

A ROBUST MULTI-ITEM AUCTION DESIGN WITH STATISTICAL LEARNING

Jiale Han
UCLA
jialehan@ucla.edu

Xiaowu Dai *
UCLA
daix@ucla.edu

ABSTRACT

We propose a novel statistical learning method for multi-item auctions that incorporates credible intervals. Our approach employs nonparametric density estimation to estimate credible intervals for bidder types based on historical data. We introduce two new strategies that leverage these credible intervals to reduce the time cost of implementing auctions. The first strategy screens potential winners' value regions within the credible intervals, while the second strategy simplifies the type distribution when the length of the interval is below a threshold value. These strategies are easy to implement and ensure fairness, dominant-strategy incentive compatibility, and dominant-strategy individual rationality with a high probability, while simultaneously reducing implementation costs. We demonstrate the effectiveness of our strategies using the Vickrey-Clarke-Groves mechanism and evaluate their performance through simulation experiments. Our results show that the proposed strategies consistently outperform alternative methods, achieving both revenue maximization and cost reduction objectives.

1 INTRODUCTION

Designing optimal-revenue auctions is a fundamental problem in economics and computer science, with applications in a wide range of fields such as e-commerce, online advertising, and spectrum auctions. Although Myerson's celebrated work (Myerson 1981) provided an optimal design for single-item auctions by assuming knowing bidders' value distributions, the problem becomes increasingly complex in multi-item scenarios. Additionally, in practical settings, bidders' value distributions are private, leaving the seller with limited information. Typically, sellers only have access to historical bidding data before an auction. Moreover, the process of querying bidders before an auction can be time-consuming, especially when the number of bidders is large.

To address the challenges in multi-item mechanism design, various methods have been proposed, including Alaei (2014), Cai & Daskalakis (2011), Alaei et al. (2012), and Cai & Huang (2013). However, these approaches often assume knowing bidder types distributions, which may not be available in practice. Recently, there has been some progress in developing mechanisms that are robust to errors in the distribution of bidder types, such as Cai & Daskalakis (2017), Brustle et al. (2020), and Cai & Daskalakis (2022). Nevertheless, there remains a lack of practical algorithms that effectively utilize data for estimating bidders' types and reducing the time cost of implementing auctions, without relying on prior knowledge of bidders' types. In this paper, we propose a novel statistical learning method for optimal multi-item auctions that leverages historical data and credible intervals. Our approach aims to minimize the computational complexity and time cost of implementing auctions, providing a more efficient and practical solution for multi-item auctions.

Our proposed method leverages nonparametric density estimation to estimate credible intervals for bidder types using their historical bidding data. We introduce two new strategies that leverage these credible intervals to reduce the time cost of implementing auctions. First, we propose a filtering method that screens potential winners by considering their value regions within the credible intervals generated by the estimated distribution. This approach ensures fairness while selecting winners. Second, we classify intervals based on their length and simplify the distribution of the type when the interval length of its credible interval falls below a threshold value.

*Corresponding author.

To demonstrate the efficacy of our proposed strategies, we implement these two strategies within the Vickrey-Clarke-Groves (VCG) mechanism, with items being sold separately. This has resulted in a new mechanism design that guarantees fairness, dominant-strategy incentive compatibility (DSIC), and dominant-strategy individual rationality (DSIR) with a high probability. We conducted simulations using both small-scale and large-scale data to evaluate the effectiveness of our proposed strategies within the VCG mechanism. Our simulation results demonstrate that our strategies can efficiently reduce the implementation costs, as evidenced by a significant reduction in the number of queries required compared to the original VCG mechanism, while maintaining a comparable level of revenue. Overall, our proposed strategies offer an efficient data preprocessing approach for implementing multi-item auctions.

Related works. The seminal work of Myerson (Myerson 1981) laid the foundation for optimal auction design in single-item settings, and the revelation principle in the paper suggests that achieving optimal design requires a round of direct communication in which bidders simultaneously and confidentially announce their value estimates to the seller. Therefore, to design a mechanism that maximizes revenue, the seller must make some queries to the bidders. However, the extension of this framework to multi-item scenarios poses significant challenges, as counterintuitive properties may arise that do not present in single-item settings (Daskalakis 2015). The complexity can become intractable as the number of items increases, as demonstrated in Dughmi et al. (2014) and Babaioff et al. (2017). In recent years, various strategies have been proposed to alleviate these analytical and computational difficulties for multi-item auctions, including Alaei (2014), Cai & Daskalakis (2011), Alaei et al. (2012), Cai & Huang (2013), Cai et al. (2012a,b), Chawla et al. (2007, 2010), Cai et al. (2013), Morgenstern & Roughgarden (2016), Goldner & Karlin (2016). However, the majority of these methods require knowing the true distribution of bidder types, which may not be available in practice. The presence of private and unknown bidder types has motivated research on robust mechanism design. Studies in this area focus on robustifying mechanisms to errors in the distribution of bidder types, without relying on Bayesian assumptions. Notable examples include Cai & Daskalakis (2017), Cai & Daskalakis (2022), Brustle et al. (2020) and Bergemann & Schlag (2011), which have all aimed to design mechanisms that achieve optimal revenue with only approximated distributions of bidder types. In contrast to these existing works, our approach does not rely on the approximated distribution of bidder types. Instead, we utilize estimated credible intervals for bidder types to design mechanisms that effectively reduce the time cost of implementing auctions and maximize revenue. Han & Dai (2025) also studies mechanism design without prior knowledge of the values, but focuses on the single-item setting. Our method is designed to be efficient and easy to implement, providing a practical solution for multi-item auction applications.

2 PRELIMINARIES

In our auction setting, a seller aims to sell $N \geq 1$ items to a group of $m > 1$ bidders, denoted as $\mathcal{B} = \{B_1, B_2, \dots, B_m\}$. Each bidder B_i has a private valuation for the N items, which is captured by their type $t_i = (t_{i,1}, t_{i,2}, \dots, t_{i,N}) \in \mathbb{R}^{1 \times N}$, and $i \in \{1, \dots, m\} \equiv [m]$. To accommodate the different preferences of bidders, we assume that each element $t_{i,j}$ is drawn from a private distribution D_i^j . We denote the set of all distributions as $D = \otimes_{i,j} D_i^j$, where the mN type distributions are independent, but not necessarily identical.

Valuation. For the set of items $S \subset [N]$, we assume a bidder B_i 's value for S is given by the additive valuation function as Goldner & Karlin (2016),

$$v_i(t_i, S) = \sum_{j \in S} t_{i,j}.$$

Auction Mechanism. An auction mechanism M is a mapping from the space of all the bidder types $t = (t_1, t_2, \dots, t_m)^T \in \mathbb{R}^{m \times N}$ to the space of allocations $Q(t) = (q_i^j(t))$ and payment $P(t) = (p_1(t), p_2(t), \dots, p_m(t))$, where $q_i^j(t) \in [0, 1]$ denotes the probability the items j is allocated to the bidder B_i and the $p_i(t) \in (-\infty, \infty)$ denotes the payment for bidder B_i . Let $p(t) = \sum_{j \in [m]} p_j(t)$ be the total payments received by the seller. Since the bidder types follow the distribution D , the goal is to design a mechanism for the seller that maximum her expected revenue,

$$\text{REV}(M, D) \equiv \mathbb{E}_{t \sim D} [p(t)].$$

Given a multi-item auction mechanism M and all bidders' reporting bids $b \in \mathbb{R}^{m \times N}$, the utility of a bidder B_i with private type t_i is formulated as

$$u_i(t_i, M(b)) = \sum_{j=1}^N t_{i,j} q_i^j(b) - p_i(b).$$

DSIC. To ensure truthfulness, we impose an incentive compatibility constraint. We adopt the definition of dominant-strategy incentive compatibility (DSIC); see, e.g., Cai et al. (2016).

Definition 1 (DSIC) An auction mechanism M is DSIC if, for every bidder B_i , $i \in [m]$, the utility of truthfully reporting type t_i is no less than the utility of reporting an alternative type t'_i . That is,

$$u_i(t_i, M(t_i, t_{-i})) \geq u_i(t'_i, M(t'_i, t_{-i})).$$

DSIR. Next, we introduce a generalization of the dominant-strategy individual rationality (DSIR), denoted as δ -DSIR, which is different from the traditional definitions in the literature (e.g., Yao 2014). The introduction of δ -DSIR is motivated by our proposed mechanism in Section 3, which utilizes a method of simplifying the distribution of bidder types based on the lower bound of their corresponding credible intervals.

Definition 2 (δ -DSIR) A mechanism M is δ -individual rationality (δ -DSIR) if for all types $t = (t_1, t_2, \dots, t_m)$ and for all bidder $i \in [m]$,

$$\mathbb{P}\{u_i(t_i, M(t)) \geq 0\} \geq 1 - \delta.$$

A mechanism is DSIR if $\delta = 0$, i.e. for all types $t = (t_1, t_2, \dots, t_m)$ and for all bidder B_i , $i \in [m]$, we have $u_i(t_i, M(t)) \geq 0$ almost surely.

