Robust Autobidding for Noisy Conversion Prediction Models

Andrey Pudovikov
AI Center & IAI MSU
Lomonosov Moscow State University
Moscow, Russia
pudovikovad@my.msu.ru

Alexandra Khirianova
Lebedev Physical Institute of the Russian Academy of Sciences & AI Center & IAI MSU, Lomonosov Moscow State University
Moscow, Russia
khirianova.alexandra@gmail.com

Ekaterina Solodneva
AI Center
Lomonosov Moscow State University
Moscow, Russia
eksolodneva@gmail.com

Gleb Molodtsov
MBZUAI
Abu Dhabi, UAE
gleb.molodcov@gmail.com

Aleksandr Katrutsa
AI Center
Lomonosov Moscow State University
Moscow, Russia
amkatrutsa@gmail.com

Yuriy Dorn
AI Center & IAI MSU
Lomonosov Moscow State University
Moscow, Russia
dornyv@my.msu.ru

Egor Samosvat
Independent Researcher
Moscow, Russia
samosvat.egor@gmail.com

ABSTRACT

Managing millions of digital auctions is essential to modern advertising auction systems. The primary approach to managing digital auctions is autobidding, which relies on Click-Through Rate and Conversion Rate metrics. While these quantities are estimated with ML models, their prediction uncertainty directly impacts advertisers' revenue and bidding strategies. To address this issue, we propose `RobustBid`, an efficient method for robust autobidding taking into account uncertainty in CTR and CVR predictions. Our approach leverages advanced, robust optimization techniques to prevent large errors in bids if the estimates of CTR/CVR are perturbed. We derive an analytical solution to the stated robust optimization problem, which improves the runtime efficiency of the `RobustBid` method. The synthetic, iPinYou, and BAT benchmarks are used in our experimental evaluation of `RobustBid`. We compare our method with the non-robust baseline and the RiskBid algorithm using total conversion volume (TCV) and average cost-per-click (CPC_{avg}) as performance metrics. The experiments demonstrate that `RobustBid` provides bids that yield larger TCV and smaller CPC_{avg} than competitors in the case of large perturbations in CTR/CVR predictions.

KEYWORDS

Autobidding problem; robust optimization; uncertainty quantification of CTR model

ACM Reference Format:
Andrey Pudovikov, Alexandra Khirianova, Ekaterina Solodneva, Gleb Molodtsov, Aleksandr Katrutsa, Yuriy Dorn, and Egor Samosvat. 2026. Robust Autobidding for Noisy Conversion Prediction Models. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 9 pages. https://doi.org/10.65109/RXYW3025

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/RXYW3025

1 INTRODUCTION

On commercial electronic platforms, a list of relevant display content, including paid options, is generated to best meet the user's needs while also maximizing revenue. This content is typically selected and placed through online auctions [4], where bids are generated by automated algorithms. These algorithms operate within advertiser-specified constraints, analyzing past bidding data, assessing previous auction performance, and predicting future ad performance for specific requests. Central to this process are the tasks of the Click-Through Rate (CTR) prediction [21, 22] and the Conversion Rate (CVR) prediction [10, 25]. CTR measures the likelihood that a user will click an ad, which is economically vital for advertisers because it indicates the potential for a purchase. In turn, CVR is the probability of conversion after a click.

The design of an auction can significantly influence bidders' behavior. Current advertising systems typically employ cost-per-impression, cost-per-click, or cost-per-conversion [12, 26] as the payment rule, which constitutes the core of the auction mechanism. Advertisers subsequently focus on their own metrics to assess the effectiveness of their advertising campaigns. A classic approach is to maximize the number of ad impressions, clicks, or conversions. However, while bidders cannot directly influence the auction design, they can strategically select their objectives based on these performance metrics. This strategic adaptation highlights that payment rules and metrics are inherently context-dependent; their efficacy and suitability must be evaluated against the specific goals and challenges of each auction problem.

The main contributions of this work are the following.

(1) We propose a robust optimization problem to determine the bid formula, treating the uncertainty in CTR/CVR estimates.

(2) Based on the analytical solution of the stated robust optimization problem, we suggest the `RobustBid` method, which provides proper bids even in the case of high uncertainty.
(3) The extensive experimental evaluation of the `RobustBid` algorithm on synthetic and industrial datasets demonstrates improvements in conversion rates and reduction in the cost per click compared to the baselines.

Related works

Autobidding algorithms automatically determine bid prices to achieve advertisers' marketing objectives, adapting to market changes and growing data complexity [1, 2, 4, 11]. These systems enable advertisers to optimize spending within budget and cost-per-click constraints [8, 9, 27, 29]. We adopt the problem from [27], which considers online bidding where bids are determined sequentially for each auction without complete market information. This problem formulation is widely used in the autobidding literature [1, 2, 14, 18], making our approach broadly applicable. In this framework, optimal bids depend on predicted CTR and CVR values from machine learning models. However, these predictions are uncertain and can affect algorithm performance.

Robustness in bidding. Bidding algorithms are affected by the uncertainty from multiple sources. For example, strategic uncertainty arises when advertisers intentionally underbid to force the platform to lower the reserve price. The lower the reserve price, the lower the bid that will win the auction. Studies [3, 13, 17] develop robust algorithms that prevent strategic manipulation and maximize platform revenue. However, these approaches maximize platform revenue, whereas our work focuses on optimizing advertiser bidding. At the same time, uncertainty in opportunity values and win rates is addressed with robust optimization techniques in [20]. Similarly, uncertainty in competing bid distributions is discussed in [15], which proposes a min-max approach. In contrast to our approach, the latter two studies require estimating ad opportunity values.

Robustness of CTR prediction. Since autobidding algorithms typically use the CTR estimation to provide bids [1], uncertainty quantification of the predicted CTR values is crucial for proper bidding. [24] propose a clustering model, which groups bid requests with similar predicted CTRs and uses reinforcement learning to train a model for bid generation. This clustering model shows empirical improvements but lacks robustness guarantees. [30] introduce a risk management framework that incorporates prediction uncertainty via modeling bid proportional to the weighted sum of CTR and its standard deviation. This approach has theoretical foundations but requires tuning of weights. At the same time, [5] introduce robustness for valuation vectors by using a central limit theorem and induced confidence intervals. However, their analysis was restricted to the offline setup, which limits its applicability to real-world systems. For the best of our knowledge, our approach is the first to handle uncertainties in both CTR and CVR values. We summarize the comparison of our work with the previous studies in Table 1.

Table 1: Comparison of the proposed approach with existing alternatives. We highlight that our approach supports uncertainties in multiple sources, includes an uncertainty set, and is evaluated on real-world datasets.

Reference	Multiple uncertainty sources	Uncertainty set	Real-world data
[27]	✗	✗	✓
[5]	✗	✓	✗
[30]	✗	✓	✓
This paper	✓	✓	✓

2 PROBLEM STATEMENT

This section presents the offline non-robust and robust optimization problems for identifying bids in cases where the CTR/CVR values are exact and perturbed, respectively.

