IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received March 28, 2019, accepted May 23, 2019, date of publication May 29, 2019, date of current version June 7, 2019.

Digital Object Identifier 10.1109/ACCESS.2019.2919450

A Revenue-Maximizing Bidding Strategy for Demand-Side Platforms

TENGYUN WANG[1], HAIZHI YANG[1], HAN YU[2], WENJUN ZHOU[3], YANG LIU[4], AND HENGJIE SONG[1]

[1]School of Software Engineering, South China University of Technology, Guangzhou 510006, China
[2]School of Computer Science and Engineering, Nanyang Technological University, Singapore 639798
[3]Department of Business Analytics and Statistics, The University of Tennessee, Knoxville, TN 37996-0532, USA
[4]WeBank, Shenzhen 518055, China

Corresponding author: Hengjie Song (sehjsong@scut.edu.cn)

This work was supported in part by the National Natural Science Foundation of under Grant 71671069, in part by the Natural Science Foundation of Guangdong Province, China, under Grant 2016A030313479, in part by the Mobile Research Foundation of Ministry of Education of China under Grant MCM20160204, and in part by the Nanyang Technological University, Nanyang Assistant Professorship (NAP).

ABSTRACT In real-time bidding (RTB) systems for display advertising, a demand-side platform (DSP) serves as an agent for advertisers and plays an important role in competing for online advertising spaces by placing proper bidding prices. A critical function of the DSP is formulating proper bidding strategies to maximize key performance indicators, such as the number of clicks and conversions. However, many small and medium-sized advertisers' main goal is to maximize revenue with an acceptable return on investment (ROI), rather than simply increase clicks or conversions. Most existing approaches are inapplicable of satisfying the revenue-maximizing goals directly. To solve this problem, we first theoretically analyze the relationships among the conversion rate, ROI, and ad cost, and how they affect revenue. By doing so, we reveal that it is a challenge to increase revenue by relying solely on improving ROI without considering the impact of the ad cost. Based on this insight, the maximal revenue (MR) bidding strategy is proposed to maximize revenue by maximizing the ad cost with a desirable ROI constraint. Unlike previous studies, the proposed MR first distinguishes bid prices from ad costs explicitly, which makes it more applicable to the real second-price auction (GSP) auction mechanism in RTB systems. Then, the winning function is empirically defined in the form of tanh that provides a promising solution for estimating ad costs by jointly considering ad costs with the winning function. The experimental results based on two real-world public datasets demonstrate that the MR significantly outperforms five state-of-the-art models in terms of both revenue and ROI.

INDEX TERMS Bid landscape forecasting, bidding strategy optimization, demand-side platform, real-time bidding.

I. INTRODUCTION

Real-time bidding (RTB) serves as an important mechanism in computational advertising. In a typical RTB system, ad publishers (e.g., search engines, or websites with display ads) sell each ad placement opportunity via a real-time auction, in which advertisers need to determine whether to bid for it and how much money to put forward after evaluating the value of the ad impression [1]. In this process, the main goal of the advertisers is to reach their target audience, who will respond to the ad in a desirable manner, such as clicking on it or making a purchase.

Since the process to achieve the above goal is very complex, advertisers are often assisted by specialized intelligent software, called demand-side platforms (DSPs). In RTB systems, DSPs work on behalf of advertisers by managing ad campaigns and optimizing bidding activities. This function requires that the DSPs process the ad information, ad placements, audiences, and campaign constraints in a real-time manner [2], [3]. To facilitate understanding of the roles and functions of DSPs, Fig. 1 shows the work-flow of a typical DSP in a RTB system. Initially (Step-0), an ad request is created once a user visits a web page containing

The associate editor coordinating the review of this manuscript and approving it for publication was Irene Amerini.

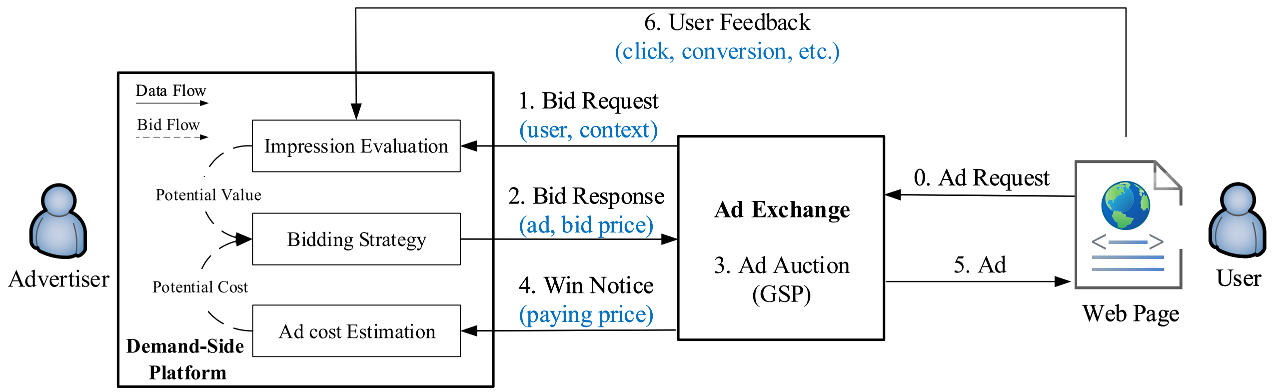

FIGURE 1. An illustration of a typical DSP's work-flow in a RTB system.

an ad placement opportunity to be sold. Then (Step-1), a bid request corresponding to the ad request, along with information about the user and web page, is triggered by the ad publisher and then is broadcasted to various DSPs via an ad exchange [4]. Based on estimating the ad impression value and possible ad cost, (Step-2) a DSP computes a bid price as a bid response to participate in the auction process for this bid request. Note that the solution to computing a bid price in this step is referred to as *bidding strategy*. Then (Step-3), the ad exchange determines the winner in terms of generalized second-price auction (GSP) mechanism [5]. In other words, the bidder with the highest bid price wins the ad auction and actually pays the second-highest bid price as ad cost for the ad impression. (Step-4) The winner is notified with the paying price and (Step-5) the winner displays ad to the user. Finally (Step-6), the user provides feedback, such as clicking on the ad or making a purchase (a.k.a. conversion), which is recorded by the DSP.

The most important challenge in a DSP is to design the bidding strategy (i.e., Step-2 in the DSP work flow) to win the desirable bid requests with appropriate bid prices. Recently, many efficient bidding strategies have been studied to maximize the number of clicks or conversions [2], [6]–[9], which are normally measured by the predicted click-through rate (CTR) or conversion rate (CVR). By doing so, these existing bidding strategies tend to bid for requests with higher CTR or CVR. *However, from the business perspective, the advertisers' ultimate goal is to maximize ad revenue rather than improve the CTR or CVR.* In this respect, most existing bidding strategies focusing on maximizing clicks or conversions are inapplicable to satisfying advertisers' requirement on ad revenue directly. On the one hand, it is essentially due to the fact that few studies in the literature fully realize the extremely vital significance of ad revenue for advertisers, which to great extent is the main motivation for advertisers to continuously invest on bidding ads. On the other hand, many existing bidding strategies overlook the potentially important impact on ad revenue by return-on-investment ratio (ROI) and ad cost, respectively.

In general, ad revenue can be described as the product of ROI and ad cost, which are closely related to the CVR. As a result, an intuitive approach to increase revenue is to improve ROI by targeting ad audiences with higher conversion rates (i.e., higher CVR). However in practice, there is a limited number of ad audiences that are more likely to purchase. This restricts the ad cost spent on the ad audiences with higher CVR. In this case, revenue cannot be increased by solely targeting ad audience with higher CVRs without considering the dynamic interactions between ROI and ad cost.

Inspired by above background, we first formulate and analyze the potential impacts on revenue that are imposed by ROI and ad cost respectively. Furthermore, we theoretically reveal that the predicted ROI associated with a bid request is not only proportional to the CVR, but also inversely proportional to the estimated ad cost. Based on this insight, we then propose maximal revenue (MR) bidding strategy in which ad cost is maximized on the premise of predefined acceptable ROI so as to satisfy the advertisers' requirements on revenue more directly and efficiently. Unlike previous studies on optimal bidding strategies in which the underlying idea is to increase CTR/CVR with the ad cost as a constraint [2], [6]–[9], the main contributions of our paper are two-fold:

1) We theoretically analyze the relationships among CVR, ROI and ad cost and how they influence revenue. By doing so, we reveal that, for most existing bidding strategies, it is a challenging task to increase revenue by relying solely on improving ROI without considering the impact of ad cost, which has never been formally discussed in existing bidding strategies [2], [9]–[11]. Based on this insight, the MR bidding strategy is proposed to satisfy the requirements on revenue by directly translating the problem of revenue maximization into maximizing ad cost with an ROI constraint. It significantly differs from previous studies in the literature that attempt to improve CTR/CVR with an ad cost constraint. To validate the effectiveness, we conducted extensive experiments on two real-world large-scale public datasets. The experimental results and analysis

demonstrate that MR bidding strategy significantly outperforms state-of-the-art methods in terms of revenue and ROI.

2) In the process of designing a bidding strategy, both the estimation of ad cost and an appropriate winning function play important roles in determining the probability to win a bid request. However, in most existing bidding strategies [2], [10], [12], the ad cost is oversimplified as the bid price provided by a bidding strategy. *In real application scenarios, this oversimplification is invalid since it is inconsistent with GSP mechanism in which the winner only pays the second highest bid price (i.e., ad cost) for the ad impression.* To deal with this issue, our proposed MR bidding strategy not only distinguishes the bid price from ad cost, but also provides a novel solution to estimating ad cost under the GSP auction mechanism by constructing the correlation between ad cost and winning function. Furthermore, in terms of Wasserstein distance and KL-divergence, the winning function is empirically described in the form of tanh so as to predict the winning probability more accurately against existing approaches in [2], [13].

The rest of the paper is organized as follows. We first provide an overview of related work in Section II. For clarity, Section III lists the preliminaries and notations. In Section IV, we define the problem of maximal revenue bidding strategy and discuss the motivations of our study. Furthermore, we describe the detailed mathematical derivations and optimization involved in the proposed MR bidding strategy, respectively. Section V discusses the experimental results and comparison analysis. Finally, Section VI concludes the paper and discusses promising future research directions.

II. RELATED WORK

This section discusses the related literatures involved in bidding strategies, which focus on the three key components, including impression evaluation, ad cost estimation, and bidding strategy optimization [11].

A. IMPRESSION EVALUATION

Impression evaluation focuses on predicting how likely the ad audiences will make desirable responses to ad impressions. The performance is commonly measured by CTR [14], [15] and CVR [16]. Since impression evaluation measures the ad audiences' interest in ad impressions, it plays a key role in real-time display advertising and is commonly used to derive a reasonable allocation of ad budget [3]. Typically, CTR/CVR predictions are modeled as probability estimation tasks in which logistic regression and relevant extensions are widely used in the literature since logistic regression is highly scalable and also offers well calibrated probability outputs [16]–[18]. To enable efficient training on large-scale data sets, an online learning algorithm for logistic regression was studied in which the parameter was updated immediately when observing new data instance [15]. However, as a linear model, logistic regression suffers from the effect of feature interactions. To solve this problem, factorization machine-based models were proposed to make better predictions by capturing the impact of feature interactions [19], [20]. Recently, a deep and cross neural network was proposed to predict number of clicks in which the cross network learned low-degree feature interactions efficiently and the deep network is trained to capture the high-degree feature interactions [21].