VCG Mechanism. In this paper, we will implement our methods upon the foundation of the Single-Stage Vickrey-Clarke-Groves (VCG) mechanism or the 2nd-price Vickrey mechanism in our situation. The VCG mechanism operates by allocating the bundle of items in accordance with the bidder who expresses the highest valuation for the bundle and charges them the value of the second highest bid. Mathematically, let $b_{i,j}$ be the reporting bid of bidder B_j to item j . The winner of the bundle S is determined as

$$i_* := \arg \max_i \sum_{j \in [S]} b_{i,j},$$

and the winner is subsequently charged with $\max_{i, i \neq i_*} \sum_{j \in [S]} b_{i,j}$. It is well-known that truthful reporting is a dominant strategy for each bidder in a VCG mechanism (e.g., Ausubel et al. 2006) and VCG mechanism has also been widely recognized for its desirable properties of DSIR and DSIC (Vickrey 1961, Clarke 1971, Che & Yoo 2001). This implies that, when using a VCG mechanism and soliciting bids from bidders, it can be expected that they will rationally provide truthful valuations, i.e. $b_{i,j} = t_{i,j}$ for any bidder B_i and item j .

The single-stage VCG mechanism simplifies cost measurement in implementation, requiring only that the seller make queries to bidders. Cost reduction can be achieved by strategically minimizing the number of queries.

3 MAIN RESULTS

In this section, we present the main results of this paper. We begin by introducing a method for estimating the distributions and constructing credible intervals for bidder types in Section 3.1. Next, we propose a novel strategy for reducing the number of bids considered in an auction in Section 3.2. We then provide a mechanism that simplifies the distributions of the bidder types in Section 3.3. Finally, in Section 3.4, we implement the VCG mechanism based on the strategy in Section 3.3 and give the lower bound of its revenue.

We consider the most challenging scenario where we lack any prior knowledge about the distributions of bidder types, and the bidder types distributions can be different for a certain item. Additionally, we assume that the distributions of the bidder types are fixed throughout the bidding process.

3.1 ESTIMATION OF DISTRIBUTION AND CREDIBLE INTERVALS FOR BIDDER TYPES

In this subsection, we make the assumption that each bidder has participated in numerous bidding instances with independent bids coming from her true value’s distribution for every item, generating a certain amount of historical data available for analysis.

Let the dataset of previous bids for bidder B_i on item j be denoted as $\Gamma_{i,j}$, which consists of $n_{i,j}$ independent and identically distributed samples, $\Gamma_{i,j} \equiv \{x_{i,j}^1, x_{i,j}^2, \dots, x_{i,j}^{n_{i,j}}\}$.

Kernel Density Estimate. With the data $\Gamma_{i,j}$, we can perform kernel density estimation of D_i^j using the function \widehat{D}_i^j , as defined by:

$$\widehat{D}_i^j(x) = \frac{1}{n_{i,j}h} \sum_{s=1}^{n_{i,j}} K\left(\frac{x - x_{i,j}^s}{h}\right).$$

Here, K represents a density function known as the kernel, and h is the bandwidth. This method, which dates back to the works of Rosenblatt (1956) and Parzen (1962), commonly utilizes a Gaussian kernel, $K(t) = \frac{1}{\sqrt{2\pi}}e^{-t^2/2}$. A choice of the bandwidth is $h = 1.06 \cdot \min\{\hat{\sigma}, \frac{IQR}{1.34}\} \cdot n_{i,j}^{-1/5}$, where $\hat{\sigma}$ is the sample standard deviation, $n_{i,j}$ denotes the sample size, and IQR represents the sample interquartile range. This choice of h is known to be robust against outliers (Scott 1992), and is minimax optimal for estimation (Zambom & Ronaldo 2013) with the mean integrated square error (MISE) of $\mathbb{E} \int [\widehat{D}_i^j(x) - D_i^j(x)]^2 dx = O(n_{i,j}^{-4/5})$.

Rejection Sampling: In order to sample from the estimated density, we can utilize the rejection sampling method. This involves selecting random samples from a given distribution, $g(x)$, and choosing a constant c such that

$$c \cdot g(x) > \widehat{D}_i^j(x), \quad \forall x.$$

Then, we accept the sample x with a probability of $\widehat{D}_i^j(x)/cg(x)$, resulting in a sampling process equivalent to that from $\widehat{D}_i^j(x)$ as described in Casella et al. (2004). In our scenario, for each item $j \in [N]$, we can define the support of the density \widehat{D}_i^j to be within the known minimum and maximum values, a_j and b_j , respectively, which may be set by the seller or obtained from the market. Thus, we can set $g(x)$ as a uniform distribution over the interval $[a_j, b_j]$.

We can repeat this sampling process \mathcal{N} times, and compute the $(\alpha/2)$ th and $(1 - \alpha/2)$ th quantiles of the samples as the lower and upper bounds of the $(1 - \alpha)$ credible interval of $t_{i,j}$. As the MISE of \widehat{D}_i^j is $O(n_{i,j}^{-4/5})$, we can assert that when $n_{i,j}$ is large enough, the difference between \widehat{D}_i^j and D_i^j is negligible, resulting in high precision of the credible interval. We provide a detailed outline of this procedure in Algorithm 1, which is deferred to Section C.1.

3.2 WINNOW DOWN POTENTIAL WINNERS

We present a novel strategy that effectively reduces the number of bids considered in an auction while maintaining fairness in the game based on credible intervals. This approach narrows down the pool of potential winners, leading to improved efficiency in the auction process. We begin by defining the fairness of the mechanism.

Definition 3 (δ -fairness) *Given a set of items $S \subset [N]$ and the winner of the bundle S , B_{i^*} , a mechanism is δ -fair if for any item j in S , with probability at least $1 - \delta$, for all $i \in [m], i \neq i^*$, it holds true that $t_{i,j} < t_{i^*,j}$. A mechanism is fair if $\delta = 0$.*

This definition draws inspiration from Joseph et al. (2016), in which fairness is defined through pairwise comparison. In other words, a fair mechanism assigns higher probabilities of winning to bidders with higher values. In our scenario, we only need to consider the winner, as the outcome is unique. Based on the definition of δ -fairness, it is clear that a smaller δ value corresponds to a fairer game.

Assume that we have already obtained the $(1 - \alpha)$ credible interval $T_{i,j} = [t_{i,j}^L, t_{i,j}^U]$ for each bidder B_i ’s type for item j , where $t_{i,j}^L$ and $t_{i,j}^U$ is the $(\alpha/2)$ th and $(1 - \alpha/2)$ th quantiles of D_i^j . This can be

achieved by employing the methodology outlined in Section 3.1, which leverages historical data to gain the interval. Then with probability more than $(1 - \alpha)$, the true value $t_{i,j}$ is within the interval $[t_{i,j}^L, t_{i,j}^U]$. For each item j , we identify the bidder $B_{i_*^j}$ with the largest *upper credible interval* among all bidders,

$$i_*^j = \arg \max_{i \in [m]} t_{i,j}^U. \quad (1)$$

This bidder has the potential to have the highest bid and is therefore more likely to win the item. However, other bidders whose estimated intervals overlap with $B_{i_*^j}$'s may also have a chance to win. To account for this scenario, we define a new interval relationship. For each item j , we consider bidders B_i and $B_{i'}$ to be *linked* if

$$[t_{i,j}^L, t_{i,j}^U] \cap (t_{i',j}^L, t_{i',j}^U] \neq \emptyset. \quad (2)$$

We note that Joseph et al. (2016) introduces the concept of "chained relation" in addition to the "linked relation" defined above. Two bidders are considered chained if they belong to the same component of the transitive closure of the linked relation. The authors then focus on all the arms chained to the arm with the highest upper confidence bound and disregard the others. Our approach is similar, but we utilize the linked relation instead of the chained relation, as only one selection is made in our setting and only one arm needs to be pulled.

In order to mitigate the computational complexity associated with determining the winner of a multi-item auction, we propose a mechanism that selectively considers bids from a subset of bidders for each item. Specifically, for each item $j \in [N]$, we only consider the bids from bidders that are linked to the bidder $B_{i_*^j}$, as defined by

$$\mathcal{B}_j = \{B_i | i \in [m], B_i \text{ is linked to } B_{i_*^j} \text{ for item } j\}. \quad (3)$$

Here i_*^j and the linked relation are defined in Eqs. (1) and (2), respectively. The algorithm of constructing the set \mathcal{B}_j is given in Section C.2. We show that if just considering the bidders in set \mathcal{B}_j in an original fair mechanism, it is still fair enough.

Proposition 1 *Let \mathcal{B}_j denote the set of bidders linked to the bidder with the highest upper confidence bound for item j , as defined by (3). Then for each item $j \in [N]$, considering only bidders in the set \mathcal{B}_j to implement a fair mechanism results in a new mechanism that is α -fair.*

The proof of this proposition is given in Appendix A.1. The reduction in the number of bidders that need to be considered results in decreased computational complexity and time cost. This is because there is no longer a need to calculate the allocation and payment for these bidders when implementing any mechanism.

3.3 LOWER CREDIBLE BOUND METHOD

In this section, we present another strategy that utilizes credible intervals to reduce the cost of the seller to implement any mechanism with good incentive properties, such as the VCG mechanism that satisfies the DSIC and DSIR.

Regarding robust auction mechanism design, since the bidder types are private, obtaining the true set of distributions D is not possible, the goal is to ensure that the distance between the estimated distribution \widehat{D} and the true distribution D is minimized so that implement any mechanism on \widehat{D} will also gain comparable revenue. Mathematically, as motivated by the characterization of the Prokhorov Metric (e.g., Strassen 1965), we want to ensure that there exists a coupling γ of D and \widehat{D} , such that

$$\mathbb{P}_{(\hat{t}, t) \sim \gamma} [\|\hat{t} - t\|_\infty > \epsilon] \leq \delta_0, \quad (4)$$

for some small values of $\epsilon, \delta_0 > 0$.

