

Bid Shading by Win-Rate Estimation and Surplus Maximization

Shengjun Pan, Brendan Kitts, Tian Zhou, Hao He, Bharatbhushan Shetty, Aaron Flores, Djordje Gligorijevic, Junwei Pan, Tingyu Mao, San Gultekin and Jianlong Zhang

Verizon Media
Sunnyvale, California

{alanpan,brendan.kitts,tian.zhou,hao.he,bharatbs,aaron.flores}@verizonmedia.com
{djordje,jwpan,tingyu.mao,sgultekin,jianlong}@verizonmedia.com

ABSTRACT

This paper describes a new win-rate based bid shading algorithm (WR) that does not rely on the minimum-bid-to-win feedback from a Sell-Side Platform (SSP). The method uses a modified logistic regression to predict the profit from each possible shaded bid price. The function form allows fast maximization at run-time, a key requirement for Real-Time Bidding (RTB) systems. We report production results from this method along with several other algorithms. We found that bid shading, in general, can deliver significant value to advertisers, reducing price per impression to about 55% of the unshaded cost. Further, the particular approach described in this paper captures 7% more profit for advertisers, than do benchmark methods of just bidding the most probable winning price. We also report 4.3% higher surplus than an industry Sell-Side Platform shading service. Furthermore, we observed 3% – 7% lower eCPM, eCPC and eCPA when the algorithm was integrated with budget controllers. We attribute the gains above as being mainly due to the explicit maximization of the surplus function, and note that other algorithms can take advantage of this same approach.

CCS CONCEPTS

• **Applied computing** → **Online auctions**; • **Information systems** → **Display advertising**; • **Computing methodologies** → **Machine learning algorithms**.

KEYWORDS

online bidding, shading, auction, advertising, bid, optimization

ACM Reference Format:

Shengjun Pan, Brendan Kitts, Tian Zhou, Hao He, Bharatbhushan Shetty, Aaron Flores, Djordje Gligorijevic, Junwei Pan, Tingyu Mao, San Gultekin and Jianlong Zhang. 2020. Bid Shading by Win-Rate Estimation and Surplus Maximization. In *AdKDD '20: SIG Conference on Knowledge Discovery and Data Mining, August 2020, San Diego, California, USA*. ACM, New York, NY, USA, 6 pages.

1 INTRODUCTION

Online Advertising auctions have been dominated by Second Priced Auctions (SPAs) since their early implementations in the 1990s.

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.
AdKDD '20, August 2020, San Diego, California, USA

© 2020 Association for Computing Machinery.
ACM ISBN 978-x-xxxx-xxxx-xx/YY/MM...\$15.00

Google famously used Second Price Auctions for its Adwords and AdSense auctions, and, in 2017, generated 90% of its revenue from Second Price Auctions [12]. However, there was a dramatic shift in online advertising between 2018 and 2019. As of 2020, almost all major display ad auctions have switched from Second to First Price Auctions (FPAs) [13, 14]. Several factors conspired to drive the industry towards the adoption of FPA, including the widespread growth of header bidding with its incompatibility with SPAs [18], increased demand for transparency and accountability [7, 11, 26, 29], and yield concerns [3, 21, 25].

Unfortunately for advertisers, First Price Auctions leave private value bidders susceptible to over-paying. For instance, if the bidder's private value of an impression was \$10.00, and the winner knew the second placed bidder's price was just \$1.00, they could bid just \$1.01 and effectively collect a \$8.99 profit. If they instead bid their private value, they would be charged the entirety of the \$10.00 and they would have \$0 profit!

The practice of strategically decreasing bid price below the buyer's private value is known as *bid shading*. Bid shading has been observed in a variety of real world auctions including FCC Spectrum [6], US Oil Deposits [5], Cattle auctions [8], US Treasury auctions [17] and others. Despite its widespread use, there has been little work done on methods to systematically exploit shading, particularly when data is available to make it possible to predict auction clearing prices.

2 THE BID SHADING PROBLEM

Given bid request b_i , and a valuation V_i , if we won the impression, which represents how much the advertiser expects to capture from the impression, how much should the advertiser discount their valuation of the dollar value that the advertiser expects to obtain, and the bid $b_i = g_i V_i$ is also in real dollars, the advertiser's financial gain, or *surplus*, is equal to:

$$surplus \stackrel{\text{def}}{=} \sum_i (V_i - g_i V_i) \mathbf{I}(g_i V_i > \hat{b}_i), \quad (1)$$

where $g_i \in (0, 1]$ is the shading factor to apply to the bidder's private value V_i , \hat{b}_i is the minimum bid price to win, and $\mathbf{I}(b_i > \hat{b}_i) = 1$ if the impression is won, and 0 otherwise. The task is to find a shading factor g_i that maximizes the surplus to the advertiser.

