Delta}^{-1}(t) \boldsymbol{\theta}(t) + \boldsymbol{\Sigma} \mathbf{b}(t)) \\ \boldsymbol{\Delta}(t+1) &= (\boldsymbol{\Delta}^{-1}(t) + \boldsymbol{\Sigma}^{-1})^{-1} \end{aligned}$$

Then with  $p(B_1, B_2)$  estimated at current time  $t$ , we can solve for a reserve price which has the maximum probability of sitting between the two bids:

$$\alpha(t)^* = \arg \max_{\alpha(t)} P(B_1 \geq \alpha(t) \geq \lambda B_2) \quad (8)$$

where  $\alpha(t)^*$  is the current reserve price and  $\frac{B_1}{B_2} \geq \lambda \geq 1$  is the risk preference parameter. It pushes the result closer to the winning bids. In experiments we set  $\lambda = 1$  to mimic a risk-averse choice. The maximum could be achieved by letting the first derivative equal zero. We use a numerical solution and approximate the result in algorithm’s implementation using the Broyden-Fletcher-Goldfarb-Shanno (BFGS)

method, the most popular one in Quasi-Newton methods class. The BFGS method uses the first derivative and has proven good performance for smooth optimizations [5] and. We refer to this algorithm as BAYESIANB in experiments.

### 3.3.2 Univariate Log-normal Distribution

The most greedy choice from Equation 8 is to let  $\lambda = \frac{B_1}{B_2}$ : regardless of the second highest bidder, the publisher could always learn and approach the highest bid directly. It is equivalent to modelling the winning bids using a univariate Log-normal distribution, which has a simpler solution

$$\begin{aligned} B_1 &\sim \text{lognorm}(\mu, \sigma^2) \\ \mu(t) &\sim \mathcal{N}(\theta(t), \delta^2(t)) \end{aligned}$$

and still we assume that  $\sigma^2$  is known. Every time we observe the highest bid  $b_1$  we obtain

$$\begin{aligned} \theta(t+1) &= \frac{\theta(t)\delta^2(t) + \sigma^2 b_1}{\sigma^2 + \delta^2(t)} \\ \delta^2(t+1) &= \frac{\delta^2(t)\sigma^2}{\sigma^2 + \delta^2(t)} \end{aligned}$$

where the priori was obtained using the training dataset.

Then the publisher simply chooses

$$\alpha(t+1) = \log(\theta(t+1)) + \beta(\sigma^2 + \delta^2(t+1)) \quad (9)$$

as the new reserve price where  $\beta$  is the risk preference for this model which could be either positive or negative. In experiment we set  $\beta = 0$  for simplicity.

We refer to this algorithm as BAYESIANU. Again, the fitting of private value distribution is beyond the scope of this paper but the ideas and algorithms always hold, allowing different models being implemented easily (it is highly likely that different models would be used for different placements, hour-of-day, and etc.).

## 3.4 Private Value Free Algorithms

There are algorithms based on common knowledge or intuition and do not require the assumption that each advertiser has their own private fixed valuation about the impression at hand. In fact, they play a major role in reserve price set up in real-world. Among those choices, we present two simple methods as our baselines.

The most basic one is the FIXED algorithm which sets a static reserve price for all time steps, regardless of outcome of auctions. Formally it is  $\alpha(t) \equiv a$  where  $a$  denotes the pre-chosen fixed reserve price.

When  $\alpha = 0$  it becomes ZERO which totally relies on the auction mechanism, in our case the second price auction, assuming that quality score (or similar factors like bid bias) is absent. It is apparently the most altruistic one in our strategy space which does not try to gain any extra benefit and could be used as a baseline to measure bidders' attrition.

The FIXED (including ZERO) is the most straightforward and easy to implement. Besides, it directly reflects the publisher's private valuation of the inventory. Before adopting the automated selling, publishers are used to negotiating contracts with advertisers or their representatives. The contract price could be easily converted to a reserve price. Also, due to simplicity this is probably the most widely adopted algorithm in today's ad marketplaces. In experiments, we consider FIXED and ZERO the most naive baselines. We

show that, under certain circumstances, although simple, they performed surprisingly well.