2.1 Non-robust maximization of expected number of conversions

Consider a sequence of T auctions involving I items, where each item participates in every auction with corresponding bids denoted as bid_t^i. In this formulation, we assume each advertiser has only one item, i.e., these terms are equivalent.

This bid depends on CTR value CTR_t^i and conversion value CVR_t^i for $i \in [1, I]$ and $t \in [1, T]$. This dependence for the i-th item is modeled by solving the specific maximization problem described below. The objective in this maximization problem is the expected number of conversions for the i-th item, which is estimated as

$$\sum_{t=1}^{T} x_t^i \cdot CTR_t^i \cdot CVR_t^i \qquad (1)$$

where $x_i^t = \begin{cases} 1, & \text{if item } i \text{ wins auction } t; \\ 0, & \text{otherwise.} \end{cases}$. Upon winning an auction t by submitting the highest bid, the advertiser pays a winning price wp_t, which depends on the auction type, e.g., first-price or second-price auctions [23]. At the same time, the i-th advertiser has a pre-defined budget B^i, which induces the natural constraint:

$$\sum_{t=1}^{T} x_t^i \cdot wp_t \le B^i, \qquad (2)$$

where the left-hand side corresponds to the expected spends from the wins of auctions. In addition, the i-th advertiser can bound the total cost per click by a constant C^i, which leads to the following CPC constraint:

$$\frac{\sum_{t=1}^{T} x_t^i \cdot wp_t}{\sum_{t=1}^{T} x_t^i \cdot CTR_t^i} \le C^i, \qquad (3)$$

where the left-hand side indicates the expected average cost per click. In (3), we estimate expected number of clicks as $\sum_{t=1}^{T} x_t^i \cdot CTR_t^i$. Thus, maximization of the objective function (1) subject to constraints (2) and (3) with binary variables x_i^t is an integer linear programming problem, which is hard to solve in a large-dimensional case [27]. To approximate the solution, studies [1, 27] propose relax variables x_i^t such that $x_i^t \in [0, 1]$. Therefore, the resulting non-robust maximization problem is stated as follows

$$\begin{aligned}\max_{0\le x_t^i\le 1} \quad & \sum_{t=1}^{T} x_t^i \cdot CTR_t^i \cdot CVR_t^i \\ \text{s.t.} \quad & \sum_{t=1}^{T} x_t^i \cdot wp_t^i \le B^i \\ & \frac{\sum_{t=1}^{T} x_t^i \cdot wp_t^i}{\sum_{t=1}^{T} x_t^i \cdot CTR_t^i} \le C^i.\end{aligned} \tag{4}$$

In problem (4), we assume that the winning price is known and corresponds to the used auction type. The main requirement for the auction is that the auction winner sets the highest bid. As a consequence, the highest bid is the upper bound for the winning price wp_t^i. If the auction mechanism satisfies this requirement, then it can be used within our framework.

To derive the dependence of bid bid_t^i on CTR/CVR, study [27] considers dual problem to (4) and obtains the following equation:

$$bid_t^i = \frac{1}{p+q} CVR_t^i \cdot CTR_t^i + \frac{q}{p+q} C^i \cdot CTR_t^i, \tag{5}$$

where $p, q > 0$ are optimal dual variables corresponding to constraints (2) and (3), respectively. See more details in the Appendix E in [19]. Equation (5) demonstrates that perturbation in CTR/CVR values significantly affects the bid value. Since CTR_i^t and CVR_i^t are typically estimated by machine learning models in practice [28], these values are not known precisely. Therefore, a natural question is how to make the bid values robust to the uncertainty in CTR_i^t and CVR_i^t. The following section proposes the robust modification of problem (4) based on a robust optimization framework.

2.2 Robust conversion maximization problem from CTR/CVR uncertainty

Accurate estimations of CTR and CVR are crucial for deriving a proper solution of (4) since even a small perturbation of these values can significantly degrade the performance of the resulting bid (5). To mitigate such effects, we apply the robust optimization framework [6] to the autobidding problem. The robust optimization framework requires introducing an uncertainty set $\mathcal{U}$ that captures feasible perturbations to the CTR/CVR values. For simplicity, we consider CTR values; CVR can be used instead, and our approach processes them correctly. Since we assume that CTR values are estimated by a pre-trained ML model f, we use the corresponding loss function L in the definition of $\mathcal{U}$. We assume that, to fit the model f, one uses training and validation datasets consisting of ground-truth CTRs, denoted by CTR.

In particular, the loss function L measures for the i-th advertiser how close the predicted CTR $\widehat{CTR}_t^i$ is to the ground-truth CTR $\mathbf{CTR}^i = [CTR_t^i]$, where $i = 1, \dots, I$ and $t = 1, \dots, T$. Therefore, we can define an uncertainty set for the i-th advertiser as follows:

$$\mathcal{U}^i = \left\{ \mathbf{a} = [a_t]_{t=1}^T, a_t \in [0,1] \mid L(\widehat{\mathbf{CTR}}^i, \mathbf{a}) \le \varepsilon \right\}, \tag{6}$$

where $L(\widehat{\mathbf{CTR}}^i, \mathbf{a}) = \sum_{t=1}^T L(\widehat{CTR}_t^i, a_t)$, $\mathbf{a}$ is a vector of feasible CTR values and $\varepsilon > 0$ is a predefined threshold. We assume that $\mathbf{CTR}^i \in \mathcal{U}^i$, $i = 1, \dots, I$ for some ε_0. According to [7], if ε_0 is empirical $(1+\frac{1}{I})(1-q)$-quantile of $L(\widehat{\mathbf{CTR}}_{val}^i, \mathbf{CTR}^i)$, where $q \in (0,1)$, then with probability at least q the following inequality holds:

$$L(\widehat{\mathbf{CTR}}_{val}^i, \mathbf{CTR}^i) \le \varepsilon_0. \tag{7}$$

To compute ε_0, we use $\widehat{\mathbf{CTR}}_{val}^i$, which is the CTR values estimated by the pre-trained model f on the validation dataset. While ε_0 is computed on the validation dataset, we will use it in the general problem (4) under the assumption of data homogeneity.

Since the loss function L and uncertainty set (6) are independent of each item i, the item index is further omitted to simplify expressions. Following the robust optimization framework, we state the robust modification of the non-robust problem (4):

$$\begin{aligned}\max_{0\le x_t\le 1} \min_{\mathbf{a}\in\mathcal{U}} & \sum_{t=1}^{T} x_t \cdot a_t \cdot CVR_t \\ \text{s.t.} \ & \sum_{t=1}^{T} x_t \cdot wp_t \le B \\ & \sum_{t=1}^{T} x_t \cdot wp_t \le C \sum_{t=1}^{T} x_t \cdot a_t\end{aligned} \tag{8}$$

The main difference between the non-robust problem (4) and the robust problem (8) is the objective function and uncertainty set for the CTR predictions. Note that this worst-case robust autobidding problem (8) depends on the choice of loss function. The following section presents the analytical solution of problem (8) and generalizes it to the case of joint uncertainty in CTR and CVR.