3 PREVIOUS WORK

3.1 Bid Shading Theory

Bid shading is a common tactic in repeated First Price Auctions. [36] found robust evidence of shading in Austrian livestock auctions ,

[8] reported shading in a Texas cattle market, and [17] found the practice in auctions for US Treasury notes.

Auctions generally need to be repeated and predictable for bid shading to be practically feasible, but under these conditions, it often occurs organically. Pownall and Wolk (2013) showed that bid shading for repeated internet auction prices increased over time; by about 26% after 10 iterations [24]. When there are enough repeated games bidders can even develop collusive shading strategies where bidders actively coordinate to have low bids [16, 23].

Although behavior varies from auction to auction, several studies have shown that the magnitude of shading tends to increase with the average price on the auction [4, 6, 17]. This is likely to occur because of the more substantial losses involved on higher priced auctions, if shading isn't sufficient. This result suggests that using a measure of the expense of the auction is valuable when trying to estimate the shading factor - a finding we revisit later in Section 6.

In situations where the supply is plentiful, and demand limited, buyers can shade deeper. In looking at this phenomenon in the US Treasury Market, Hortacsu et. al. (2017) find that large institutional buyers on average shade more aggressively than small indirect buyers [17]. This seems to be because these large buyers effectively control a large percentage of bidders, and so it is almost like they are able to coordinate the buying of multiple buyers. They can therefore drive the bid prices for a large percentage of bidders down, whilst still meeting their goals.

3.2 Previous Algorithms

In 2018 and 2019, Rubicon [1, 25], AppNexus [3] and Google [14, 15, 28, 30] all released Sell-Side bid shading services. Leading up to this, there had been reports of dramatically lower ROI from the new First Price Auctions [18, 21]. Never-the-less, this is a surprising move as Sell-Side Platforms are potentially decreasing their yield, and they clearly have a different incentive from buyers. The sell-side algorithms seem to reflect this incentive difference. The descriptions of these services suggest that they try to keep bid prices high enough to maintain a set win-rate, but preventing the bid price from becoming too extreme; which might risk an advertiser to halt their bidding due to poor Return on Investment. Rubicon released data suggesting that their service decreases First Price CPMs by a modest 5% over 4 months [25]. AppNexus reported that prices under their service were 25% lower over 100 days [3]. We tried one of the services, and recorded the shading distribution in Figure 2. Most of the bid shades were about 90%, which is conservative for our problem. Further analysis on Sell-Side "Bid Shaders" are in Section 7.

On the Demand Side, a variety of algorithms have been explored, although generally not exactly for bid shading applications. [35] developed a censored winning bid probability estimator. They observed that when a bidder submitted a bid and lost, the information gained is that the winning price is somewhere above the submitted price, and when a bidder submits a bid and wins, the minimum bid to win is at a price somewhere below their submitted bid. Using these two cases, the authors developed a Maximum Likelihood procedure to estimate the probability of the winning bid being any of the bid prices. This created a distribution of the probable winning bids, with the most likely winning price being used for bidding.

[34] extended their work to using a neural network to estimate the parameters of the win probability distribution.

An unpublished implementation [19] used Logistic Regression to predict the optimal bid shading factor using features in the request. The predicted factor was then used as a multiplier on the unshaded bid price.

The approaches described above [19, 34, 35] all focus on predicting the probable winning bid price. However, the surplus maximum is very different from the minimum bid to win. An accurate (unbiased, symmetric noise) win probability estimator will be below the winning bid price about 50% of the time - this means that 50% of the surplus won't be captured *by design*. If the change in new impressions captured at a higher bid price, over-weights the marginal decrease in profitability per impression, the optimum for surplus can be higher than the most probable bid.

Unpublished work [20] is one of the few that we know of to attempt to explicitly maximize the surplus function. These authors estimate shading factors for a set of fixed segments based on three bid samples taken in real-time to estimate the local surplus landscape. However the approach has many drawbacks: the segments have to be predetermined and finding a suitable segment definition requires substantial analysis. The information across segments is not shared, which is a problem for segments that do not have enough traffic. Further, the set of possible segments quickly explode as the number of variables used to define them increases. The approach taken in this paper uses a model to estimate the surplus function, and so a very large number of features can be used, and model induction is also automated, easy to maintain, and improve.

In order to compare the method we used to prior work, we have included an implementation of the Logistic Regression algorithm from [19], the Distribution Estimator algorithm from [35], and the Segment-based Surplus maximizer [20] in the benchmarks which we use to analyze algorithm performance in Section 7.