Another family of private-value-free algorithms is based on historical payoff of auctions. For example, the AVERAGE sets the reserve price to the average of past payoff. Formally it is

$$\alpha(t) = \frac{1}{M} \sum_{i=t-M}^{t-1} r(i) \quad (10)$$

where  $r(t)$  is the payoff at time step  $t$  and  $M$  is the averaging window.

A natural extension is the weighted average variation, which values more of recent payoff. Formally it is

$$\alpha(t) = \frac{1}{M} \sum_{i=t-M}^{t-1} w(i, t)r(i) \quad (11)$$

where  $w(\cdot)$  is the weighting function, which could take various forms. In our experiments we used linear weights and denoted it as WEIGHTEDL.

AVERAGE and WEIGHTEDL share the similar insights with FIXED, except recognizing the fluctuation of demand and supply in the market. This fluctuation could greatly affect the reserve price setting especially at hour-of-day level. It is discussed in details in Section 4.1.

## 4. AN EMPIRICAL STUDY

Understanding the bidder behaviour in RTB auctions is crucial to the reserve price problem. In this section we first report our findings from an empirical study of bidders' behaviours in RTB auctions, focusing on the daily bidding pattern, the bidders' lifetime, the change of bidders and winners, and bids' distribution. Based on real-world statistics, we discuss their impact on setting reserve prices. Then we report a large scale online experiment in the same production platform to test various algorithms proposed before.

### 4.1 Live Experiment Setups and Dataset

First we highlight the dataset used for this empirical study:

- The dataset is obtained from a production platform in UK;
- The dataset contains 13,478,591 impressions from Dec 2012 to Feb 2013;
- These impressions are sampled from 39 placements belonging to 16 websites of different categories (e.g., finance, sports, pc & console games, news, technology & gadgets).

On average, these websites receive 1.8m impressions daily. We roughly sampled one impression every 10 seconds to construct the dataset due to the limit of computing power and storage. For each auction we have logged the following information: the venue (URL, position and size of the placement, above or below fold), date and time, demographics of visitor (geographical location, IP, browser agent, language, and segments if available), bidding activities, and current reserve price.

Parameters of tested algorithms were obtained during the training stage (14 Dec 2012 to 18 Jan 2013) then they were evaluated during the testing stage (19 Jan to 21 Feb 2013) in the live system. There were 6 placements having few bidders (less than 5) or small traffic (less than 10k impressions per day). They were dropped from the online experiment but were compared with in bidders' attrition analysis.

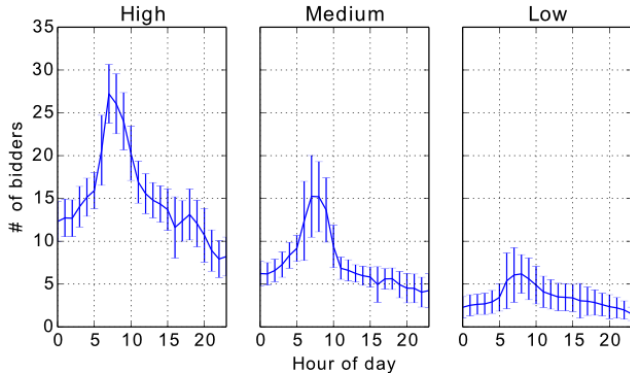


Figure 4: The average number of bidders varies across hour-of-day and placements. The error bars show standard deviation. A clear pattern emerges that for all placements the number of bidders peaks at 8-10am.

Besides, we selected 8 placements as the control group to which the optimisation had never been applied. Although we evaluated the ZERO in the experiment, having a zero reserve price at some hours is not equivalent to no optimisation at all. Having the control group makes it easier and more accurate to understand the buyers’ reaction or attrition. Offline evaluation had been performed on these placements too (not reported here due to space limit).

## 4.2 Bidders’ Behaviours

Statistically speaking, the more bidders, the higher winning bids would be. Thus, the competition level (the demand) is the first important factor we use to group the placements. In Figure 4 we illustrated the number of bidders varies across a day on different placements. By using the basic mean, std, max, and min statistics we divided placements into groups:

- Group High, 10 to 30 bidders, 6 placements;
- Group Medium, 5 to 20 bidders, 16 placements;
- Group Low, 2 to 10 bidders, 11 placements;
- Dropped, 1 to 5 bidders, 6 placements.