For any bidder $B_i, i \in [m]$ and item $j \in [N]$, $T_{i,j} = [t_{i,j}^L, t_{i,j}^U]$ is the $(1 - \alpha)$ credible interval for $t_{i,j}$, as described in Sections 3.2. Thus $\mathbb{P}(t_{i,j}^L \leq t_{i,j} \leq t_{i,j}^U) = 1 - \alpha$. Additionally, we define $d_{i,j} = t_{i,j}^U - t_{i,j}^L$ as the length of the interval. Then, we have

$$\begin{aligned} \mathbb{P}(t_{i,j}^L - d_{i,j} \leq t_{i,j} \leq t_{i,j}^L + d_{i,j}) &= \mathbb{P}(-d_{i,j} \leq t_{i,j} - t_{i,j}^L \leq d_{i,j}) \\ &= \mathbb{P}(|t_{i,j}^L - t_{i,j}| \leq d_{i,j}) \geq 1 - \alpha, \end{aligned}$$

where $t_{i,j}$ is a random variable and $t_{i,j}^L$ is a fixed value.

If we choose \widehat{D}_i^j to be the one-point distribution, denoted by $\mathbb{P}(\hat{t}_{i,j} = t_{i,j}^L) = 1$ for all i and j , and define the norm $\|\hat{t}\|_\infty \equiv \max_{i,j} \hat{t}_{i,j}$, then due to the independence of all mN type distributions, we can deduce that:

$$\begin{aligned} \mathbb{P}_{(\hat{t},t) \sim (\widehat{D},D)} \left[\|\hat{t} - t\|_\infty \leq \max_{i,j} d_{i,j} \right] &= \prod_{i \in [m], j \in [N]} \mathbb{P}(|t_{i,j}^L - t_{i,j}| \leq \max_{i,j} d_{i,j}) \\ &\geq \prod_{i \in [m], j \in [N]} \mathbb{P}(|t_{i,j}^L - t_{i,j}| \leq d_{i,j}) \geq (1 - \alpha)^{mN}, \end{aligned}$$

so it satisfies Eq. (4) with $\epsilon = \max_{i,j} d_{i,j}$ and $\delta_0 = 1 - (1 - \alpha)^{mN} \leq mN\alpha$. Thus, the length of the intervals $d_{i,j}$ becomes a crucial factor in determining the accuracy of the degenerated \widehat{D}_i^j as one point $t_{i,j}^L$.

By simplifying all \widehat{D}_i^j to one-point distributions, we can significantly reduce the computational complexity and time cost involved. This is because we eliminate the need to deal with complex probability distributions when computing the reserve price (e.g. implementing the optimal auction design for a single item, Myerson 1981) or designing a mechanism. However, it is important to acknowledge that this approach may lead to a considerable reduction in revenue, particularly if certain credible intervals have a large bandwidth. Therefore, it becomes necessary to carefully manage the set of bids utilizing the one-point distribution in order to minimize complexity while still maintaining revenue reductions within acceptable tolerances.

To deal with this trade-off dilemma, for every bidder $B_i, i \in [m]$, we first define two types of credible intervals based on a specified threshold value $d > 0$. Let

$$Q_i = \{j | t_{i,j}^U - t_{i,j}^L > d, j \in [N]\}, \quad L_i = \{j | t_{i,j}^U - t_{i,j}^L \leq d, j \in [N]\}. \quad (5)$$

We let the distribution of the types for items in L_i degenerate to one-point distribution, i.e. the distribution of the types in $L' \equiv \{t_{i,j} | i \in [m], j \in L_i\}$ will be replaced by a one-point distribution. Then let

$$Q = \{t_{i,j} | i \in [m], j \in Q_i\}, \quad L = \{\hat{t}_{i,j} | i \in [m], j \in L_i\}. \quad (6)$$

be the sets of types that the seller will use to do analysis, where $t_{i,j} \sim D_i^j$ and $\hat{t}_{i,j} \sim [\mathbb{P}(\hat{t}_{i,j} = t_{i,j}^L) = 1]$ are all random variables. We can let the degenerated type distribution be \widehat{D} and $\hat{t} = Q \cup L$.

Let $|L_i| \equiv n_i$ so that $|Q_i| = N - n_i$, and

$$n \equiv \sum_{i \in [m]} n_i = |L|, \quad |Q| = mN - n \quad (7)$$

The following theorem demonstrates the effectiveness of our strategy in maintaining both the DSIC and δ -DSIR properties for any DSIC and DSIR mechanism. However, our mechanism is specifically trained using a simpler distribution \widehat{D} and has query access limited to the types in Q . This approach reduces the computational complexity and time cost involved.

Theorem 1 *Suppose we implement a DSIR mechanism \widehat{M} w.r.t. \widehat{D} and get the allocation $(\hat{q}_i^j(\hat{t}))$ and payment rule $(\hat{p}_i(\hat{t}))$. We can construct a mechanism M using only query access to \widehat{M} (which equals to only use query access to the types in Q), such that the mechanism M is δ -DSIR with $\delta = \alpha n/2$ w.r.t. D . If \widehat{M} is also DSIC, and we assume that a change in type for a specific item only affects the allocation and payment of that item, then M is DSIC w.r.t. D as well.*

The proof of this theorem is given in Appendix A.2. Theorem 1 implies that our proposed mechanism is practically viable. The only challenge is to keep αn as small as possible to ensure that bidders are willing to participate in the auction. To make α small, we can use historical data or sample more extensively from the estimated distribution to obtain more precise credible intervals for t . Additionally, we can control n by carefully choosing the value of d .

3.4 DEGENERATED VCG MECHANISM

The VCG mechanism is both a DSIC and DSIR mechanism, and it meets the definition of fairness as outlined in Definition 3. Moreover, when there is a change in type for a specific item, it only affects the allocation and payment for that particular item. Therefore, the VCG mechanism satisfies all the assumptions stated in Proposition 1 and Theorem 1. Consequently, we can utilize the VCG mechanism to validate the properties of our strategies in Section 3.2 and Section 3.3. In this section, our attention will be directed towards the method discussed in Section 3.3.

To make the validation process easier, we consider that the seller sells each item separately, which gives a practical solution to alleviate the computational complexity associated with determining the allocation of items in a bundle. As an example, we calculate the lower bound of the complexity of allocation when each bundle has at most 2 items in Appendix B and show the complexity of the allocation grows exponentially as a function of the number of items.

Based lower credible bound method in Section 3.3, we can give a degenerated VCG mechanism, this approach involves keeping the types in set L fixed while only querying the types in Q . For any bidder $B_i, i \in [m]$, we have the set of estimated types L_i , which is defined in Eq. (5). Let $L_i^* := \{j | t_{i,j}^U - t_{i,j}^L > 0, j \in [N]\}$, and

$$K = \cup_{i \in [m]} (L_i \cap L_i^*). \quad (8)$$

Let $k = |K|$ be the number of items in set K . The revenue reduction is limited to the items in K since only the distributions of bidder B_i 's valuation for items in K are altered and degenerate into one-point distributions. We have the following result.

Theorem 2 *Let \widehat{M} be the VCG mechanism implemented on \widehat{D} constructed in Section 3.3, and M_0 be the VCG mechanism implemented directly on D . M is the degenerated VCG mechanism using only query access to \widehat{M} . Then with probability $(1 - \alpha/2)^n$,*

$$REV(M, D) \geq REV(M_0, D) - kd.$$

The proof of this theorem is given in Appendix A.3. We observe that when the number of items, N , is limited, it is likely that the number of items, k , in the set K in Eq. (8) will also be limited due to the constraint $k < N$. Consequently, kd will also be relatively small, indicating the potential effectiveness of the proposed mechanism when N is small. Moreover, our proposed mechanism holds a distinct advantage over the robust mechanism design described in Cai & Daskalakis (2022) in situations where the number of bidders, m , is exceptionally large. While their revenue lower bound demonstrates a negative correlation with m , our revenue lower bound remains largely unaffected by m . This is due to the fact that k is strictly smaller than N , and d is a hyper-parameter that is typically reduced to control the number of n as m increases. Therefore, our mechanism is particularly well-suited for environments with a high number of bidders and a limited number of items.

Theorem 2 shows that the performance of the revised VCG mechanism based on Section 3.3 depends on kd . The value of d is a hyper-parameter that is chosen based on the specific characteristics of the data set at hand. The density function, and subsequently the length of the confidence intervals, can vary from dataset to dataset, thus requiring the selection of an appropriate value of d . In Appendix C.3, we give the algorithms of the implemented VCG mechanism and the details of how to explore the appropriate range of values for d .

4 NUMERICAL EXPERIMENTS

In this section, we present our experimental results. More details and results are presented in the supplementary materials appendix D. We utilize a truncated re-normalized Gaussian distribution to model the private valuations of each bidder, employing distinct means and variances. Once the distribution D is obtained, bidder types t can be generated from D , enabling the simulation of historical data sets $\Gamma_{i,j}$ as described in Section 3.1. To validate the effectiveness of our proposed methods, we compare three revised VCG Mechanisms with the original VCG Mechanism.

Mechanism 1 (M1) employs the method outlined in Section 3.1 to obtain credible intervals, followed by the utilization of the approach described in Section 3.2 to narrow down the dataset. Additionally,

the method explained in Section 3.3 is used to exclude queries to types whose distributions have degenerated into one-point distributions. In comparison to M1, Mechanism 2 (M2) does not employ the winnow-down method, while Mechanism 3 (M3) constructs a credible interval using the minimum and maximum values from the original data, without utilizing the distribution estimation method. The benchmark is the revenue generated by implementing the original VCG mechanism without employing any methods from our paper.

