Robust Auto-Bidding Strategies for Online Advertising

Qilong Lin
Shanghai Jiao Tong University
Shanghai, China
edersnow@sjtu.edu.cn

Zhenzhe Zheng*
Shanghai Jiao Tong University
Shanghai, China
zhengzhenzhe@sjtu.edu.cn

Fan Wu
Shanghai Jiao Tong University
Shanghai, China
fwu@cs.sjtu.edu.cn

ABSTRACT

In online advertising, existing auto-bidding strategies for bid shading mainly adopt the approach of first predicting the winning price distribution and then calculating the optimal bid. However, the winning price information available to the Demand Side Platforms (DSPs) is extremely limited, and the associated uncertainties make it challenging for DSPs to accurately estimate winning price distribution. To address this challenge, we conducted a comprehensive analysis of the process by which DSPs obtain winning price information, and abstracted two types of uncertainties from it: known uncertainty and unknown uncertainty. Based on these uncertainties, we proposed two levels of robust bidding strategies: Robust Bidding for Censorship (RBC) and Robust Bidding for Distribution Shift (RBDS), which offer guarantees for the surplus in the worst-case scenarios under uncertain conditions. Experimental results on public datasets demonstrate that our robust bidding strategies consistently enable DSPs to achieve superior surpluses, both on test sets and under worst-case conditions.

CCS CONCEPTS

• **Applied computing → Online auctions**; • **Information systems → Display advertising**; • **Mathematics of computing → Mathematical optimization**.

KEYWORDS

Auto-Bidding, Bid Shading, Robust Optimization

ACM Reference Format:

Qilong Lin, Zhenzhe Zheng, and Fan Wu. 2024. Robust Auto-Bidding Strategies for Online Advertising. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24), August 25–29, 2024, Barcelona, Spain.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3637528.3671729

1 INTRODUCTION

In recent years, the landscape of auto-bidding in online advertising has undergone a significant transformation, with the sale of vast quantities of ad impressions shifting from the traditional second-price auction to the first-price auction [8, 26, 33]. This shift has altered the bidding strategies of the auction participants, namely Demand Side Platforms (DSPs), and has given rise to the issue of bid shading [32, 34, 35] in the first-price auction context.

From the perspective of auction theory, in contrast to traditional second-price auctions, first-price auctions do not possess the incentive compatibility property [17, 18]. For DSPs, this indicates that truthfully revealing their value of winning the auction may not necessarily yield the maximal surplus for themselves. Therefore, DSPs need to develop bidding strategies tailored to the first-price auction environment in order to maximize their own surplus.

A natural bidding approach in first-price auction is firstly predicting the distribution of winning price for each auction, where winning price is the minimal bid that could win the auction. Based on this distribution, one can solve for the bid that maximizes the expected surplus. This approach has been adopted by many works on bid shading, and has led to the development of specialized works that focus on the prediction of winning prices [19, 30, 31].

However, the uncertainty of winning price is a substantial challenge. In reality, the real distribution of winning prices is unattainable, and we can only estimate it from the sampled winning price data within each auction. Moreover, a significant issue for DSPs is that they can only access partial information from this sampled data. For instance, a DSP may know the exact winning price only upon winning an auction, whereas in the cases of not winning, DSP would merely know that winning price is higher than her bid. Additionally, the proportion of auctions that a DSP wins constitutes a minor segment of all auctions. This dilemma is commonly referred to as the censorship problem [2, 27, 32]. The impediments in accessing winning price information result in uncertainties that cannot be disregarded in the modeling of winning price.

Existing research endeavors to predict the distribution of winning prices in the context of censorship, which inherently necessitates the introduction of certain assumptions as criteria for evaluating the quality of the predicted winning price distributions. For example, the method of censored regression assumes that the distribution of winning prices should remain consistent across auctions that a DSP loses and those her wins. Such an assumption may not correspond to actual conditions [32], thereby introducing an intrinsic bias into the predicted distribution.

In our work, to address the challenges posed by uncertainties, we propose two levels robust bidding strategies. More specifically, we categorize the uncertainty within the winning price into two levels: known uncertainty and unknown uncertainty, corresponding respectively to the issue of censorship and the sampling process of winning price. These two types of uncertainties collectively describe the limited information on the winning price distribution that the DSP can obtain.

Given that both types of uncertainties pertain to winning price distribution, We draw inspiration from the idea of Distributionally

*Z.Zheng is the corresponding author.

KDD '24, August 25–29, 2024, Barcelona, Spain

ACM ISBN 979-8-4007-0490-1/24/08
https://doi.org/10.1145/3637528.3671729

Robust Optimization (DRO) [5, 16, 21] and model such uncertainties on distribution by the concept of ambiguity sets. Subsequently, we solve the distributionally robust optimization problem to select the optimal bids. Following this approach, we propose two levels robust bidding strategies—Robust Bidding for Censorship (RBC) and Robust Bidding for Distributional Shift (RBDS)—for known and unknown uncertainties, respectively, and design corresponding algorithms to solve the bidding problem. Such robust bidding strategies aim to optimize the worst-case surplus, thereby providing lower bound guarantees for the DSP's surplus in first-price auction environments, where the winning price is fraught with uncertainty. Our contributions in this work can be summarized as follows:

- We model the uncertainties of winning price, and propose the corresponding distributionally robust bidding strategies. Our robust strategy aims to optimize the worst-case surplus, thereby achieving advanced performance in the bidding environments characterized by significant uncertainty.
- From the technical perspective, we designed the construction of the ambiguity set and the algorithm for the RBC and RBDS strategies. This approach can provide insights into solving specialized distributionally robust optimization problems under discrete distribution scenarios.
- We conduct comprehensive experiments on public datasets. The experimental results show that our robust bidding strategies outperform the existing strategies, especially exhibiting better performance in the worst-case situations.

The rest of this work is organized as follows. In Section 2, we provide an overview of the existing works on bid shading and robust optimization. In Section 3, we conduct a detailed analysis of the uncertainty in bid shading, and introduce the robust bidding problem. In Section 4, in correspondence to the two types of uncertainties abstracted in the analysis, we specifically propose two levels robust bidding strategies. In Section 5, we present the experimental results on public datasets, and demonstrate the effectiveness of our robust strategies. Finally, in Section 6, we briefly summarize the content of this work and potential future works.

2 RELATED WORK

In this section, we provide an overview of the existing auto-bidding strategies in bid shading, and introduce the closely related series of works on winning price prediction. Additionally, we briefly review the relevant works on distributionally robust optimization, which is the primary technique employed in this work.

2.1 Bid Shading

In bid shading, mainstream works adopt the idea of distribution-based bidding strategies. According to whether the distribution type of winning price is pre-assumed, we can roughly divide existing works into two categories. One kind of work assumes that the winning price follows a certain type of distribution, which is called the parametric method. Existing work has tried various distribution types to model winning price distribution [11, 20, 29]. Among them, [35] compares several basic distribution types, and finds that the lognormal distribution fits best in actual business.

Although the research on the parametric method is comprehensive, the actual winning price distribution contains lots of detailed information, which is difficult to be described by a certain type of distribution. Based on this observation, another series of works tries to directly fit the original distribution. This type of method is called the non-parametric method. [34] uses a table-based algorithm to record historical surplus and make bidding decisions based on these records. At present, the main discussion of non-parametric methods lies in the series of works about winning price prediction.

2.2 Winning Price Prediction

To deal with censorship problem, a series of works focused on the winning price prediction model. [32] is the first to consider the censorship problem in distribution prediction. In their work, they introduce the traditional censored regression method and model the winning price prediction as a linear fitting problem with normal noise. [7] further models winning price as a mixture distribution and uses the mixture density network corresponding. [14] focuses on the design of feature engineering, and improves the loss function of the network. Besides these works related to censored regression, there is also a series of works that use survival analysis to deal with censorship problem. [30] is the first to introduce survival analysis into winning price prediction, and they consider combining it with decision trees to apply to bid shading problem. [22, 28] further consider combining survival analysis with recurrent neural network and Markov network. In general, this series of works gradually develops towards complex non-parametric methods.

2.3 Distributionally Robust Optimization

Distribution robust optimization is a burgeoning robust optimization method that we mainly refer to. Its principal idea is to optimize the worst-case performance within an uncertain environment, where the optimized variables are distributions. This method can probably be traced back to a study on the inventory problem [24], and is widely known as distributionally robust optimization after a more recent study on moment uncertainty [4]. Nowadays, this optimization idea has been applied to various scenarios, such as machine learning [13] and auction mechanism design [12].

In our work, since the bidding decision function of bid shading is not a convex function, we need to consider non-convex distributionally robust optimization. Recently, [10] and [9] have expanded the inner optimization problem in DRO to rewrite the DRO problem as a stochastic optimization problem. With this transformation, they can naturally use existing algorithms to solve the original DRO problem. However, the number of studies on this topic is small and further exploration is still needed.

3 PRELIMINARIES

In this section, we conduct a comprehensive analysis of the process by which DSPs acquire winning price information, and abstract two types of uncertainties from it: known uncertainty and unknown uncertainty. Based on this uncertain environment, we introduce the robust bidding problem, which encapsulates the objectives of robust strategies in the next section.

3.1 Uncertain Bidding Environment

The environment we consider is the context of first-price auction, in which a DSP aims to devise a bidding strategy that could maximize her surplus. Formally speaking, from the DSP's perspective, we define her private value as v, and she can submit a bid $b \in \mathcal{B}$ to the auction platform to compete for the advertisement slot, where $\mathcal{B}$ denote the discrete bidding space with size M. Let w be the winning price of the auction. The DSP wins the auction when her bid b is greater than or equal to w, at which point she incurs a cost equal to her bid b, and gets the surplus $v-b$. If DSP does not win the auction, her surplus is 0. Hence, the objective of DSP can be expressed as optimizing her surplus:

$$h(b, w) = (v - b) \cdot \mathbb{I}\{b \geq w\}, \tag{1}$$

where $\mathbb{I}\{b \geq w\}$ indicates whether DSP wins or not, and it converts the logical judgment result to 0 or 1 accordingly. In practical business scenarios, the private value v is generally given. Hence, similar to other bid shading studies, we focus our attention on the modeling of winning price w.