The experiment results confirm that for different competition levels the same policy would perform differently. Another finding is that almost on all placements the number of bidders peaks at 8-10am. With low level of competition it is still noticeable (c.f Figure 4, right). This is hard to justify from the publishers’ perspective; however, we assume this is due to daily budget settings of campaigns and hour-of-day targeting.

Note that the number of bidders does not always align with the number of impressions across hour of day. It is commonly known that the number of impressions peaks at evening hours (6-11pm) as reported in [34]. Having more bidders at the morning hours (8-10am) when there are less impressions results in even higher winning bids. It implies the possibility of optimisation from both demand and supply side. In this paper, we argue it is necessary to perform at least hourly optimisation and report the results later.

Another observation is on the lifetime of bidders, sometimes referred to as the *flight time* of campaigns. There are 755 bidders (Note again that we consider each campaign a bidder) in total in the dataset. As illustrated in Figure 5 we can see two clear patterns: “learners” and “4-week campaigns”, as well as a less clear pattern “6-week campaigns”. This plot of lifetime is a reflection of how online advertising work: most of the campaigns are relative short, limited to




Figure 5: The histogram of 755 bidders’ lifetime across all placements. We consider each campaign a bidder. From the figure we observe three patterns: two created by common flight-time of campaigns (4 weeks and 6 weeks) and one created by exploration of learning algorithms (0-2 days). Due to limited time of data collection there are some instances lying between 0-25 days because these bidders started their activities before the start date of our experiments.

4 weeks or 6 weeks by advertisers. Meanwhile, we observed spikes around days 0-3 that indicates the activities of *learners*: advertisers spend budget on a placement to learn its performance against a specific campaign, which is essential in demand-side optimisation.

The pattern between days 6-20 is less clear due to the limited time span of data collection. There were good quantity of campaigns already running at the beginning of the collection, and they ended randomly several days later. This could be due to the *holiday effect*, too (the data collection started on 14 Dec 2012). It is hard to identify and remove these campaigns from the dataset; however, they do not affect the result of auctions with reserve prices, and would become less significant when more data is available. The figure was created by aggregating logs across all placements. Although not reported, placements with different levels of competition showed the similar pattern; when aggregated they become easier to identify.

The impact of bidders’ lifetime on reserve price algorithms is mostly indirect. However, since the optimal auction theory models auctions as iterated and dynamic games, it is preferable that bidders keep bidding for a while so that both players could learn and act properly. On the other hand, the publisher might want to set different reserve prices with respect to ages of bidders (or at least winners), esp. when lots of learners are present. We leave this discussion to the future work.

Another factor that would affect reserve prices is the change of winners of auctions, as illustrated in Figure 6. The instances of such change are in fact not negligible. Combining with the bidders’ lifetime plot, a significant impact emerges. Since the winner of auctions on the same placement is constantly changing, it is difficult for buyers to detect if a reserve price is place, esp. when the reserve price is changing, too. This in turn favours publishers in the repeated and dynamic game setting, since advertisers may have only incomplete information set.

## 4.3 Bids’ Distribution Test

An important assumption made by optimal auction theory is that bidders draw their private values from a monotonically increasing distribution. In second price auction they bid on their private values; thus the distribution of bids




Figure 6: The change of winners for placements with different levels of competition in 4 days. The fact that a bidder does not always win could add difficulty to reserve price detection if undisclosed. The result also implies the change rate does not necessary relate to the competition level.

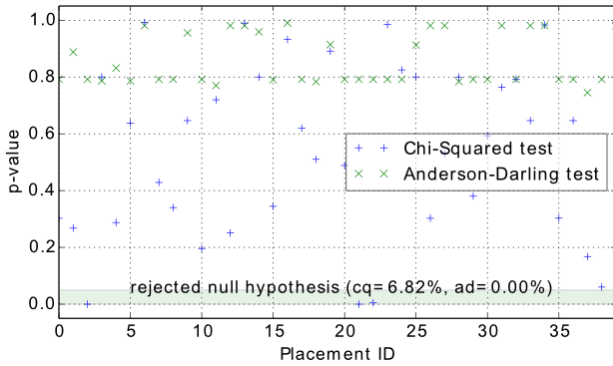


Figure 7: Only bids from 3 out of 44 placements (6.82%) accept the Uniform distribution hypothesis. The Uniform distribution is tested by Chi-Squared test and the Log-normal distribution is tested by Anderson-Darling test.

reflects their individual valuations of an impression. As introduced in Section 2 researchers often assume bidders are symmetric, and bids follow Log-normal distribution or Uniform distribution. This assumption then becomes the core of optimal auction theory based algorithms of computing reserve prices [24, 26, 31].