3 ROBUSTBID METHOD

This section presents the RobustBid method, which is based on the analytical solution of the robust optimization problem (8) for a particular uncertainty set. The explicit form of uncertainty set $\mathcal{U}$ depends on the choice of the loss function L, see (6). The straightforward loss function for the CTR prediction problem is cross-entropy; however, it yields an analytically intractable optimization problem. To make the problem analytically tractable and still estimate the perturbation of the predicted CTR values $\widehat{\mathbf{CTR}}$, we select MSE loss to construct the uncertainty set

$$\mathcal{U}_{mse} = \left\{ \mathbf{a} = [a_t]_{t=1}^T, a_t \in [0,1] \mid \frac{1}{2} \left\| \widehat{\mathbf{CTR}} - \mathbf{a} \right\|_2^2 \le \varepsilon_a \right\}. \tag{9}$$

Incorporating the uncertainty set $\mathcal{U}_{mse}$ in problem (8) makes it possible to solve this problem analytically. If instead of uncertainty in CTR estimates one considers uncertainty in CVR estimates, then the derived formulas can be used, see details in Appendix C in [19]. Moreover, the RobustBid method supports uncertainty in *both* CTR and CVR estimates. The analytical solution to the corresponding robust optimization problem is presented in this section. Thus, the uncertainty set $\mathcal{U}_{mse}$ provides a reasonable trade-off between the tractability of the resulting robust optimization problem and the correctness of the uncertainty measurements.

3.1 Bidding formulas for individual CTR/CVR uncertainties

This section presents the analytical bidding formulas for cases of uncertainty in CTR and CVR values, separately. We derive the bidding formula for the case of uncertainty in CTR from the analytical solution of (8), where $\mathcal{U} = \mathcal{U}_{mse}$. To solve this problem, we introduce a slack variable s such that the original problem is rewritten as follows

$$\begin{aligned} &\max_{0 \le x_t \le 1, s} s \\ \text{s.t.} \ &\sum_{t=1}^{T} x_t \cdot wp_t \le B, \quad \sum_{t=1}^{T} x_t \cdot wp_t \le C \sum_{t=1}^{T} x_t \cdot a_t \\ &s \le \min_{\mathbf{a} \in \mathcal{U}_{mse}} \sum_{t=1}^{T} x_t \cdot a_t \cdot CVR_t. \end{aligned} \tag{10}$$

The minimization problem in the last constraint assumes that x_t and CVR_t are known and only the vector $\mathbf{a}$ is unknown. Therefore, this problem has a linear objective function in the form $\mathbf{m}^\top \mathbf{a}$, where $m_t = x_t \cdot CVR_t$, and quadratic inequality constraint $\frac{1}{2}\|\widehat{\mathbf{CTR}} - \mathbf{a}\|_2^2 \le \varepsilon_a$. We eliminate inequality constraints of the form $0 \le a_t \le 1$ to make this problem analytically tractable and further discuss the cases in which they will be satisfied automatically. Since the resulting problem is convex, we use KKT optimality conditions and derive the optimal solution $\tilde{\mathbf{a}}$ such that:

$$\tilde{\mathbf{a}} = \widehat{\mathbf{CTR}} - \frac{\mathbf{m}}{\|\mathbf{m}\|_2}\alpha, \quad \alpha = \sqrt{2\varepsilon_a}. \tag{11}$$

The detailed derivation of (11) is presented in Appendix A in [19]. Note that $\tilde{a}_t \le 1$ since $CTR_t \le 1$ and $\alpha, m_t \ge 0$. To satisfy constraint $\tilde{a} \ge 0$ we set ε_a such that

$$\varepsilon_a \le \varepsilon_m = \min_t \frac{\widehat{CTR_t} \cdot \|\mathbf{CVR}\|_2}{CVR_t}. \tag{12}$$

The estimated $\varepsilon_a = \varepsilon_m$ can lead to excessive pessimistic probabilistic guarantees according to (7).

After that, we substitute $\tilde{\mathbf{a}}$ in problem (10) and denote $\alpha = \sqrt{2\varepsilon_a}$ to get the following optimization problem with linear objective function and convex quadratic constraints:

$$\begin{aligned} &\max_{\mathbf{x},\mathbf{y},s} s \\ \text{s.t.} \ &y_t = x_t \cdot CVR_t \quad \forall t \\ &s \le \mathbf{y}^\top \widehat{\mathbf{CTR}} - \alpha\|\mathbf{y}\|_2 \\ &\sum_{t=1}^{T} x_t \cdot wp_t \le B \\ &\sum_{t=1}^{T} x_t \cdot wp_t \le C \cdot [\mathbf{x}^\top \widehat{\mathbf{CTR}} - \alpha\|\mathbf{x}\|_2] \\ &0 \le x_t \le 1, \forall t \end{aligned} \tag{13}$$

Note that quadratic constraints from this maximization problem, including the 2-norm of vectors $\mathbf{x}$ and $\mathbf{y}$, can be expressed as cone constraints. In particular, denote by $K_2 = \{(\mathbf{u}, v) \mid ||\mathbf{u}||_2 \le v\}$ the cone induced by 2-norm, then we have the following equivalence:

$$\alpha||\mathbf{y}||_2 \le -s + \mathbf{y}^\top \widehat{\mathbf{CTR}} \Leftrightarrow \begin{pmatrix} \alpha \mathbf{y} \\ -s + \mathbf{y}^\top \widehat{\mathbf{CTR}} \end{pmatrix} \in K_2$$

$$C\alpha\|\mathbf{x}\|_2 \le \mathbf{x}^\top (C \cdot \widehat{\mathbf{CTR}} - \mathbf{wp}) \Leftrightarrow \begin{pmatrix} C\alpha \mathbf{x} \\ \mathbf{x}^\top (C \cdot \widehat{\mathbf{CTR}} - \mathbf{wp}) \end{pmatrix} \in K_2,$$

where $\mathbf{wp} = [wp_1, \dots, wp_T]$ is a vector of winning prices. The resulting cone maximization problem can be solved analytically with KKT optimality conditions. Moreover, the dual problem can also be derived and solved in closed form. Thus, we can compute the optimal dual variables p^*_{ctr} and q^*_{ctr} associated with the budget and CPC constraints. Similar to the approach from [27], we use p^*_{ctr}, q^*_{ctr} to construct the bid formula:

$$bid_t = \frac{1}{p^*_{ctr} + q^*_{ctr}} CVR_t \cdot \widehat{CTR_t} + \frac{q^*_{ctr}}{p^*_{ctr} + q^*_{ctr}} C \cdot \widehat{CTR_t} + \delta^{ctr}_t, \tag{14}$$

where $\delta^{ctr}_t = \begin{cases} -\frac{\alpha}{p^*_{ctr}+q^*_{ctr}} \left(\frac{Cq^*_{ctr}}{\sqrt{|\mathcal{T}|}} + \frac{CVR_t^2}{\sqrt{\sum_{l \in \mathcal{T}} CVR_l^2}} \right) & t \in \mathcal{T}, \\ 0 & t \notin \mathcal{T}, \end{cases}$ is the additive perturbation to (5) and a set of timestamps $\mathcal{T}$ is the following

$$\mathcal{T} = \{t \mid CTR_t \cdot CVR_t + q^*_{ctr} \cdot C \cdot CTR_t \le (p^*_{ctr} + q^*_{ctr}) wp_t\}. \tag{15}$$