B. AD COST ESTIMATION

Ad cost estimation is to predict the actual payment for winning an impression in DSP systems, which can be utilized for more reasonable allocation of budget in bidding strategies. In many existing bidding strategies [2], [10], [12], the ad cost is oversimplified as the bid price. However in real application scenarios, this oversimplification is invalid since it is inconsistent with GSP mechanism in which the winner only pays the second highest bid price (i.e.,ad cost) for the ad impression, instead of bid price. Aiming at the problem, [11] modeled ad cost as the mathematical expectation of market price when winning the ad auction. Rather than fitting a popular log-normal market price distribution, [13] and [11] obtained the distribution of market price by taking the derivative of a pre-specified winning function. In these approaches, several hypothetical winning function forms had been studied to predict the winning probability for a given bid price. These include logistic regression form [22], and long tail form [2], [11]. In addition, since the market price of each ad auction can be observed only if the advertiser wins the corresponding bid, the observed market prices for a bidding strategy are censored. To provide appropriate ad cost estimations under the censored context, [23] utilized a censored linear regression model to jointly fit the likelihood of observed prices in winning cases and the likelihood of censored ones in losing cases.

C. BIDDING STRATEGY OPTIMIZATION

Recently, many studies have been devoted to designing and optimizing bidding strategies in RTB systems so as to meet the advertisers' diverse requirements. In terms of different optimization objectives, most existing bidding strategies can be categorized into two types: *KPI-oriented bidding* and *budget-pacing-oriented bidding*.

KPI-oriented bidding strategies attempt to deliver the appropriate ads to the ad audience with reasonable bid price so as to satisfy advertisers' requirements on certain KPIs, such as ad audience's CTR or CVR [2], [6], [9], [10]. From the perspective of the second price auction underlying GSP mechanism, [24] proved that the truthful bidding (i.e., bid price is determined as the click/conversion value times CTR/CVR) is the optimal strategy in RTB system. Inspired by truthful bidding, a linear bidding strategy was further designed and widely used in industry due to the simplicity and scalability in which the bid prices is described as the predicted CTR/CVR weighted with a pre-specified constant parameter [6]. However, this approach overlooked the actual effects on ad audiences that are imposed by the displayed ad, which should be measured by the conversion probability lift after a certain ad being exposed to the ad audience [25]. To deal with this issue, a lift-based bidding strategy was proposed to determine the bid prices by assuming bid prices to be proportional to the estimated CVR lifts [25]. Considering practical constraints such as the limited auction volume and budget, [2] proposed a general bid optimization framework to maximize the number of clicks or conversions with limited budget and auction volume constraints. In this approach, the parameters underlying bidding function were tuned by integrating ad-reward function with ad cost function.

To further maximize the number of clicks, [9] formulated the bidding strategy as a reinforcement learning problem in which Markov Decision Process was utilized to determine the bid prices according to the remaining bid requests and budget. However in practice, this approach may result in unfavorable KPI (e.g., relatively high cost per click and low winning rate, etc.) due to the instability of RTBs [8]. To deal with this issue, [8] explored a proportional-integral-derivative control function to approximate the bid price in which KPI errors against the preferred KPI values and corresponding adjustment of bid price served as the input signal and output signal, respectively. Furthermore, [26] highlighted the data censorship problem in bidding process. It is mainly due to the fact that advertisers only know the statistics (market price, user clicks etc.) when winning ad impressions. To address this problem, [26] designed a non-parametric survival model to describe winning probability so as to remedy the biased market price learned from censored data.

More recently, a joint learning framework [11] was proposed to maximize the overall profit of an ad campaign. To this end, the authors attempt to jointly optimize user response prediction, bid landscape forecasting, and bidding function. In this process, the bidding strategies favor bidding requests with relatively high user response rates but incur lower ad costs. In addition, the profit maximization bidding strategies [10] are proposed to bid requests with relatively higher profits (i.e., revenues minus ad costs). Nevertheless, as analyzed in Section IV-A, these two profit maximization models may overlook the potential impact on profit by ad cost and ROI.

Budget-pacing-oriented bidding approaches aim to satisfy advertisers' requirements on brand promotion by spending the ad budget smoothly to reach as wide an ad audience as possible [7]. In particular, [27] proposed a budget pacing model that attempts to evenly spread out the ad opportunity by controlling bid probability according to the information about traffic patterns of eligible impressions. To more efficiently balance the optimization of ad clicks/conversions and smooth budget spending, the model proposed by [28] tends to bid the ad impressions with high quality based on the prior click/conversion distributions while adjusting the bid price to distribute ad budget optimally. Similarly, [7] optimizes budget pacing control and campaign performance simultaneously. In particular, [7] first assumed that the requests in the same group share the same group pacing rate. Based on this assumption, it divided bid requests into different groups by designing an offline response prediction model which groups similarly responses to ad requests together. Furthermore, a novel control-based method was designed to dynamically adjust the pacing of budget spending for each group so as to disburse the budget smoothly.

To the best of our knowledge, no study in the literature highlights the trade-off between ROI and ad cost and analyzes their potential impact on maximizing revenue. Inspired by this background, we propose the MR bidding strategy to maximize revenue more efficiently by maximizing ad cost with a pre-specified acceptable ROI constraint.

III. PRELIMINARIES AND NOTATIONS

Typically, a bid request in an RTB scenario is described as a random feature vector $\mathcal{X}$ that consists of various information about the ad display opportunity, such as the location of the audience, timestamp of this visit, and contextual information about the web page. Following previous work [2], [16], we assume the bid requests are generated i.i.d. from a prior distribution $p_{\mathcal{X}}(\boldsymbol{x})$. Correspondingly, the advertising effect is commonly measured by the audience's historical feedbacks, such as ad audience's CTR/CVR [2], [6], desired action rate lift [25] or memory retention of displayed ads [29]. Without loss of generality, in this paper, we adopt ad audience's CVR to evaluate the advertising effect since it is more relevant to revenue.

TABLE 1. A summary of our symbols and descriptions.

Symbol	Description
$\boldsymbol{x}$	The feature vector representing a bid request.
$p_{\mathcal{X}}(\boldsymbol{x})$	The prior *p.d.f.* of bid requests.
$\theta(\boldsymbol{x})$	The predicted CVR of $\boldsymbol{x}$.
$p_{\Theta}(\theta)$	The prior *p.d.f.* of CVRs.
$b(\theta(\boldsymbol{x});\omega)$	The bidding strategy represented by a bid function with parameter ω. We use b to denote a specific bid price.
$\mathcal{V}$	The commercial value of each conversion (i.e., the unit price of the product being advertised).
$ROI(\boldsymbol{x})$	The estimated ROI of $\boldsymbol{x}$.
$W_{\mathcal{Z}}(b;\beta)$	The winning function with parameter β, which indicates the probability of winning the ad auction with bid price b and random market price $\mathcal{Z}$.
$p_{\mathcal{Z}}(z)$	The *p.d.f.* of the market price $\mathcal{Z}$.
$\mathbb{E}[Cost(b)]$	The expected ad cost corresponding to the bid price b.

For clarity, all the symbols used throughout this paper are listed in Table 1. We define function $\theta(\boldsymbol{x})$ to perform CVR prediction, which maps a bid request $\boldsymbol{x}$ to the probability of accomplishing a conversion. Similar to $\boldsymbol{x} \sim p_{\mathcal{X}}(\boldsymbol{x})$, we assume θ be subject to a prior distribution $p_{\Theta}(\theta)$.

The constant $\mathcal{V}$ denotes the commercial value of each conversion (i.e., the unit price of the product being advertised). Once receiving a bid request $\boldsymbol{x}$, the main task of bidding strategy $b(\theta(\boldsymbol{x}); \omega)$ with parameter ω is to calculate a bid price b for this ad auction, which essentially maps a bid request to a bid price. In this process, many factors have potential impacts on calculating the bid price, such as the impression value (i.e., $\mathcal{V} * \theta(\boldsymbol{x})$) [2], [6] and the ad cost of the ad impression being auctioned [10], [11], [30]. To further analyze the potential impact of these factors on revenue, $ROI(\boldsymbol{x})$ indicates the estimated ROI for a given bid request. Given the bid price b and the market price $\mathcal{Z}$, $W_{\mathcal{Z}}(b; \beta)$ indicates the probability of winning the auction with parameter β where the market price $\mathcal{Z}$ is regarded as a random variable $\mathcal{Z} \sim p_{\mathcal{Z}}(z)$; and the corresponding ad cost is estimated by $\mathbb{E}[Cost(b)]$.

IV. MAXIMAL REVENUE BIDDING

In this section, we first formulate revenue, ROI and ad cost. By revealing that the predicted ROI is not only directly proportional to CVR, but also inversely proportional to the estimated ad cost through theoretical analysis, we prove that it is challenging to improve revenue by solely targeting ad audience with higher CVRs without considering the dynamics between ROI and ad cost. Based on this insight, the MR bidding strategy is proposed to address the problem of maximizing ad cost given a pre-defined acceptable ROI.

A. PROBLEM FORMULATION AND MOTIVATION

Given N eligible bid requests during the campaign's lifetime, the objective to maximize the overall revenue for the campaign can be described as:

$$\text{Overall\_Revenue} = \text{Overall\_ROI} * \text{Overall\_Cost}. \tag{1}$$

Equation (1) indicates that the overall revenue is affected by both overall ROI and overall ad cost simultaneously.

For a given bid request $\boldsymbol{x}$ in the campaign, the corresponding ad cost is described as:

$$Cost(\boldsymbol{x}) = \mathbb{E}[Cost(b(\theta(\boldsymbol{x}); \omega))]. \tag{2}$$

In this case, the estimated $ROI(\boldsymbol{x})$ is expressed as:

$$ROI(\boldsymbol{x}) = \frac{W_{\mathcal{Z}}(b(\theta(\boldsymbol{x}); \omega); \beta)\theta(\boldsymbol{x})\mathcal{V}}{Cost(\boldsymbol{x})}. \tag{3}$$

Note that since the occurrence of conversion is subject to a Bernoulli distribution [31] with a success probability of $\theta(\boldsymbol{x})$, the numerator in (3) is essentially the mathematical expectation of revenue associated with the given bid request.

From (2), we note that the ad cost is influenced by the corresponding bid request $\boldsymbol{x}$ through $\theta(\boldsymbol{x})$, i.e., $\boldsymbol{x} \to \theta(\boldsymbol{x}) \to Cost(\boldsymbol{x})$. Therefore, (2) can be rewritten as

$$Cost(\theta) = \mathbb{E}[Cost(b(\theta; \omega))]. \tag{4}$$

Similarly, (3) can be rewritten as

$$ROI(\theta) = \frac{W_{\mathcal{Z}}(b(\theta; \omega); \beta)\theta\mathcal{V}}{\mathbb{E}[Cost(b(\theta; \omega))]}. \tag{5}$$

From (4) and (5), we can see that both ad cost and ROI are closely related to CVR. In this case, it is necessary to analyze the correlation between ROI and CVR as well as ad cost and CVR.

Proposition 1: Given a fixed bid price, ROI is monotonically increasing w.r.t. CVR. It indicates that targeting bid requests with higher CVR will result in higher ROI.