The *regret* is defined as the discrepancy in revenue between the original VCG mechanism and the modified VCG mechanism.




Figure 1: Comparison of the revenue and regret when $m = 30$, $N = 10$.

Figure 1 illustrates a comparison of revenue and regret when the number of items m is 30 and when the number of items N is 10. The revenue of the original VCG mechanism is represented by a green dashed line at the top of Figure 1(a) and serves as a benchmark. Mechanism 1 is depicted by a red solid line, Mechanism 2 by a black dashed line, and Mechanism 3 by a blue long dash line. To further analyze the results, we present a modified version of Figure 1(a) in Figure 1(b), focusing on regret. In Figure 1(b), each line’s position represents the discrepancy with the green line, allowing for a clearer understanding of the regret associated with each mechanism.

It can be observed that the revenue of M1 and M2 overlap across all values of d , indicating that the winnow-down method described in Section 3.2 does not have an impact on the revenue. This finding validates Proposition 1, which asserts that the winnow-down method preserves fairness in the auction, as none of the winners are omitted. Moreover, it is worth noting that the revenue generated by M3 consistently remains lower than that of M1 and M2. This demonstrates the advantage of utilizing the method described in Section 3.1 to obtain credible intervals. Additionally, we conducted a verification of the interval coverage and found that all of the bidders’ bids fall within our constructed credible intervals. This implies that utilizing the lower bound as the estimated type still ensures the mechanism is DSIR, which aligns with the findings of Theorem 1 in Section 3.3 and further affirms the feasibility of employing the method outlined in Section 3.1. Furthermore, upon comparing Figure 1(b) and Figure 1(c), it can be observed that the practical regret consistently remains smaller than the upper bound of the theoretical regret, which is kd . This observation serves to validate the result presented in Theorem 2. Further refinement is possible. Currently, when calculating the upper bound of regret theoretically, the revenue for items in set K defined in Eq. (8) is reduced. In practice, a more precise method would be to sort the intervals’ bandwidths and deduce the maximum bandwidth for each item. Specifically, we can define

$$d_j = \max_i \{t_{i,j}^U - t_{i,j}^L | i \in [m], 0 < t_{i,j}^U - t_{i,j}^L \leq d\}.$$

By utilizing this approach, the expected revenue derived from the degenerated distribution \hat{D} would be lower than that based on the true distribution D by a maximum of $\sum_{j \in K} d_j$, which is actually smaller than kd . This would yield a more accurate result than Theorem 2.

To further compare M1 and M2, Figure 2(a) presents a comparison of the proportion of bidder types without queries. It can be observed that after implementing the winnow-down method, approximately half of the queries are eliminated. When examined alongside Figure 1, it becomes apparent that these eliminated queries are redundant, underscoring the significant advantage of the winnow-down method in Section 3.2.

Figure 2(b) illustrates the comparison of the confidence rate $(1 - \alpha/2)^n$, where n represents the total number of degenerated types distributions. According to Theorem 2, we can only guarantee that the

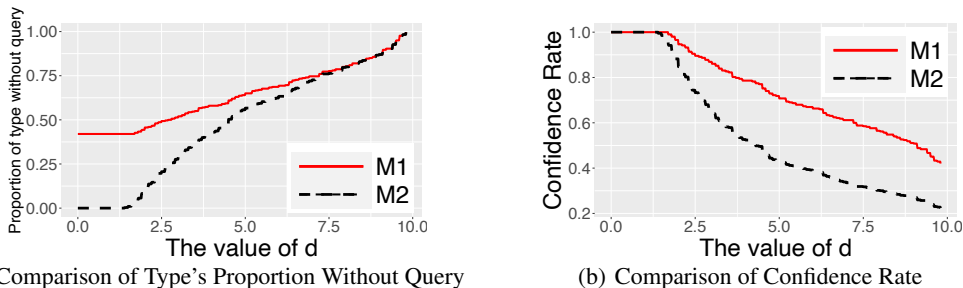


Figure 2: Comparison of the proportion of bidder types without queries and comparison of the confidence rate when $m = 30$, $N = 10$.

regret is smaller than kd with a probability of $(1 - \alpha/2)^n$. We demonstrate in Section C.3.1 that, when the confidence rate $\eta = 0.9$ and $\alpha = 0.01$, at most $n = 21$ types distributions can degenerate to achieve a high level of precision. Nevertheless, our experimental results demonstrate that in practice, a greater number of distributions can be degenerated without significantly affecting the revenue. As observed in Figure 2(b), even with a low confidence rate, the obtained regret in Figure 1(b) is still smaller than the upper bound kd shown in Figure 1(c). This finding provides support for the effectiveness of our proposed method, which employs the lower bound to construct one-point distributions. Moreover, it underscores the robustness of our approach in handling variations in the confidence rate.

We also conducted experiments with different parameter settings, specifically when $m = 50$ and $N = 30$, as well as when $m = 50$ and $N = 100$. Remarkably, in all of the cases, we arrived at the same conclusions and observed consistent results. The analysis of each of them can be found in Appendix D of the supplementary material.

5 CONCLUSION

This paper proposes a statistical learning method that leverages credible intervals to effectively reduce computation complexity and time costs in auction implementations. Our method utilizes kernel density estimation with historical data to estimate credible intervals. We propose two easily implementable strategies that ensure fairness, dominant-strategy incentive compatibility, and dominant-strategy individual rationality, while also reducing implementation costs. Through simulation experiments using the Vickrey-Clarke-Groves mechanism, we demonstrate the effectiveness of our proposed strategies. We provide the lower bound of the expected revenue when using the lower credible bound method, highlighting the advantages of our approach. Our suggested mechanism offers a highly effective solution for sellers in multi-item auctions. There are several directions for future research. For example, it is of interest to explore distribution-free methods for obtaining credible intervals, and we are currently working on this direction.

REFERENCES

- Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014.
- Saeed Alaei, Hu Fu, Nima Haghpanah, Jason D Hartline, and Azarakhsh Malekian. Bayesian optimal auctions via multi-to single-agent reduction. In *13th ACM Conference on Electronic Commerce, EC’12*, pp. 17, 2012.
- Lawrence M Ausubel, Paul Milgrom, et al. The lovely but lonely vickrey auction. *Combinatorial auctions*, 17: 22–26, 2006.
- Moshe Babaioff, Yannai A Gonczarowski, and Noam Nisan. The menu-size complexity of revenue approximation. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 869–877, 2017.
- Dirk Bergemann and Karl Schlag. Robust monopoly pricing. *Journal of Economic Theory*, 146(6):2527–2543, 2011.

- David Alexander Brannan. *A first course in mathematical analysis*. Cambridge University Press, 2006.
- Johannes Brustle, Yang Cai, and Constantinos Daskalakis. Multi-item mechanisms without item-independence: Learnability via robustness. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pp. 715–761, 2020.
- Yang Cai and Constantinos Daskalakis. Extreme-value theorems for optimal multidimensional pricing. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pp. 522–531. IEEE, 2011.
- Yang Cai and Constantinos Daskalakis. Learning multi-item auctions with (or without) samples. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 516–527. IEEE, 2017.
- Yang Cai and Constantinos Daskalakis. Recommender systems meet mechanism design. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pp. 897–914, 2022.
- Yang Cai and Zhiyi Huang. Simple and nearly optimal multi-item auctions. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 564–577. SIAM, 2013.
- Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pp. 459–478, 2012a.
- Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pp. 130–139. IEEE, 2012b.
- Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. Understanding incentives: Mechanism design becomes algorithm design. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pp. 618–627. IEEE, 2013.
- Yang Cai, Nikhil R Devanur, and S Matthew Weinberg. A duality based unified approach to bayesian mechanism design. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pp. 926–939, 2016.
- George Casella, Christian P Robert, and Martin T Wells. Generalized accept-reject sampling schemes. *Lecture Notes-Monograph Series*, pp. 342–347, 2004.
- Shuchi Chawla, Jason D Hartline, and Robert Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, pp. 243–251, 2007.
- Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *42nd ACM Symposium on Theory of Computing, STOC 2010*, pp. 311–320, 2010.
- Yeon-Koo Che and Seung-Weon Yoo. Optimal incentives for teams. *American Economic Review*, 91(3):525–541, 2001.
- Edward H Clarke. Multipart pricing of public goods. *Public choice*, pp. 17–33, 1971.
- Constantinos Daskalakis. Multi-item auctions defying intuition? *ACM SIGecom Exchanges*, 14(1):41–75, 2015.
- Shaddin Dughmi, Li Han, and Noam Nisan. Sampling and representation complexity of revenue maximization. In *International Conference on Web and Internet Economics*, pp. 277–291. Springer, 2014.
- Kira Goldner and Anna R Karlin. A prior-independent revenue-maximizing auction for multiple additive bidders. In *International Conference on Web and Internet Economics*, pp. 160–173. Springer, 2016.
- Jiale Han and Xiaowu Dai. Online auction design using distribution-free uncertainty quantification with applications to e-commerce. *Journal of the American Statistical Association*, pp. 1–12, 2025.
- Matthew Joseph, Michael Kearns, Jamie H Morgenstern, and Aaron Roth. Fairness in learning: Classic and contextual bandits. *Advances in neural information processing systems*, 29, 2016.
- Jamie Morgenstern and Tim Roughgarden. Learning simple auctions. In *Conference on Learning Theory*, pp. 1298–1318. PMLR, 2016.
- Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- DW Scott. Multivariate density estimation. *Multivariate Density Estimation*, 1992.
- Volker Strassen. The existence of probability measures with given marginals. *The Annals of Mathematical Statistics*, 36(2):423–439, 1965.
- William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1): 8–37, 1961.
- Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pp. 92–109. SIAM, 2014.
- Adriano Z Zambom and Dias Ronaldo. A review of kernel density estimation with applications to econometrics. *International Econometric Review*, 5(1):20–42, 2013.