If only uncertainty in CVR is assumed, then the corresponding uncertainty set $\mathcal{V}_{mse}$ can be defined as

$$\mathcal{V}_{mse} = \left\{ \mathbf{a} = [a_t]_{t=1}^{T}, a_t \in [0, 1] \mid \frac{1}{2} \left\| \widehat{\mathbf{CVR}} - \mathbf{a} \right\|_2^2 \le \varepsilon_b \right\}. \tag{16}$$

In this case, the robust version of problem (4) becomes easier to solve since CVR uncertainty affects only the objective function and does not modify constraints. The corresponding dual problem also has an analytical solution, which includes optimal dual variables p^*_{cvr} and q^*_{cvr}. Following the same steps as in the derivation of (5) and (14), these dual variables can be used to construct the bid formula robust to CVR perturbation:

$$bid_t = \frac{1}{p^*_{cvr} + q^*_{cvr}} \widehat{CVR_t} \cdot CTR_t + \frac{q^*_{cvr}}{p^*_{cvr} + q^*_{cvr}} C \cdot CTR_t + \delta^{cvr}_t, \tag{17}$$

where $\delta^{cvr}_t = \begin{cases} -\frac{\alpha'}{p^*_{cvr}+q^*_{cvr}} \frac{CTR_t}{\sqrt{\sum_{l \in \mathcal{T}} CTR_l^2}}, & t \in \mathcal{T} \\ 0, & t \notin \mathcal{T}. \end{cases}$, $\alpha' = \sqrt{2\varepsilon_b}$ and $\mathcal{T}$ is defined in (15). The detailed derivaton of (17) is presented in Appendix C in [19]. The final bid formula for CVR uncertainties (17) looks similar to the bid formula for CTR uncertainties (14) except for the form of the perturbation item, where the CVR is replaced with CTR and the term related to the CPC constraint is dropped.

3.2 Bidding formula for joint CTR and CVR uncertainties

This section extends a robust optimization approach to address uncertainties in *both* CTR and CVR predictions simultaneously. This setup represents a more comprehensive modeling of uncertainties in the bidding process. Since we have introduced uncertainty sets for CTR $\mathcal{U}_{mse}$ and for CVR $\mathcal{V}_{mse}$ in Section 3.1, we can state the robust optimization problem which treats the Cartesian product of these uncertainty sets:

$$\begin{aligned} &\max_{0 \le x_t \le 1} \min_{(\mathbf{a},\mathbf{b}) \in \mathcal{U}_{mse} \times \mathcal{V}_{mse}} \sum_{t=1}^{T} x_t \, a_t \, b_t; \\ \text{s.t.} \ &\sum_{t=1}^{T} x_t \, wp_t \le B \\ &\frac{\sum_{t=1}^{T} x_t \, wp_t}{\sum_{t=1}^{T} x_t \, a_t} \le C. \end{aligned} \tag{18}$$

This problem is more complex than the single-uncertainty case, as it involves interactions between uncertainties in CTR and CVR.

To solve problem (18), we use advanced techniques from convex optimization theory, including the S-lemma and Schur complement theorem. Below, we provide the main steps to derive a robust optimal bid formula (for more details, see Appendix D in [19]).

To simplify notations, we denote $a_t^0 = \widehat{\mathrm{CTR}}_t, b_t^0 = \widehat{\mathrm{CVR}}_t$, then introduce perturbations $\delta a_t = a_t - a_t^0,\ \delta b_t = b_t - b_t^0$ and obtain new form for inequality in uncertainty set

$$\|\delta\mathbf{a}\|_2^2 \le r_a^2 = 2\varepsilon_a, \|\delta\mathbf{b}\|_2^2 \le r_b^2 = 2\varepsilon_b. \tag{19}$$

The objective is rewritten as follows:

$$\sum_{t=1}^{T} x_t\, a_t\, b_t = \sum_t x_t\, (a_t^0 + \delta a_t)(b_t^0 + \delta b_t) = \mathbf{z}^\top \mathbf{Q}\mathbf{z} + 2\mathbf{q}^\top \mathbf{z} + c,$$

where

$$\mathbf{z} = \begin{pmatrix} \delta\mathbf{a} \\ \delta\mathbf{b} \end{pmatrix}, \quad \mathbf{D} = \mathrm{diag}(x_1, \dots, x_T), \quad \mathbf{Q} = \begin{pmatrix} 0 & \frac{1}{2}\mathbf{D} \\ \frac{1}{2}\mathbf{D} & 0 \end{pmatrix},$$

$$\mathbf{q} = \frac{1}{2}\begin{pmatrix} \mathbf{D}\,\mathbf{b}^0 \\ \mathbf{D}\,\mathbf{a}^0 \end{pmatrix}, \quad c = \sum_{t=1}^{T} x_t\, a_t^0\, b_t^0.$$

Then we introduce a slack scalar variable s similar to Section 3.1:

$$s \le \min_{\|\delta\mathbf{a}\|_2 \le r_a, \|\delta\mathbf{b}\|_2 \le r_b} (\mathbf{z}^\top \mathbf{Q}\mathbf{z} + 2\mathbf{q}^\top \mathbf{z} + c)$$

which is equivalent to

$$\mathbf{z}^\top \mathbf{Q}\mathbf{z} + 2\mathbf{q}^\top \mathbf{z} + (c - s) \ \ge 0 \quad \forall z.$$

To reformulate constraints (19) in the quadratic forms with positive semi-definite matrices, we use the S-lemma. For this purpose, we rewrite inequalities:

$$g_1(\mathbf{z}) \equiv r_a^2 - \|\delta\mathbf{a}\|_2^2 = r_a^2 - \delta\mathbf{a}^T \mathbf{I}\delta\mathbf{a} \ \ge 0,$$

$$g_2(\mathbf{z}) \equiv r_b^2 - \|\delta\mathbf{b}\|_2^2 = r_b^2 - \delta\mathbf{b}^T \mathbf{I}\delta\mathbf{b} \ge 0$$

and apply the inhomogeneous S-lemma [6]. Then,

$$g_1(\mathbf{z}) \ge 0,\ g_2(\mathbf{z}) \ge 0 \ \Rightarrow\ \mathbf{z}^\top \mathbf{Q}\mathbf{z} + 2\mathbf{q}^\top \mathbf{z} + (c - s) \ge 0 \quad \forall \mathbf{z} \tag{20}$$

holds if and only if there exist $\lambda_a, \lambda_b \ge 0$ such that

$$\mathbf{M} = \begin{pmatrix} \mathbf{Q} + \mathrm{diag}(\lambda_a \mathbf{I}, \lambda_b \mathbf{I}) & \mathbf{q} \\ \mathbf{q}^\top & (c - s - \lambda_a r_a^2 - \lambda_b r_b^2) \end{pmatrix} \succeq 0.$$