Proof: Giving a fixed bid price that bid price does not vary with CVR, we obtain $\frac{\partial b(\theta;\omega)}{\partial\theta} = 0$, $\frac{\partial W_{\mathcal{Z}}(b(\theta;\omega);\beta)}{\partial\theta} = 0$ and $\frac{\partial \mathbb{E}[Cost(b(\theta;\omega))]}{\partial\theta} = 0$. According to (5), by taking the first derivative of $ROI(\theta)$ w.r.t. θ, we obtain:

$$\frac{\partial ROI(\theta)}{\partial\theta} = \mathcal{V}W_{\mathcal{Z}}(b(\theta; \omega); \beta)\mathbb{E}^{-1}[Cost(b(\theta; \omega))]. \tag{6}$$

According to the definition of each item involved in (6), we conclude that $\frac{\partial ROI(\theta)}{\partial\theta} \geq 0$. It indicates that the ROI of a bid request is monotonically increasing with respect to its θ. Therefore, targeting bid requests with higher θ will improve the ROI associated with every bid request for a campaign, and the overall ROI of this campaign will be improved. □

Proposition 2: Targeting bid requests with higher predicted CVR will decrease the overall ad cost.

Proof: Let ϑ be the acceptable minimum CVR on the bid requests. The truncated bidding strategy is expressed as

$$b^{truncated}(\theta; \omega) = \begin{cases} b(\theta; \omega), & \text{if } \theta \geq \vartheta, \\ 0, & \text{if } \theta < \vartheta, \end{cases} \tag{7}$$

where $b(\theta; \omega)$ represents the bid price to compete for bid requests with CVR $\theta \geq \vartheta$; otherwise for the case $\theta < \vartheta$, the bidding strategy will not participate in bidding. By substituting the truncated bidding strategy described by (7) into (4), the overall ad cost of the campaign can be expressed as:

$$\begin{aligned} \text{Overall\_Cost}(\vartheta) &= N \int_{\theta} \mathbb{I}[\theta \geq \vartheta]\mathbb{E}[Cost(b(\theta; \omega))]p_{\Theta}(\theta)d\theta \\ &= N \int_{\vartheta}^{1} \mathbb{E}[Cost(b(\theta; \omega))]p_{\Theta}(\theta)d\theta, \end{aligned} \tag{8}$$

where $\mathbb{I}[\theta \geq \vartheta] = 1$ if $\theta \geq \vartheta$; otherwise, $\mathbb{I}[\theta \geq \vartheta] = 0$. According to (8), we calculate the first partial derivative w.r.t. ϑ:

$$\frac{\partial \text{Overall\_Cost}(\vartheta)}{\partial\vartheta} = -N\mathbb{E}[Cost(b(\theta; \omega))]p_{\Theta}(\theta). \tag{9}$$

Taking into account the definitions of each term involved in (9), it is obvious that $\frac{\partial \text{Overall\_Cost}(\vartheta)}{\partial\vartheta} \leq 0$. Therefore, the overall ad cost is monotonically decreasing with respect to ϑ, which indicates that targeting bid requests with higher CVR will reduce the overall ad cost. □

Proposition 1 shows that, on the premise of a fixed bid price $b(\theta; \omega)$, targeting bid requests with higher CVR will lead to higher overall ROI for the ad campaign. In this case, according to (1), the overall revenue of a campaign will increase if the overall ad cost remains unchanged. Essentially, it is the main reason to inspire the studies in KPI-oriented bidding strategies [7], [28]. However, Proposition 2 shows that the higher CVR (i.e., higher ROI) results in decreasing overall ad cost. It is mainly due to the fact that the number of converted bid requests (i.e., with higher CVR or ROI) is always limited in practice, which results in severe competitions for these bid requests. Such intense competitions always restrict the overall ad cost for the bid requests with high CVR or ROI.

Therefore, the above analysis reveals that the overall ROI and overall ad cost move in opposite trends when the bidding strategy targets the bid requests with higher CVRs. In this case, considering their potential impact on overall revenue as described in (1), it is necessary to further analyze the dynamics between overall revenue and CVR under Proposition 3.

Proposition 3: Targeting the bid requests with higher predicted CVR cannot guarantee the maximization of revenue.

Proof: By substituting the truncated bidding strategy described by (7) into (5), the overall ROI of the campaign can be described as:

$$\text{Overall\_ROI}(\vartheta) = \frac{N\int_{\vartheta}^{1} W_{\mathcal{Z}}(b(\theta;\omega);\beta)\theta\mathcal{V}p_{\Theta}(\theta)d\theta}{N\int_{\vartheta}^{1}\mathbb{E}[Cost(b(\theta;\omega))]p_{\Theta}(\theta)d\theta}. \quad (10)$$

By substituting the overall ROI definition in (10) and overall ad cost definition in (8) into (1), the overall revenue of the campaign can be rewritten as:

$$\text{Overall\_Revenue}(\vartheta) = N\int_{\vartheta}^{1} W_{\mathcal{Z}}(b(\theta;\omega);\beta)\theta\mathcal{V}p_{\Theta}(\theta)d\theta. \quad (11)$$

According to (11), we calculate the first partial derivative w.r.t. ϑ:

$$\frac{\partial\text{Overall\_Revenue}(\vartheta)}{\partial\vartheta} = -NW_{\mathcal{Z}}(b(\theta;\omega);\beta)\vartheta\mathcal{V}p_{\Theta}(\theta) \quad (12)$$

Taking into account the definitions of each term involved in (12), it is obvious that $\frac{\partial\text{Overall\_Revenue}(\vartheta)}{\partial\vartheta} \leq 0$. Therefore, the overall revenue is monotonically decreasing with respect to ϑ, which indicates that targeting bid requests with higher CVR will restrict the increase in revenue. □

Proposition 3 reveals the fact that it is challenging to increase revenue by solely targeting bid requests with higher CVRs without considering the dynamics between ROI and ad cost.

B. THE MR BIDDING STRATEGY

Based on the above analysis, we can see that many existing bidding strategies [2], [7], [9], [10] which focus on improving CTR/CVR are inapplicable of maximizing revenue directly. To address this problem, we propose the MR bidding strategy.

Definition 1 (Maximal Revenue Problem): The maximal revenue problem is formulated as:

$$\begin{aligned}\mathcal{B}^* &= \text{argmax} \;\; Overall\_Cost \\ &\quad subject\ to\ \ Overall\_ROI \geqslant \mathcal{R}.\end{aligned} \quad (13)$$

In (13), $\mathcal{B}^*$ represents the proposed MR bidding strategy that translates revenue maximization into maximizing overall ad cost for a given ROI $\mathcal{R}$. Moreover, we assume that if the ROI associated with each bid request is no less than $\mathcal{R}$, then the overall ROI of the campaign will be no less than $\mathcal{R}$. In this case, we substitute the formulations about overall ad cost (8) and ROI (5) into (13), and obtain:

$$\begin{aligned}\mathcal{B}^* &= \underset{b(\theta;\omega)}{\text{argmax}} \;\; N\int_{\theta}\mathbb{E}[Cost(b(\theta;\omega))]p_{\Theta}(\theta)d\theta \\ &\quad \text{subject to } \frac{\mathcal{V}}{\mathcal{R}} * \frac{W_{\mathcal{Z}}(b(\theta_i;\omega);\beta)\theta_i}{\mathbb{E}[Cost(b(\theta_i;\omega))]} \geqslant 1, \quad i \in N,\end{aligned} \quad (14)$$

where θ_i denotes the predicted CVR associated with the i^{th} bid request. Since the ROI constraint in (14) is too strict to be satisfied in practice, we define $\xi_{\theta_i} = \max\left(0, 1 - \frac{\mathcal{R}_{\theta_i}}{\mathcal{R}}\right)$ to relax the constraint where $\mathcal{R}_{\theta_i} = \frac{\mathcal{V}W_{\mathcal{Z}}(b(\theta_i;\omega);\beta)\theta_i}{\mathbb{E}[Cost(b(\theta;\omega))]}$ indicates the ROI corresponding to θ_i. As a result, (14) is rewritten as,

$$\begin{aligned}\mathcal{B}^* &= \underset{b(\theta;\omega)}{\text{argmax}}\, N\int_{\theta}\mathbb{E}[Cost(b(\theta;\omega))]p_{\Theta}(\theta)d\theta \\ &\quad \text{subject to } \frac{\mathcal{V}}{\mathcal{R}} * \frac{W_{\mathcal{Z}}(b(\theta_i;\omega);\beta)\theta_i}{\mathbb{E}[Cost(b(\theta_i;\omega))]} \geqslant 1-\xi_{\theta_i}, \quad i \in N.\end{aligned} \quad (15)$$

By introducing a regularization parameter $\mathcal{C}$ to control the trade-off between overall ad cost and ROI constraint (i.e., a larger $\mathcal{C}$ makes the bidding strategy more inclined towards satisfying the ROI constraint), we define the loss function of $b(\theta;\omega)$ according to (15) as follows:

$$\mathcal{L}(b(\theta;\omega)) = -N\int_{\theta}(1-\mathcal{C}\xi_{\theta})\mathbb{E}[Cost(b(\theta;\omega))]p_{\Theta}(\theta)d\theta, \quad (16)$$

where ξ_{θ} is described as

$$\xi_{\theta} = \max\left(0, 1 - \frac{\mathcal{V}}{\mathcal{R}} * \frac{W_{\mathcal{Z}}(b(\theta;\omega);\beta)\theta}{\mathbb{E}[Cost(b(\theta;\omega))]}\right). \quad (17)$$

To facilitate understanding the loss function defined in (16) and (17), we further discuss as follows. Given a bid request with CVR θ:

- For the case $\mathcal{R}_{\theta} \geqslant \mathcal{R}$, it means that the ROI associated with the given bid request satisfies the constrain in (15) and $\xi_{\theta} = 0$. In this case, the loss value associated with a bid request is calculated as $-\mathbb{E}[Cost(b(\theta;\omega))]$, which is identical to maximize the ad cost in (15).
- For the case $\mathcal{R}_{\theta} < \mathcal{R}$, only if $\xi_{\theta} = 1 - \frac{\mathcal{V}}{\mathcal{R}} * \frac{W_{\mathcal{Z}}(b(\theta;\omega);\beta)\theta}{\mathbb{E}[Cost(b(\theta;\omega))]}$, it is possible to satisfy the ROI constraint in (15). In this case, the expected minimum revenue associated with the given bid request is described as $\mathcal{R} * \mathbb{E}[Cost(b(\theta;\omega))]$. As a result, the difference between the expected minimum revenue and the actual revenue is described as $\Omega_{\theta} = \mathcal{R} * \mathbb{E}[Cost(b(\theta;\omega))] - \mathcal{V}W_{\mathcal{Z}}(b(\theta;\omega);\beta)\theta$. Furthermore by substituting $\xi_{\theta} = 1 - \frac{\mathcal{V}}{\mathcal{R}} * \frac{W_{\mathcal{Z}}(b(\theta;\omega);\beta)\theta}{\mathbb{E}[Cost(b(\theta;\omega))]}$ into (16), we derive the loss value associated with the given bid request as $-\left(\mathbb{E}[Cost(b(\theta;\omega))] - \frac{\mathcal{C}}{\mathcal{R}}\Omega_{\theta}\right)$, which is also the same as maximizing ad cost with the acceptable ROI constrain in (15).

From (16), (17) and above discussion, we note that the predicted CVR θ, expectation of ad cost $\mathbb{E}[Cost(b(\theta;\omega))]$

and winning function $W_{\mathcal{Z}}(b(\theta;\omega);\beta)$ play important roles in designing MR bidding strategies. Therefore, in the following sections, we provide the details to calculate CVR θ, $\mathbb{E}[Cost(b(\theta;\omega))]$ and $W_{\mathcal{Z}}(b(\theta;\omega);\beta)$ respectively.