To model winning price, we analyze its origins and the associated uncertainty issues from the perspective of DSP. First, we assume that winning price w follows a distribution $\boldsymbol{p}$. In this work, for ease of discussion, we utilize this kind of probability vector $\boldsymbol{p}$ to discretely represent the winning price distribution. Specifically, we define $b_0 = 0$, and set the bids in $\mathcal{B}$ to satisfy $b_0 < b_1 < ... < b_M$ and form an arithmetic sequence. Consequently, the probability vector can be denoted as $\boldsymbol{p} = (p^1, p^2, ..., p^M)$, where for any $j \in [1, M]$, p^j corresponds to the probability of the winning price being in the interval $[b_{j-1}, b_j)$.

In reality, the DSP cannot directly access the real winning price distribution $\boldsymbol{p}$. Instead, it indirectly acquires information about this distribution through the following process:

- **Sampling**. Historical auctions contain sampled winning price data $\{\hat{w}_i | i \in \mathcal{A}\}$, where i is used to identify different auctions in historical auctions set $\mathcal{A}$. These sampled data constitutes an sampling distribution $\hat{\boldsymbol{p}} = \mathcal{F}_{\mathcal{B}}(\{\hat{w}_i | i \in \mathcal{A}\})$. In our work, we use the function $\mathcal{F}_{\mathcal{B}}$ to denote the function that map the dataset to a distribution $\hat{\boldsymbol{p}}$ on $\mathcal{B}$.
- **Censorship**. DSP can only obtain the sampled winning price $\{\hat{w}_i | i \in \mathcal{W}\}$ when her wins, where $\mathcal{W}$ is the set of auctions DSP won. For those losing auctions $i \in \mathcal{L} = \mathcal{A} \setminus \mathcal{W}$, DSP normally only knows that the sampled winning price exceeds her own bid. We abstract the situation of these auctions as cases where the DSP only knows the interval $[\hat{l}_i, \hat{r}_i]$ in which $\hat{w}_i$ is located.

Here, the censorship problem has already been widely discussed in related work [32], and we further explain the meaning of the interval $[\hat{l}_i, \hat{r}_i]$ in it. In the most general case, when DSP does not win the auction, she only knows that the winning price is higher than her bid, so the left end of the interval $\hat{l}_i$ is DSP's bid in that auction, and the right end $\hat{r}_i$ is infinite. However, some studies have shown that the interval in which the winning price resides can be empirically narrowed down [14]. Therefore, in order to make the modeling of our problem more universal, we abstract the accessible information when the DSP does not win as the winning price being within a known interval $\hat{w}_i \in [\hat{l}_i, \hat{r}_i]$.

Since the sampling and censorship process introduces considerable uncertainty to the bid shading problem, for further robust strategies design, we divide the uncertainty into two types, corresponding to the sampling and censorship process respectively:

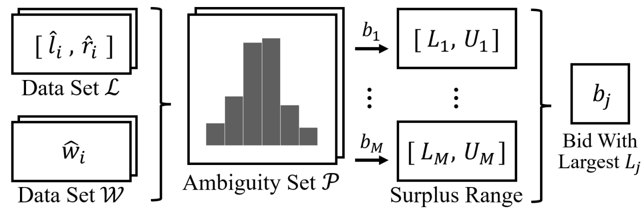

Figure 1: The idea of uncertainty analysis and distributionally robust optimization

- **Known uncertainty**. For any losing auction $i \in \mathcal{L}$, the actual winning price $\hat{w}_i$ could be any value within the interval $[\hat{l}_i, \hat{r}_i]$. Since the sampling distribution is comprised of winning price data from the dataset, DSP can determine that the sampling distribution $\hat{\boldsymbol{p}}_i$ belongs to an ambiguity set $\mathcal{P}_{kn}$, where the uncertainty can be specified by the intervals.
- **Unknown uncertainty**. The difference between the sampling distribution $\hat{\boldsymbol{p}}$ and the real distribution $\boldsymbol{p}$, as well as the potential changes of the real distribution $\boldsymbol{p}$ over time, constitutes the unknown uncertainty that DSP cannot be certain about due to the limited information. Corresponding to this uncertainty, we denote the ambiguity set in which the real distribution $\boldsymbol{p}$ could resides as $\mathcal{P}_{un}$.

The term "known uncertainty" here refers to the situation where DSP is aware of the set $\mathcal{P}_{kn}$ but does not know which specific element from the set is the actual sampling distribution $\hat{P}$. In other words, the DSP knows the range of possible distributions but cannot pinpoint the exact one within that range. On the other hand, "unknown uncertainty" implies that DSP cannot even accurately grasp or define the ambiguity set $\mathcal{P}_{un}$ itself.

3.2 Robust Bidding Strategies

In order to achieve the satisfied surplus of DSP in uncertain environments, we introduce the robust optimization [1] into the design of bidding strategies. Specifically, since uncertainty in bidding problem arising from the winning price distribution, the bidding decision can be formulated as the following distributionally robust optimization problem:

$$\max_{b \in \mathcal{B}} \min_{\boldsymbol{p} \in \mathcal{P}} \mathbb{E}_{w \sim \boldsymbol{p}}[h(b, w)], \tag{2}$$

where the decision function h is defined in (1). This minimax optimization process can be understood as follows: as the real distribution could be any element within the ambiguity set $\mathcal{P}$, for each bid b, there exists a range for the expected surplus obtained by DSP, and we can denote its lower and upper bound as L and U respectively. In this case, distributionally robust hopes to take this uncertainty into account and chooses the bid b_j that maximizes the minimal surplus L_j. Figure 1 provides a visual representation of this process, where we adopt the known uncertainty and the corresponding ambiguity set $\mathcal{P} = \mathcal{P}_{kn}$ for illustrating.

To facilitate subsequent discussions, we denote the cumulative distribution function (CDF) corresponding to the distribution $\boldsymbol{p}$

as $P_{\boldsymbol{p}}$, and present a more tractable form for the distributionally robust optimization problem (2):

$$\max_{b \in \mathcal{B}} \min_{\boldsymbol{p} \in \mathcal{P}} \ (v - b) \cdot P_{\boldsymbol{p}}(b). \tag{3}$$

The distributionally robust problem (2) and (3), along with the ambiguity set $\mathcal{P}_{kn}$ and $\mathcal{P}_{un}$ considered within it, abstractly formulates the mathematical problem in bid shading, and provides a general description of the objective for our robust strategies design. Subsequently, we will further elaborate on the design details of our robust bidding strategies.

4 DESIGN

In this section, we further elaborate on the design of robust bidding strategies. In response to the two types of uncertainties summarized in the previous section, we propose two levels robust bidding strategies: Robust Bidding for Censorship and Robust Bidding for Distribution Shift, and further provide solutions for the robust optimization problems within these strategies respectively.

4.1 Robust Bidding for Censorship

We first discuss the design of robust bidding strategy considering only the known uncertainty. We refer to this as the Robust Bidding for Censorship (RBC), which addresses a problem of the same form as (2), but with the ambiguity set defined as follows:

$$\mathcal{P}_{kn} = \{\mathcal{F}_{\mathcal{B}}(\{\hat{w}_i | i \in \mathcal{W}\} \cup \{w_i | w_i \in [\hat{l}_i, \hat{r}_i], i \in \mathcal{L}\})\}, \tag{4}$$

where we use the sampled winning price set $\{\hat{w}_i\}$ to represent the distribution. In this set, the data for subscript $i \in \mathcal{W}$ is known, while the data for subscript $i \in \mathcal{L}$ is uncertain, where the data w_i can take any value within the interval $[\hat{l}_i, \hat{r}_i]$.

Problem (2) with the ambiguity set $\mathcal{P}_{kn}$ appears complex, but it can be transformed into a simple maximization problem, which is tractable to some extent. We arrive at the following conclusion:

Remark 4.1. *Given the ambiguity set $\mathcal{P}_{kn}$, problem (2) is equivalent to the following problem:*

$$\max_{b \in \mathcal{B}} \mathbb{E}_{w \sim \boldsymbol{p}_0} [h(b, w)], \tag{5}$$

with the distribution $\boldsymbol{p}_0 = \hat{\boldsymbol{p}}_c$, where $\hat{\boldsymbol{p}}_c \in \mathcal{P}_{kn}$ represents the worst-case winning price distribution, and it can be expressed as:

$$\hat{\boldsymbol{p}}_c = \mathcal{F}_{\mathcal{B}}(\{\hat{w}_i | i \in \mathcal{W}\} \cup \{\hat{r}_i | i \in \mathcal{L}\}). \tag{6}$$

This result is quite intuitive. Similar to the equivalence between problems (2) and (3), problem (5) can be equivalently transformed into the following problem:

$$\max_{b \in \mathcal{B}} (v - b) \cdot P_{\boldsymbol{p}_0}(b). \tag{7}$$

By analyzing problem (3), for any given bid b, the worst-case scenario corresponds to the smallest cumulative probability $P_{\hat{\boldsymbol{p}}_c}(b)$. Within the ambiguity set $\mathcal{P}_{kn}$, the sampled data in $\hat{\boldsymbol{p}}_c$ should be as large as possible, meaning that for each auction $i \in \mathcal{L}$, selecting $\hat{r}_i$ to form the sampling distribution $\hat{\boldsymbol{p}}_c$. Thus problem (3) with the ambiguity set $\mathcal{P}_{kn}$ is equivalent to problem (7) with the worst-case distribution $\hat{\boldsymbol{p}}_c$. A detailed proof can be found in Appendix A.1.