A PROOFS

In this section, we present the technical proofs for the main results outlined in our paper. We begin by introducing Lemma 1, which serves as a foundational component for our subsequent proofs. Following that, in Section A.1, we demonstrate the proof of Proposition 1, which validates the fairness of our method in the process of winnowing down potential winners. Moving on to Section A.2, we provide the proof of Theorem 1, establishing the incentive compatibility of our proposed mechanism. Finally, we delve into Section A.3 where we prove the lower bound of the revenue generated by the degenerated VCG mechanism as outlined in Theorem 2.

Lemma 1 (Bernoulli’s Inequality) *For any real number $x \geq -1$ and any natural number n , $(1 + x)^n \geq 1 + nx$.*

Proof A.1 (Proof of Lemma 1) *See Brannan (2006). A First Course in Mathematical Analysis. Cambridge University Press. p. 20.*

A.1 PROOF OF PROPOSITION 1

Proof A.2

For each item $j \in S \subset [N]$ and every bidder B_k , the revised mechanism ensures α -fairness as follows:

1. *If B_k is not in \mathcal{B}_j defined in Eq. (3), according to the definition of \mathcal{B}_j , we have $t_{k,j}^U \leq t_{i_*^j,j}^L$. Based on the definition of the credible intervals, with probability $1 - \alpha/2$, $t_{k,j}$ is smaller than $t_{k,j}^U$. Similarly, with probability $1 - \alpha/2$, the winner’s bid $t_{i_*^j,j}$ is bigger than $t_{i_*^j,j}^L$. Since the bidders’ type distributions are independent, we can write:*

$$\begin{aligned} \mathbb{P}(t_{k,j} < t_{i_*^j,j}) &\geq \mathbb{P}(t_{k,j} < t_{k,j}^U \leq t_{i_*^j,j}^L < t_{i_*^j,j}) \\ &= \mathbb{P}(t_{k,j} < t_{k,j}^U) \mathbb{P}(t_{i_*^j,j}^L < t_{i_*^j,j}) \\ &= (1 - \alpha/2)^2 \geq 1 - \alpha. \end{aligned}$$

The last inequality holds because $0 \leq \alpha \leq 1$ and Lemma 1.

2. *If B_k is in \mathcal{B}_j , because the original mechanism is fair, which means it satisfies the fairness definition 3 with $\delta = 0$, so that:*

$$\mathbb{P}(t_{k,j} < t_{i_*^j,j}) = 1 \geq 1 - \alpha.$$

In conclusion, the revised mechanism guarantees α -fairness.

A.2 PROOF OF THEOREM 1

Proof A.3 *First, we prove the δ -DSIR property of the mechanism. For any bid b , the utility of a bidder B_i with a private type t_i in mechanism \widehat{M} , with allocations $(\hat{q}_i^j(b))$ and payments $(\hat{p}_1(b), \dots, \hat{p}_m(b))$, is formulated as:*

$$u_i(t_i, \widehat{M}(b)) = \sum_{j=1}^N t_{i,j} \hat{q}_i^j(b) - \hat{p}_i(b).$$

Because \widehat{M} is DSIR with respect to \widehat{D} , it guarantees the following property for any private type \hat{t} sampled from the distribution \widehat{D} :

$$\begin{aligned} u_i(\hat{t}_i, \widehat{M}(\hat{t})) &= \sum_{j \in Q_i} t_{i,j} \hat{q}_i^j(\hat{t}) + \sum_{j \in L_i} \hat{t}_{i,j} \hat{q}_i^j(\hat{t}) - \hat{p}_i(\hat{t}) \\ &= \sum_{j \in Q_i} t_{i,j} \hat{q}_i^j(\hat{t}) + \sum_{j \in L_i} t_{i,j}^L \hat{q}_i^j(\hat{t}) - \hat{p}_i(\hat{t}) \\ &\geq 0. \end{aligned} \tag{C.1}$$

Since the mechanism M only relies on query access to \widehat{M} , it ensures that if the bidder's total type is t , \widehat{M} will solely make queries on types in set Q and eventually operate on \hat{t} . Moreover, M preserves the same allocation and payment functions as \widehat{M} after the querying process, denoted as $M(t) = \widehat{M}(\hat{t})$.

Given that $\mathbb{P}(t_{i,j} \geq t_{i,j}^L) = 1 - \alpha/2$, we can conclude that with a probability of $(1 - \alpha/2)^{n_i}$, for every j in L_i , $t_{i,j} \geq t_{i,j}^L$. Consequently, with a probability of $(1 - \alpha/2)^{n_i}$, for all types t sampled from the distribution \widehat{D} ,

$$\begin{aligned} u_i(t_i, M(t)) &= u_i(t_i, \widehat{M}(\hat{t})) \\ &= \sum_{j \in Q_i} t_{i,j} \hat{q}_i^j(\hat{t}) + \sum_{j \in L_i} t_{i,j} \hat{q}_i^j(\hat{t}) - \hat{p}_i(\hat{t}) \\ &\geq \sum_{j \in Q_i} t_{i,j} \hat{q}_i^j(\hat{t}) + \sum_{j \in L_i} t_{i,j}^L \hat{q}_i^j(\hat{t}) - \hat{p}_i(\hat{t}) \\ &= u_i(\hat{t}_i, \widehat{M}(\hat{t})) \geq 0. \end{aligned} \tag{C.2}$$

Then, for every i , we have

$$\mathbb{P}(u_i(t_i, M(t)) \geq 0) \geq \left(1 - \frac{\alpha}{2}\right)^{\sum_{i=1}^m n_i} = \left(1 - \frac{\alpha}{2}\right)^n \geq 1 - \frac{\alpha n}{2}.$$

The last inequality holds because $0 \leq \alpha \leq 1$ and Lemma 1. Therefore, our mechanism is M is DSIR w.r.t. D with $\delta = \frac{\alpha n}{2}$.

Next, we prove the DSIC property of the mechanism. Consider any bidder B_i with type distribution \widehat{D} . Only the types for items in Q_i can take different values due to the definition of the degenerated distribution \widehat{D} . Since the mechanism \widehat{M} is DSIC with respect to \widehat{D} , untruthful reporting of types by bidder B_i will result in a reduction in her utility in \widehat{M} . Thus, we have:

$$u_i(\hat{t}_i, \widehat{M}(\hat{t}_i, \hat{t}_{-i})) \geq u_i(\hat{t}'_i, \widehat{M}(\hat{t}'_i, \hat{t}_{-i})), \tag{C.3}$$

where the difference between \hat{t}_i and \hat{t}'_i occurs only in the items in Q_i .

Based on the assumption of the theorem, untruthful reporting of the types of items in Q_i will only influence the allocation and payment for the items in Q_i . Let $\hat{p}_i(\hat{t}) = \hat{p}_i^1(\hat{t}) + \dots + \hat{p}_i^N(\hat{t})$, where $\hat{p}_i^j(\hat{t})$ is the payment that bidder B_i ultimately pays for item j in the Mechanism \widehat{M} given the type $\hat{t} \sim \widehat{D}$. For any item j in L_i , we have $\hat{q}_i^j(\hat{t}_i, \hat{t}_{-i}) = \hat{q}_i^j(\hat{t}'_i, \hat{t}_{-i})$ and $\hat{p}_i^j(\hat{t}_i, \hat{t}_{-i}) = \hat{p}_i^j(\hat{t}'_i, \hat{t}_{-i})$.

Based on Eq.(C.1), we can rewrite Eq.(C.3) as follows:

$$\sum_{j \in Q_i} t_{i,j} \hat{q}_i^j(\hat{t}_i, \hat{t}_{-i}) - \sum_{j \in Q_i} \hat{p}_i^j(\hat{t}_i, \hat{t}_{-i}) \geq \sum_{j \in Q_i} t_{i,j} \hat{q}_i^j(\hat{t}'_i, \hat{t}_{-i}) - \sum_{j \in Q_i} \hat{p}_i^j(\hat{t}'_i, \hat{t}_{-i}).$$

This equation is equivalent to:

$$u_i(t_i, M(t_i, t_{-i})) \geq u_i(t_i, M(t'_i, t_{-i})),$$

based on Eq. (C.2). Hence, we can conclude that M is DSIC with respect to D .

A.3 PROOF OF THEOREM 2

Proof A.4 If the bidders have total types $t \sim D$, both \widehat{M} and M will generate the payment and allocation based on the revised types $\hat{t} \sim \widehat{D}$ as defined in Section 3.3, because M uses only query access to \widehat{M} . On the other hand, M_0 directly generates the payment and allocation based on t . The difference between t and \hat{t} only occurs for the types in the set L defined in Eq. (6).

Let $p_i(t) = p_i^1(t) + \dots + p_i^N(t)$, where $p_i^j(t)$ is the payment that bidder B_i ultimately pays for item j in Mechanism M_0 given the type $t \sim D$. Similarly, we can define $\hat{p}_i(\hat{t}) = \hat{p}_i^1(\hat{t}) + \dots + \hat{p}_i^N(\hat{t})$, where $\hat{p}_i^j(\hat{t})$ is the payment that bidder B_i ultimately pays for item j in Mechanism \widehat{M} given the type $\hat{t} \sim \widehat{D}$.