1) PREDICTION OF CVR

From the above discussion, we can see that the accurate prediction of CVR θ associated with bid request $\boldsymbol{x}$ is indispensable for identifying the bidding strategy $\mathcal{B}^*$. Regarding this issue, many studies have been proposed to predict CVR as we have reviewed in Section II.

For simplicity and without loss of generality, we follow the study in [32] and apply the popular Logistic Regression to predict CVR, which is described as:

$$\theta(\boldsymbol{x}) = \frac{1}{1+e^{-\phi^T \boldsymbol{x}}}, \tag{18}$$

where ϕ is the parameter of the model. Specifically, given a training dataset $D = \{(\boldsymbol{x}_1, y_1), \ldots, (\boldsymbol{x}_i, y_i), \ldots, (\boldsymbol{x}_T, y_T)\}$ where T is the number of bid requests in this training dataset; $\boldsymbol{x}_i$ denotes the features of the i^{th} bid request; and $y_i \in \{0, 1\}$ is a binary variable indicating whether the conversion occurred (1) or not (0). The cross entropy loss with a regularization term associated with $\theta(\boldsymbol{x})$ is described as,

$$\mathcal{L}(\theta(\boldsymbol{x})) = -\frac{1}{T}\sum_{(\boldsymbol{x},y)\in D} y\log\theta(\boldsymbol{x}) + (1-y)\log(1-\theta(\boldsymbol{x})) + \frac{1}{2}\lambda_\phi \|\phi\|^2. \tag{19}$$

To determine the parameter ϕ involved in (18), we employ gradient descent to update ϕ as follows:

$$\phi \leftarrow \phi - \eta_\phi \frac{\partial \mathcal{L}(\theta(\boldsymbol{x}))}{\partial \phi}, \tag{20}$$

where η_ϕ is the learning rate, and the gradient of $\mathcal{L}(\theta(\boldsymbol{x}))$ regarding to ϕ is calculated as:

$$\frac{\partial \mathcal{L}(\theta(\boldsymbol{x}))}{\partial \phi} = \frac{1}{T}\sum_{(\boldsymbol{x},y)\in D}\left(\frac{1}{1+e^{-\phi^T \boldsymbol{x}}} - y\right)\boldsymbol{x} + \lambda_\phi \|\phi\|. \tag{21}$$

2) MODELLING AD COST

As mentioned, ad cost of a bid (i.e. the actual consumption) in many existing approaches is oversimplified as the bid price [2], [12]. However, in real RTB systems, this oversimplification is inconsistent with the GSP mechanism used to determine the winners. In the GSP mechanism, a bidder wins the bid request only if the bid price $b(\theta;\omega) > \mathcal{Z}$ where $\mathcal{Z}$ indicates the market price. Then, the winner pays $\mathcal{Z}$ as the actual ad cost. Otherwise, the DSP loses this bid without any payment.

Corresponding to the above GSP mechanism, the ad cost associated with a bidding price $b(\theta;\omega)$ can be formulated as:

$$Cost(b(\theta;\omega)) = \begin{cases} \mathcal{Z} & if\ b(\theta;\omega) > \mathcal{Z}, \\ 0 & if\ b(\theta;\omega) \le \mathcal{Z}. \end{cases} \tag{22}$$

Due to $\mathcal{Z} \sim p_{\mathcal{Z}}(z)$, the mathematical expectation of the ad cost described in (22) can be derived as:

$$\begin{aligned}\mathbb{E}[Cost(b(\theta;\omega))] &= \underbrace{\int_{z=0}^{b(\theta;\omega)} p_{\mathcal{Z}}(z) * z dz}_{when\ b(\theta;\omega) > z} + \underbrace{\int_{z=b(\theta;\omega)+1}^{+\infty} p_{\mathcal{Z}}(z) * 0 dz}_{when\ b(\theta;\omega) \le z} \\ &= \int_{z=0}^{b(\theta;\omega)} p_{\mathcal{Z}}(z) * z dz.\end{aligned} \tag{23}$$

However, in real-world scenarios, the competition for ad audiences as a result of different bidding strategies is very intense. It leads to the market price associated with massive ads changing frequently, which makes it challenging to model $p_{\mathcal{Z}}(z)$ directly.

To solve this problem, we propose an alternative solution to obtain $p_{\mathcal{Z}}(z)$ by defining $\mathcal{P}_{\mathcal{Z}}(z)$ as the cumulative distribution function (i.e., *c.d.f.*) of a random variable $\mathcal{Z}$. It describes the probability of a given bid price b being greater than the market price $\mathcal{Z}$ that is expressed as follows:

$$\mathcal{P}_{\mathcal{Z}}(b) = \int_{z=0}^{b} p_{\mathcal{Z}}(z) dz. \tag{24}$$

Equation (24) is the probability to win the bid request with bid price b and market price $\mathcal{Z}$. Considering the definition of the winning function $W_{\mathcal{Z}}(b;\beta)$, it is obvious that $\mathcal{P}_{\mathcal{Z}}(b)$ is identical to $W_{\mathcal{Z}}(b;\beta)$. Thus, we can derive:

$$W_{\mathcal{Z}}(b;\beta) \equiv \mathcal{P}_{\mathcal{Z}}(b) = \int_{z=0}^{b} p_{\mathcal{Z}}(z) dz. \tag{25}$$

Then, we can obtain $p_{\mathcal{Z}}(z)$ by taking the first derivative of $W_{\mathcal{Z}}(b;\beta)$ with respect to b:

$$\frac{\partial W_{\mathcal{Z}}(b;\beta)}{\partial b} \equiv \frac{\partial \mathcal{P}_{\mathcal{Z}}(b)}{\partial b} = p_{\mathcal{Z}}(b). \tag{26}$$

In (26), $p_{\mathcal{Z}}(b)$ represents the *p.d.f.* of a random variable $\mathcal{Z}$ and the estimation of $p_{\mathcal{Z}}(b)$ depends on the specific formulation of the winning function $W_{\mathcal{Z}}(b;\beta)$ which will be discussed in the next section.

3) MODELLING THE WINNING FUNCTION

In the above discussion, the specific expression of a winning function $W_{\mathcal{Z}}(b;\beta)$ is necessary to estimate the expectation of ad cost (as described in (23) and (26)) and to identify the optimal bidding strategy. In previous studies, different mathematical expressions about the winning function $W_{\mathcal{Z}}(b;\beta)$ have been proposed. For example, [13] assumed that the logarithm of market price $\log(z)$ follows the Normal distribution $\mathcal{N}(\mu_z, \sigma_z^2)$ where μ_z and σ_z describe the mean and the standard deviation, respectively. In this case, the winning function is defined as $W_{\mathcal{Z}}(b;\mu_z,\sigma_z) = \Phi((\log(b)-\mu_z)/\sigma_z)$ where Φ denotes the *c.d.f.* of a standard Normal distribution. In addition, the method in [2] formulates the winning function as $W_{\mathcal{Z}}(b;\beta) = b/(b+\beta)$.

Unlike previous studies, we empirically define $W_{\mathcal{Z}}(b;\beta)$ in the form of tanh to model winning function based on analyzing the iPinYou [1] and the YOYI datasets[2] as follows:

$$W_{\mathcal{Z}}(b;\beta)=\tanh(\beta b). \tag{27}$$

To validate the accuracy, we further compared the winning probability that are calculated by different winning functions against the ground truth values in the iPinYou dataset and the YOYI dataset, respectively. The detailed comparison results reported in Section V-D show that our proposed winning function in (27) predicts winning probability with high accuracy in terms of Wasserstein distance and KL-divergence.

Note that given the winning function (27), we employ regularized regression to further determine the parameter β. Specifically, we denote the training dataset as $Z=\{z_1,\ldots,z_i,\ldots,z_T\}$ where T is the number of bid requests in this training dataset and z_i indicates the market price associated with the i^{th} bid request. Since the market price z_i is commonly represented as a positive integer [26], the minimum market price and maximum market price in Z can be described as 1 and z^*, respectively. By grouping the dataset Z, we further define $Z'=\{(1,n_1),\ldots,(z_i,n_{z_i}),\ldots,(z^*,n_{z^*})\}$ in which (z_i,n_{z_i}) means that there exists n_{z_i} bid requests with same market price z_i. In this case, the mean squared error loss about the winning function $W_{\mathcal{Z}}(b;\beta)$ is described as follows:

$$\mathcal{L}(W_{\mathcal{Z}}(b;\beta))=\frac{1}{2\,|Z'|}\sum_{b=1}^{z^*}(W_{\mathcal{Z}}(b;\beta)-\mathcal{P}_{\mathcal{Z}}(b))^2+\frac{1}{2}\lambda_\beta\,\|\beta\|^2, \tag{28}$$

where $\frac{1}{2}\lambda_\beta\,\|\beta\|^2$ is a regularization term. By substituting the winning function $W_{\mathcal{Z}}(b;\beta)=\tanh(\beta b)$ into (28), the gradient of $\mathcal{L}(W_{\mathcal{Z}}(b;\beta))$ regarding to β is calculated as:

$$\frac{\partial\mathcal{L}(W_{\mathcal{Z}}(b;\beta))}{\partial\beta}=\frac{1}{|Z'|}\sum_{b=1}^{z^*}b\,(\tanh(\beta b)-\mathcal{P}_{\mathcal{Z}}(b))\times(1-\tanh^2(\beta b))+\lambda_\beta\,\|\beta\|. \tag{29}$$

Since Z' is composed of discrete values, the corresponding $\mathcal{P}_{\mathcal{Z}}(b)$ defined in Equation (24) is calculated by,

$$\mathcal{P}_{\mathcal{Z}}(b)=\frac{1}{N}\sum_{\{n_{z_i}|(z_i,n_{z_i})\in Z',z_i<b\}}n_{z_i}. \tag{30}$$

Then, we apply gradient descent to update β as follows:

$$\beta\leftarrow\beta-\eta_\beta\frac{\partial\mathcal{L}(W_{\mathcal{Z}}(b;\beta))}{\partial\beta} \tag{31}$$

where η_β is the learning rate. The detailed procedure to learn the parameter β in the winning function is summarized in Algorithm 1. The computational complexity of the learning algorithm depends on the size of all possible bid prices in Z'.