In fact, problem (5) is a highly versatile optimization problem, which can be used in conjunction with the winning price distribution prediction model to formulate a bidding strategy, and we refer to this expectation-based optimization as Stochastic Optimization (SO) in the remaining parts. Broadly speaking, the crux of solving problem (5) lies in obtaining the distribution $\boldsymbol{p}_0$ that the winning price w follows. In existing work, $\boldsymbol{p}_0$ is the estimated winning price distribution, which can be predicted using existing methods, such as the works presented in Section 2.2; in our RBC strategy, $\boldsymbol{p}_0$ is the worst-case distribution, but it can also be predicted by emulating existing methods. The detailed approach of predicting the worst-case distribution is as follows.

Considering that DSP can obtain feature data $\hat{\boldsymbol{x}}_i$ for each auction, we design a non-parametric distribution estimation model, denoted as f_c, which employs a simple two-layer fully connected network and utilizes the Softmax function to process the outputs. This network maps the input feature data $\hat{\boldsymbol{x}}_i$ to a discrete distribution $\hat{\boldsymbol{p}}_i = f_c(\hat{\boldsymbol{x}}_i)$, whereby the output dimension corresponds to the dimension M of the discrete distribution, with the i-th output corresponding to the value of $\hat{p}_i^M$. This distribution $\hat{\boldsymbol{p}}_i$ estimated will serve as the worst-case distribution $\boldsymbol{p}_0 = \hat{\boldsymbol{p}}_c$ within the SO problem (5) to obtain the robust bids. We denote the probability density function (PDF) associated with the output distribution $\hat{\boldsymbol{p}}_i$ as $p_{\hat{\boldsymbol{p}}_i}$, and thus the loss function utilizing the concept of maximum likelihood can be written as:

$$L_s = -\sum_{i \in \mathcal{W}} \log p_{\hat{\boldsymbol{p}}_i}(\hat{w}_i) - \sum_{i \in \mathcal{L}} \log p_{\hat{\boldsymbol{p}}_i}(\hat{r}_i). \tag{8}$$

In practice, the right end of the interval can be calculated based on some empirical circumstances, such as the method in [14]. However, in this work, we wish to consider a more general case, where the DSP only knows that the winning price is higher than her own bid, thus assuming that the upper bound of the winning price is infinity in each auction. Under this assumption, the worst-case distribution in (6) can be further rewritten as $\hat{\boldsymbol{p}}_c = \mathcal{F}_{\mathcal{B}}(\{w_i | i \in \mathcal{W}\})$, while keeping the solution of problem (5) unchanged. The loss function (8) now can be simplified to:

$$L_s' = -\sum_{i \in \mathcal{W}} \log p_{\hat{\boldsymbol{p}}_i}(\hat{w}_i). \tag{9}$$

This result is natural. Since all known uncertainty resides within the auctions in $\mathcal{L}$, the most conservative approach is to refrain from utilizing the uncertain data from these auctions, that is, only using data in winning auctions $\mathcal{W}$ like loss function (9).

We would like to further discuss the significance of RBC from the perspective of existing works. Since censorship issue was raised by [32], the academic discussion on bid shading has mostly been limited to the distribution prediction methods under censorship scenarios, such as the censored regression and survival analysis mentioned in Section 2.2. However, the results of RBC suggest that using only the data in winning auctions $\mathcal{W}$ has a certain degree of robustness, and therefore, it may yield favorable results in an environment with considerable uncertainty such as bid shading. Our subsequent experiments have validated this. Moreover, since RBC considers the worst-case distributions that are only composed of deterministic data, RBC has the potential to bring in research on conditional density estimation [23] beyond censorship, an area that is currently lacking in the discussion of bid shading. We will leave this potential direction for future work.

4.2 Robust Bidding for Distribution Shift

In contrast to the RBC strategy, where the ambiguity set $\mathcal{P}_{kn}$ is known, the Robust Bidding for Distribution Shift (RBDS) contemplates scenarios in which even the information about the ambiguity set $\mathcal{P}_{un}$ is unattainable. Analogous to the preceding analysis, the DSP is restricted to obtaining sampled data of the winning prices, and there may be a disparity between sampling distribution and the real distribution of winning prices. Furthermore, the real distribution of winning prices might undergo shifts over time, thereby introducing an element of inevitable bias into the DSP's predictions that are based on the sampled data.

To address this kind of uncertainty, we consider employing the conventional distributionally robust optimization approach, and adopting the discrepancy-based ambiguity set. This approach assumes that the discrepancy between the real distribution and the sampling distribution lies within a certain range, and it ensures a lower bound on the performance of the real distribution by optimizing for the worst-case scenario among all distributions within this range. The optimization problem can still be expressed in the form of problem (2), where the ambiguity set is represented as:

$$\mathcal{P}_{un}(\boldsymbol{p}_0, \epsilon_0) = \{\boldsymbol{p} \mid d(\boldsymbol{p}, \boldsymbol{p}_0) \leq \epsilon_0\}, \tag{10}$$

where $\boldsymbol{p}_0$ is the winning price distribution estimated by DSP, ϵ_0 is the upper bound of the discrepancy between the real distribution and estimated distribution, and d is a function that measures the discrepancy between distributions. Therefore, $\mathcal{P}_{un}$ mathematically describes our method of considering all distributions that differ from the estimated $\boldsymbol{p}_0$ within a certain range. Conversely, problem (2) with the ambiguity set $\mathcal{P}_{un}$ aims to optimize the worst-case scenario within this distribution range.

After introducing the complete strategic framework, we will further provide the definition of the ambiguity set $\mathcal{P}_{un}$ and the algorithm for problem (2) with $\mathcal{P}_{un}$ in the following parts.

4.2.1 Ambiguity Set in RBDS. We firstly introduce the estimated distribution $\boldsymbol{p}_0$, the discrepancy function d, and the discrepancy upper bound ϵ within the ambiguity set $\mathcal{P}_{un}$. The distribution $\boldsymbol{p}_0$ is the distribution estimated based on the sampled data. Similar to the distribution $\boldsymbol{p}_0$ in SO problem (5), $\boldsymbol{p}_0$ in $\mathcal{P}_{un}$ could be the estimated winning price distribution, in which case problem (2) can be simply interpreted as ensuring a lower bound on the bidding surplus by optimizing the worst-case outcome when there is a discrepancy between the estimated distribution and the real distribution. Alternatively, it could also be the worst-case distribution $\hat{\boldsymbol{p}}_c$ estimated in our RBC strategy (6). In this case, problem (2) can similarly be understood as acknowledging that discrepancies exist between the sampling worst-case distribution and the real worst-case distribution, and we likewise ensure a lower bound on bidding surplus by optimizing for the "worst of the worst-case" distribution. From this perspective, the RBDS strategy can be regarded as a robust form of the SO problem (5).

For the discrepancy function d, we adopt the Wasserstein distance function [25] to measure the discrepancy between distributions. This is a commonly used discrepancy function in distributionally robust optimization. Furthermore, when employing the Wasserstein distance, the form of the worst-case winning price

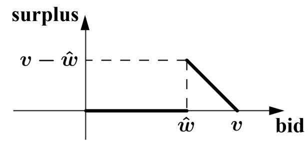

Figure 2: Decision function for the bid shading problem

distribution has an intuitive interpretation, aligning with observations in some existing works [34]. This intuitive interpretation is provided in the subsequent Section 4.2.2. Here, we first present the method of calculating this discrepancy function. Specifically, given any two distributions $\boldsymbol{p}$ and $\boldsymbol{q}$ with discrete representation $(p^1,\ p^2,\ \dots,\ p^M)$ and $(q^1,\ q^2,\ \dots,\ q^M)$, the Wasserstein distance of $\boldsymbol{p}$ and $\boldsymbol{q}$ can be defined as the solution of problem:

$$\begin{array}{ll} \min_{\boldsymbol{d}} & \sum_i \sum_j d_{ij} |i-j| \\ \text{s.t.} & 0 \leq \sum_{j=1}^N d_{ij} \leq p^i \\ & p^i + \sum_{j=1}^N d_{ji} - \sum_{i=1}^N d_{ij} = q^i \\ & \boldsymbol{d} = \{d_{ij}\}, i \in [1, M], j \in [1, M], \end{array} \tag{11}$$

where the matrix $\boldsymbol{d}$ is the optimized variable, and each d_{ij} represents the probability quantity transferred from p^i to p^j, which is non-negative and cannot be greater than the original probability quantity p^i in total. In addition, after all the transfer, the probability distribution $\boldsymbol{p}$ should become $\boldsymbol{q}$. Then Wasserstein distance is the minimum value of the sum of the product of the probability transferred and the distance under these constraints, and we denote the result of problem (11) as the Wasserstein distance $d(\boldsymbol{p}, \boldsymbol{q})$.

Finally, ϵ_0 defines the distance between the estimated distribution $\boldsymbol{p}_0$ and real distribution. In practice, DSP can never know the real distribution of winning prices, but can only obtain sampled data. Therefore, it is impossible for DSP to accurately determine the value of ϵ_0, which is a specific manifestation of the unknown uncertainty associated with P_{un}. Given this, we adopt a second-best approach, considering the integration of the DSP's objective of maximizing surplus to choose ϵ_0. Of course, this method of selection dilutes its physical meaning, making it more akin to a hyperparameter. Hence, we will not delve into an extensive discussion on the selection of ϵ_0, but will only report the relationship between the DSP's surplus and ϵ_0 in experiments, as well as the maximum surplus that can be achieved by adjusting ϵ_0.

4.2.2 Algorithm for RBDS. To solve problem (2) with the ambiguity set $\mathcal{P}_{un}$, we first consider whether we can use traditional distributionally robust optimization algorithms, which requires the analysis of the decision function h. Given winning price value $\hat{w}$, the image of the function $h(\cdot, \hat{w})$ is shown in Figure 2. It can be seen that it is a non-continuous and non-convex function, and its mathematical properties are relatively poor.