4 CANONICAL ALGORITHM

Given a bid request for First Price Auction, let x_1, x_2, \dots, x_k be the set of publisher and user attributes that we will use to find the best bid price b^* . Let \hat{b} be the highest bid price from other competing bidders, which value is unknown. Note that \hat{b} depends on both attributes x_i s which represent the item that is being auctioned, and external competing bidder behavior. \hat{b} follows an unknown distribution $\mathcal{D}_{\hat{b}|x_1, x_2, \dots, x_k}$ with cumulative probability distribution $\text{cdf}_{\hat{b}|x_1, x_2, \dots, x_k}$. When the context is clear, we use $\mathcal{D}_{\hat{b}}$ and $\text{cdf}_{\hat{b}}$ for simplicity.

If the distribution $\mathcal{D}_{\hat{b}}$ is known, we can calculate the optimal bid price b^* directly as follows. Let $\mathbf{I}(b > \hat{b})$ be 1 if $b > \hat{b}$ and 0 otherwise, which indicates if the submitted price b wins the auction. Then the surplus when the submitted price is b would be

$$\text{surplus} = (V - b)\mathbf{I}(b > \hat{b}) = \begin{cases} V - b, & \text{if } b > \hat{b}, \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

The optimal bid price can be calculated as the price that maximizes the expected surplus

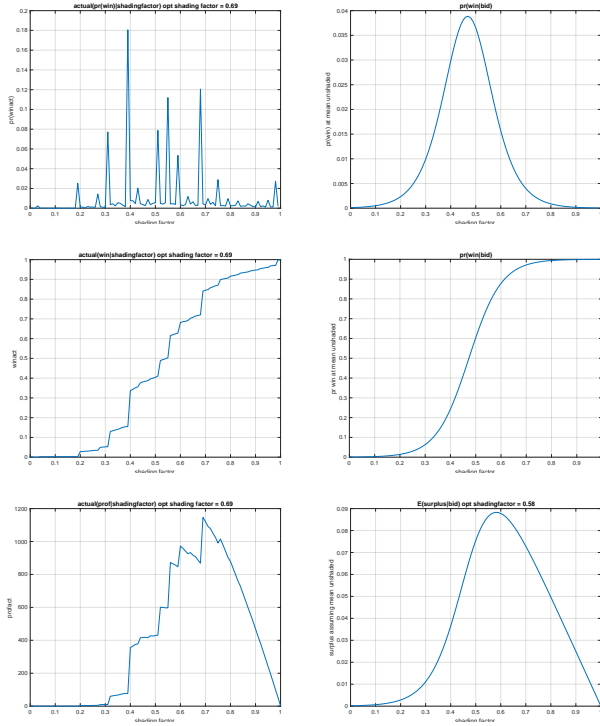


Figure 1: Top: Actual PDF for \hat{b} (left) versus estimate (right); middle: CDF actual versus estimate; bottom: Surplus distribution actual versus estimate.

$$\begin{aligned}
 b^* &= \arg \max_{b>0} \mathbb{E}[\text{surplus}] \\
 &= \arg \max_{b>0} \mathbb{E} \left[(V - b) \mathbf{I}(b > \hat{b}) \right] \\
 &= \arg \max_{b>0} (V - b) \text{cdf}_{\hat{b}}(b). \tag{3}
 \end{aligned}$$

For simple forms of $\text{cdf}_{\hat{b}}(b)$, the optimization problem (3) can be solved analytically. For example, suppose \hat{b} distributes uniformly over the interval $[B_0, B_1]$, where $0 \leq B_0 < B_1$. This produces a $\text{cdf}_{\hat{b}}(b)$ that is piece-wise linear, with a flat region of 0.0 from $[0, B_0]$, a constant slope from $[B_0, B_1]$, and another flat region of 1.0 above B_1 . The bid price b^* that maximizes the surplus can be calculated as below

$$\mathbb{E}[\text{surplus}] = (V - b) \text{cdf}_{\hat{b}}(b) = \begin{cases} 0, & \text{if } b < B_0, \\ (V - b)(b - B_0)/(B_1 - B_0), & \text{if } B_0 \leq b \leq B_1, \\ V - b, & \text{if } b > B_1. \end{cases}$$

It is straightforward to see that

$$\max \mathbb{E}[\text{surplus}] = \begin{cases} \frac{(V - B_0)^2}{4(B_1 - B_0)} \text{ at } b^* = \frac{V - B_0}{2}, & \text{if } V \leq 2B_1 - B_0, \\ V - B_1 \text{ at } b^* = B_1, & \text{if } V > 2B_1 - B_0. \end{cases}$$

However, in practice, we rarely see such simple form of distributions. Figure 1 shows an example of the empirical PDF of \hat{b} , including the derived surplus distribution.