[1]Dataset link: http://data.computational-advertising.org
[2]Dataset link: http://apex.sjtu.edu.cn/datasets/7

Algorithm 1 Learning β in the Winning Function

INPUT: Training set Z'
OUTPUT: Winning function $W_{\mathcal{Z}}(b;\beta)$

```
Initially set parameter β of the winning function
for number of training rounds do
    for all samples in Z' do
        Calculate the gradient of β via (29)
        Update parameter β via (31)
    end for
end for
return W_Z(b; β)
```

Finally, for the winning function $W_{\mathcal{Z}}(b;\beta)$, the *p.d.f.* of the market price $\mathcal{Z}$ is computed following (26) and (27):

$$p_{\mathcal{Z}}(b)=\frac{\partial W_{\mathcal{Z}}(b;\beta)}{\partial b}=\beta(1-\tanh^2(\beta b)). \tag{32}$$

4) OPTIMIZATION ALGORITHM

Based on above derivations, this section focuses on optimizing the parameter ω involved in the MR bidding strategy $b(\theta;\omega)$. Without loss of generality, we apply the gradient decent method to update the parameter ω in the bid functions as follows:

$$\omega\leftarrow\omega-\eta_\omega\frac{\partial\mathcal{L}(b(\theta;\omega))}{\partial b(\theta;\omega)}\frac{\partial b(\theta;\omega)}{\partial\omega}, \tag{33}$$

where η_ω is the learning rate and $\mathcal{L}(b(\theta;\omega))$ is the loss function. It is obvious that the update in (33) depends on the calculation of $\frac{\partial\mathcal{L}(b(\theta;\omega))}{\partial b(\theta;\omega)}$ and the specific expression of $b(\theta;\omega)$. Given a training dataset D described in Section IV-B.1, we transform it into $D'=\{\theta_1,\ldots,\theta_i,\ldots,\theta_T\}$ where θ_i denotes the predicted CVR associated with the i^{th} bid request. According to (16) and (17), the loss function $\mathcal{L}(b(\theta;\omega))$ can be re-written as

$$\mathcal{L}(b(\theta;\omega))=-\sum_{\theta\in D'}(1-\mathcal{C}\xi_\theta)\mathbb{E}[Cost(b(\theta;\omega))],$$
$$\text{and }\xi_\theta=\max\left(0,1-\frac{\mathcal{V}}{\mathcal{R}}*\frac{W_{\mathcal{Z}}(b(\theta;\omega);\beta)\theta}{\mathbb{E}[Cost(b(\theta;\omega))]}\right). \tag{34}$$

By differentiating (34) w.r.t. $b(\theta;\omega)$, we obtain:

$$\begin{aligned}\frac{\partial\mathcal{L}(b(\theta;\omega))}{\partial b(\theta;\omega)}&\\=-\sum_{\theta\in D'}\Bigg\{&\frac{\partial\mathbb{E}[Cost(b(\theta;\omega))]}{\partial b(\theta;\omega)}\\&*\left(1-\mathbb{I}[\xi_\theta>0]\mathcal{C}\left(1-\frac{\mathcal{V}\theta}{\mathcal{R}}\frac{W_{\mathcal{Z}}(b(\theta;\omega);\beta)}{\mathbb{E}[Cost(b(\theta;\omega))]}\right)\right)\\&+\mathbb{I}[\xi_\theta>0]\mathcal{C}\frac{\mathcal{V}\theta}{\mathcal{R}}*\Bigg(\frac{\partial W_{\mathcal{Z}}(b(\theta;\omega);\beta)}{\partial b(\theta;\omega)}\\&-\frac{W_{\mathcal{Z}}(b(\theta;\omega);\beta)}{\mathbb{E}[Cost(b(\theta;\omega))]}\frac{\partial\mathbb{E}[Cost(b(\theta;\omega))]}{\partial b(\theta;\omega)}\Bigg)\Bigg\},\end{aligned} \tag{35}$$

where $\mathbb{I}[\xi_\theta > 0]$ is an indicator function (i.e., $\mathbb{I}[\xi_\theta > 0] = 1$ if $\xi_\theta > 0$; otherwise $\mathbb{I}[\xi_\theta > 0] = 0$). In terms of (23), the $\frac{\partial \mathbb{E}[Cost(b(\theta;\omega))]}{\partial b(\theta;\omega)}$ in (35) is derived as follows:

$$\frac{\partial \mathbb{E}[Cost(b(\theta;\omega))]}{\partial b(\theta;\omega)} = b(\theta;\omega) * p_{\mathcal{Z}}(b(\theta;\omega)). \tag{36}$$

By further substituting (26) and (36) into (35), we can obtain:

$$\begin{aligned}
&\frac{\partial \mathcal{L}(b(\theta;\omega))}{\partial b(\theta;\omega)} \\
&\quad = -\sum_{\theta \in D'} \Bigg\{ b(\theta;\omega) p_{\mathcal{Z}}(b(\theta;\omega)) \\
&\qquad * \left(1 - \mathbb{I}[\xi_\theta > 0]\mathcal{C}\left(1 - \frac{\mathcal{V}\theta}{\mathcal{R}} \frac{W_{\mathcal{Z}}(b(\theta;\omega);\beta)}{\mathbb{E}[Cost(b(\theta;\omega))]}\right)\right) \\
&\qquad + \mathbb{I}[\xi_\theta > 0]\mathcal{C}\frac{\mathcal{V}\theta}{\mathcal{R}} * \Bigg(p_{\mathcal{Z}}(b(\theta;\omega)) \\
&\qquad - \frac{W_{\mathcal{Z}}(b(\theta;\omega);\beta)}{\mathbb{E}[Cost(b(\theta;\omega))]} b(\theta;\omega) p_{\mathcal{Z}}(b(\theta;\omega)) \Bigg) \Bigg\}.
\end{aligned} \tag{37}$$

Moreover, we need to specify the expression of the bidding strategy $b(\theta;\omega)$. To facilitate comparisons against existing approaches in [2], [10], [11], we employ two popular bidding functions that are defined as follows:

Definition 2 (MR1): The first one is a widely used linear bidding function [6], [11]:

$$b^{MR1}(\theta;\omega) = \omega * \theta, \tag{38}$$

which is referred to as the MR1 bidding strategy.

Definition 3 (MR2): The second one is similar to the square root bid function defined in [2], [10],

$$b^{MR2}(\theta;\omega) = \omega_1 \sqrt{\theta + \omega_2^2} - \omega_2, \tag{39}$$

which is referred to as the MR2 bidding strategy.

Given the two bidding strategies $b^{MR1}(\theta;\omega)$ and $b^{MR2}(\theta;\omega)$, we can calculate the corresponding bid price $b(\theta;\omega)$ for each input θ. In addition, all the terms in (37), including $p_{\mathcal{Z}}(b(\theta;\omega))$ (32), ξ_θ (34), $W_{\mathcal{Z}}(b(\theta;\omega);\beta)$ (27) and $\mathbb{E}[Cost(b(\theta;\omega))]$ (23) can be calculated. They serve as the inputs to optimize the parameter ω involved in the bidding strategies. Finally, the detailed learning procedure to optimize the parameter ω is summarized in Algorithm 2.

Algorithm 2 Learning the Parameter ω in the MR Bidding Strategy

INPUT: Training set D', winning function $W_{\mathcal{Z}}(b;\beta)$,
ad cost estimator $\mathbb{E}[Cost(b)]$,
market price distribution $p_{\mathcal{Z}}(z)$,
OUTPUT: Bidding strategy $b^{MR}(\theta;\omega)$

```
Pre-specify the bidding function b^{MR}(θ; ω) as MR1 (38) or MR2 (39)
Initializing the parameter ω involved in the bidding function MR1 and MR2
for number of training rounds do
    for all samples θ ∈ D' do
        Calculate the gradient of loss function L(b(θ; ω)) w.r.t. b^{MR}(θ; ω) via (37)
        Calculate the gradient of pre-specified b^{MR}(θ; ω) in (38) and (39) w.r.t. ω
        Update parameter ω via (33)
    end for
end for
return b^{MR}(θ; ω)
```

V. EXPERIMENTS

In this section, we first briefly introduce the datasets, experimental setup with evaluation metrics and baseline bidding strategies. Secondly, to validate our proposed winning function which affects the estimation of ad cost and the optimization of bidding strategy (as discussed in section IV-B.3), we calculate and compare the Wasserstein distance and KL-divergence against those of other popular winning functions. Thirdly, we compare the performance of the proposed MR bidding strategy with the baselines and discuss the experimental results. Last but not least, the impact of the regularization parameter $\mathcal{C}$ and different winning functions on bidding performance are studied to validate our research motivation.

A. DATASETS

To study the effectiveness of the proposed MR bidding strategy, we conduct extensive offline experiments on two real-world benchmark datasets: iPinYou and YOYI.

1) iPinYou

This dataset is used for the global RTB algorithm competition [33]. The available dataset consists of 9 different display ad campaigns during 10 days in 2013, including 64.75 million bids, 19.50 million impressions, 14.79 thousand clicks and 1,253 conversions. For each campaign, the first 7 days of data are used as training data while the rest are used as test data.

2) YOYI

This dataset recorded multi-device display advertising during 8 days in January 2016, which contains 5.13 million impressions and 428.27 thousand clicks. The first 7 days of data are used as training data while the rest are used as test data. Since the campaign information is unavailable, we treat all records as a single campaign.

In contrast to the YOYI dataset, which does not include conversion data, the conversion data are available for 4 out of 9 campaigns in the iPinYou dataset. In order to have more data for the experiments, we follow the methods in [10], [11] and regard the number of clicks as a proxy for conversions. In the above datasets, each record is described as a triple

$(\boldsymbol{x}, y, z)$ where $\boldsymbol{x}$ represents the feature vector of a bid request, y is the binary user response (i.e., click or not) and z is the corresponding market price.

B. EXPERIMENT SETUP

1) LEARNING PROCEDURE

The whole learning procedure of our proposed method in all experiments is as follows. Since clicks are regarded as the proxy for conversions, click data are first utilized to predict CVR $\theta(\boldsymbol{x})$ by training a Logistic Regression model as described in Section IV-B.1. Then, we use the market price data to determine the parameter β involved in the pre-specified winning function $W_{\mathcal{Z}}(b; \beta)$ as described in Algorithm 1. Furthermore, the expected ad cost $\mathbb{E}[Cost(b))]$ is derived according to (23) in which $p_{\mathcal{Z}}(z)$ is given by (26). Finally, the parameter ω underlying bidding strategies MR1 and MR2 is optimized by performing Algorithm 2. In this process, to facilitate quick convergence, we apply BFGS [34] that is a state-of-the-art optimization algorithm to update ω according to (33) $\sim$ (39). The source code for our experiment is available on Github via the following link: https://github.com/wty9391/maximal-revenue-rtb.

2) EVALUATION PROCEDURE

To ensure the comparability, we employed the evaluation procedure similar to the previous work in [2] as follows. The original impression log is used as the full volume bid request data and delivered to the bidding strategy with the original logged sequence. Once a bid request is received, the bidding strategy calculates a bid price to participate in the real-time bidding auction. If this bid price is higher than the recorded market price, the bidding strategy wins this ad impression and pays the recorded market price; otherwise, the bidding strategy loses the ad impression without any payment. Then, for the winning impression, the bidding strategy collects the corresponding user feedback (i.e. clicked or not). After bidding all received bid requests, we record the total number of winning impressions as I, the total number of clicks as C, and the amount of ad cost as A.

Since the primary task of our study is to maximize the revenue of a campaign on the premise of acceptable ROI, we adopt **revenue** and **ROI** w.r.t. the corresponding ad cost as our primary evaluation metrics. To ensure fairness of evaluation and comparisons, we assign the value of each click $\mathcal{V}$ to be 0.1 CNY. Thus, the **revenue** is calculated as $C * \mathcal{V}$ and the **ROI** is calculated as $C * \mathcal{V}/A$. It is worth mentioning that if the bidding strategy bids ads with very high prices each time, the ad cost and revenue will be same as the original test log. Therefore, the budget constraints play a key role in the evaluation [2]. To deal with this issue, we follow the previous work [2], [11] and set a budget (i.e. the upper limit of the ad cost) on the bidding strategy. More specifically, we perform the evaluation by using $1/8$, $1/16$ and $1/32$ of the original total cost in the test log as the budget constraints, respectively following [2]. In addition, we also take the related metrics, including the ad cost associated with different bidding strategies and the number of winning impressions, to demonstrate the bidding performance.