In the existing works, the latest researches [10] and [9] have studied the non-convex distributionally robust optimization algorithm, but they still require the decision function to satisfy certain continuous assumptions. Besides, since the decision variable b is discrete, the equivalence problem (3) is also similar to the form of discrete minimax problem [36], but in this series of research, the set of distributions is finite, while our ambiguity set is infinite. If we

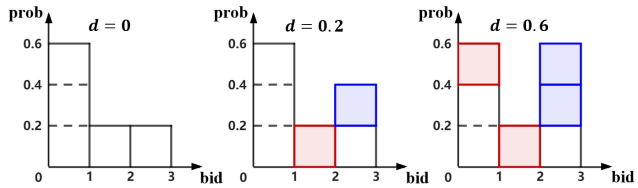

Figure 3: An example for the constructive approach

hope to apply these algorithms to our problem, we need to modify our decision function to meet the requirements of their algorithms.

In our work, due to the particularity of our problem, we consider designing an algorithm without changing the decision function. Since the bid space $\mathcal{B}$ in bid shading problem is finite, we can further write the problem (3) as:

$$\max_{b\in\mathcal{B}}\{(v-b)\min_{\boldsymbol{p}\in\mathcal{P}}P_{\boldsymbol{p}}(b)\}, \tag{12}$$

where the key point is to solve the min term on discrete domain $\mathcal{B}$, and we denote it as $f_{\mathcal{P}}$:

$$f_{\mathcal{P}}(b)=\min_{\boldsymbol{p}\in\mathcal{P}}P_{\boldsymbol{p}}(b). \tag{13}$$

If we can obtain the values of function $f_{\mathcal{P}}$ on $\mathcal{B}$, problem (12) becomes easy to solve, like problem (7) in previous section. Next, given $\mathcal{P}=\mathcal{P}_{un}(\boldsymbol{p}_0,\epsilon_0)$, we will elaborate on the approach to solving $f_{\mathcal{P}}$ and provide a specific algorithm for it.

Let's start with a simple example. Assuming $\mathcal{B}=\{1,2,3\}$ and $\boldsymbol{p}_0$ can be discretely represented by (p_0^1,p_0^2,p_0^3), where $p_0^1=0.6$ and $p_0^2=p_0^3=0.2$. In this case, distribution $\boldsymbol{p}_0$ can be represented by the left histogram in Figure 3. Assuming that given $\epsilon_0=0.6$, we are calculating $f_{\mathcal{P}}(2)=\min_{\boldsymbol{p}\in\mathcal{P}}P_{\boldsymbol{p}}(2)$, that is, finding a distribution $\boldsymbol{q}_0\in\mathcal{P}_{un}(\boldsymbol{p}_0,\epsilon_0)$ with the minimal cumulative probability $P_{\boldsymbol{q}_0}(2)$.

We consider a constructive approach, devising a certain procedure to construct $\boldsymbol{q}_0$ from $\boldsymbol{p}_0$. Let's start with $\boldsymbol{q}_0=\boldsymbol{p}_0$. To make $P_{\boldsymbol{q}_0}(2)=q_0^1+q_0^2$ as small as possible, we need to decrease q_0^1 and q_0^2, and increase q_0^3. But which one should we decrease first? Given our constraint $d(\boldsymbol{p}_0,\boldsymbol{q}_0)\le\epsilon_0=0.6$, it would be sensible to reduce q_0^2 first because reducing q_0^2 by the same amount will lead to a smaller $d(\boldsymbol{p}_0,\boldsymbol{q}_0)$ compared to reducing q_0^1. We decrease q_2 to 0, which results in an increase of q_3 to 0.2. At this point, the distribution can be represented by the middle histogram in Figure 3, where the red shaded area indicates the probability reduced relative to $\boldsymbol{p}_0$, and the blue shaded area indicates the probability increased relative to $\boldsymbol{p}_0$, with the computed $d(\boldsymbol{p}_0,\boldsymbol{q}_0)$ being 0.2. We continue to reduce q_0^1 and increase q_0^3, and we find that after reducing q_0^1 by 0.2, $d(\boldsymbol{p}_0,\boldsymbol{q}_0)$ becomes $\epsilon_0=0.6$; this new distribution can be represented by the histogram on the right side of Figure 3. At this point, $P_{\boldsymbol{q}_0}(2)=0.2$ is the function value $f_{\mathcal{P}}(2)$ that we are looking for.

An issue that this example doesn't cover is, assuming we allow for a bid amount of 4, and considering the same initial distribution $\boldsymbol{p}_0$ above, should we increase the probability at q_0^3 or q_0^4? This is actually similar to the choice of decreasing q_0^1 or q_0^2 first. Since increasing the same amount of probability, increasing q_0^3 results in a smaller distance $d(\boldsymbol{p}_0,\boldsymbol{q}_0)$, we similarly choose to increase q_0^3. In fact, summarizing the previous example, we find that given

Algorithm 1 Algorithm for $f_{\mathcal{P}}$ on Discrete Domain

Require: Distribution $\boldsymbol{p}_0$, discrepancy upper bound ϵ_0

for $j=1$ to M **do**
  $\boldsymbol{q}_0\leftarrow\boldsymbol{p}_0$;
  $k\leftarrow j$;
  while $d(\boldsymbol{p}_0,\boldsymbol{q}_0)\le\epsilon$ and $k>0$ **do**
    $q^{j+1}\leftarrow q^{j+1}+q^k$;
    $q^k\leftarrow 0$;
    $k\leftarrow k-1$;
  end while
  $q^{k+1}\leftarrow(d(\boldsymbol{p}_0,\boldsymbol{q}_0)-\epsilon_0)/(j-k)$;
  $f_{\mathcal{P}}(b_j)\leftarrow\sum_{k=1}^{j}q^k$;
end for

bid b, the method to construct $\boldsymbol{q}_0$ is to transfer the probabilities from $p_0^b,p_0^{b-1},...,p_0^1$ to p_0^{b+1} in descending order of index, until the distance from the original distribution reaches the upper bound ϵ_0. Based on this approach, we propose a procedural method as shown in Algorithm 1, which provides the pseudo code for solving the values of $f_{\mathcal{P}}$ on the discretely defined domain. We have the following conclusion:

Remark 4.2. *Algorithm 1 provide the strictly optimal solution for problem* $\min_{\boldsymbol{p}\in\mathcal{P}}P_{\boldsymbol{p}}(b)$ *on discrete domain* $\mathcal{B}$*, hence provide the exact values of function* $f_{\mathcal{P}}$ *on* $\mathcal{B}$ *defined in (13).*

A detailed proof can be found in Appendix A.2. In addition to outlining the algorithm, the example in Figure 3 vividly demonstrates the worst-case scenario considered in distributionally robust optimization problems. It can be seen that given bid $b=2$ and the initial distribution $\boldsymbol{p}_0$, in the worst-case scenario, the distribution $\boldsymbol{q}_0$ forms a spike right after the bid $b=2$, specifically at q_0^3. Existing literature indicates that such spikes are common in real-world settings [34], therefore, employing Wasserstein distance in distributionally robust optimization can be interpreted as a robust bidding strategy that accounts for the occurrence of these spikes.

After obtaining the algorithm for solving $f_{\mathcal{P}}(b_j)$, the optimization problem (12) of RBDS becomes as easy to solve as the SO problem (5). In summary and comparison, the RBDS method is equivalent to performing a robust treatment on the original distribution, and then solving the optimization problem (5) to obtain a robust bid against unknown uncertainty.

5 EXPERIMENT RESULTS

5.1 Experimental Setup

Firstly, we introduce the experimental datasets we used, the experimental procedures, and the models utilized in the experiments.

5.1.1 Datasets. Our experiments use two public datasets including the iPinYou dataset [15] and the Criteo dataset [6]:

- The iPinYou dataset is a widely used dataset in bid shading related research. It contains 10 days of real-world data from auctions in which the iPinYou DSP participates. We follow the existing work [30] and use the data in the first 7 days as training data, and the remaining data as test data.

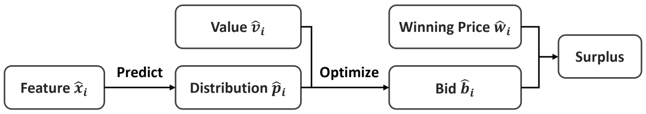

Figure 4: The main process of the experiments

- The Criteo dataset contains 30 days of real-world data from auctions that the Criteo DSP participates. We use the first 24 days' data as the training set, and the remaining 6 days' data as the test set.

The datasets above have provided the feature $\hat{x}_i$ and winning price $\hat{w}_i$ in each auction i. However, there is no censored data in the original datasets, so we need to simulate the censorship problem on the training set. In this work, we adopt the same method of simulating censored data as in [28], which firstly simulates the DSP bid $\hat{b}_i$, and then divides the training set into the set $\mathcal{W}$ and $\mathcal{L}$ according to the relationship between $\hat{b}_i$ and $\hat{w}_i$.

In addition, we also need the estimated value data $\hat{v}_i$ of the DSP. However, since the value is confidential for DSP, public datasets do not include this data, and we can only obtain it through simulation. In our paper, we assume that the DSP uses a common strategy, where in each auction, the DSP uses the value multiplied by a shading rate as her bid [8]. We further assume that the shading rate is randomly selected by DSP, that is, the bid $\hat{b}_i$ is generated in random proportion according to the value $\hat{v}_i$. From these assumptions, we are able to simulate the data $\hat{v}_i$ for our experiments.