To simplify this constraint, we apply the Schur complement theorem [6], so $\mathbf{M} \succeq 0$ is equivalent to the two convex conditions:

$$\begin{pmatrix} \lambda_a \mathbf{I}_T & \frac{1}{2}\mathbf{D} \\ \frac{1}{2}\mathbf{D} & \lambda_b \mathbf{I}_T \end{pmatrix} \succeq 0$$

$$(c - s - \lambda_a r_a^2 - \lambda_b r_b^2) - \mathbf{q}^\top (\mathbf{Q} + \mathrm{diag}(\lambda_a \mathbf{I}, \lambda_b \mathbf{I}))^{-1} \mathbf{q} \ge 0.$$

By combining the fact that $(\mathbf{Q} + \mathrm{diag}(\lambda_a \mathbf{I}, \lambda_b \mathbf{I}))^{-1} \succeq 0$, since inversion preserves positive semi-definiteness, and applying the definitions of $\mathbf{Q}$ and $\mathbf{q}$, we obtain:

$$\begin{pmatrix} \lambda_a \mathbf{I_T} & \frac{1}{2}\mathbf{D} \\ \frac{1}{2}\mathbf{D} & \lambda_b \mathbf{I_T} \end{pmatrix} \succeq 0 \quad \text{and} \quad -\lambda_a r_a^2 - \lambda_b r_b^2 + c - s \ge 0.$$

Here, the first is a $2T \times 2T$ linear matrix inequality, and the second is an affine bound. However, the positive semi-definite constraint can be simplified by the Schur complement theorem and the diagonal structure of D, we obtain:

$$\begin{cases} \lambda_a \ge 0, \\ \lambda_a \lambda_b \ge \frac{1}{4} x_t^2 \quad \forall t \in 1, \dots, T. \end{cases}$$

Therefore, the resulting simplified problem can be expressed as:

$$\begin{aligned} \max_{x_1,\dots,x_T,\lambda_a,\lambda_b} \quad & -\lambda_a\, r_a^2 - \lambda_b\, r_b^2 + \sum_{t=1}^{T} x_t\, a_t^0\, b_t^0 - \\ & - \sum_{t=1}^{T} \frac{2x_t^2[\lambda_b (b_t^0)^2 + \lambda_a (a_t^0)^2] - 2x_t^3 a_t^0 b_t^0}{4\lambda_a \lambda_b - x_t^2} \\ \text{s.t.} \quad & \sum_{t=1}^{T} x_t\, wp_t \le B, \\ & \sum_{t=1}^{T} x_t\, wp_t - C \sum_{t=1}^{T} x_t\, a_t^0 + C\, r_a\, \|x\|_2 \le 0, \\ & \lambda_a \lambda_b \ge \frac{1}{4} x_t^2, \\ & \lambda_a \ge 0, \lambda_b \ge 0, 0 \le x_t \le 1 \end{aligned} \tag{21}$$

The solution of problem (21) and the solution of the corresponding dual problem lead to the optimal bidding formula in the case of joint CTR and CVR uncertainties:

$$bid_t = \frac{1}{p_{joint}^* + q_{joint}^*} \widehat{CVR}_t \cdot \widehat{CTR}_t + \frac{q_{joint}^*}{p_{joint}^* + q_{joint}^*} C \cdot \widehat{CTR}_t + \delta_t^{joint}, \tag{22}$$

where bid perturbation is expressed as

$$\delta_t^{joint} = \begin{cases} -\frac{\alpha}{p_{joint}^* + q_{joint}^*} \left(\frac{q_{joint}^*}{\sqrt{|\mathcal{T}|}} + A(\lambda_a^*, \lambda_b^*, \widehat{CTR}_t, \widehat{CVR}_t) \right) & t \in \mathcal{T}, \\ 0 & t \notin \mathcal{T}, \end{cases},$$

where p_{joint}^* and q_{joint}^* are optimal dual variables corresponding to the CPC and budget constraints, λ_a^*, λ_b^* are optimal dual variables corresponding to the uncertainty bounds, and

$$\mathcal{T} = \{t \mid CTR_t CVR_t + Cq\mathrm{CTR}_t - A(\lambda_a, \lambda_b, \widehat{CTR}_t, \widehat{CVR}_t) \le (p+q) wp_t\} \tag{23}$$

is the set of active indices and helpful $A(\lambda_a, \lambda_b, \widehat{CTR}_t, \widehat{CVR}_t)$ is defined as:

$$\begin{aligned} A(\lambda_a^*, \lambda_b^*, \widehat{CTR}_t, \widehat{CVR}_t) = & \frac{4[\lambda_a^* \widehat{CTR}_t^2 + \lambda_b^* \widehat{CVR}_t^2] - 6\widehat{CTR}_t \widehat{CVR}_t}{(4\lambda_a^* \lambda_b^* - 1)} + \\ & + \frac{4(\lambda_a^* \widehat{CTR}_t^2 + \lambda_b^* \widehat{CVR}_t^2 - \widehat{CTR}_t \widehat{CVR}_t)}{(4\lambda_a^* \lambda_b^* - 1)^2} \end{aligned} \tag{24}$$

The detailed derivation of this formula is in Appendix D in [19].

Equation (22) requires the exact values of the $p_{joint}^*, q_{joint}^*, \lambda_a^*, \lambda_b^*$; therefore, the similar approach from [27] is used: determine these variables on the history and apply them to the next steps. Note that this bid formula has a similar structure to the bid formulas (14) and (17) corresponding to the individual uncertainties in CTR and CVR. However, equation (22) includes an additional term that captures the interaction between CTR and CVR uncertainties.

4 NUMERICAL EXPERIMENTS

This section evaluates the performance of the `RobustBid` algorithm through experiments on both synthetic and real-world datasets. We compare our approach against two baselines: the `NonRobustBid` algorithm [27] based on (5), and the RiskBid algorithm [30]. All results are reproducible using the code available in repository[1].

4.1 Datasets and environment design

To simulate real-world uncertainty in CTR predictions, various levels of noise ε^a and ε^b are introduced into CTR and CVR data. The variation of ε is naturally carried out in permissible normalization boundaries, upper bounds on budget B and cost per click C are chosen so that equation (4) has a finite solution (see values of B and C in Appendix E [19]).

Advertisers have access to their historical data, including bids, winning prices, wins, accumulated clicks, and spending. Before placing a bid, they can also access the current values $\widehat{CTR_t^i}$ and $\widehat{CVR_t^i}$. The auction outcomes and click events are determined by the ground-truth values CTR_t^i and CVR_t^i. Moreover, the winning prices are not allowed before the auction ends. Therefore, an advertiser can obtain winning prices from historical data (all auctions he has participated in). For the autobidding task, parameters are updated at each step by solving a constrained nonlinear optimization problem (8) using the advertiser's full historical data. When a bid is placed, the advertiser is charged the minimum of the full bid amount and the remaining budget. The experimental evaluation was carried out using synthetic and industrial datasets. The $\varepsilon_a, \varepsilon_b$ lie in range $[10^{-6}, 10^{-2}]$ for all experiments.