C. BASELINES

Since the proposed MR bidding strategy aims at maximizing revenue with ROI constraint, it can be considered as a typical KPI-oriented bidding strategy. Therefore, five state-of-the-art KPI-oriented bidding strategies are employed as baseline models, which are:

- **Max-eCPC bidding (Mcpc):** This bidding strategy is widely used in existing works [2], [9] as a baseline. Mcpc is given as $b^{Mcpc}(\theta) = eCPC * \theta$ where $eCPC$ is calculated by dividing the total ad cost by the total number of clicks in the training data.
- **Optimal real-time bidding (ortb):** This bidding strategy [2] is designed to maximize the number of clicks by combining CTR prediction and ad cost estimation for each bid request.
- **Statistical arbitrage mining (sam):** This bidding strategy [10] is designed to maximize the profits (i.e., subtracting ad cost from revenue) of campaign based on the estimated CVR and ad cost for each bid request.
- **Reinforcement learning based bid (RLB):** To maximize the number of clicks, RLB [9] formulates the bid decision process as a reinforcement learning problem in which a Markov Decision Process is utilized to determine bid prices according to the remaining bid request volume and the remaining budget.
- **Bidding machine (BM):** This bidding strategy [11] is presented to maximize ad campaign profit by jointly optimizing CTR prediction, ad cost estimation and the bidding function.

D. EFFECTS OF THE WINNING FUNCTIONS ON PREDICTING WINNING PROBABILITY

As aforementioned in Section IV-B.3, the winning function plays an important role in bidding strategy optimization. To validate the effectiveness of the proposed winning function, Fig. 2 depicts the winning probabilities that are calculated by different winning functions $\tanh(\beta b)$, $b/(b+\beta)$ [13] and $\Phi((\log(b) - \mu_z)/\sigma_z)$ [2] respectively, and compares them against the ground truth value in the iPinYou dataset and YOYI dataset. From Fig. 2, we can see that the curve generated by our proposed winning function $\tanh(\beta b)$ is much closer to the truth value.

Furthermore, we quantitatively assess the winning functions in terms of Wasserstein distance and KL-divergence since they are commonly used to measure the differences between probability distributions. Table 2 presents a detailed comparisons among different fine-tuned winning functions for each campaign. The results show that, for most campaigns, the winning function $W_{\mathcal{Z}}(b; \beta) = \tanh(\beta b)$ results in the least Wasserstein distance (7 out of 11 campaigns) and KL-divergence (6 out of 11 campaigns). It indicates that $W_{\mathcal{Z}}(b; \beta) = \tanh(\beta b)$ predicts winning probability more

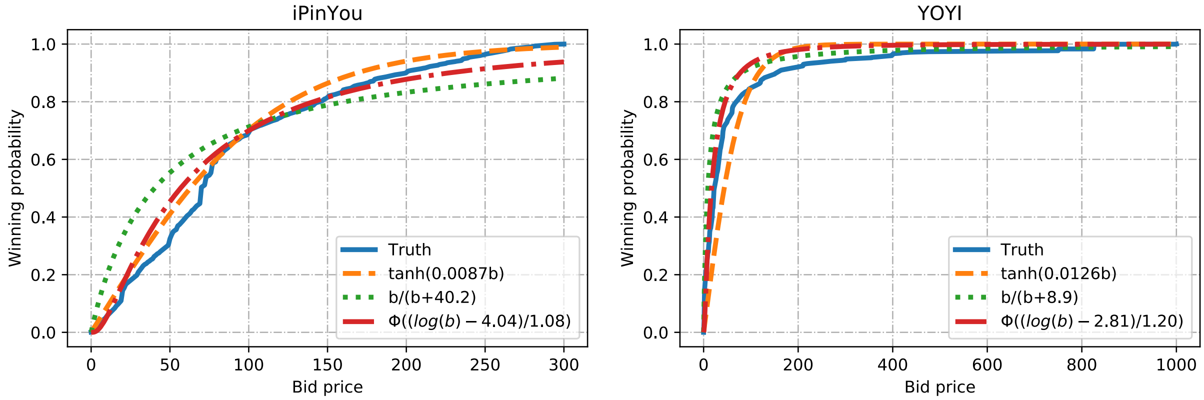

(a) (b)

FIGURE 2. Winning probability vs. bid price based on the iPinYou dataset and YOYI dataset.

TABLE 2. Comparisons among three fine-tuned winning functions.

iPinYou campaign	Wasserstein distance			KL-divergence		
	$\tanh(\beta b)$	$\frac{b}{b+\beta}$	$\Phi(\frac{\log(b)-\mu_z}{\sigma_z})$	$\tanh(\beta b)$	$\frac{b}{b+\beta}$	$\Phi(\frac{\log(b)-\mu_z}{\sigma_z})$
1458	0.0353	0.1035	**0.0255**	0.0045	0.0251	**0.0029**
2259	**0.0378**	0.1056	0.0631	**0.0015**	0.0220	0.0065
2261	**0.0388**	0.0580	0.1471	**0.0066**	0.0067	0.0084
2821	**0.0478**	0.1070	0.0681	**0.0017**	0.0229	0.0074
2997	0.0448	0.0510	**0.0274**	0.0082	**0.0028**	0.0037
3358	0.1007	0.1415	**0.0865**	0.0161	0.0574	**0.0145**
3386	**0.0327**	0.1163	0.0347	**0.0038**	0.0278	0.0049
3427	0.0383	0.1139	**0.0321**	0.0065	0.0355	**0.0060**
3476	**0.0376**	0.1164	0.0454	0.0080	0.0362	**0.0065**
All campaigns	**0.0335**	0.0936	0.0399	**0.0022**	0.0226	0.0039
YOYI	**0.0390**	0.0403	0.0418	**0.0033**	0.0037	0.0041

accurately than the existing approaches in [13] and [2]. Therefore, it is utilized in the MR bidding strategy to maximize revenue in Section V-E.

E. ADVERTISING PERFORMANCE

In this section, we compare and report the advertising performance achieved by MR1/MR2 with that achieved by other baselines. To conduct more extensive experiments, we regard all campaigns in the iPinYou dataset as a single campaign to perform comparisons and report the results in Table 3. In addition, the detailed experimental results on the YOYI dataset are reported in Table 4. For fairness of comparison, we pre-define the acceptable ROI $\mathcal{R}$ as 1.0 and make the ad cost associated with MR1/MR2 similar to that of other baseline models by tuning $\mathcal{C}$. From Table 3 and Table 4, we can observe that:

1) For both iPinYou and YOYI datasets, MR1 and MR2 consistently achieve higher revenue than other state-of-the-art baselines. Regarding ROI metric, MR bidding strategy gains less ROI than ortb and sam bidding strategy in the iPinYou dataset with 1/8 budget setting; the MR bidding strategy gains less ROI than sam and BM bidding strategies in the YOYI dataset with 1/8 budget setting. With budget declining form 1/8 to 1/32, MR1 and MR2 provide more stable improvements on ROI. These results demonstrate the effectiveness of the proposed MR bidding strategy to maximize revenue with acceptable ROI.
2) Similar to Mcpc that utilizes a linear bidding function, MR1 also employs a linear bidding function as shown in (38). By contrast, MR1 provides significant improvements on the revenue and ROI. Analogously, in contrast with sam which adopts similar bidding functions as MR2 (39), MR2 achieves more revenue. From this perspective, those results demonstrate that MR1 and MR2 provide more stable improvements on revenue.
3) From Table 3 we can see that, given similar ad cost, ortb wins more impressions than MR1 and MR2 for iPinYou dataset. Nevertheless, MR1 and MR2 gain more revenue with less impressions. The possible reason is that ortb considers each bid request equally. In contrast, MR1 and MR2 tend to spend the budget on the bid requests with higher ROI. These results suggest that our proposed MR bidding strategy is much more cost-effective.
4) In contrast with BM, MR1 and MR2 significantly improve revenue. The possible reason is that the market price distribution $p_{\mathcal{Z}}(z)$ in BM is derived by marginalizing the estimation of market price density $p_{\mathcal{Z}}(z, \boldsymbol{x})$

TABLE 3. Performance comparisons on the iPinYou dataset under various budget settings.

Strategy	Budget setting	Ad cost	Revenue	ROI	Impressions (K)	MR1 improvements on		MR2 improvements on	
						Revenue	ROI	Revenue	ROI
Mcpc	1/8	455.20	178.40	0.392	848.61	13.90%	19.64%	13.79%	16.07%
RLB	1/8	397.24	162.90	0.410	760.87	24.74%	14.39%	24.62%	10.96%
ortb	1/8	410.70	190.10	0.463	968.84	6.89%	1.30%	6.79%	-1.73%
sam	1/8	346.89	189.90	**0.547**	739.77	7.00%	-14.26%	6.90%	-16.82%
BM	1/8	430.77	193.80	0.450	751.22	4.85%	4.22%	4.75%	1.11%
MR1	1/8	433.26	**203.20**	0.469	863.67	-	-	-	-
MR2	1/8	445.71	203.00	0.455	879.23	-	-	-	-
Mcpc	1/16	227.60	92.20	0.405	407.22	89.91%	125.43%	87.85%	121.48%
RLB	1/16	197.35	136.30	0.691	458.33	28.47%	32.13%	27.07%	29.81%
ortb	1/16	215.78	156.80	0.727	611.59	11.67%	25.58%	10.46%	23.38%
sam	1/16	194.40	162.30	0.835	618.25	7.89%	9.34%	6.72%	7.43%
BM	1/16	194.88	148.70	0.763	433.83	17.75%	19.66%	16.48%	17.56%
MR1	1/16	191.84	**175.10**	**0.913**	490.55	-	-	-	-
MR2	1/16	193.11	173.20	0.897	497.93	-	-	-	-
Mcpc	1/32	113.80	48.10	0.423	214.26	214.97%	220.33%	212.06%	214.89%
RLB	1/32	108.98	119.30	1.095	281.52	26.99%	23.74%	25.82%	21.64%
ortb	1/32	113.78	128.80	1.132	377.24	17.62%	19.70%	16.54%	17.67%
sam	1/32	113.59	145.40	1.280	401.85	4.20%	5.86%	3.23%	4.06%
BM	1/32	113.80	134.20	1.180	271.90	12.89%	14.83%	11.85%	12.88%
MR1	1/32	111.83	**151.50**	**1.355**	304.69	-	-	-	-
MR2	1/32	112.63	150.10	1.332	306.28	-	-	-	-

TABLE 4. Performance comparisons on the YOYI dataset under various budget settings.