Our approach breaks into two steps:

- Step 1** Estimate the distribution $\mathcal{D}_{\hat{b}|x_1, x_2, \dots, x_k}$;
- Step 2** Solve the maximization problem (3).

4.1 Distribution Estimation

Given publisher and user attributions x_1, \dots, x_k and bid price b , we first train a classification model with historical data:

$$\text{Pr}(\text{win}) = \text{cdf}_{\hat{b}}(b) = F \left(w_0 + \sum_{i=1}^k w_i x_i + \beta g(b) \right), \tag{4}$$

where F is a fitting function that outputs a value between 0 and 1, which must be monotonically increasing in b (higher bid price leads to higher winning rate), and $g(b)$ is a bid transformation function such that $F \rightarrow 0$ as $b \rightarrow 0$, that is, as bid price goes to 0, the winning probability also goes to 0, and the weights to be learned are w_0, w_1, \dots, w_k and β .

For $g(b)$, we use the logarithm of bid price $g(b) = \log(b)$ so that $g(b) \rightarrow -\infty$, as $b \rightarrow 0$. For F , we use the logistic function [9, 27] so that $F(x) \rightarrow 0$ as $x \rightarrow -\infty$, with the constraint that $\beta > 0$.

Other forms $g(b)$ and F can be explored, but our choices of simple forms, besides satisfying mentioned constraints, allow the maximization problem (3) in Step 2 to be solved efficiently. More details will follow later in Subsection 4.2. With our choice of functions F and $g(b)$, we have the following win-rate classification model:

$$\text{Pr}(\text{win}) = \left(1 + e^{-(w_0 + \sum_{i=1}^k w_i x_i + \beta \log b)} \right)^{-1}, \tag{5}$$

which can be trained by gradient descent [9]. Note that the training should be constrained such that $\beta > 0$. In practice we found that it's not necessary since our learned β without constraint turns out always positive.

4.2 Surplus Maximization

With a trained win-rate model from (5), the optimal bid price b^* can now be found by solving the optimization (3):

$$\begin{aligned}
 b^* &= \arg \max_{b>0} (V - b) \text{logistic} \left(w_0 + \sum_{i=1}^k w_i x_i + \beta \log b \right) \\
 &= \arg \max_{b>0} (V - b) \left(1 + e^{-w_0 - \sum_{i=1}^k w_i x_i - \beta \log b} \right)^{-1} \\
 &= \arg \max_{b>0} \frac{V - b}{1 + e^{-\alpha b - \beta}}, \tag{6}
 \end{aligned}$$

where $\alpha = w_0 + \sum_{i=1}^k w_i x_i$.

We show below that, for $b > 0$, there is a single optimum bid price b^* which can be bounded from above and below. These bounds make it possible to implement a fast bisection search.

THEOREM 1. For any $\beta > 0$,

$$f(b) = \frac{V - b}{1 + e^{-\alpha b - \beta}}$$

is maximized at some unique b^* such that

$$\frac{\beta}{\beta+1+e^{\alpha V\beta}}V \leq b^* < \frac{\beta}{\beta+1}V.$$

PROOF. Taking the derivative, we have

$$f'(b) = \frac{\beta V - (\beta+1)b - e^{\alpha b\beta+1}}{(1+e^{-\alpha b-\beta})^2 e^{\alpha b\beta+1}}.$$

Note that the denominator is always positive. Thus to find b^* that maximizes $f(b)$ it's sufficient to consider the sign of

$$h(b) \stackrel{\text{def}}{=} \beta V - (\beta+1)b - e^{\alpha b\beta+1}.$$

Since $h(b)$ is a decreasing function in b , for any $b \in (0, V]$, $h(b)$ can be bounded as

$$\beta V - (\beta+1)b - e^{\alpha V\beta} b \leq h(b) < \beta V - (\beta+1)b,$$

which implies that

$$h\left(\frac{\beta}{\beta+1+e^{\alpha V\beta}}V\right) \geq 0 \text{ and } h\left(\frac{\beta}{\beta+1}V\right) < 0.$$

Therefore there is a unique value $b^* \in \left[\frac{\beta}{\beta+1+e^{\alpha V\beta}}V, \frac{\beta}{\beta+1}V\right)$ such that $h(b^*) = 0$, and hence $f'(b^*) = 0$. In other words, $f(b)$ is maximized at $b = b^*$. \square

Theorem 1 allows us to implement a fast bisection search algorithm 4.1 for the optimal bid price. Starting with the minimum and

Algorithm 4.1 Bisection Algorithm Surplus Maximization

Require:

- 1: • Model weights: $w_0, w_1, \dots, w_k, \beta$;
- Feature values x_1, x_2, \dots, x_k ;
- V : expected value of the current ad opportunity
- $\epsilon > 0$: minimum valid interval length
- N : maximum number of search steps