Synthetic dataset. This dataset is used to evaluate performance under controlled uncertainties in CTR and CVR values. The first-price auction is simulated over $T = 100$ auctions with $n = 10$ participating advertisers. The ground-truth CTR values CTR_t^i and CVR values CVR_t^i are random from the interval $(0.01, 0.1]$. To verify the stability and consistency of results, we use a large number of auctions $T = 500, 1000$ and average results across 100 random seeds.

iPinYou dataset. The first-price auction environment consists of $n = 10$ advertisers competing over $T = 100$ auctions. For the first 9 advertisers, bid distributions are derived from the iPinYou dataset. Since the original dataset contains only 2-3 discrete bid values, kernel density estimation is applied to generate a continuous bid distribution that better reflects real-world behavior. The ground-truth CTR_t^i and CVR_t^i values and their perturbations are sampled in the same way as in the synthetic dataset. The 10th advertiser follows the proposed bidding strategy, with parameters adjusted based on bid statistics from the iPinYou dataset.

BAT dataset. The study [16] proposed the new BAT dataset for benchmarking autobidding algorithms. The dataset consists of thousands of First-Price Auction records from real advertising campaigns. Unlike the synthetic and iPinYou datasets, where auctions occur simultaneously, auctions in BAT are distributed across time in line with real-world patterns. The range of CTR_t^i values in this dataset is $[0.0017, 0.63]$, and the range of CVR_t^i values is $[0.001, 0.3]$.

4.2 Baselines

NonRobustBid. To address the issue of non-robustness, the dual problem formulation proposed in [27] is adopted (see Appendix E in [19] for details). Following the approach outlined in the original paper, there are two adjusted dual variables corresponding to the budget and CPC constraints. The baseline solution is obtained by solving the original optimization problem exactly. Thus, we compare the robust solution against the non-robust one to demonstrate the specific benefits of incorporating uncertainty handling.

RiskBid [30]. This method is designed to enhance bidding robustness by incorporating an estimate of the CTR standard deviation, CTR_{std}, derived from historical data, into the bidding formula (5): $CTR_{risk} = CTR - \alpha \cdot CTR_{std}$, where α is a hyperparameter.

4.3 Metrics

To evaluate the performance of the considered autobidding algorithms, we use the following metrics. *Total Conversion Value* (TCV) quantifies the aggregate conversion performance across all advertising campaigns and auctions, computed as:

$$\text{TCV} = \sum_{t=1}^{T}\sum_{i=1}^{I} CTR_t^i \cdot CVR_t^i \cdot x_t^i,$$

where CTR_t^i and CVR_t^i represent the click-through and conversion rates respectively for campaign i at time t, and x_t^i a binary indicator of auction win. *Average Cost-Per-Click* (CPC_{avg}) represents the average cost per click across all advertising campaigns. It represents the ratio of total spending to total clicks across all campaigns:

$$CPC_{\text{avg}} = \frac{\sum_{i=1}^{I}\sum_{t=1}^{T} x_t^i \cdot bid_t^i}{\sum_{i=1}^{I}\sum_{t=1}^{T} x_t^i \cdot CTR_t^i}.$$

We have focused on the metrics explicitly used in the considered optimization problems (total conversion and cost per click). Additional metrics, such as return on investment (ROI) and budget utilization, can be used to support the observed results, while total conversion and cost per click can serve as proxy metrics for them.

4.4 Individual uncertainty

The dependencies of the metrics on the CTR uncertainty level, assuming CVR is known exactly, are presented in Figure 1. Note that the results of the RobustBid and NonRobustBid naturally converge to the same values if epsilon is close to zero. RiskBid performs better than NonRobustBid, and the RobustBid outperforms both baselines. The dependencies of the metrics on the CTR and CVR uncertainty levels are presented in Figure 1, where the root-square standard deviation is shown as gradient shading.

4.5 Joint uncertainty

We analyzed the performance of our RobustBid approach against NonRobustBid across three datasets with varying ε_a and ε_b, denoting errors in CTR and CVR, respectively. Figure 2 illustrates the dependence of conversion and average cost-per-click metrics on the noise in CTR and CVR values for the Synthetic dataset. We observe that NonRobustBid provides similar values of these metrics independent of the uncertainties ε_a and ε_b. In contrast, RobustBid gives higher values for conversions and lower average cost-per-click for a broad range of ε_a and ε_b. Thus, on the Synthetic dataset, we confirm that our approach makes bidding more robust to noise in both CTR and CVR values.

[1] https://github.com/IAIOnline/RobustBid

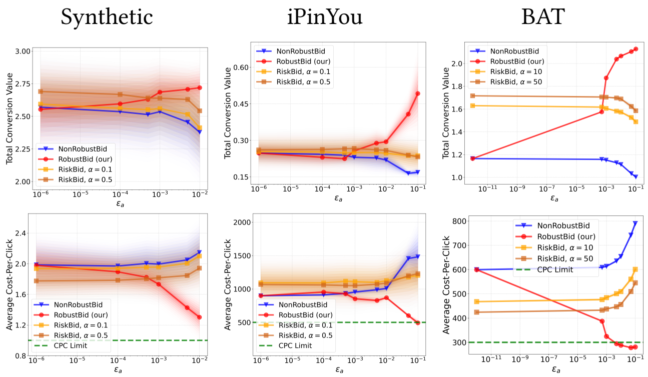

Figure 1: Experimental results with CTR uncertainty and CVR accurate estimate. The green dotted line represents the upper bound on the CPC constraint. Root mean squared deviation is specified in the gradient shading form. α is a risk-averse parameter, the larger α corresponds to less risky strategies.

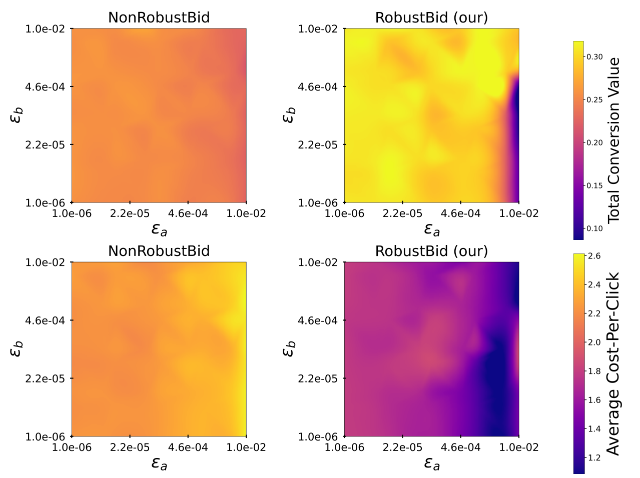

Figure 2: Heatmaps with comparison TCV and CPC_{avg} for Synthetic dataset. NonRobustBid metrics remain approximately the same, slightly decreasing at large values of ε_a and ε_b. RobustBid performs better overall, with TCV decreasing only at large ε_a, while maintaining lower CPC_{avg}.