Strategy	Budget setting	Ad cost	Revenue	ROI	Impressions (K)	MR1 improvements on		MR2 improvements on	
						Revenue	ROI	Revenue	ROI
Mcpc	1/8	54.60	26.60	0.487	165.41	27.82%	38.19%	29.70%	33.88%
RLB	1/8	52.73	32.20	0.611	173.93	5.59%	10.15%	7.14%	6.71%
ortb	1/8	45.67	29.90	0.655	162.50	13.71%	2.75%	15.38%	-0.46%
sam	1/8	41.36	28.80	**0.696**	115.70	18.06%	-3.30%	19.79%	-6.32%
BM	1/8	48.60	32.80	0.675	163.13	3.66%	-0.30%	5.18%	-3.41%
MR1	1/8	50.50	34.00	0.673	167.81	-	-	-	-
MR2	1/8	52.90	**34.50**	0.652	170.13	-	-	-	-
Mcpc	1/16	27.30	14.20	0.520	98.20	84.51%	84.51%	83.10%	83.10%
RLB	1/16	26.92	25.00	0.929	114.68	4.80%	3.34%	4.00%	2.48%
ortb	1/16	22.48	20.80	0.925	114.56	25.96%	3.78%	25.00%	2.92%
sam	1/16	20.60	19.10	0.927	102.26	37.17%	3.56%	36.13%	2.70%
BM	1/16	27.30	23.90	0.875	111.65	9.62%	9.71%	8.79%	8.80%
MR1	1/16	27.30	**26.20**	**0.960**	112.93	-	-	-	-
MR2	1/16	27.30	26.00	0.952	114.28	-	-	-	-
Mcpc	1/32	13.65	8.80	0.645	62.74	130.68%	132.56%	126.14%	143.57%
RLB	1/32	13.43	19.00	1.415	78.19	6.84%	6.01%	4.74%	11.02%
ortb	1/32	12.48	17.60	1.410	90.55	15.34%	6.38%	13.07%	11.42%
sam	1/32	11.71	17.50	1.494	87.34	16.00%	0.40%	13.71%	5.15%
BM	1/32	13.65	19.80	1.451	80.12	2.53%	3.38%	0.51%	8.27%
MR1	1/32	13.53	**20.30**	1.500	79.41	-	-	-	-
MR2	1/32	12.67	19.90	**1.571**	78.43	-	-	-	-

for each bid request $\boldsymbol{x}$, which may be overfitting to the ground truth value. Different from BM, MR obtains $p_{\mathcal{Z}}(z)$ in terms of the pre-defined winning function $W_{\mathcal{Z}}(b;\beta)$ as described in (26). By doing so, MR predicts the expectation of the ad cost more accurately, which is indispensable to optimizing the bidding strategy as discussed in Section IV-B.4.

To further demonstrate the possible advantage of our approach intuitively, Fig. 3 depicts the real market price distribution on iPinYou dataset and compares it with the distributions of bid price that are generated by MR1, MR2, sam and BM respectively. Since the maximal market price is 300 in the dataset, we cut off the figure for price >300. In Fig. 3, when budget setting decreasing from 1/8 to 1/32, MR1/MR2 significantly decreases the bidding number for the ad request with high market price (i.e., price >100). In contrast with sam and BM, it shows that the proposed MR1/MR2 bidding strategies provide more flexible bidding responses for tackling different budget settings, which enables the MR strategy to bid for ads in a more productive manner.

The above discussion shows that both MR1 and MR2 significantly improve the revenue and ROI against other state-of-the-art bidding strategies. Moreover, we also note that the linear bidding function (i.e. MR1) performs better than the quadratic bidding function (i.e. MR2), since the linear bidding has been proved to be the optimal bidding function in GSP [11].

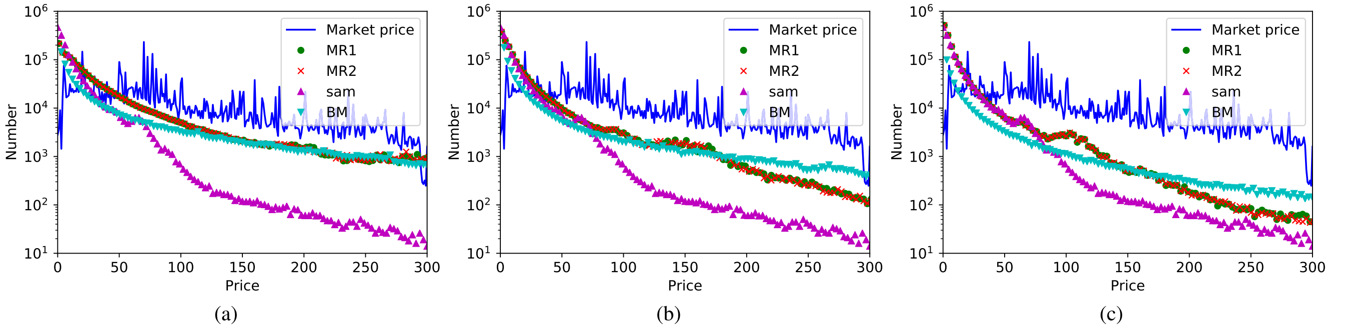

FIGURE 3. Analysis of bid price and market price distributions on the iPinYou dataset under various budget settings. (a) 1/8 budget setting. (b) 1/16 budget setting. (c) 1/32 budget setting.

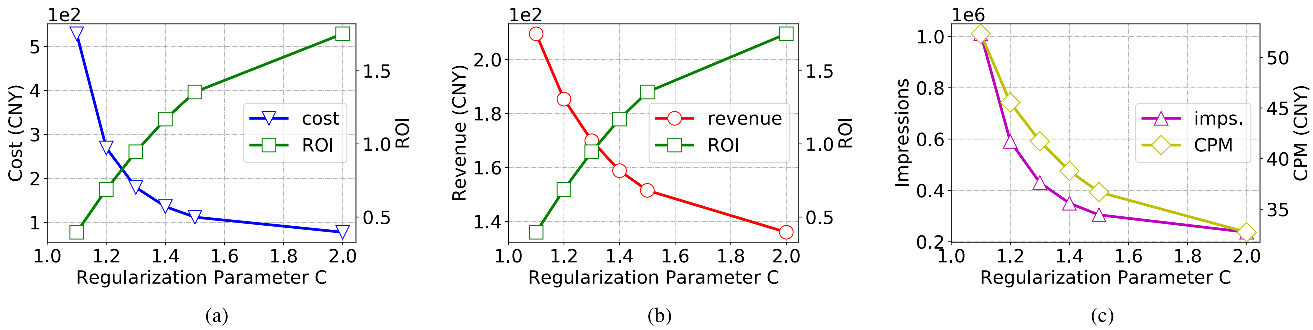

FIGURE 4. Performance comparison under different $\mathcal{C}$ values on the iPinYou dataset.

TABLE 5. Performance comparisons among different winning functions on the iPinYou dataset under 1/32 budget setting.

Winning function	$\mathcal{C}$	Ad cost	Revenue	ROI	Impressions (K)	tanh(·) improvements on	
						Revenue	ROI
$b/(b+\beta)$	1.40	113.80	123.60	1.086	272.18	17.80%	17.80%
$\Phi((\log(b)-\mu_z)/\sigma_z)$	1.40	113.80	131.30	1.154	275.55	10.89%	10.89%
$\tanh(\beta b)$	1.40	113.80	**145.60**	**1.280**	293.60	-	-
$b/(b+\beta)$	1.50	113.80	133.10	1.170	276.63	11.34%	19.45%
$\Phi((\log(b)-\mu_z)/\sigma_z)$	1.50	113.80	145.70	1.280	293.17	1.72%	9.12%
$\tanh(\beta b)$	1.50	106.08	**148.20**	**1.397**	297.15	-	-
$b/(b+\beta)$	1.60	113.80	143.30	1.259	288.68	0.07%	22.56%
$\Phi((\log(b)-\mu_z)/\sigma_z)$	1.60	103.33	142.00	1.374	285.13	0.99%	12.30%
$\tanh(\beta b)$	1.60	92.94	**143.40**	**1.543**	271.34	-	-

F. POTENTIAL IMPACT OF REGULARIZATION PARAMETER AND THE RATIONALITY OF THE RESEARCH MOTIVATION

To further examine the potential impact of regularization parameter $\mathcal{C}$ in (16) on the desirable KPIs (i.e., ad cost, ROI, revenue, impressions and CPM) and validate the rationality of research motivation, we apply MR1 with different $\mathcal{C} = \{1.10, 1.20, 1.30, 1.40, 1.50, 2.00\}$ on iPinYou dataset without budget constraint. The experimental results are illustrated in Fig. 4.

In Fig. 4a, we observe that the regularization parameter $\mathcal{C}$ balances the trade-off between ROI and overall ad cost. The higher $\mathcal{C}$ makes bidding strategy MR1 more inclined to bid the requests with higher ROI while decreasing ad cost. From this perspective, it verifies Proposition 1 and Proposition 2. Furthermore, Fig. 4b shows that the overall ROI and the overall ad cost follow opposing trends. In Fig. 4b, the higher $\mathcal{C}$ results in decreasing overall revenue, which validates the reasonableness of our research motivation (i.e. it is challenging to increase revenue by solely targeting bid requests with higher CVR/ROI without considering the impact of the overall ad cost). In contrast, Fig. 4c shows that both the winning impressions and the CPM decrease with increasing $\mathcal{C}$ (i.e. preferring the ad audiences with higher CVR/ROI). It is mainly due to the fact that the stiff competition on the limited ad audiences with higher CVR/ROI decreases the chance to win impressions and the corresponding CPM.

G. POTENTIAL IMPACT OF THE WINNING FUNCTION

In this section, we examine the potential impact on the performance due to different winning functions involved in our proposed bidding strategy. To this end, we equip MR1 with fine tuned winning functions $b/(b+\beta)$, $\Phi((\log(b)-\mu_z)/\sigma_z)$ and $\tanh(\beta b)$, respectively, and perform the experiments on the iPinYou dataset under the 1/32 budget setting.

For different values of the regularization parameter $\mathcal{C}$, Table 5 reports the corresponding experimental results about ad cost, number of clicks, revenue, ROI and impressions. In Table 5, for a given $\mathcal{C}$, we can see that the ad cost corresponding to each winning function form is comparable. In this case, the revenue and ROI associated with the proposed tanh winning function are much better than those of the other two winning functions. Combining these results with the discussions about (23) and (26), it validates that the winning function in tanh form provides more accurate expectation of ad cost, which makes it possible to maximize the revenue with acceptable ROI directly.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we propose a novel bidding strategy to satisfy the advertisers' requirements on revenue more efficiently. To this end, we first theoretically analyze the relations among CVR, ROI and ad cost that respectively impose the potential impacts on revenue. Inspired by this insight, our proposed MR translates and formulates the revenue maximization problem into maximizing ad cost with desirable ROI constraint. On the one hand, the MR first distinguishes bid price from ad cost explicitly, which makes it more applicable to the real GSP auction mechanism in real-time bidding systems. On the other hand, we empirically define the winning function underlying the bidding process as tanh form so as to estimate the expectation of ad cost more accurately, which plays an important role in optimizing bidding strategy. The experiments on public datasets show that, compared to five popular state-of-the-art methods, our approach significantly boosts the campaign performance in terms of revenue and ROI. In addition, we also perform the experiments to validate the reasonableness of our research motivation.

In the future, we plan to learn more realistic and more expressive winning functions and bidding functions so as to improve the efficiency of MR bidding strategy. Additionally, we will also study the effects of censored data on winning probability estimation [23] and impression evaluation [26] in order to improve the robustness of the MR bidding strategy. As user privacy and data confidentiality becomes an important consideration for online businesses relying on AI technologies, it has become imperative for future DSPs to put in place privacy by default mechanisms to ensure that RTS approaches operate in compliance with privacy preservation regulations such as GDPR [35]. A new paradigm of Federated Machine Learning (FML) in which a group of data owners collectively train a model while keeping their data stored locally [36] has emerged to tackle this challenge. Secure protocols such as homogeneous encryption, differential privacy, and secret sharing are used to ensure that the information actually shared among data federation participants (e.g., model loss and gradients) does not breach user privacy [37]. Open source codes for FML are available [38], [39] and an FML Standard with the IEEE (IEEE P3652.1) is being developed. We plan to carry out further research in order to enable the proposed MR bidding strategy to operate under the FML paradigm to offer better privacy protection for users.