5.1.2 Experimental Procedure. Our experiment simulates the real process of the DSP participating in auctions, as shown in Figure 4. In this process, DSP first predicts the distribution $\hat{\boldsymbol{p}}_i$ based on feature $\hat{\boldsymbol{x}}_i$ using the prediction model. This distribution could be the winning price distribution or the worst-case distribution in RBC strategy, corresponding to different prediction models utilized. After obtaining the distribution, DSP derives the bid $\hat{b}_i$ according to the optimization problem where the optimization corresponds to the aforementioned distribution and could be the SO problem (5) or the RBDS problem (12). Once the bid is determined, the DSP's surplus is calculated based on the actual winning price $\hat{w}_i$ and the decision function h.

5.1.3 Prediction Models. In our experiments, we primarily employ the following bidding strategies:

STM is a winning price distribution prediction method that combines survival analysis with decision trees [30]. We integrate it with SO problem (5) as a traditional bidding strategy that utilizes survival analysis.

MCN is a parametric distribution prediction method where the winning prices are modeled as a mixture of Gaussian distributions [7]. By integrating this approach with SO problem (5), it is utilized as a traditional parametric method bidding strategy.

KMMN is a winning price distribution prediction method that combines survival analysis with Markov network [28]. We integrate it with SO problem (5) to serve as an enhanced survival analysis-based bidding strategy.

NPM is a non-parametric distribution prediction method that employs the same network as depicted in RBC strategy. It is similar to the method described in [14], but we utilize a simpler network architecture and loss function. We integrate it with SO problem (5) to serve as a non-parametric bidding strategy.

RBC is the robust bidding strategy we have proposed for known uncertainty. Unless otherwise stated, we employ the network depicted in previous section and train our model using the loss function (9) to predict the worst-case distribution. This prediction is then incorporated into SO problem (5) as the robust bidding strategy. The selection of the distribution prediction model and loss function will be discussed in subsequent experimental sections.

RBDS is our proposed robust bidding method designed for unknown uncertainty. It necessitates the integration with the aforementioned distribution prediction model and supersedes the original SO problems (5). In our experiments, we denote the combination of this bidding strategy with any distribution prediction method M as M+R.

5.2 Performance of RBC Strategy

5.2.1 Overall Performance. We initially conduct experiments on the surplus performance of each strategy without employing RBDS strategy, with the results displayed in Table 1. The distribution prediction model is trained on the training set, and the resulting bidding strategy is tested on the test set. The first column of Table 1 lists different Campaigns, comprising nine campaigns from the iPinYou dataset and the overall Criteo dataset. Within the iPinYou dataset, the winning price distributions vary among different campaigns, hence existing works provide independent results for each campaign, a practice we also continue. The main body of the table presents the surplus obtained through different strategies, and it can be observed that the RBC strategy we propose achieves the highest surplus in the majority of the campaigns.

Table 1: Overall Surplus of Different Strategies (10^6)

Camp.	STM	MCN	KMMN	NPM	RBC
1458	11.29	11.61	11.03	12.06	**12.31**
2259	4.988	5.442	5.278	5.436	**5.590**
2261	4.959	5.602	5.802	6.021	**6.152**
2821	8.191	8.803	8.548	8.784	**9.095**
2997	2.460	2.939	2.832	2.876	**3.005**
3358	3.309	3.327	2.961	**3.610**	3.451
3386	9.253	9.659	9.491	10.34	**10.42**
3427	7.295	7.767	7.783	8.441	**8.486**
3476	7.202	7.470	7.389	7.899	**8.051**
Criteo	101.8	104.0	103.6	109.2	**109.3**

A natural question arises as to why the robust bidding strategy RBC, being a conservative approach, can perform better than the strategy based on direct prediction of winning price distribution. The primary reason is likely due to the fact that direct prediction of winning price distribution requires the introduction of assumptions that may not necessarily hold in practice, leading to an inherent bias in the predicted winning price distribution. From the experimental results, it appears that the surplus loss caused by these biases is greater than the loss due to robustness.

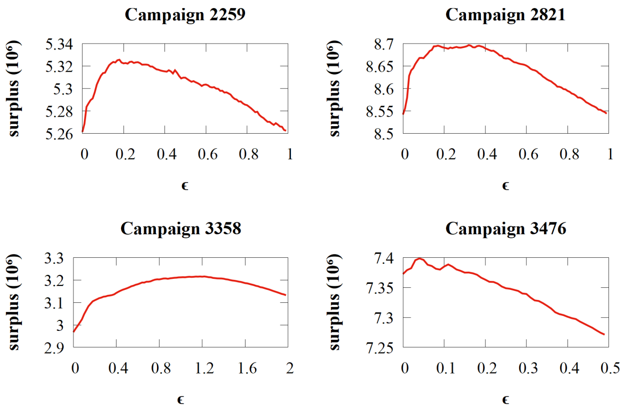

Figure 5: Relation between the upper bound parameter ϵ_0 and surplus on some campaigns

5.2.2 *Worst-case Performance.* Subsequently, we compare the performance of different strategies under the worst-case scenario. Here, the worst-case refers to the most adverse situation within the known uncertainty, corresponding to the case where $\hat{w}_i = \hat{r}_i$ for $i \in \mathcal{L}$. We set the winning price in the auctions lost by the DSP in the training set to infinity, and test the surplus of different models under this worst-case training set. The results are shown in Table 2. It can be observed that our robust approach exhibits the best worst-case performance across all campaigns.

Table 2: Worst-case Surplus of Different Strategies (10^7)

Camp.	STM	MCN	KMMN	NPM	RBC
1458	5.952	5.996	6.107	6.324	**6.626**
2259	1.145	1.264	1.304	1.277	**1.425**
2261	0.973	1.076	1.103	1.117	**1.215**
2821	1.947	2.074	2.125	2.137	**2.310**
2997	0.479	0.558	0.529	0.540	**0.570**
3358	2.156	2.273	2.267	2.391	**2.584**
3386	5.223	5.262	5.475	5.679	**5.967**
3427	3.704	3.815	3.951	4.100	**4.339**
3476	2.862	2.907	3.002	3.085	**3.271**
Criteo	41.05	39.70	40.15	41.68	**41.72**

5.3 Performance of RBDS Strategy

5.3.1 *Relationship between Upper Bound Parameter and Surplus.* We first give a certain relationship between the upper bound ϵ_0 and the surplus on the test set. The results on the test set of some campaigns with KMMN+R strategy are shown in Figure 5. We can observe that with the increase of ϵ_0, the surplus on the test set basically increases first and then decreases. But in different campaigns, the specific shape of the relation curve between surplus and ϵ_0 is different. In practice, we combine this insight with binary search to find the maximum surplus that can be obtained on the test set by adjusting ϵ_0.

5.3.2 *Overall Performance.* In this section, we report on the upper bound of surplus attainable through the adjustment of ϵ_0 in the RBDS strategy. Taking RBC as an example, we compare the surplus generated by the RBC strategy with the upper bound of surplus that can be obtained through the RBC+R strategy, as illustrated in Table 3. It can be observed that RBDS has delivered a notable increase in surplus for DSP in parts of campaigns. At this juncture, a question similar to that regarding RBC may naturally arise: why does the employment of the robust strategy RBDS lead to an enhancement in surplus? This is primarily due to the divergence in the winning price distribution between the training and testing datasets, which leads to the model derived from the training set failing to accurately predict the winning price distribution on the testing set. This may result in an increased surplus when an appropriate robustness parameter ϵ_0 is selected.

Table 3: Surplus of RBC and RBC+R Strategies (10^6)

Camp.	RBC	RBC+R
1458	12.31	12.40
2259	5.590	5.596
2261	6.152	6.161
2821	9.095	9.175
2997	3.005	3.008
3358	3.451	3.484
3386	10.42	10.44
3427	8.486	8.490
3476	8.051	8.130
Criteo	109.3	110.5

5.3.3 *Worst-case Performance.* To verify the robustness of our RBDS strategy, we simulate the case where the real winning price distribution is different from the predicted distribution, and then compare the expected surplus of the SO strategy and our RBDS strategy. In this context, the SO strategy refers to obtaining optimal bids by solving problem (5), whereas RBDS strategy derives optimal bids by solving problem (12). Specifically, we assume that the value of DSP is evenly distributed on [50, 200]. We take the overall winning price distribution on some campaigns as the predicted distribution $\boldsymbol{p}_0$, and construct the worst-case distributions as the real distributions subject to different Wasserstein distances. For the RBDS strategy, we set the upper bound $\epsilon_0 = 0.02$. Under these settings, we calculate the expected surplus for SO and the worst-case expected surplus for RBDS in some of the campaigns, and the results are shown in Figure 6. We can observe that the performance of SO strategy is generally better than RBDS strategy when the distance is small. However, as the distance increases, our RBDS strategy gradually outperforms the SO strategy. Compared to SO strategy, it is easy to find that the surplus of our RBDS strategy decreases less as the distance increases, which means that the surplus of RBDS is less affected by the distribution shift in the worst-case situations. This verifies the robustness of our RBDS strategy.