In the original VCG mechanism M_0 , the payment for item j is defined as $P_j(t) \equiv \sum_{i=1}^m p_i^j(t)$. This payment is determined as the second highest bid among $t_{i,j}$ for all $i \in [m]$. In mechanism M or \widehat{M} , the payment for item j is defined as $P_j(\hat{t}) \equiv \sum_{i=1}^m \hat{p}_i^j(\hat{t}) = \sum_{i=1}^m p_i^j(\hat{t})$. This payment is determined as the second highest bid among $\hat{t}_{i,j}$ for all $i \in [m]$.

For every $i \in [m]$, $\mathbb{P}(\hat{t}_{i,j} \geq t_{i,j} - d) = (1 - \alpha/2)$ holds for each item $j \in K$ as defined in Eq. (8). Let $n_{j,1}$ be the number of elements in the set $\{i | 0 < t_{i,j}^U - t_{i,j}^L \leq d, i \in [m]\}$, given by:

$$n_{j,1} = \left| \{i | 0 < t_{i,j}^U - t_{i,j}^L \leq d, i \in [m]\} \right|.$$

1. If item $j \in K$, with probability $(1 - \alpha/2)^{n_{j,1}}$, $\hat{t}_{i,j} \geq t_{i,j} - d$ holds for all $i \in [m]$. This implies:

$$P_j(\hat{t}) > P_j(t) - d.$$

2. When $j \notin K$, for any $i \in [m]$, if $t_{i,j}^U - t_{i,j}^L = 0$, then with probability $(1 - \alpha)$, $\hat{t}_{i,j} > t_{i,j}$. If $t_{i,j}^U - t_{i,j}^L \neq 0$, then $t_{i,j}$ is in Q , so $\hat{t}_{i,j} = t_{i,j}$. Let $n_{j,2}$ be the number of elements in the set $\{i | t_{i,j}^U - t_{i,j}^L = 0, i \in [m]\}$, given by:

$$n_{j,2} = \left| \{i | t_{i,j}^U - t_{i,j}^L = 0, i \in [m]\} \right|.$$

Then with probability $(1 - \alpha/2)^{n_{j,2}}$, we have:

$$P_j(\hat{t}) \geq P_j(t).$$

Altogether, we have that with probability $(1 - \alpha/2)^{\sum_{j=1}^N (n_{j,1} + n_{j,2})} = (1 - \alpha/2)^n$, the following inequality holds:

$$\begin{aligned} \sum_{j \in [N]} P_j(\hat{t}) &\geq \sum_{j \in [N]} P_j(t) - \sum_{j \in K} d \\ &\geq \sum_{j \in [N]} P_j(t) - kd. \end{aligned} \tag{C.4}$$

By taking expectations on both sides of Eq. (C.4), we have:

$$\mathbb{E}_{\hat{t} \sim \widehat{D}} \left[\sum_{j \in [N]} P_j(\hat{t}) \right] \geq \mathbb{E}_{t \sim D} \left[\sum_{j \in [N]} P_j(t) \right] - kd, \tag{C.5}$$

where we used the fact that only the term $\sum_{j \in [N]} P_j(\hat{t})$ is random over the change of \hat{t} , and the term $\sum_{j \in [N]} P_j(t)$ is random over the change of t .

Because for any bidders' total bid b , we have:

$$\sum_{j \in [N]} P_j(b) = \sum_{j=1}^N \sum_{i=1}^m p_i^j(b) = \sum_{i=1}^m \sum_{j=1}^N p_i^j(b) = \sum_{i=1}^m p_i(b),$$

we can rewrite Eq. (C.5) as:

$$\mathbb{E}_{\hat{t} \sim \widehat{D}} \left[\sum_{i=1}^m p_i(\hat{t}) \right] \geq \mathbb{E}_{t \sim D} \left[\sum_{i=1}^m p_i(t) \right] - kd,$$

which can be further simplified to:

$$REV(M, D) \geq REV(M_0, D) - kd.$$

B ALLOCATION COMPLEXITY ANALYSIS

In this section, we establish the lower bound for the number of bundle selections in scenarios where each bound is constrained to a maximum of 2 items. This analysis serves to emphasize the inherent intractability of the allocation problem, particularly when confronted with a considerable number of items.

Proposition 2 *When each bundle has at most 2 items, the number of bundle selections grows exponentially as a function of the number of items N . Specifically, the number of potential bundle selections exceeds $(2N/3)^{\lfloor N/3 \rfloor}$.*

Proof B.1 *When each bundle just has one item, We only have 1 bundle selection rule:*

$$\{\{1\}, \{2\}, \{3\}, \dots, \{N\}\}.$$

When at least a bundle has 2 items, We use $\lfloor N/2 \rfloor$ to denote the maximum integer which less than or equal to $N/2$. Then if we have i ($i \leq \lfloor N/2 \rfloor$) 2-item bundles in total, the number of bundle selections is

$$\frac{\binom{N}{2i} \binom{2i}{2} \binom{2i-2}{2} \cdots \binom{4}{2}}{i!} = \frac{N!}{(N-2i)!i!2^i}.$$

When $N \geq 3i$, $N - i \geq 2i$, we have $(N - 2i + 1)(N - 2i) \cdots (N - i) \geq 1 \times 2 \times \cdots \times 2i$, and $N - i \leq \frac{2N}{3}$, so

$$\begin{aligned} \frac{N!}{(N-2i)!i!2^i} &= \frac{(N-2i+1)(N-2i) \cdots (N)}{1 \times 2 \times \cdots \times i \times 2^i} \\ &= \frac{(N-2i+1)(N-2i) \cdots (N)}{2 \times 4 \times \cdots \times 2i} \\ &\geq (N-i+1)(N-i+2) \cdots N \\ &\geq (N-i)^i \geq \left(\frac{2N}{3}\right)^i. \end{aligned} \tag{C.6}$$

The total number of bundle selections is the sum of the number of bundle selections for i ($i \leq \lfloor N/2 \rfloor$), which is

$$\begin{aligned} \sum_{i=1}^{\lfloor N/2 \rfloor} \frac{N!}{(N-2i)!i!2^i} &\geq \sum_{i=1}^{\lfloor N/3 \rfloor} \frac{N!}{(N-2i)!i!2^i} \geq \sum_{i=1}^{\lfloor N/3 \rfloor} \left(\frac{2N}{3}\right)^i \\ &= \frac{2N}{2N-1} \left(\left(\frac{2N}{3}\right)^{\lfloor N/3 \rfloor} - 1 \right) \geq \left(\frac{2N}{3}\right)^{\lfloor N/3 \rfloor} - 1, \end{aligned}$$

where the first inequality uses $N/2 > N/3$, the second inequality uses inequality (C.6), and the third inequality uses $\frac{2N}{2N-1} < 1$. Hence the total number of bundle selections is

$$1 + \sum_{i=1}^{\lfloor N/2 \rfloor} \frac{N!}{(N-2i)!i!2^i} \geq \left(\frac{2N}{3}\right)^{\lfloor N/3 \rfloor}.$$

This completes the proof.

C ALGORITHM

In this section, we give the three algorithms based on the methods proposed in Section 3. A flow chart of the proposed mechanism is shown in Figure 3. First, we employ kernel density estimation to estimate the distribution of the data set $\Gamma_{i,j}$ and obtain an estimated distribution \hat{D}_i^j . Then, we

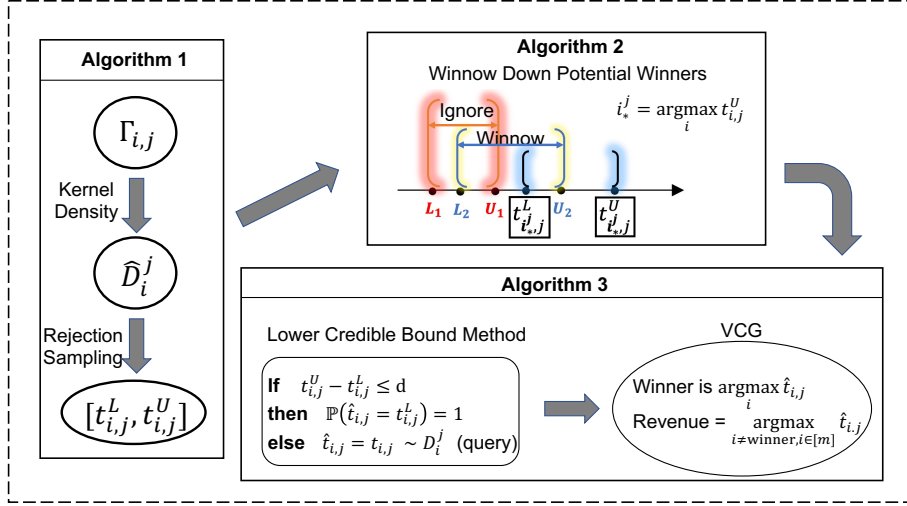


Figure 3: Flow chart of the proposed mechanism.

use rejection sampling to calculate a $(1 - \alpha)$ credible interval, denoted by $[t_{i,j}^L, t_{i,j}^U]$. This procedure is presented as Algorithm 1. Second, we implement a winnow-down procedure for each item j by considering only the intervals that intersect with the interval with the largest upper bound $t_{i_*^j}^U$, while disregarding the others. For example, as shown in Figure 3, the interval $[L_2, U_2]$ is winnowed down, while $[L_1, U_1]$ is disregarded. This procedure is presented as Algorithm 2. Finally, we employ the lower credible bound method on the VCG mechanism and obtain the implemented VCG mechanism using the degenerated types distribution. This procedure is presented as Algorithm 3.