Ensure: $\beta > 0, V > 0$

```

2:  $\alpha \leftarrow w_0 + \sum_{i=1}^k w_i x_i.$ 
3:  $b_{\min} \leftarrow \frac{\beta}{\beta+1+e^{\alpha V\beta}}V$ 
4:  $b_{\max} \leftarrow \frac{\beta}{\beta+1}V$ 
5: for  $i = 1, 2, \dots, N$  do
6:    $\text{fp}_{\min} \leftarrow \beta V - (\beta+1)b_{\min} - e^{\alpha b_{\min}^{\beta+1}}$ 
7:    $\text{fp}_{\max} \leftarrow \beta V - (\beta+1)b_{\max} - e^{\alpha b_{\max}^{\beta+1}}$ 
8:    $r \leftarrow -\text{fp}_{\min} / (\text{fp}_{\max} - \text{fp}_{\min})$ 
9:    $b \leftarrow (1-r)b_{\min} + r b_{\max}$ 
10:   $\text{fpb} \leftarrow \beta V - (\beta+1)b - e^{\alpha b^{\beta+1}}$ 
11:  if  $\text{fpb} < 0$  then
12:     $b_{\min} \leftarrow b$ 
13:  else
14:     $b_{\max} \leftarrow b$ 
15:  end if
16:  if  $b_{\max} - b_{\min} < \epsilon$  then
17:    break
18:  end if
19: end for
return  $b$ 

```

maximum bounds on the surplus optimum, $b_{\min} = \frac{\beta}{\beta+1+e^{\alpha V\beta}}V$ and $b_{\max} = \frac{\beta}{\beta+1}V$, per Theorem 1, we know that the lower bound for optimum has positive derivative, and the high bound has negative. Bisection can divide the range and find the zero point for the derivative in at most $\log_2[(b_{\max} - b_{\min})/\epsilon]$ steps; this logarithmic time is extremely desirable since the maximization search must run in real-time in the ad-server.

We found in practice that we could use the gradient information to speed up the search further. Rather than cutting the range in half each time ($r = 0.5$; step 8), after testing the gradient of the minimum and maximum bid points, we use our knowledge that the surplus function is convex and so derivatives shorten close to the optimum. We calculate the ratio between the surplus derivative at min and max bid locations, and then use that estimate for the relative distance to the optimum in bid space. Steps 8 and 9 of the pseudo-code show this modification to r . Empirically we observed that the bisection ends in less than 10 iterations to achieve a sufficient precision, which is controlled by ϵ .

5 IMPLEMENTATION

The features used for predicting win probability comprise 12 variables extracted from the HTTP of an incoming bid request, along with $\log(\text{bid price})$ and $\log(\text{bid price before shading})$. The HTTP attributes include the requesting page (e.g., yahoo.com), device type (e.g., desktop), hour of day; day of week, country, user segment, and other variables. All of the HTTP features are encoded to be binary variables.

For production we use one week of historical data for training so that weekly patterns are captured. The training data typically contain over a billion of bid requests with less than 100K encoded features. For the curve fit, we used the LogisticRegression method that is part of the PySpark `pyspark.ml.classification` library [2], which is distributed. The training time depends on the number of allocated machines. With less than 100 machines the training can finish within a few hours.

At run-time, the shading algorithm needs to respond to millions of requests per second peak load, within 100 milliseconds for all systems. In order to meet these speed constraints, bid shading has to minimize the number of computations that it performs. In terms of memory, by using a single global model, memory consumption is kept to less than 100K floating point numbers. In terms of time, shading optimization averages just less than 20 floating-point operations per request.

6 SHADING INSIGHTS

Here we describe a few features that we observed to be predictive in the win-rate model. The numerical features $\log(\text{bid price before shading})$ and $\log(\text{bid price})$ are both highly predictive¹ ($w = -0.39$ and 0.565 ; McFadden $R^2 = 0.24$ and 0.20 respectively [10, 32]). The high predictiveness of bid price before shading - and yet negative sign when included with bid price - is consistent with previous observations that bid shading tends to be deeper in auctions with higher valuations [4, 6, 17].

¹In the following, the regression coefficient is labeled w and PR is the observed positive rate of the binary variable

The top binary feature in terms of impact on win probability is `is_new_user` ($w = 0.831$; $PR=0.52$), which is associated with an increase in chance of winning the auction (since bid prices are lower). `hour_of_day=6am` (user local hour) ($w = -0.267$; $PR=0.01$) is associated with a drop in the probability of winning, likely due to the reduction in supply [22]. `country=US` ($w = -0.110$; $PR=0.84$) decreases the chance of winning; and the largest 768x1024 ads also are less likely to be won ($w = -0.267$; $PR=0.01$).