In our empirical study, bids are fitted to these two distributions. For the Log-normal distribution we check their normality using Anderson-Darling test [1] after taking the logarithm. We do this test at both auction level and placement level. For the Uniform distribution we use the Chi-squared test [17] only at the placement level. We use  $p\text{-value} < 0.05$  to reject the null hypothesis.

The results are reported in Figure 8 and 7. Clearly only the acceptance of Uniform distribution assumption at placement level is noticeable yet still at a low ratio of 6.82%. The poor results are probably due to the bursts and randomness of bids as illustrated in Figure 9. However, to be consistent with the research literature, we still adopt the Log-normal assumption when implementing the optimal auction theory based algorithm in the experiment. The undesirable fitting could be the main reason that the OPTAUC does not perform the best.

#### 4.4 Live Test Results and Discussions

We carried out the online experiment in the same production platform. The experiment includes 25 placements as the treatment group and additional 8 as the control group. For




Figure 8: Only bids from less than 0.1% of all auctions accept the Log-normal distribution hypothesis. The plot shows a randomly sampled of 1000 auctions. Only the Log-normal distribution is tested by Anderson-Darling test.

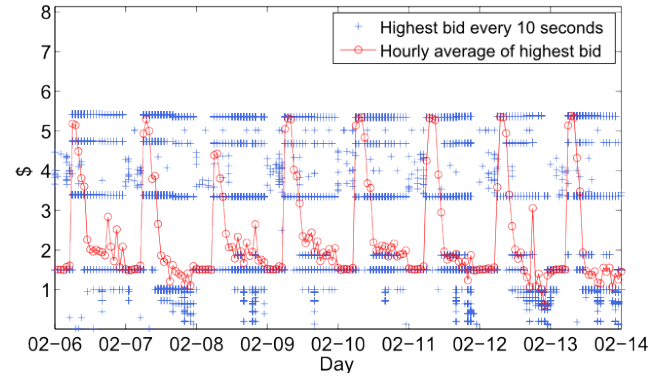


Figure 9: The plot of hourly average winning bids and bursting instances on a placement of high level of competition. Due to these bursts the model-based algorithms fitted poorly and did not perform well.

each placement we have collected RTB auction logs which record bidders and their bids in every auction. The live experiments lasted from 19 Jan to 21 Feb 2013. The parameters of algorithms were trained using logs from 14 Dec 2012 to 18 Jan 2013.

##### 4.4.1 Algorithms and Scheduling

The following algorithms and configurations were evaluated.

- GAME, with publisher value set to \$1.0;
- ZERO, with a fixed reserve price set to \$0.0;
- AVERAGE;
- WEIGHTEDL, with a looking back window size of 5;
- ONESHOT, with  $\epsilon = 1$  and  $\lambda_h, \lambda_m,$  and  $\lambda_l$  learned individually for each placement;
- BAYESIANB, with  $\theta, \Sigma,$  and  $\Delta$  learned individually for each placement;
- BAYESIANU, with  $\theta, \sigma,$  and  $\delta$  learned individually for each placement.

Note that the private value of \$1.0 of the publisher was chosen based on the publisher's suggestion.