5.3.4 *Further Discussion of RBDS strategy.* To better understand the robustness of RBDS strategy, we specifically construct an example to analyze how the SO and RBDS strategies bid for a given

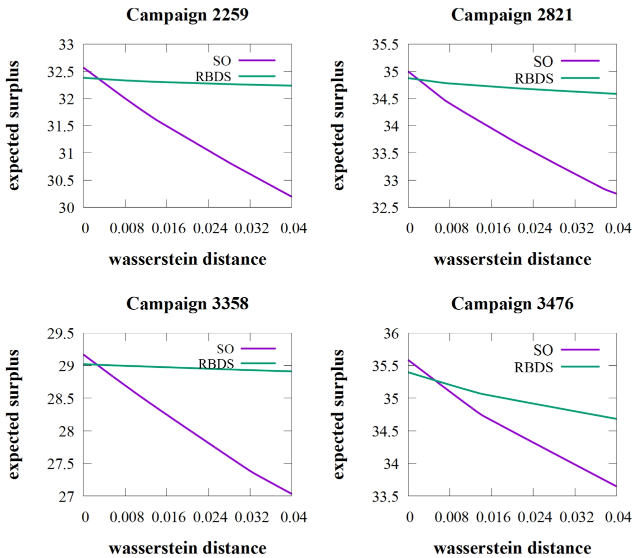

Figure 6: Worst-case performance of RBDS on some campaigns

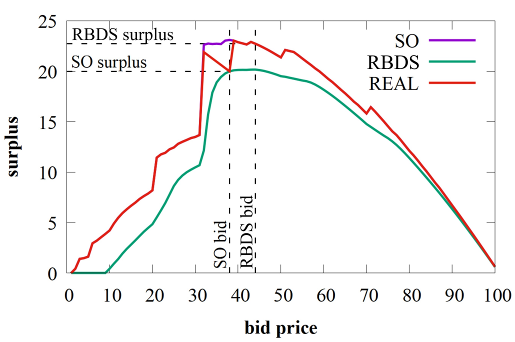

Figure 7: An example for the bid selection process in SO and RBDS strategies

value v. We use the overall winning price distribution on the 2821 campaign as the predicted distribution of SO problem (5) and RBDS problem (12) respectively. Assuming that the Wasserstein distance between the real (worst-case) distribution and the predicted distribution, as well as the upper bound ϵ_0 in RBDS are both 0.2, and we set the DSP's value $v = 101$. At this point, for different bids, the predicted surplus considered by the SO strategy, the worst-case surplus considered by the RBDS strategy, and the real surplus are shown in Figure 7. Among them, the real surplus is corresponding to the specially constructed worst-case distribution, and it overlaps with the predicted surplus for most bids, as shown in Figure 7. For better illustration, we mark the bid prices selected by the SO and RBDS strategies with dashed lines, which are the values corresponding to the highest points in the two curves. Based on these selected bids, we further mark out the surplus of these two strategies under the real distribution to give an intuitive comparison.

As can be observed from Figure 7, there is a sharp increase in the predicted surplus when the bid price is around 30, which means that the probability of winning price at this point is very high, forming a spike in the probability distribution. In this case, compared with the steep predicted surplus curve of SO strategy, RBDS strategy considers a smoother worst-case surplus curve in order to prevent this spike from moving within a small range. Therefore the bid of RBDS strategy is farther from the spike than SO strategy. In the worst case, the probability around the spike is shifted, and RBDS can deal with this situation more robustly and obtain a better surplus.

6 CONCLUSION

In this work, we model the uncertain environment inherent in the design of auto-bidding strategies within the context of bid shading, and propose two levels robust bidding strategies to achieve better surplus in such environments with considerable uncertainty. The experimental results on public dataset validate the effectiveness and robustness of our robust bidding strategies.

Since this work is the first to consider the uncertainty issue in bid shading, there are still many aspects to explore in both theory and experiment. For instance, one could attempt using the chance-constrained [3] surplus instead of worst-case surplus as the optimization objective, introduce techniques of advanced conditional density estimation [23] into RBC, and validate the results of robust strategies in real bidding environments, etc. We leave these promising directions to future works.

ACKNOWLEDGMENTS

This work was supported in part by National Key R&D Program of China (No. 2022ZD0119100), in part by China NSF grant No. 62322206, 62132018, U2268204, 62025204, 62272307, 62372296. The opinions, findings, conclusions, and recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agencies or the government.

REFERENCES

[1] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. 2009. *Robust Optimization.* Princeton University Press.
[2] Jonathan Buckley and Ian James. 1979. Linear regression with censored data. *Biometrika* 66, 3 (1979), 429–436.
[3] A. Charnes and W. W. Cooper. 1959. Chance-Constrained Programming. *Management science* 6, 1 (1959), 73–79.
[4] Erick Delage and Yinyu Ye. 2010. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations research* 58, 3 (2010), 595–612.
[5] Erick Delage and Yinyu Ye. 2010. Wiesemann, Wolfram and Kuhn, Daniel and Sim, Melvyn. *Operations research* 62, 6 (2010), 1358–1376.
[6] Eustache Diemert, Julien Meynet, Pierre Galland, and Damien Lefortier. 2017. Attribution Modeling Increases Efficiency of Bidding in Display Advertising. In *Proceedings of ADKDD.* 1–6.
[7] Aritra Ghosh, Saayan Mitra, Somdeb Sarkhel, Jason Xie, Gang Wu, and Viswanathan Swaminathan. 2020. Scalable Bid Landscape Forecasting in Real-Time Bidding. In *Machine Learning and Knowledge Discovery in Databases.* 451–466.
[8] Djordje Gligorijevic, Tian Zhou, Bharatbhushan Shetty, Brendan Kitts, Shengjun Pan, Junwei Pan, and Aaron Flores. 2020. Bid Shading in The Brave New World of First-Price Auctions. In *Proceedings of CIKM.* 2453–2460.
[9] Mert Gürbüzbalaban, Andrzej Ruszczyński, and Landi Zhu. 2022. A Stochastic Subgradient Method for Distributionally Robust Non-convex and Non-smooth Learning. *Journal of Optimization Theory and Applications* 194, 3 (2022), 1014–1041.
[10] Jikai Jin, Bohang Zhang, Haiyang Wang, and Liwei Wang. 2021. Non-convex Distributionally Robust Optimization: Non-asymptotic Analysis. *Advances in Neural Information Processing Systems* 34 (2021), 2771–2782.

[11] Niklas Karlsson and Qian Sang. 2021. Adaptive Bid Shading Optimization of First-Price Ad Inventory. In *2021 American Control Conference (ACC)*. 4983–4990.
[12] Çağıl Koçyiğit, Garud Iyengar, Daniel Kuhn, and Wolfram Wiesemann. 2020. Distributionally Robust Mechanism Design. *Management Science* 66, 1 (2020), 159–189.
[13] Daniel Kuhn, Peyman Mohajerin Esfahani, Viet Anh Nguyen, and Soroosh Shafieezadeh-Abadeh. 2019. Wasserstein Distributionally Robust Optimization: Theory and Applications in Machine Learning. arXiv:1908.08729
[14] Xu Li, Michelle Ma Zhang, Zhenya Wang, and Youjun Tong. 2022. Arbitrary Distribution Modeling with Censorship in Real-Time Bidding Advertising. In *Proceedings of SIGKDD*. 3250–3258.
[15] Hairen Liao, Lingxiao Peng, Zhenchuan Liu, and Xuehua Shen. 2014. IPinYou Global RTB Bidding Algorithm Competition Dataset. In *Proceedings of the Eighth International Workshop on Data Mining for Online Advertising*. 1–6.
[16] Fengming Lin, Xiaolei Fang, and Zheming Gao. 2022. Distributionally Robust Optimization: A review on theory and applications. *Numerical Algebra, Control and Optimization* 12, 1 (2022), 159–212.
[17] Roger B Myerson. 1981. Optimal auction design. *Mathematics of operations research* 6, 1 (1981), 58–73.
[18] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani. 2007. *Algorithmic Game Theory*. Cambridge University Press.
[19] Weitong Ou, Bo Chen, Yingxuan Yang, Xinyi Dai, Weiwen Liu, Weinan Zhang, Ruiming Tang, and Yong Yu. 2023. Deep Landscape Forecasting in Multi-Slot Real-Time Bidding. In *Proceedings of SIGKDD*. 4685–4695.
[20] Shengjun Pan, Brendan Kitts, Tian Zhou, Hao He, Bharatbhushan Shetty, Aaron Flores, Djordje Gligorijevic, Junwei Pan, Tingyu Mao, San Gultekin, and Jianlong Zhang. 2020. Bid Shading by Win-Rate Estimation and Surplus Maximization. arXiv:2009.09259
[21] Hamed Rahimian and Sanjay Mehrotra. 2022. Frameworks and results in distributionally robust optimization. *Open Journal of Mathematical Optimization* 3 (2022), 1–85.
[22] Kan Ren, Jiarui Qin, Lei Zheng, Zhengyu Yang, Weinan Zhang, and Yong Yu. 2019. Deep Landscape Forecasting for Real-Time Bidding Advertising. In *Proceedings of SIGKDD*. 363–372.
[23] Jonas Rothfuss, Fabio Ferreira, Simon Walther, and Maxim Ulrich. 2019. Conditional Density Estimation with Neural Networks: Best Practices and Benchmarks. arXiv:1903.00954
[24] Herbert E Scarf, KJ Arrow, and S Karlin. 1957. *A min-max solution of an inventory problem*. Rand Corporation Santa Monica.
[25] SS Vallender. 1974. Calculation of the Wasserstein Distance Between Probability Distributions on the Line. *Theory of Probability & Its Applications* 18, 4 (1974), 784–786.
[26] Jun Wang, Weinan Zhang, Shuai Yuan, et al. 2017. Display advertising with real-time bidding (RTB) and behavioural targeting. *Foundations and Trends® in Information Retrieval* 11, 4-5 (2017), 297–435.
[27] Ping Wang, Yan Li, and Chandan K. Reddy. 2019. Machine Learning for Survival Analysis: A Survey. *ACM Comput. Surv.* 51, 6 (2019), 36 pages.
[28] Tengyun Wang, Haizhi Yang, Siyu Jiang, Yueyue Shi, Qianyu Li, Xiaoli Tang, Han Yu, and Hengjie Song. 2022. Kaplan–Meier Markov network: Learning the distribution of market price by censored data in online advertising. *Know.-Based Syst.* 251 (2022), 11 pages.
[29] Tengyun Wang, Haizhi Yang, Han Yu, Wenjun Zhou, Yang Liu, and Hengjie Song. 2019. A revenue-maximizing bidding strategy for demand-side platforms. *IEEE Access* 7 (2019), 68692–68706.
[30] Yuchen Wang, Kan Ren, Weinan Zhang, Jun Wang, and Yong Yu. 2016. Functional Bid Landscape Forecasting for Display Advertising. In *Machine Learning and Knowledge Discovery in Databases*. 115–131.
[31] Wush Wu, Mi-Yen Yeh, and Ming-Syan Chen. 2018. Deep Censored Learning of the Winning Price in the Real Time Bidding. In *Proceedings of SIGKDD*. 2526–2535.
[32] Wush Chi-Hsuan Wu, Mi-Yen Yeh, and Ming-Syan Chen. 2015. Predicting Winning Price in Real Time Bidding with Censored Data. In *Proceedings of SIGKDD*. 1305–1314.
[33] Yong Yuan, Feiyue Wang, Juanjuan Li, and Rui Qin. 2014. A survey on real time bidding advertising. In *Proceedings of 2014 IEEE International Conference on Service Operations and Logistics, and Informatics*. 418–423.
[34] Wei Zhang, Brendan Kitts, Yanjun Han, Zhengyuan Zhou, Tingyu Mao, Hao He, Shengjun Pan, Aaron Flores, San Gultekin, and Tsachy Weissman. 2021. MEOW: A Space-Efficient Nonparametric Bid Shading Algorithm. In *Proceedings of SIGKDD*. 3928–3936.
[35] Tian Zhou, Hao He, Shengjun Pan, Niklas Karlsson, Bharatbhushan Shetty, Brendan Kitts, Djordje Gligorijevic, San Gultekin, Tingyu Mao, Junwei Pan, Jianlong Zhang, and Aaron Flores. 2021. An Efficient Deep Distribution Network for Bid Shading in First-Price Auctions. In *Proceedings of SIGKDD*. 3996–4004.
[36] Shen Zuhe, Huane Zhen Yu, and MA Wolfe. 1997. An Interval Maximum Entropy Method for a Discrete Minimax Problem. *Applied mathematics and computation* 87, 1 (1997), 49–68.