The predictability of time, user, and other features, for estimating auction clearing prices, suggest that shading should be effective, as noted in work [24] on the preconditions for shading described in Section 3.

7 COMPARISON TO BENCHMARKS

We ran several of the algorithms in Section 3 as benchmarks. These included: (1) Sell-Side Shading Service (S4) [1, 3, 14, 15, 25], (2) Non-linear Segment-Based (SEG) [20], Distribution estimator with Normal (NRML), Exponential (EXP) Distributions [35], Logistic Regression (LR) [19] and Unshaded (Uns). The win-rate based algorithm in this paper is labeled *WR* in the tables to follow.

The prior work benchmarks aren't ideal - the win distribution approaches [35] don't explicitly maximize surplus and so we expect them to not perform as well. The S4 algorithm seems to be geared towards maintaining win rate. Nevertheless, we have included them not only to compare to prior work but also to quantify the gain that surplus maximization approaches can deliver in practice.

Unlike the other benchmarks, the SEG algorithm does maximize surplus [20]. Under a favorable selection of segments, the Segment-Based algorithm might even be tuned to perform as well or better than the current method, despite the scaling problem with using more features. Our purpose in showing these benchmarks isn't to claim that this particular algorithm outperforms the others in all metrics, but rather to show that surplus maximizers have an advantage, to quantify the gain, and to note that WR, which is fully automated, uses all available features to estimate the surplus landscape, and has excellent memory and speed properties, performs comparable to other reported approaches.

The experiments below (except ones with S4) were run on auctions for which the minimum bid prices to win were known. Using this data it was possible to calculate surplus performance as a percentage of the total optimal surplus:

$$\% \text{ opt surplus} \stackrel{\text{def}}{=} \frac{\sum_i (V_i - b_i) \mathbf{I}(b_i > \hat{b}_i)}{\sum_i (V_i - b_i)},$$

i.e., the surplus achieved by a particular algorithm out of total available surplus by bidding optimally. Spend and impression performances can be defined similarly.

The algorithms were tested on one day of auction data. For fair comparison training is done on data from the previous day, since not all algorithms are designed to be trained on multiple days of data. All the bid requests are scored by each algorithm, so all algorithms operate on the same set of records. The results are shown in Table 1.

The distribution estimator methods (NRML, EXP) estimate the minimum bid to win and so are not expected to do well in maximizing surplus. As a group they were about 7% below WR. The Nonlinear Segment method generated the second highest surplus

Metric	WR	SEG	NRML	LR	EXP	Uns
%opt surplus	50.6%	49.0%	48.0%	47.3%	46.0%	0%
%opt spend	41.7%	56.0%	42.7%	39.8%	31.1%	176%
%opt imps	56.6%	49.1%	53.1%	50.3%	42.6%	100%
avg shading factor	0.6	0.55	0.62	0.61	0.42	1.00

Table 1: Benchmark Algorithms




Figure 2: Shading factor distributions for three algorithms. S4 has more shallow shading factors.

Metric	WR	LR	SEG	Uns
% opt surplus	46.7%	44.8%	38.2%	0.0%
% opt spend	79.1%	72.6%	89.9%	410%
% opt imps	60.3%	51.4%	56.0%	100%
avg shading factor	0.55	0.53	0.59	1.00

Table 2: Production Results

besides WR. This makes sense given that it is a legitimate surplus maximizer. WR generates the highest surplus (50.6%). In sum, the surplus maximizers produced the most surplus, which was expected.

We also compared an S4 algorithm from an anonymous SSP. We had to separate this analysis due to a service issue. When using the S4 for real-time bidding, the service disabled the minimum bid to win functionality. As a result, we were unable to do an optimality analysis.

Overall, the S4 delivered about 15% more impressions than WR - as noted SSPs have an incentive to try to monetize more traffic. However it delivered about 4.3% lower surplus. The bidding distribution from the S4 is shown in Figure 2. Whereas the SSP's shading distribution is right-skewed, with most shading at 90% and above, the WR distribution - which generates more surplus - is left-skewed, with most shades below 72%. It seems likely that the S4 is geared towards generating high sales, but not necessarily high advertiser surplus.

8 PRODUCTION RESULTS

After rolling out the WR algorithm, we were able to monitor its online performance by maintaining a percentage of traffic that was randomly allocated to each algorithm. The analysis shown in Table 2 spans about two months, during which time all algorithms were automatically updated at daily basis. Overall WR captured 46.7% of the maximum possible surplus, whereas Non-linear captured 38%. Bid prices on WR were about 45% lower than their unshaded prices.