We evaluated algorithms as the following: impression-level logs were constantly sampled then fed to algorithms. We ran the experiment for 30 days. The experiment used a Round Robin scheduler at the hour level. For example, at Hour=1 and Day=1 we used Algorithm=OPTAUC on Placement=1, Algorithm=ZERO on Placement=2, etc.; on Hour=2 and Day=1 we used Algorithm=BAYESIANU on



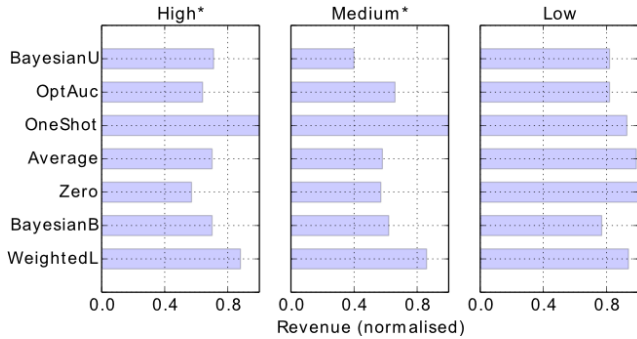


Figure 10: A snapshot of main results of the online experiment, on all placements, from 6am to 10am across all days. The \* indicates the OneShot performed significantly better than others (except the WeightedL) using Wilcoxon signed-rank test. Considering the percentage of placements of high, medium, low level of competition, the OneShot performed significantly better in about 70% cases. On placements having fewer bidders it is hard to distinguish the results. Although not reported, at other hours the results are similar.

Placement=1, Algorithm=OPTAUC on Placement=2, etc. With this scheduling method the day-of-week factor does not influence the evaluation of performance of algorithms. In this way for every (hour-of-day, algorithm, level of competition) tuple we obtained 10 sets of observations. We split the results into 6 non-overlapping chunks and performed Wilcoxon signed-rank test [9] to check significance.

#### 4.4.2 Results

We report the performance of algorithms with respect to levels of competition in Figure 10, where we find the private-value-free algorithms performed better than private-value-based ones. Note that the figure only includes data from 6-10am across all days but for other hours the results are similar. Considering the percentage of placements of high, medium, low level of competition, the ONESHOT performed significantly better in about 70% cases. This is very consistent with the findings made in [26].

Looking at the results and various analysis made before, we believe the undesirable fitting of Log-normal distribution led to the poor performance of private-value-based algorithms on placements with many bidders. Although there is a steady hourly average winning bids pattern on these placements, there are also lots of short periods of burst (both upward and downward) around the curve as shown in Figure 9, which could have dragged the models away from the curve easily. On the contrary the private-value-free algorithms had better chance of capturing these bursts.

The ONESHOT shares the same intuition as the WEIGHTEDL: using the observations of recent future to infer the reserve price. More specifically, we replayed the auction logs in training dataset against various parameter sets and adopted the best performing one.

### 4.5 Bidders' Attrition

As we argued before, a key difference of reserve price optimisation in RTB is that the inventory is not as limited as ordinary auctions or sponsored search. Thus, it is also interesting to study whether the reserve price optimisation could result in bidders' attrition [25] in long term (e.g., reduce their bid values or volume). Figure 11 compares winning bids and number of bidders before and after the reserve price optimisation experiment. It suggests that it may not

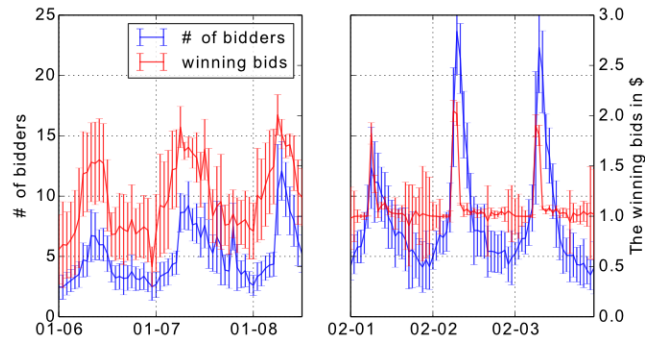


Figure 11: The comparison of winning bids and number of bidders before and after the experiment. Both factors increased after the experiment, which contradicts the attrition hypothesis that bidders would reduce their bid or volume after reserve price optimisation.