A PROOFS

A.1 Proof of Remark 4.1

We firstly show that for any $b \in \mathcal{B}$:

$$\min_{\boldsymbol{p} \in \mathcal{P}_{kn}} P_{\boldsymbol{p}}(b) = P_{\hat{\boldsymbol{p}}_c}(b), \tag{14}$$

where $\mathcal{P}_{kn}$ and $\hat{\boldsymbol{p}}_c$ are defined in (4) and (6). Since $\mathcal{B}$ is finite:

(a) If $b = 0$, $P_{\boldsymbol{p}}(b) = 0$ for any $\boldsymbol{p}$, thus (14) holds.

(b) If $b \neq 0$, we denote $b = b_J \in \mathcal{B}$. For rigor, we formally provide the definition of the probability vector $\boldsymbol{p}$ in the paper. For a probability vector constructed from a winning price dataset $\boldsymbol{p} = \mathcal{F}_{\mathcal{B}}(\{w_i | i \in \mathcal{A}\})$, the j-th element p^j is defined as:

$$p^j = \frac{1}{|\mathcal{A}|} \sum_{i \in \mathcal{A}} \mathbb{I}\{b_{j-1} \le w_i < b_j\}, \tag{15}$$

where $|\mathcal{A}|$ denotes the size of set $\mathcal{A}$. Then the cumulative distribution function (CDF) value for any b_J is:

$$P_{\boldsymbol{p}}(b_J) = \sum_{j=1}^{J} p_j = \frac{1}{|\mathcal{A}|} \sum_{i \in \mathcal{A}} \mathbb{I}\{w_i < b_j\}. \tag{16}$$

Defined in (4), $\mathcal{P}_{kn}$ contains all $\boldsymbol{p} = \mathcal{F}_{\mathcal{B}}(\{w_i | i \in \mathcal{A}\})$ that satisfies $w_i \in \mathcal{T}_i$, where $\mathcal{T}_i = \{\hat{w}_i\}$ for $i \in \mathcal{W}$ and $\mathcal{T}_i = \{w | \hat{l}_i \le w \le \hat{r}_i\}$ for $i \in \mathcal{L}$. Hence the left-hand side in (14) equals to:

$$\begin{aligned} \min_{\boldsymbol{p} \in \mathcal{P}_{kn}} P_{\boldsymbol{p}}(b_J) &= \min_{\{w_i \in \mathcal{T}_i | i \in \mathcal{A}\}} \frac{1}{|\mathcal{A}|} \sum_{i \in \mathcal{A}} \mathbb{I}\{w_i < b_J\} \\ &= \frac{1}{|\mathcal{A}|} \sum_{i \in \mathcal{A}} \min_{w_i \in \mathcal{T}_i} \mathbb{I}\{w_i < b_J\} \\ &= \frac{1}{|\mathcal{A}|} \sum_{i \in \mathcal{W}} \min_{w_i \in \mathcal{T}_i} \mathbb{I}\{w_i < b_J\} + \\ &\quad \frac{1}{|\mathcal{A}|} \sum_{i \in \mathcal{L}} \min_{w_i \in \mathcal{T}_i} \mathbb{I}\{w_i < b_J\} \end{aligned} \tag{17}$$

For $i \in \mathcal{W}$, $\mathcal{T}_i = \{\hat{w}_i\}$ and $\min_{w_i \in \mathcal{T}_i} \mathbb{I}\{w_i < b_J\} = \mathbb{I}\{\hat{w}_i < b_J\}$; for $i \in \mathcal{L}$, $\mathcal{T}_i = \{w | \hat{l}_i \le w \le \hat{r}_i\}$ and $\min_{w_i \in \mathcal{T}_i} \mathbb{I}\{w_i < b_J\}) = \mathbb{I}\{\hat{r}_i < b_J\}$. Hence from (17) we have:

$$\min_{\boldsymbol{p} \in \mathcal{P}_{kn}} P_{\boldsymbol{p}}(b_J) = \frac{1}{|\mathcal{A}|}(\sum_{i \in \mathcal{W}} \mathbb{I}\{\hat{w}_i < b_J\} + \sum_{i \in \mathcal{L}} \mathbb{I}\{\hat{r}_i < b_J\}), \tag{18}$$

which is exactly $P_{\hat{\boldsymbol{p}}_c}(b_J)$ for $\hat{\boldsymbol{p}}_c$ defined in (6), thus (14) holds.

We have thus proven that equation (14) holds for any $b \in \mathcal{B}$. Now, we further provide the proof of Remark 4.1. Since problems (2) and (3) are equivalent, we only need to prove that given v, problem (3) with ambiguity set (4) is equivalent to problem (5) with distribution (6). According to (14), the problem (3) with ambiguity set (4) can be written as:

$$\begin{aligned} \max_{b \in \mathcal{B}} \min_{\boldsymbol{p} \in \mathcal{P}_{kn}} (v - b) \cdot P_{\boldsymbol{p}}(b) &= \max_{b \in \mathcal{B}} \{(v - b) \min_{\boldsymbol{p} \in \mathcal{P}_{kn}} \{P_{\boldsymbol{p}}(b)\}\} \\ &= \max_{b \in \mathcal{B}} (v - b) \cdot P_{\hat{\boldsymbol{p}}_c}(b), \end{aligned} \tag{19}$$

which is equivalent to problem (5) with distribution (6). Hence, the conclusion of Remark 4.1 is established.

A.2 Proof of Remark 4.2

For convenience, we introduce some additional notations. Firstly, since both $\boldsymbol{p}_0$ and ϵ_0 are given values, we will use $\mathcal{P}_{un}$ to denote $\mathcal{P}_{un}(\boldsymbol{p}_0, \epsilon_0)$ in the subsequent proof. Besides, we additionally appended an element of zero to the end of the probability vector, that is, $\boldsymbol{p}_0 = (p_0^1, p_0^2, ..., p_0^M, p_0^{M+1} = 0)$. This addition does not alter the probability distribution of $\boldsymbol{p}_0$, hence it does not affect the correctness of our conclusions. Next, we will prove that Algorithm 1 obtains the optimal solution to problem (12) on domain $\mathcal{B}$.

Firstly, since $\mathcal{B}$ is a finite set, we only need to prove that for any $b \in \mathcal{B}/\{0\}$, Algorithm 1 can provide the accurate value of $f_{\mathcal{P}_{un}}(b)$ (clearly $f_{\mathcal{P}_{un}}(0) = 0$). Therefore, in the subsequent proof, we fix b and denote it as $b = b_J \in \mathcal{B}$, where $J \in [1, M]$.

Next, for the problem of $f_{\mathcal{P}_{un}}(b_J) = \min_{\boldsymbol{p} \in \mathcal{P}_{un}} P_{\boldsymbol{p}}(b_J)$, we provide the form of the probability vector that could attains the optimal value, that is, the form of $\boldsymbol{p}_* \in \mathcal{P}_{un}$ that could satisfies $P_{\boldsymbol{p}_*}(b_J) = \min_{\boldsymbol{p} \in \mathcal{P}_{un}} P_{\boldsymbol{p}}(b_J)$ for any $\boldsymbol{p}_0$, ϵ_0 and J. We subsequently refer to $\boldsymbol{p}_*$ as the "worst-case form".

We show that for any $\boldsymbol{p}_0$, ϵ_0 and J, there exists some $t \in [1, J]$ and $p^t \le p_0^t$ such that the worst-case probability vector $\boldsymbol{p}_*$ can be constructed as:

- $p_*^t = p^t$
- $p_*^k = 0, \forall k \in [t+1, J]$ (if $t < J$)
- $p_*^{J+1} = p_0^{J+1} - p^t + \sum_{k=t}^{J} p_0^k$
- $p_*^k = p_0^k, \forall k \notin [t, J+1]$ (that is, for the rest element of $\boldsymbol{p}_*$)

This is consistent with the construction results considered in Fig. 3 of our paper. For the proof of correctness regarding the worst-case form, we conduct a classification discussion:

(a) If for $t = 1$ and $p^t = 0$ we have $d(\boldsymbol{p}_0, \boldsymbol{p}_*) \le \epsilon_0$, then since $P_{\boldsymbol{p}_*}(b_J) = \sum_{k=1}^{J} p_*^k = 0$, it attains the optimal value (since the probability value should not be negative).