A/B Testing	Spend	Surplus	eCPM	eCPC	eCPA
WR v.s. LR	+1.3%	+1.4%	-7.4%	-4.5%	-2.7%
WR v.s. SEG	+1.2%	+2.5%	-5.4%	-5.5%	-3.9%

Table 3: Improvements on Business Metrics

Note that in a real-time bidding system, campaigns usually have finite budgets, and budget controller is a necessary component in such a system. The production performance of a bid shading algorithm relies on how well it works together with the budget controller. In a simplified view, a reasonable controller is expected to spend all the daily budget, and hence the budget *saved* by a bid shading algorithm, that is, the surplus, would be spent again to buy more impressions, thus leads to lower eCPM, eCPC, and eCPA [31]. Indeed, as shown in Table 3, with similar spend WR achieved significant improvements on these business metrics.

9 CONCLUSION

There is evidence that First Price Auctions have created problems for advertisers. Average traffic prices are higher, with estimates ranging between 5% and 50% [3, 18, 21, 25]. [21] also reported that after their SSP switched to First Price, 10% of advertisers actually discontinued bidding. Our experiments confirm these findings; without a shading solution, CPM would approximately double.

DSPs are required to compute the private value of impressions based on advertiser parameters, and they also execute a large number of trades, and so can build up an ability to predict auction prices. This makes it possible to implement rational shading similar to other industries [16, 17, 23]. Advertiser bids follow the value of traffic, and this follows daily, hourly, and site patterns. As a result, auction prices will always have structure that can be used by some advertisers with other advertisers have less flexibility.

The surplus maximization approach of this paper delivered about 7% higher surplus than naive methods just designed to submit the probable clearing price. Furthermore, when integrated with budget controllers, it significantly reduced eCPM, eCPC and eCPA by 3% – 7%. Publicly available data shows medium sized DSPs managing between 260 to 1 billion US dollars in advertiser spend [33]. The Shading gains reported in this paper therefore represent 18 to 100 million US dollars in additional yield that is provided to advertisers. Shading has an enormous impact on advertiser profitability. Now that the online ad industry has increasingly shifted to First Price Auctions, it seems likely that the new advertising technology arms race will be in the domain of bid shading.