REFERENCES

[1] J. Shen, B. Orten, S. C. Geyik, D. Liu, S. Shariat, F. Bian, and A. Dasdan, "From 0.5 million to 2.5 million: Efficiently scaling up real-time bidding," in *Proc. IEEE Int. Conf. Data Mining (ICDM)*, Nov. 2015, pp. 973–978.

[2] W. Zhang, S. Yuan, and J. Wang, "Optimal real-time bidding for display advertising," in *Proc. 20th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2014, pp. 1077–1086.

[3] J. Wang and S. Yuan, "Real-time bidding: A new frontier of computational advertising research," in *Proc. 8th ACM Int. Conf. Web Search Data Mining*, 2015, pp. 415–416.

[4] J. Wang, W. Zhang, and S. Yuan, "Display advertising with real-time bidding (RTB) and behavioural targeting," 2016, *arXiv:1610.03013*. [Online]. Available: https://arxiv.org/abs/1610.03013

[5] B. Lucier, R. P. Leme, and E. Tardos, "On revenue in the generalized second price auction," in *Proc. 21st Int. Conf. World Wide Web*, 2012, pp. 361–370.

[6] C. Perlich, B. Dalessandro, R. Hook, O. Stitelman, T. Raeder, and F. Provost, "Bid optimizing and inventory scoring in targeted online advertising," in *Proc. 18th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2012, pp. 804–812.

[7] J. Xu, K.-C. Lee, W. Li, H. Qi, and Q. Lu, "Smart pacing for effective online ad campaign optimization," in *Proc. 21th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2015, pp. 2217–2226.

[8] W. Zhang, Y. Rong, J. Wang, T. Zhu, and X. Wang, "Feedback control of real-time display advertising," in *Proc. 9th ACM Int. Conf. Web Search Data Mining*, 2016, pp. 407–416.

[9] H. Cai, K. Ren, W. Zhang, K. Malialis, J. Wang, Y. Yu, and D. Guo, "Real-time bidding by reinforcement learning in display advertising," in *Proc. 10th ACM Int. Conf. Web Search Data Mining*, 2017, pp. 661–670.

[10] W. Zhang and J. Wang, "Statistical arbitrage mining for display advertising," in *Proc. 21th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2015, pp. 1465–1474.

[11] K. Ren, W. Zhang, K. Chang, Y. Rong, Y. Yu, and J. Wang, "Bidding machine: Learning to bid for directly optimizing profits in display advertising," *IEEE Trans. Knowl. Data Eng.*, vol. 30, no. 4, pp. 645–659, Nov. 2018.

[12] H. Zhu, J. Jin, C. Tan, F. Pan, Y. Zeng, H. Li, and K. Gai, "Optimized cost per click in taobao display advertising," in *Proc. 23rd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2017, pp. 2191–2200.

[13] Y. Cui, R. Zhang, W. Li, and J. Mao, "Bid landscape forecasting in online ad exchange marketplace," in *Proc. 17th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2011, pp. 265–273.

[14] D. Agarwal, A. Z. Broder, D. Chakrabarti, D. Diklic, V. Josifovski, and M. Sayyadian, "Estimating rates of rare events at multiple resolutions," in *Proc. 13th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2007, pp. 16–25.

[15] H. B. McMahan, G. Holt, D. Sculley, M. Young, D. Ebner, J. Grady, L. Nie, T. Phillips, E. Davydov, D. Golovin, S. Chikkerur, D. Liu, M. Wattenberg, A. M. Hrafnkelsson, T. Boulos, and J. Kubica, "Ad click prediction: A view from the trenches," in *Proc. 19th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2013, pp. 1222–1230.

[16] K.-C. Lee, B. Orten, A. Dasdan, and W. Li, "Estimating conversion rate in display advertising from past erformance data," in *Proc. 18th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2012, pp. 768–776.

[17] D. Agarwal, R. Agrawal, R. Khanna, and N. Kota, "Estimating rates of rare events with multiple hierarchies through scalable log-linear models," in *Proc. 16th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2010, pp. 213–222.

[18] O. Chapelle, E. Manavoglu, and R. Rosales, "Simple and scalable response prediction for display advertising," *ACM Trans. Intell. Syst. Technol.*, vol. 5, no. 4, p. 61, 2015.

[19] A.-P. Ta, "Factorization machines with follow-the-regularized-leader for CTR prediction in display advertising," in *Proc. IEEE Int. Conf. Big Data (Big Data)*, Oct. 2015, pp. 2889–2891.

[20] J. Pan, J. Xu, A. L. Ruiz, W. Zhao, S. Pan, Y. Sun, and Q. Lu, "Field-weighted factorization machines for click-through rate prediction in display advertising," in *Proc. World Wide Web Conf.*, 2018, pp. 1349–1357.

[21] R. Wang, B. Fu, G. Fu, and M. Wang, "Deep & cross network for ad click predictions," in *Proc. ADKDD*, 2017, p. 12.

[22] X. Li and D. Guan, "Programmatic buying bidding strategies with win rate and winning price estimation in real time mobile advertising," in *Proc. Pacific–Asia Conf. Knowl. Discovery Data Mining*. Springer, 2014, pp. 447–460.
[23] W. C.-H. Wu, M.-Y. Yeh, and M.-S. Chen, "Predicting winning price in real time bidding with censored data," in *Proc. 21th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2015, pp. 1305–1314.
[24] B. Edelman, M. Ostrovsky, and M. Schwarz, "Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords," *Amer. Econ. Rev.*, vol. 97, no. 1, pp. 242–259, 2007.
[25] J. Xu, X. Shao, J. Ma, K.-C. Lee, H. Qi, and Q. Lu, "Lift-based bidding in ad selection," in *Proc. AAAI*, 2016, pp. 651–657.
[26] W. Zhang, T. Zhou, J. Wang, and J. Xu, "Bid-aware gradient descent for unbiased learning with censored data in display advertising," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2016, pp. 665–674.
[27] D. Agarwal, S. Ghosh, K. Wei, and S. You, "Budget pacing for targeted online advertisements at linkedin," in *Proc. 20th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2014, pp. 1613–1619.
[28] K.-C. Lee, A. Jalali, and A. Dasdan, "Real time bid optimization with smooth budget delivery in online advertising," in *Proc. 7th Int. Workshop Data Mining Online Advertising*, 2013, p. 1.
[29] T. Maehara, A. Narita, J. Baba, and T. Kawabata, "Optimal bidding strategy for brand advertising," in *Proc. 27th Int. Joint Conf. Artif. Intell. (IJCAI)*, 2018, pp. 424–432. doi: 10.24963/ijcai.2018/59.
[30] K. Ren, W. Zhang, Y. Rong, H. Zhang, Y. Yu, and J. Wang, "User response learning for directly optimizing campaign performance in display advertising," in *Proc. 25th ACM Int. Conf. Inf. Knowl. Manage.*, 2016, pp. 679–688.
[31] E. Ikonomovska, S. Jafarpour, and A. Dasdan, "Real-time bid prediction using thompson sampling-based expert selection," in *Proc. 21th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2015, pp. 1869–1878.
[32] M. Richardson, E. Dominowska, and R. Ragno, "Predicting clicks: Estimating the click-through rate for new ads," in *Proc. 16th Int. Conf. World Wide Web*, 2007, pp. 521–530.
[33] H. Liao, L. Peng, Z. Liu, and X. Shen, "iPinYou global RTB bidding algorithm competition dataset," in *Proc. 8th Int. Workshop Data Mining Online Advertising*, 2014, pp. 1–6.
[34] R. Fletcher, *Practical Methods of Optimization*. Hoboken, NJ, USA: Wiley, 2013.
[35] EU. (2016). *REGULATION (EU) 2016/679 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)*. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT
[36] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, "Practical secure aggregation for privacy-preserving machine learning," in *Proc. CCS*, 2017, pp. 1175–1191.
[37] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Trans. Intell. Syst. Technol.*, vol. 10, no. 2, pp. 12:1–12:19, 2019.
[38] (2019). *Federated AI Technology Enabler (FATE)*. [Online]. Available: https://www.fedai.org/
[39] (2019). *Tensorflow Federated*. [Online]. Available: https://www.tensorflow.org/federated

TENGYUN WANG is currently pursuing the Ph.D. degree with the School of Software Engineering, South China University of Technology. His research interests are data mining, machine learning, and applied data science including computational advertising.

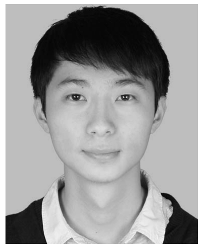

HAIZHI YANG is currently pursuing the Ph.D. degree with the School of Software Engineering, South China University of Technology. His research interests include machine learning, data mining, and applied data science including computational advertising.

HAN YU received the Ph.D. degree from the School of Computer Science and Engineering, Nanyang Technological University (NTU), Singapore, in 2014, where he is currently a Nanyang Assistant Professor (NAP). He held the prestigious Lee Kuan Yew Postdoctoral Fellowship (LKY PDF), from 2015 to 2018. He has published over 100 research papers in book chapters and leading international conferences and journals, including AAAI, IJCAI, AAMAS, MM, ACM/IEEE TRANSACTIONS, and *Scientific Reports*. His research focuses on real-time artificial intelligence (AI)-powered data-driven algorithmic management for complex collaborative networks. His research work has received 11 conference and journal awards (including from IJCAI and AAAI).

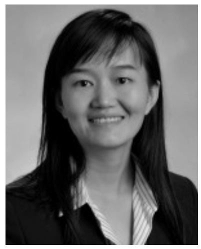

WENJUN ZHOU received the B.Sc. degree from the University of Science and Technology of China (USTC), in 2004, the M.Sc. degree from the University of Michigan, Ann Arbor, in 2006, and the Ph.D. degree from Rutgers University, in 2011. She is currently an Assistant Professor with the Department of Business Analytics and Statistics, Haslam College of Business, The University of Tennessee, Knoxville. Her general research interests include data mining, business analytics, and statistical computing. She was nominated as the R. Stanley Bowden II Faculty research fellow, in 2016, and the Roy & Audrey Fancher Faculty Research Fellow, in 2017. She was a recipient of the Best Paper Award at WAIM 2013, the Best Student Paper Runner-Up Award at KDD 2008, and the Best Paper Runner-Up Award at ICTAI 2006. She was among the top five finalists for the ACM SIGKDD Doctoral Dissertation Award, in 2011.

YANG LIU received the B.Sc. degree from Tsinghua University, in 2007, and the Ph.D. degree from Princeton University, in 2012. She is currently a Senior Researcher with the AI Department, WeBank, China. Her research has been published in leading scientific journals, such as ACM TIST and *Nature*. Her research interests include machine learning, federated learning, transfer learning, multi-agent systems, statistical mechanics, and applications of these technologies in the financial industry.

HENGJIE SONG received the B.Sc. degree in electrical engineering and the M.Sc. degree in aerospace computer design from the Harbin Institute of Technology, in 1998 and 2002, respectively, and the Ph.D. degree from Nanyang Technological University, Singapore. He is currently a Professor with the School of Software Engineering, South China University of Technology. He has published several high-quality papers on the best journals and conference proceedings, including the *ACM Transactions on Web*, the IEEE CIMs, the IEEE TRANSACTIONS ON FUZZY SYSTEMS, *Neural Networks* (Elsevier), AAAI, and ICDM. His research interest focuses on artificial intelligence and the applications of AI in commercial search engines.