The similar heatmaps for the iPinYou dataset are presented in Figure 3. They emphasize the drawbacks of the NonRobustBid method, noting that resulting conversions are very low and the average cost-per-click is excessively high across the considered range of ε_a and ε_b. Moreover, these metrics are almost constant for all considered ε_a and ε_b. Conversely, RobustBid better addresses the noise in CTR and CVR values. In particular, it increases the conversion metric and decreases the average cost-per-click metric for a medium level of noise. Hence, the RobustBid method outperforms the NonRobustBid in the iPinYou dataset, too.

Figure 4 compares TCV and average cost-per-click metrics generated by NonRobustBid and RobustBid algorithms for the BAT dataset. For Average Cost-Per-Click and Total Conversion Value,

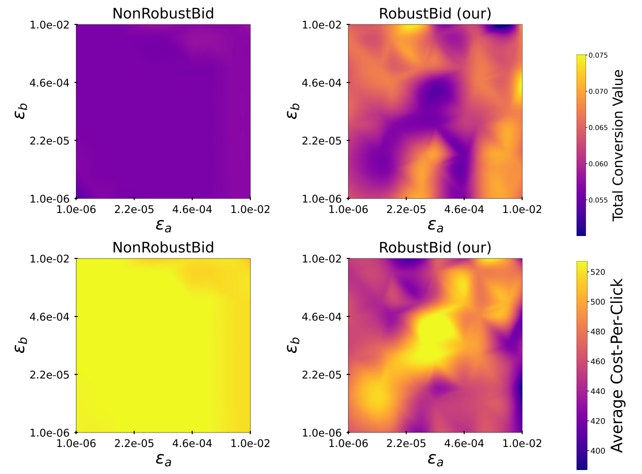

Figure 3: Heatmaps with comparison TCV and CPC_{avg} for iPinYou dataset. NonRobustBid metrics remain approximately the same across all ε_a and ε_b. In contrast, RobustBid metrics show significant variation while demonstrating better performance overall.

NonRobustBid performs poorly on both metrics. The Total Conversion Value of NonRobustBid remains almost constant, while the Cost-Per-Click varies with changes in ε_a. At the same time, RobustBid outperforms NonRobustBid, while showing the more complex dynamics of changes in the ε_a and ε_b.

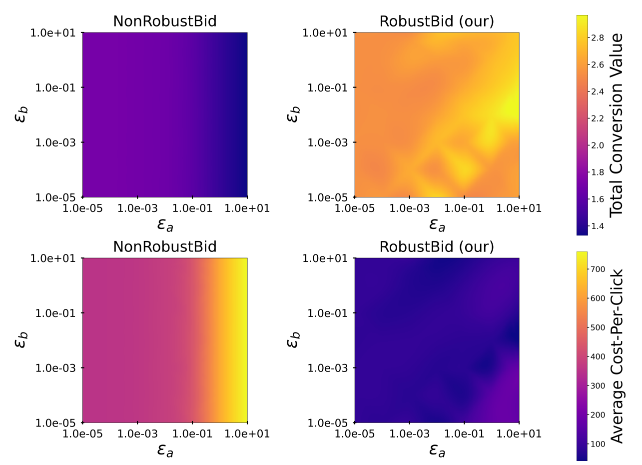

Figure 4: Heatmaps with comparison TCV and CPC_{avg} for BAT dataset. While ε_a grows, NonRobustBid shows lower TCV and higher CPC_{avg}. At the same time, RobustBid metrics are nearly independent of ε_a and ε_b.

The overall comparison of heatmap patterns highlights that BAT (see Figure 4) shows more stable behavior than iPinYou (see Figure 3). This stability correlates with budget diversity: BAT has a wide budget range, whereas iPinYou uses fixed budgets (see Table 2 in Appendix F [19]). The BAT dataset (Figure 4) demonstrates strong results for RobustBid, maintaining stable metrics and consistently outperforming NonRobustBid. This is especially significant as BAT most closely represents real-world advertising scenarios. Across all three datasets and for the majority of ε_a and ε_b parameter pairs, RobustBid consistently achieves higher total conversions (TCV) and lower average cost-per-click (CPC_{avg}) than NonRobustBid, demonstrating its effectiveness.

Table 2 presents the standard deviation of joint CTR and CVR uncertainties. For Synthetic data, the standard deviation for CPC_{avg} of RobustBid is higher. On more complex IPinYou and real-world BAT datasets, RobustBid exhibits lower standard deviations for the CPC_{avg} and TCV metrics, indicating more stable behavior.

Table 2: Standard deviation of target metrics with MSE loss under joint uncertainty (CTR and CVR).

Metric	Synthetic, $\varepsilon_a = \varepsilon_b = 10^{-2}$		IPinYou, $\varepsilon_a = 5 \cdot 10^{-3}, \varepsilon_b = 10^{-3}$		BAT, $\varepsilon_a = \varepsilon_b = 10^{-3}$	
	Robust	Non-robust	Robust	Non-robust	Robust	Non-robust
TCV	0.09	0.17	0.001	0.007	0.001	0.01
CPC_{avg}	0.36	0.05	8.28	11.11	0.04	3.01

5 LIMITATIONS AND FUTURE WORK

Our approach works in the offline setup, where a lot of data is available and can not be used in the online setup. Extending the proposed method to the online setup with controllers or RL frameworks is a promising research direction. The incorporation of uncertainty sets based on cross-entropy loss in our method remains a challenge that could improve the interpretability of the uncertainty bounds.

6 CONCLUSION

This paper proposes the RobustBid algorithm which generates bids that are robust to uncertainties in CTR and CVR. This algorithm is based on the robust modifications of the conversion maximization problem, incorporating uncertainty sets for individual CTR and CVR values and joint ones. The analytical solutions of the corresponding problems lead to the explicit formulas for bid, which naturally generalize the non-robust approach. Extensive experiments on synthetic and industrial datasets confirm that RobustBid generates bids that are robust to perturbations in CTR/CVR. The performance of the considered methods is evaluated through the total conversion rate and average cost per click. According to these metrics, RobustBid gives a larger total conversion rate and a lower average cost per click than non-robust and robust competitors.

ACKNOWLEDGMENTS

This work was supported by the The Ministry of Economic Development of the Russian Federation in accordance with the subsidy agreement (agreement identifier 000000C313925P4H0002; grant No 139-15-2025-012).