C.1 ALGORITHM OF DISTRIBUTION AND CREDIBLE INTERVALS ESTIMATION

Algorithm 1 shows the details of estimating bidder types and constructing credible intervals, as outlined in Section 3.1. The inputs to the algorithm include the number of bidders, m , the number of items, N , a historical dataset, $\Gamma_{i,j}$, consisting of $n_{i,j}$ independently and identically distributed samples of the distribution D_i^j for $i \in [m], j \in [N]$, the minimum and maximum values, a_j and b_j , for item j , a critical value, α , and the sampling number \mathcal{N} . The output of the algorithm is a $(1 - \alpha)$ credible interval for each type, $t_{i,j}$, with lower bound, $t_{i,j}^L$, and upper bound, $t_{i,j}^U$. The algorithm begins by utilizing a kernel density function to estimate the distribution, \hat{D}_i^j . Then, a uniform distribution on $[a_j, b_j]$ is employed through rejection sampling. We sample \mathcal{N} results from \hat{D}_i^j , and take the $\alpha/2$ th quantile as the lower bound, $t_{i,j}^L$, and take the $(1 - \alpha/2)$ th quantile as the upper bound, $t_{i,j}^U$.

C.2 ALGORITHM OF WINNOWING DOWN POTENTIAL WINNERS

In Algorithm 2, we present a method for identifying potential winners among bidders based on the linkage of their confidence intervals, as outlined in Section 3.2. The input to the algorithm includes the number of bidders, m , the number of items, N , and confidence intervals, $[t_{i,j}^L, t_{i,j}^U]$, for each $t_{i,j}$, where $i \in [m], j \in [N]$. The output is a set of potential winners, \mathcal{B}_j , for each item j . For each item j , we identify the bidder, i_*^j , with the highest upper bound, $t_{i_*^j}^U$, of the credible interval. We then consider only those bidders whose credible intervals for item j are linked to the interval of bidder i_*^j . These bidders are then added to the set of potential winners, \mathcal{B}_j . The number of bidders that we will disregard for item j can be determined by $N - |\mathcal{B}_j|$. Furthermore, the total number of bids that will be neglected in the multi-item auction is given by

$$m^* \equiv mN - \sum_{j=1}^N |\mathcal{B}_j|. \quad (\text{C.7})$$

Algorithm 1: Estimation of distribution and credible intervals for bidder types

Input: m, N , historical dataset $\Gamma_{i,j} = \{x_{i,j}^1, x_{i,j}^2, \dots, x_{i,j}^{n_{i,j}}\}$, minimum value a_j , maximum value b_j , $a_j < b_j$, $i \in [m], j \in [N]$, $\alpha \in [0, 1]$, $\mathcal{N} \in \mathbb{N}$

Output: Lower bound $t_{i,j}^L$ and upper bound $t_{i,j}^U$ for $i \in [m], j \in [N]$

```

1 for  $i \in [m], j \in [N]$  do
2    $t=0$ ;
3   repeat
4      $K(t) = \frac{1}{\sqrt{2\pi}}e^{-t^2/2}$ ;  $h = 1.06 \cdot \min\{\hat{\sigma}, \frac{IQR}{1.34}\} \cdot n_{i,j}^{-1/5}$ ;
5      $\hat{D}_i^j(x) = \frac{1}{n_{i,j}h} \sum_{s=1}^{n_{i,j}} K(\frac{x-x_{i,j}^s}{h})$ ;  $g(x) = \frac{1}{b_j-a_j} \mathbb{I}\{x \in [a_j, b_j]\}$ ;
6     sample a point  $x \sim \text{Unif}[a_j, b_j]$ ;
7     choose  $c$  big enough s.t.  $\hat{D}_i^j(x_0)/cg(x_0) < 1$  for  $\forall x_0 \in [a_j, b_j]$ ;
8     sample a point  $u \sim \text{Unif}[0,1]$ ;
9     if  $u < \hat{D}_i^j(x)/cg(x)$  then
10      | accept  $x$ ;
11      |  $t \leftarrow t+1$ 
12   end
13   until  $t = \mathcal{N}$ ;
14    $t_{i,j}^L = \alpha/2$ th quantile of the accepted  $\mathcal{N}$  results;
15    $t_{i,j}^U = (1 - \alpha/2)$ th quantile of the accepted  $\mathcal{N}$  results
16 end

```

Algorithm 2: Winnow down the potential winners

Input: m, N , $[t_{i,j}^L, t_{i,j}^U]$ for $i \in [m], j \in [N]$

Output: For each $j \in [N]$, output a bidders' set \mathcal{B}_j

```

1 for  $j \in [N]$  do
2    $\mathcal{B}_j = \emptyset$ ;
3    $i_* = \text{argmax}_i t_{i,j}^U$ ;
4   for  $i \in [m]$  do
5     | if  $t_{i,j}^U > t_{i_*,j}^L$  then
6     | |  $\mathcal{B}_j \leftarrow \mathcal{B}_j \cup B_i$ 
7     | end
8   end
9 end

```

C.3 ALGORITHM OF THE IMPLEMENTED VCG MECHANISM

In this section, we begin by providing a detailed explanation of how to determine the suitable range of values for d . Subsequently, we present the algorithms for the implemented VCG mechanism.

C.3.1 TUNING PARAMETERS FOR THE DEGENERATED VCG MECHANISM

Theorems 1 and 2 provide a theoretical foundation for our mechanism. We consider two key tuning parameters, n in Eq. (7) and d in Eq. (5).

First, we consider the tuning of n . In order to achieve the desired IR and expected revenue in high confidence, the value of $(1 - \alpha/2)^n$ needs to be as large as possible. For any confidence rate $\eta \in [0, 1]$, let

$$(1 - \alpha/2)^n \geq \eta.$$

This yields the inequality $n \leq \log(\eta)/\log(1 - \alpha/2)$. Additionally, as defined in Section 3.3, n must be less than mN . Therefore, we can set

$$n \leq n^* \equiv \min \left\{ \frac{\log(\eta)}{\log(1 - \alpha/2)}, mN \right\}.$$

Second, we consider the tuning of d . We desire to control d such that $n \leq n^*$. Formally, this can be achieved by selecting d as the maximum value that satisfies

$$n = \left| \{(i, j) | t_{i,j}^U - t_{i,j}^L \leq d, i \in [m], j \in [N]\} \right| \leq n^*. \quad (\text{C.8})$$

In practice, we begin by selecting d as the theoretical maximum value given by Eq. (C.8). We then calculate the corresponding value of kd . If this value is too large, we reduce the value of d , whereas if the seller is willing to tolerate a reduced revenue and take on a small risk, we may increase the value of d to degenerate more types' distributions, thereby reducing the number of queries required in the VCG mechanism. This process of adjusting d can be performed by sorting the length of intervals in increasing order and selecting the n^* th length as the initial value of d , and iteratively adjusting until a satisfactory value is found. However, when the product mN is large, this method may encounter limitations. This can be observed when we consider the values $\eta = 0.9$ and $\alpha = 0.01$. In this scenario, we observe that $n^* = 21$ when the product mN exceeds 21. This implies that we can degenerate a maximum of 21 types' distributions. In such cases, we can use the strategy outlined in Section 3.2 to filter out unnecessary intervals, thus reducing the number of intervals considered and allowing us to utilize the mechanism proposed in Section 3.3 more effectively. In practice, we recommend starting with the value of d given by Eq. (C.8) for a specified value of η , and adjusting as necessary.

C.3.2 ALGORITHM OF THE IMPLEMENTED VCG MECHANISM

The proposed algorithm, outlined in Algorithm 3, leverages the results of Algorithms 1 and 2 to provide a final implementation for determining the potential winners in an auction setting. Algorithm 1 yields the estimated confidence intervals $[t_{i,j}^L, t_{i,j}^U]$, while Algorithm 2 yields the set of potential bidders \mathcal{B}_j for each item j and the number of neglected bids m^* defined in Eq. (C.7). Additionally, the algorithm introduces a hyper-parameter q , which represents the maximum loss the seller is willing to tolerate. Algorithm 3 first filters the credible intervals based on the potential bidders $\mathcal{B}_j, j \in [N]$, by setting the lower and upper bounds of the interval to 0 for bidders not in \mathcal{B}_j . The remaining intervals are sorted in increasing order according to the values of $t_{i,j}^U - t_{i,j}^L$. The value of \hat{t} is then determined as described in Section 3.3, with the value of d chosen to minimize the query number while ensuring that the maximum reduced revenue is less than q , as discussed in Section C.3.1. The algorithm begins by assuming an initial value of d , and repeatedly increases or decreases it until it fluctuates between two numbers d_- and d_+ , such that $kd_- \leq q < kd_+$, and the final value of d is set to d_- . The complete procedure is presented in Algorithm 3.

Algorithm 3 will eventually terminate, and it does not loop indefinitely. This is because the decrease in the value of d ensures that k becomes less than or equal to its original value, resulting in a strictly decreasing value of kd . When $M = 0$, kd becomes 0, and when $M = mN + 1$, kd becomes $+\infty$. Hence, the loop will eventually terminate since q is within the range of $(0, +\infty)$.