(b) If for $t = 1$ and $p^t = 0$ we have $d(\boldsymbol{p}_0, \boldsymbol{p}_*) > \epsilon_0$:

(b.1) We first need to demonstrate that there exists $t = t_0$ and $p^t = p^{t_0}$ such that the worst-case form can satisfy $d(\boldsymbol{p}_0, \boldsymbol{p}_*) = \epsilon_0$. This is because the distribution distance can be represented as:

$$d(\boldsymbol{p}_0, \boldsymbol{p}_*) = (p_0^t - p^t)(J+1-t) + \sum_{k=t+1}^{J} p_0^k (J+1-k), \tag{20}$$

in which the summation term equals 0 if $t = J$. Note that when $t = J$ and $p^t = p_0^t$, we have $\boldsymbol{p}_* = \boldsymbol{p}_0$ such that $d(\boldsymbol{p}_0, \boldsymbol{p}_*) = 0$. Besides, this expression is continuous with respect to t and p^t. Therefore, there must exist some t_0 and p^{t_0} such that $d(\boldsymbol{p}_0, \boldsymbol{p}_*) = \epsilon_0$.

(b.2) Fix t_0 and p^{t_0}, among all $\boldsymbol{p}$ within the constraint $d(\boldsymbol{p}_0, \boldsymbol{p}) \le \epsilon_0$, we show that $\boldsymbol{p}_*$ attains the optimal value of $P_{\boldsymbol{p}}(b_J)$. Note that:

$$P_{\boldsymbol{p}}(b_J) = \sum_{k=1}^{J} p^k, \tag{21}$$

we show that if there is some $\hat{\boldsymbol{p}}$ such that $P_{\hat{\boldsymbol{p}}}(b_J) < P_{\boldsymbol{p}_*}(b_J)$, we have $d(\boldsymbol{p}_0, \hat{\boldsymbol{p}}) > \epsilon_0$ such that $\hat{\boldsymbol{p}} \notin \mathcal{P}_{un}$. We first construct $\hat{\boldsymbol{p}}_*$ from $\hat{\boldsymbol{p}}$ that satisfies:

- $\hat{p}_*^k = \hat{p}^k, \forall k \in [1, J]$
- $\hat{p}_*^{J+1} = p_0^{J+1} + \sum_{k=J+1}^{M+1} (\hat{p}^k - p_0^k)$
- $\hat{p}_*^k = p_0^k, \forall k \in [J+2, M+1]$ (if $J < M$)

We have $P_{\hat{\boldsymbol{p}}_*}(b_J) = P_{\hat{\boldsymbol{p}}}(b_J)$ and $d(\boldsymbol{p}_0, \hat{\boldsymbol{p}}) \ge d(\boldsymbol{p}_0, \hat{\boldsymbol{p}}_*)$. The proof of the latter result, although intuitive, is somewhat verbose. Intuitively, for the parts with indexes greater than J, the probability of $\hat{\boldsymbol{p}}$ being different from $\boldsymbol{p}_0$ is concentrated at $\hat{p}_*^{J+1}$. This reduces the probability transfer within the indexes greater than J and decreases the distance for the remaining parts to transfer probabilities to this portion. Here, for brevity, we skip the intricate proof.

Hence we only needs to prove that $d(\boldsymbol{p}_0, \hat{\boldsymbol{p}}_*) > \epsilon_0$. Note that $P_{\hat{\boldsymbol{p}}_*}(b_J) = P_{\hat{\boldsymbol{p}}}(b_J) < P_{\boldsymbol{p}_*}(b_J)$, and $\hat{p}_*^k = p_*^k = p_0^k$ for all $k \in [J+2, M+1]$ (if $J < M$), we have $\hat{p}_*^{J+1} > p_*^{J+1}$. Assuming that the optimal transfer quantities in problem (11) are $\hat{\boldsymbol{d}}$ and $\boldsymbol{d}$ for $\hat{\boldsymbol{p}}_*$ and $\boldsymbol{p}_*$ respectively. Then according to the definition of the Wasserstein distance, by expanding the expression of the Wasserstein distance between distributions $\boldsymbol{p}_0$ and $\hat{\boldsymbol{p}}_*$, we can obtain:

$$\begin{aligned} d(\boldsymbol{p}_0, \hat{\boldsymbol{p}}_*) &= \sum_{i=1}^{M+1} \sum_{j=1}^{M+1} \hat{d}_{ij} |i-j| \\ &\ge \sum_{i=1}^{J} \hat{d}_{i(J+1)} (J+1-i) \\ &\ge \sum_{i=t+1}^{J} \hat{d}_{i(J+1)} (J+1-i) + \sum_{i=1}^{t} \hat{d}_{i(J+1)} (J+1-t), \end{aligned} \tag{22}$$

where the first $\ge$ hold since we only consider a subset of the transfer quantities, that is, the probabilities transferred from indexes in [1,J] to index J+1, and the second $\ge$ hold for $J+1-i \ge (J+1-t), \forall i \le t$. For the first term, according to the definition of $\hat{\boldsymbol{p}}_*$, we can expand it and obtain:

$$\begin{aligned} &\sum_{i=t+1}^{J} \hat{d}_{i(J+1)} (J+1-i) \\ &= \sum_{i=t+1}^{J} \hat{p}_0^i (J+1-i) + \sum_{i=t+1}^{J} (\hat{d}_{i(J+1)} - p_0^i)(J+1-i) \\ &\ge \sum_{i=t+1}^{J} \hat{p}_0^i (J+1-i) + \sum_{i=t+1}^{J} (\hat{d}_{i(J+1)} - p_0^i)(J+1-t), \end{aligned} \tag{23}$$

where the $\ge$ holds for $\hat{d}_{i(J+1)} - p_0^i \le 0, \forall i$. Organizing the derivation of (22) and (23), We now obtain a lower bound for the Wasserstein distance between distributions $\boldsymbol{p}_0$ and $\hat{\boldsymbol{p}}_*$:

$$d(\boldsymbol{p}_0, \hat{\boldsymbol{p}}_*) \ge \sum_{i=t+1}^{J} p_0^i (J+1-i) + (\sum_{i=1}^{J} \hat{d}_{i(J+1)} - \sum_{i=t+1}^{J} p_0^i)(J+1-t), \tag{24}$$

where the remaining transfer quantities to consider is only those from indexes in [1,J] to index J+1, and our goal is to compare the size of the right-hand expression in (24) with that of $d(\boldsymbol{p}_0, \boldsymbol{p}_*)$. Note that we only need to consider the sum of these transfer quantities, which is easy to compare:

$$\sum_{i=1}^{J} \hat{d}_{i(J+1)} = \hat{p}_*^{J+1} - p_0^{J+1} > p_*^{J+1} - p_0^{J+1} = \sum_{i=t}^{J} d_{i(J+1)}, \tag{25}$$

and according to the definition of $\boldsymbol{p}_*$, we have $d_{i(J+1)} = p_0^i$ for any $i \in [t+1, J]$. We now can further derive the right-hand expression in inequality (24) as follows:

$$
\begin{aligned}
&\sum_{i=t+1}^{J} p_0^i (J+1-i) + (\sum_{i=1}^{J} \hat{d}_{i(J+1)} - \sum_{i=t+1}^{J} p_0^i)(J+1-t) \\
&> \sum_{i=t+1}^{J} p_0^i (J+1-i) + (\sum_{i=t}^{J} d_{i(J+1)} - \sum_{i=t+1}^{J} p_0^i)(J+1-t) \\
&= \sum_{i=t}^{J} d_{i(J+1)} (J+1-i) \\
&= d(\boldsymbol{p}_0, \boldsymbol{p}_*) \\
&= \epsilon_0,
\end{aligned}
\tag{26}
$$

where the penultimate equation is because only when $j = J+1$ and $i \in [t, J]$ we have $d_{ij} \neq 0$. Combining (24) and (26) we have $d(\boldsymbol{p}_0, \hat{\boldsymbol{p}}_*) > \epsilon_0$. Hence the "worst-case form" $\boldsymbol{p}_*$ attains the optimal value of $P_{\boldsymbol{p}}(b_J)$ in case (b).

Combining (a) and (b), we have proven that for any $\boldsymbol{p}_0$, ϵ_0 and J, the "worst-case form" $\boldsymbol{p}_*$ satisfies $P_{\boldsymbol{p}_*}(b_J) = \min_{\boldsymbol{p} \in \mathcal{P}_{un}} P_{\boldsymbol{p}}(b_J)$.

Finally, we can prove the correctness of Remark 4.2 by demonstrating that for each b_J, Algorithm 1 provides the correct t_0 and p^{t_0} in the previous discussion.

For t_0, note that at each end of the step in the "while" loop, $\boldsymbol{q}_0$ satisfies the "worst-case form" with $t = k$ and $p^t = 0$, and the loop stops for the first time $d(\boldsymbol{p}_0, \boldsymbol{q}_0) > \epsilon$, hence $t_0 = k$ upon exiting the "while" loop because $d(\boldsymbol{p}_0, \boldsymbol{q}_0)$ increases as $t = k$ decreases.

For p^{t_0}, since we have got t_0, we can calculate its value through the Wasserstein distance, as shown in line 9 in Algorithm 1.

Therefore, at this point, we have the accurate value of $f_{\mathcal{P}_{un}}(b_J) = \sum_{k=1}^{t_0-1} q_0^k + p^{t_0}$, as shown in line 10 in Algorithm 1 (note that $q_0^k = 0$ for $k \in [t_0+1, J]$).

Thus, Algorithm 1 obtains the optimal solution to problem (12) on domain $\mathcal{B}$, and the conclusion in Remark 4.2 have been proven.