REFERENCES

- [1] AdExchanger. 2017. Rubicon Joins First-Price Auction Club; Diageo Is Latest Brand To Demand More Transparency. <https://adexchanger.com/ad-exchange-news/tuesday-12122017/>
- [2] Apache. PySpark. Source code for pyspark.ml.classification. https://spark.apache.org/docs/2.1.1/api/python/_modules/pyspark/ml/classification.html
- [3] AppNexus. 2018. Demystifying Auction Dynamics for Digital Buyers and Sellers. <https://www.appnexus.com/sites/default/files/whitepapers/49344-CM-Auction-Type-Whitepaper-V9.pdf>
- [4] Pierpaolo Battigalli and Marciano Siniscalchi. 2003. Rationalizable bidding in first-price auctions. *Games and Economic Behavior* 45, 1 (2003), 38–72.
- [5] R.V. Clapp Capen, E.C. and W.M. Campbell. 1971. Auctions and Bidding. *Journal of Petroleum Technology* 23, 6 (1971), 641–643.
- [6] Bhaskar Chakravorti, William W Sharkey, Yossef Spiegel, and Simon Wilkie. 1995. Auctioning the airwaves: the contest for broadband PCS spectrum. *Journal of Economics & Management Strategy* 4, 2 (1995), 345–373.
- [7] V Chari and Robert Weber. 1992. How the US Treasury should auction its debt. *Federal Reserve Bank of Minneapolis Quarterly Review* 16, 4 (1992), 1–12. <http://kylewoodward.com/blog-data/pdfs/references/chari+weber-quarterly-review-1992A.pdf>
- [8] John M Crespi and Richard J Sexton. 2005. A Multinomial logit framework to estimate bid shading in procurement auctions: Application to cattle sales in the Texas Panhandle. *Review of industrial organization* 27, 3 (2005), 253–278.
- [9] Julian Faraway. 2006. *Extending the Linear Model with R: Generalized Linear, Mixed Effects and Nonparametric Regression Models*. Chapman Hall/CRC Press, Boca Raton.
- [10] Jeremy Freese and J. Scott Long. 2006. *Regression Models for Categorical Dependent Variables Using Stata*. Stata Press, College Station, Texas.
- [11] Getint. 2017. RTB Auctions: Fair Play? <https://blog.getint.com/rtb-auctions-fair-play-3b372d505089>
- [12] Google. 2018. Form 10K for Alphabet Inc. https://abc.xyz/investor/static/pdf/20180204_alphabet_10K.pdf?cache=11336e3
- [13] Google. 2019. Real Time Bidding Protocol Protocol Buffer v.167. <https://developers.google.com/authorized-buyers/rtb/downloads/realtime-bidding-PROTO>
- [14] Google. 2019. Real Time Bidding Protocol, Release Notes. <https://developers.google.com/authorized-buyers/rtb/relnotes#updates-2019-03-13>
- [15] Google. 2020. Real-Time Bidding Protocol Buffer v.173. <https://developers.google.com/authorized-buyers/rtb/downloads/realtime-bidding-PROTO>
- [16] Kenneth Hendricks and Robert H Porter. 1989. Collusion in auctions. *Annales d’Economie et de Statistique*, 15/16 (1989), 217–230.
- [17] Ali Hortaçsu, Jakub Kastl, and Allen Zhang. 2018. Bid shading and bidder surplus in the us treasury auction system. *American Economic Review* 108, 1 (2018), 147–69.
- [18] B. Hovaness. 2018. Sold for more than you should have paid. <https://www.hearts-science.com/sold-for-more-than-you-should-have-paid/>
- [19] Verizon Media internal report. 2019. Predicting Optimal Bid Shading Factor Using Logistic Regression.
- [20] Niklas Karlsson and Qian Sang. 2020. Adaptive bid shading optimization of first price ad inventory.
- [21] Brendan Kitts. 2019. Bidder Behavior after Shifting from Second to First Price Auctions in Online Advertising. http://www.appliedaisystems.com/papers/FPA_Effects33.pdf
- [22] Brendan Kitts and Yongbo Zeng. 2019. Lookahead Algorithms for Online Bidding.
- [23] Yvan Lengwiler and Elmar Wolfstetter. 2010. Auctions and corruption: An analysis of bid rigging by a corrupt auctioneer. *Journal of Economic Dynamics and control* 34, 10 (2010), 1872–1892.
- [24] Rachel AJ Pownall and Leonard Wolk. 2013. Bidding behavior and experience in internet auctions. *European Economic Review* 61 (2013), 14–27.
- [25] Rubicon. 2018. Bridging the Gap to First-Price Auctions: A Buyer’s Guide. http://go.rubiconproject.com/rs/958-XBX-033/images/Buyers_Guide_to_First_Price_Rubicon_Project.pdf
- [26] Rubicon. 2018. Principles for a Better Programmatic Marketplace, Open Letter from Rubicon, SpotX, OpenX, Pubmatic, Telaris, and Sovr. <https://rubiconproject.com/insights/thought-leadership/principles-better-programmatic-marketplace-open-letter-advertisers-publishers/>
- [27] Cosma Shalizi. 2020. Logistic Regression Lecture 12. <https://www.stat.cmu.edu/~shalizi/uADA/12/lectures/ch12.pdf>
- [28] R. Shields. 2019. Google Ad Manager to Offer First-Price Auctions, Simplifying Programmatic Buying. <https://www.adweek.com/programmatic/google-ad-manager-to-offer-first-price-auctions-simplifying-programmatic-buying/>
- [29] S. Sluis. 2017. Explainer: More On The Widespread Fee Practice Behind The Guardian’s Lawsuit Vs. Rubicon Project. <https://adexchanger.com/ad-exchange-news/explainer-widespread-fee-practice-behind-guardians-lawsuit-vs-rubicon-project/>
- [30] S. Sluis. 2019. Google Switches To First-Price Auction. <https://www.adweek.com/programmatic/google-ad-manager-to-offer-first-price-auctions-simplifying-programmatic-buying/>
- [31] Automatad Team. 2018. How to Calculate CPM, CPC, CPA, CR, eCPM, eCPC, eCPA, and ROI. <https://headerbidding.co/calculate-cpm-cpc-cpa-ecpm-ecpc-ecpa-roi/>
- [32] UCLA. 2011. What are Pseudo R Squares? <https://stats.idre.ucla.edu/other/mult-pkg/faq/general/faq-what-are-pseudo-r-squareds/>
- [33] Karsten Weide. 2017. Worldwide Digital Advertising Software Market Shares 2017. https://www.criteo.com/wp-content/uploads/2018/09/US44240218e_Criteo.pdf
- [34] Wush Wu, Mi-Yen Yeh, and Ming-Syan Chen. 2018. Deep censored learning of the winning price in the real time bidding. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. KDD ’18, London, United Kingdom, 2526–2535.
- [35] Wush Chi-Hsuan Wu, Mi-Yen Yeh, and Ming-Syan Chen. 2015. Predicting winning price in real time bidding with censored data. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD ’15, Sydney, 1305–1314.
- [36] Christine Zulehner. 2009. Bidding behavior in sequential cattle auctions. *International Journal of Industrial Organization* 27, 1 (2009), 33–42.