REFERENCES

[1] Gagan Aggarwal, Ashwinkumar Badanidiyuru, Santiago R Balseiro, Kshipra Bhawalkar, Yuan Deng, Zhe Feng, Gagan Goel, Christopher Liaw, Haihao Lu, Mohammad Mahdian, et al. 2024. Auto-bidding and Auctions in Online Advertising: A Survey. *ACM SIGecom Exchanges* 22, 1 (2024), 159–183.
[2] Gagan Aggarwal, Ashwinkumar Badanidiyuru, and Aranyak Mehta. 2019. Autobidding with constraints. In *Web and Internet Economics: 15th International Conference, WINE 2019, New York, NY, USA, December 10–12, 2019, Proceedings 15*. Springer, 17–30.
[3] Santiago Balseiro, Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. 2021. Robust auction design in the auto-bidding world. *Advances in Neural Information Processing Systems* 34 (2021), 17777–17788.
[4] Santiago R Balseiro, Yuan Deng, Jieming Mao, Vahab S Mirrokni, and Song Zuo. 2021. The landscape of auto-bidding auctions: Value versus utility maximization. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 132–133.
[5] Chaithanya Bandi and Dimitris Bertsimas. 2014. Optimal design for multi-item auctions: A robust optimization approach. *Mathematics of Operations Research* 39, 4 (2014), 1012–1038.
[6] Aharon Ben-Tal, Laurent Ghaoui, and Arkadi Nemirovski. 2009. *Robust Optimization*. https://doi.org/10.1515/9781400831050
[7] Dimitris Bertsimas and Benjamin Boucher. 2025. From Data to Uncertainty Sets: a Machine Learning Approach. *arXiv preprint arXiv:2503.02173* (2025).
[8] Han Cai, Kan Ren, Weinan Zhang, Kleanthis Malialis, Jun Wang, Yong Yu, and Defeng Guo. 2017. Real-time bidding by reinforcement learning in display advertising. In *Proceedings of the tenth ACM international conference on web search and data mining*. 661–670.
[9] Qinyi Chen, Phuong Ha Nguyen, and Djordje Gligorijevic. 2024. Optimization-Based Budget Pacing in eBay Sponsored Search. In *Companion Proceedings of the ACM on Web Conference 2024*. 328–337.
[10] Chao Deng, Hao Wang, Qing Tan, Jian Xu, and Kun Gai. 2021. Calibrating User Response Predictions in Online Advertising. In *Machine Learning and Knowledge Discovery in Databases: Applied Data Science Track*, Yuxiao Dong, Dunja Mladenić, and Craig Saunders (Eds.). Springer International Publishing, Cham, 208–223.
[11] Yuan Deng, Jieming Mao, Vahab Mirrokni, Hanrui Zhang, and Song Zuo. 2024. Efficiency of the first-price auction in the autobidding world. *Advances in Neural Information Processing Systems* 37 (2024), 139270–139293.
[12] Kristin Fridgeirsdottir and Sami Najafi-Asadolahi. 2018. Cost-per-impression pricing for display advertising. *Operations Research* 66, 3 (2018), 653–672.
[13] Negin Golrezaei, Adel Javanmard, and Vahab Mirrokni. 2019. Dynamic incentive-aware learning: Robust pricing in contextual auctions. *Advances in Neural Information Processing Systems* 32 (2019).
[14] Yue He, Xiujun Chen, Di Wu, Junwei Pan, Qing Tan, Chuan Yu, Jian Xu, and Xiaoqiang Zhu. 2021. A Unified Solution to Constrained Bidding in Online Display Advertising. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining* (Virtual Event, Singapore) *(KDD '21)*. Association for Computing Machinery, New York, NY, USA, 2993–3001. https://doi.org/10.1145/3447548.3467199
[15] Bernhard Kasberger and Karl H Schlag. 2024. Robust bidding in first-price auctions: How to bid without knowing what others are doing. *Management Science* 70, 7 (2024), 4219–4235.
[16] Alexandra Khirianova, Ekaterina Solodneva, Andrey Pudovikov, Sergey Osokin, Egor Samosvat, Yuriy Dorn, Alexander Ledovsky, and Yana Zenkova. 2025. BAT: Benchmark for Auto-bidding Task. In *Proceedings of the ACM on Web Conference 2025*. 2657–2667.
[17] Rachitesh Kumar, Jon Schneider, and Balasubramanian Sivan. 2024. Strategically-robust learning algorithms for bidding in first-price auctions. *arXiv preprint arXiv:2402.07363* (2024).
[18] Yewen Li, Shuai Mao, Jingtong Gao, Nan Jiang, Yunjian Xu, Qingpeng Cai, Fei Pan, Peng Jiang, and Bo An. 2025. GAS: Generative Auto-bidding with Post-training Search. In *Companion Proceedings of the ACM on Web Conference 2025*. 315–324.
[19] Andrey Pudovikov, Alexandra Khirianova, Ekaterina Solodneva, Gleb Molodtsov, Aleksandr Katrutsa, Yuriy Dorn, and Egor Samosvat. 2025. Robust autobidding for noisy conversion prediction models. *arXiv preprint arXiv:2510.08788* (2025).
[20] Yanlin Qu, Ravi Kant, Yan Chen, Brendan Kitts, San Gultekin, Aaron Flores, and Jose Blanchet. 2024. Double Distributionally Robust Bid Shading for First Price Auctions. *arXiv preprint arXiv:2410.14864* (2024).
[21] Matthew Richardson, Ewa Dominowska, and Robert Ragno. 2007. Predicting clicks: estimating the click-through rate for new ads. In *Proceedings of the 16th international conference on World Wide Web*. 521–530.
[22] Helen Robinson, Anna Wysocka, and Chris Hand. 2007. Internet advertising effectiveness: the effect of design on click-through rates for banner ads. *International journal of advertising* 26, 4 (2007), 527–541.
[23] Tim Roughgarden. 2010. Algorithmic game theory. *Commun. ACM* 53, 7 (2010), 78–86.
[24] Wen-Yueh Shih, Hsu-Chao Lai, and Jiun-Long Huang. 2023. A Robust Real Time Bidding Strategy Against Inaccurate CTR Predictions by Using Cluster Expected Win Rate. *IEEE Access* 11 (2023), 126917–126926.
[25] Alex Shtoff, Yohay Kaplan, and Ariel Raviv. 2023. Improving conversion rate prediction via self-supervised pre-training in online advertising. In *2023 IEEE International Conference on Big Data (BigData)*. IEEE, 1835–1842.
[26] Pingzhong Tang, Xun Wang, Zihe Wang, Yadong Xu, and Xiwang Yang. 2020. Optimized cost per mille in feeds advertising. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*. 1359–1367.

[27] Xun Yang, Yasong Li, Hao Wang, Di Wu, Qing Tan, Jian Xu, and Kun Gai. 2019. Bid optimization by multivariable control in display advertising. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 1966–1974.

[28] Yanwu Yang and Panyu Zhai. 2022. Click-through rate prediction in online advertising: A literature review. *Information Processing & Management* 59, 2 (2022), 102853.

[29] Shuai Yuan, Jun Wang, and Xiaoxue Zhao. 2013. Real-time bidding for online advertising: measurement and analysis. In *Proceedings of the seventh international workshop on data mining for online advertising*. 1–8.

[30] Haifeng Zhang, Weinan Zhang, Yifei Rong, Kan Ren, Wenxin Li, and Jun Wang. 2017. Managing risk of bidding in display advertising. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. 581–590.