Table 2: The change of number of bidders after the experiment. On most of placements, especially ones with high level of competition, the increment is observed.

|         | $\leq -10\%$ | $-5\% \sim -10\%$ | $5\% \sim 10\%$ | $\geq 10\%$ |
|---------|--------------|-------------------|-----------------|-------------|
| High    | 3            | 0                 | 0               | 3           |
| Medium  | 4            | 2                 | 4               | 6           |
| Low     | 3            | 0                 | 0               | 8           |
| Control | 1            | 0                 | 0               | 5           |




Figure 12: Bidders' attrition test after the optimisation, comparing with the number of bidders before the experiment. The t-test suggests the attrition is significant on 7 placements. 3 are placement of low level of competition and 4 of medium level. However, it is also worth noting that many more placements see an increment of number of buyers after the experiment, c.f. Table 2.

be the case in RTB. To test the statistical significance of the change, we take the mean and standard deviation of number of bidders for each placement before (12-18 Jan 2013) and after (15-21 Feb 2013) the optimisation experiment. The null hypothesis of the buyers' attrition could be denoted as  $H_0: \mu_1 - \mu_2 > 0$  where  $N_1(\mu_1, \sigma_1)$  and  $N_2(\mu_2, \sigma_2)$  are the Normal distribution fitted from number of bidders before and after the experiment respectively. The rejection of this null hypothesis implies no significance of bidders' attrition.

The result of t-test with p-value 0.05 is summarised in Figure 12. There are only 7 placements out of 39 showed significant drop of number of bidders after the experiment. 3 of them are placements of low level of competition and 4 of medium level, where a slight change of would be significant. Interestingly, from Table 2 we see an opposite pattern emerges: the publisher was actually seeing more bidders on

most of placements. This could be due to three reasons: 1) during holidays (the beginning of the experiment) there are generally less bidders because of reduced amount of Internet traffic; 2) bidders are not aware of the reserve price optimisation so no reaction is taken; 3) the campaigns' lifetime introduces fluctuation.

## 5. CONCLUSIONS

In this paper, we discussed a specific problem for publisher revenue optimisation in Real-Time Bidding: to find the optimal placement level reserve price for single-item display ad auctions. It is a problem of significant importance, which has been studied extensively in the game theory, but is relatively new in online advertising research literature, especially in RTB which has many unique features. In the paper we analysed drawbacks of the optimal auction theory in the RTB practice, derived dominant strategies from simplified version of the auction game, and compared it with other commonly adopted algorithms in a production platform. We also discussed the challenges of optimising reserve prices in RTB and analysed bidders' attrition in the experiment. Due to the complexity of the problem and practical constraints there are unsolved issues in this paper, including the fitting of bids, parameters tuning, and more comprehensive attrition analysis. We would like to address these in the future work.

## 6. REFERENCES

- [1] ANDERSON, T. W., AND DARLING, D. A. A Test of Goodness of Fit. *Journal of the American Statistical Association* 49 (1954), 765–769.
- [2] AUSUBEL, L. M., AND CRAMTON, P. Vickrey auctions with reserve pricing, 2004.
- [3] BAPNA, R., GOES, P., GUPTA, A., AND OTHERS. Insights and analyses of online auctions. *Communications of the ACM* 44, 11 (2001), 42–50.
- [4] BHARADWAJ, V., MA, W., SCHWARZ, M., SHANMUGASUNDARAM, J., VEE, E., XIE, J., AND YANG, J. Pricing guaranteed contracts in online display advertising. In *Proceedings of the ACM CIKM* (2010).
- [5] BONNANS, J. F., GILBERT, J. C., LEMARÉCHAL, C., AND SAGASTIZÁBAL, C. A. *Numerical optimization: theoretical and practical aspects*. Springer, 2006.
- [6] BRODER, A., FONTOURA, M., JOSIFOVSKI, V., AND RIEDEL, L. A semantic approach to contextual advertising. In *Proceedings of the ACM SIGIR* (2007).
- [7] CAVALLO, R., MCAFEE, R. P., AND VASSILVITSKII, S. Display Advertising Auctions with Arbitrage. 1–21.
- [8] CHAKRABORTY, T., EVEN-DAR, E., GUHA, S., MANSOUR, Y., AND MUTHUKRISHNAN, S. Selective call out and real time bidding. *Internet and Network Economics* (2010), 145–157.
- [9] CRICHTON, N. Wilcoxon signed rank test. *Journal of Clinical Nursing, Blackwell Science Ltd* (2000).
- [10] EDELMAN, B., OSTROVSKY, M., AND SCHWARZ, M. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. Tech. rep., National Bureau of Economic Research, 2005.
- [11] EDELMAN, B., AND SCHWARZ, M. Optimal auction design in a multi-unit environment: The case of sponsored search auctions. *Unpublished manuscript, Harvard Business School* (2006).
- [12] EMARKETER. RTB Ad Spend Continues Robust Growth - eMarketer.
- [13] EVEN-DAR, E., FELDMAN, J., MANSOUR, Y., AND MUTHUKRISHNAN, S. Position auctions with bidder-specific minimum prices. *Internet and Network Economics* (2008), 577–584.
- [14] FELDMAN, J., MIRROKNI, V., MUTHUKRISHNAN, S., AND PAI, M. M. Auctions with intermediaries: extended abstract. In *Proceedings of the ACM EC 2010* (2010).
- [15] GHOSH, A., MCAFEE, P., PAPANENI, K., AND VASSILVITSKII, S. Bidding for representative allocations for display advertising. *Internet and Network Economics* (2009), 208–219.
- [16] GHOSH, A., RUBINSTEIN, B. I. P., VASSILVITSKII, S., AND ZINKEVICH, M. Adaptive bidding for display advertising. In *Proceedings of the WWW 2009* (2009).
- [17] GREENWOOD, P. E., AND NIKULIN, M. S. *A guide to chi-squared testing*, vol. 280. Wiley-Interscience, 1996.
- [18] KAUFFMAN, R. J., AND WOOD, C. A. Running up the bid: detecting, predicting, and preventing reserve price shilling in online auctions. In *Proceedings of the ACM EC 2003* (2003).
- [19] KRISHNA, V. *Auction theory*. Academic press, 2009.
- [20] LEVIN, D., AND SMITH, J. L. Optimal reservation prices in auctions. *The Economic Journal* 106, 438 (1996), 1271–1283.
- [21] MATTHEWS, S. *A technical primer on auction theory I: Independent private values*. No. 1096. 1995.
- [22] MCAFEE, R. P., AND VINCENT, D. Sequentially optimal auctions. *Games and Economic Behavior* 18, 2 (1997), 246–276.
- [23] MUTHUKRISHNAN, S. Ad exchanges: Research issues. *Internet and network economics* 8, 2 (2009), 1–12.
- [24] MYERSON, R. B. Optimal auction design. *Mathematics of operations research* 6, 1 (1981), 58–73.
- [25] OSBORNE, M. *A Course in Game Theory*. 2011.
- [26] OSTROVSKY, M., AND SCHWARZ, M. Reserve prices in internet advertising auctions: A field experiment. *Search* (2009), 1–18.
- [27] PERLICH, C., DALESSANDRO, B., HOOK, R., STITELMAN, O., RAEDER, T., AND PROVOST, F. Bid optimizing and inventory scoring in targeted online advertising. In *Proceedings of the ACM SIGKDD 2012* (2012), p. 804.
- [28] RILEY, J. G., AND SAMUELSON, W. F. Optimal auctions. *The American Economic Review* (1981), 381–392.
- [29] ROELS, G., AND FRIDGERSDOTTIR, K. Dynamic revenue management for online display advertising. *Journal of Revenue & Pricing Management* 8, 5 (May 2009), 452–466.
- [30] STEIN, C. M. Estimation of the mean of a multivariate normal distribution. *The annals of Statistics* (1981), 1135–1151.
- [31] XIAO, B., YANG, W., AND LI, J. Optimal Reserve Price for the Generalized Second-Price Auction in Sponsored Search Advertising. *Journal of Electronic Commerce Research* 10, 3 (2009), 114.
- [32] XU, H., BATES, C. K., AND SHATZ, S. M. Real-time model checking for shill detection in live online auctions. In *Proc. Int. Conf. Software Engineering Research and Practice (SERP'09)* (2009), pp. 134–140.
- [33] YAN, J., LIU, N., WANG, G., AND ZHANG, W. How much can behavioral targeting help online advertising? *Proceeding WWW '09 Proceedings of the 18th international conference on World wide web* (2009), 261–270.
- [34] YUAN, S., WANG, J., AND ZHAO, X. Real-time Bidding for Online Advertising: Measurement and Analysis. In *Proceedings of the ADKDD workshop 2013* (2013).