D EXPERIMENTS RESULTS

We consider a scenario where there are N items and m bidders. For each item $j \in [N]$, the minimum and maximum values are set at $a_j = 10(j - 1)$ and $b_j = 10j$ respectively. The private valuations of each bidder $B_i, i \in [m]$, are modeled by a truncated Gaussian distribution D_i^j with means chosen uniformly from the range $[10(j - 1), 10j]$ and variances of 10^x , where x is chosen uniformly from the interval $[-1, 1]$. The distribution is then normalized to 1 over the range $[10(j - 1), 10j]$. Once the distribution D is obtained, we can generate the types of bidders t from D . Historical data sets, $\Gamma_{i,j}$, can then be simulated for each item j by drawing $n_{i,j}$ samples from the corresponding true distribution D_i^j . In this simulation, we set $n_{i,j} = |\Gamma_{i,j}| = 50$ for all i and j . Additionally, we set $\mathcal{N} = 1000$, $\alpha = 0.01$, and $\eta = 0.9$ for our experiments. The mechanism can be implemented once the values of m and N are determined.

In our experiments, we do not specify a certain value for the hyper-parameter q described in Algorithm 3 as we wish to analyze the revenue, regret, and number of queries for various values of d . To accomplish this, we start by initializing $d = 0$ and increment it gradually in a continuous manner while simultaneously keeping track of the revenue generated and the corresponding value of n . In other words, we implement a modified version of Algorithm 3 by excluding the execution of lines 7,

Algorithm 3: Degenerated VCG mechanism incorporates winnow down data and a regret tolerance threshold of q .

Input: $m, N, [t_{i,j}^L, t_{i,j}^U]$ for $i \in [m], j \in [N], \alpha \in [0, 1]$, accepted confidence rate $\eta \in [0, 1]$, regret tolerance threshold $q \in (0, +\infty), \mathcal{B}_j$ for $j \in [N], m^*$

Output: The winners set and the revenue

```

1 for  $j \in [N]$  do
2   | if  $i \notin \mathcal{B}_j$  then  $t_{i,j}^L = t_{i,j}^U = 0$ ;
3 end
4  $n^* = \min \left\{ \frac{\log(\eta)}{\log(1-\alpha/2)}, mN - m^* \right\}$ ;
5 sort the value of  $(t_{i,j}^U - t_{i,j}^L), i \in [m], j \in [N]$  in an increasing sequence  $(\ell_1, \ell_2, \dots, \ell_{mN})$ ;
6  $M = n^* + m^*, \ell_0 = 0, \ell_{mN+1} = +\infty$ ;
7 repeat
8   |  $d = \ell_M, K = \emptyset$ ;
9   | for  $i \in [m]$  do
10    | | if  $j \in L_i \cap L_i^*$  then  $K = K \cup j$ ;
11    | end
12    |  $k = |K|$ ;
13    | if  $kd > q$  then  $M \leftarrow M - 1$  else  $M \leftarrow M + 1$ ;
14 until  $d$  fluctuates between the two values  $d_-$  and  $d_+$ , where  $d_- < d_+$ ;
15  $d = d_-$ ;
16 for  $i \in [m]$  do
17   | if  $j \in L_i$  then  $\hat{t}_{i,j} = t_{i,j}^L$  else  $\hat{t}_{i,j} = t_{i,j}$  (make a query);
18 end
19 revenue = 0;
20 for  $j \in [N]$  do
21   | winner for item  $j = \arg \max_i \hat{t}_{i,j}$ ; revenue  $\leftarrow$  revenue +  $\arg \max_{i \neq \text{winner}, i \in [m]} \hat{t}_{i,j}$ ;
22 end

```

13-15. Additionally, we expand the range of possible values for d to include all values in the interval $[0, \ell_{mN}]$, instead of limiting it to ℓ_M alone.

We use three different revised mechanisms to verify our proposed strategies.

Mechanism 1 (Kernel density & Winnow down data): This mechanism implements our entire proposed strategies, comprising of Algorithms 1, Algorithm 2, and Algorithm 3, in sequence.

Mechanism 2 (Kernel density & Full data): This mechanism is adopted to avoid potential compromise of auction fairness that may arise from winnowing down bidders and neglecting some bids. The revenue may be reduced if potential winners with higher types than the upper bound of their confidence intervals are neglected. In mechanism 2, Algorithm 1 and Algorithm 3 are implemented, omitting the use of Algorithm 2 by setting $\mathcal{B}_j = [m], \forall j \in [N]$, and $m^* = 0$ in the input of Algorithm 3.

Mechanism 3 (Ordinary method & Winnow down data): This mechanism does not estimate the distribution of D , and instead rely on the historical data $\Gamma_{i,j} = \{x_{i,j}^1, x_{i,j}^2, \dots, x_{i,j}^{n_{i,j}}\}, i \in [m], j \in [N]$. In this scenario, the minimum type of bidder B_i for item j is determined by the smallest value of $\Gamma_{i,j}$, while the maximum type of bidder B_i for item j is determined by the largest value of $\Gamma_{i,j}$. Specifically, let $a_{i,j} = \min\{\Gamma_{i,j}\}$ and $b_{i,j} = \max\{\Gamma_{i,j}\}$. In mechanism 3, we just implement Algorithm 2 and Algorithm 3, without using Algorithm 1. To do this, let $t_{i,j}^L = a_{i,j}$ and $t_{i,j}^U = b_{i,j}$, then we implement the Algorithm 2 and Algorithm 3.

D.1 SMALL-SCALE DATA

In order to evaluate the efficacy of our proposed methods, we conduct an initial experiment using a small dataset. Specifically, we consider a scenario in which there are $m = 30$ bidders and $N = 10$ items to be sold. Using the Vickrey-Clarke-Groves (VCG) mechanism to sell items individually,

we calculate the revenue of the original VCG mechanism based on type t , represented by the green dashed line in the bottom-left plot of Figure 4. This value, Revenue= 543, serves as a benchmark against which we compare the performance of our proposed methods. Our goal is to obtain revenue values as close as possible to this benchmark through the application of our methods. In other words, we want to minimize the regret.

D.1.1 BENEFITS OF DENSITY ESTIMATION AND SAMPLING REJECTION




Figure 4: Comparison of the revenue and regret when $m = 30, N = 10$.

In Figure 4, we present a comparison of revenue and regret across the three mechanisms as the value of d varies. Mechanism 1 is represented by a red solid line, Mechanism 2 by a black dashed line, and Mechanism 3 by a blue long dashed line. As shown in Figure 4(a) and 4(b), which represent revenue and regret comparison respectively, all three mechanisms exhibit revenue exceeding 515 and regret less than 25, indicating overall good performance. However, we note that Mechanism 3 performs worse than Mechanism 1 and 2 in most cases. This is due to the absence of kernel density function in Mechanism 3, resulting in intervals of lower confidence compared to those obtained from Mechanism 1 and 2. Additionally, Mechanisms 1 and 2 overlap, indicating that the neglected data does not significantly impact the final outcome. This is in line with the α -fair mechanism as proposed in Proposition 1, where $\alpha = 0.01$, a value small enough to ensure fairness of the game.

In Figure 4(c), Mechanism 1 is represented by a red solid line, Mechanism 2 by a black dotted line, and Mechanism 3 by a blue long dashed line. Figure 4(c) presents the comparison of theoretical regret and reveals that the lines of the three mechanisms overlap in most cases. This is because when d is too small, the set K , as defined in Eq. (8), is empty, resulting in $k = 0$. On the other hand, when d is sufficiently large, K tends towards $[N]$, resulting in $k = N$, and subsequently $kd = Nd$, a line with slope N . However, within a small interval, we find that for a fixed value of d , kd of Mechanism 3 is the largest, and kd of Mechanism 1 is the smallest. This is because, in the ordinary method, the length of the types' confidence intervals tends to be smaller, as it does not incorporate density function estimation, thus being limited to the historical data. As a result, for a fixed d , Mechanism 3 will have more intervals with a bandwidth smaller than d , resulting in a larger value of k . Conversely, Mechanism 1, by virtue of its sampling rejection method, already neglects some intervals, leading to a smaller number of intervals with a bandwidth smaller than d , and thus a smaller value of k .

In conclusion, our analysis of Figure 4 suggests that Mechanism 1 and Mechanism 2 perform better than Mechanism 3 in practice, highlighting the benefits of utilizing kernel density estimation and sampling rejection in Algorithm 1.

D.1.2 PERFORMANCE OF MECHANISM 1

In Figure 5, we compare the proportion of bidder types without queries and the confidence rate of Mechanism 1 and Mechanism 2 as the value of d varies. Mechanism 1 is represented by a red solid line, while Mechanism 2 is represented by a black dotted line. Figure 5(a) illustrates that the proportion of types without query in Mechanism 1 is consistently higher or equal to that of Mechanism 2. This is because all types with ignored credible intervals in Mechanism 1 are dropped, and the seller does not make a query to those types. In our experiment, there were $m^* = 126$ such types, resulting in an initial proportion of $126/300 = 0.42$ in Mechanism 1, while Mechanism 2 had a proportion of 0 as L , defined in Eq. (6), was an empty set at the beginning. As d increases, the number of elements in L also increases, but since Mechanism 1 has dropped some credible




Figure 5: Comparison of the proportion of bidder types without queries and comparison of the confidence rate when $m = 30, N = 10$.

intervals, the number of intervals whose bandwidths are smaller than a certain d increases more slowly, resulting in a line with a smaller slope for Mechanism 1 compared to Mechanism 2. Figure 5(b) illustrates the confidence rate, which is $(1-\alpha)^{n-m^*}$ in Mechanism 1 and $(1-\alpha)^n$ in Mechanism 2. The confidence rate of Mechanism 1 is consistently higher than that of Mechanism 2.

The results presented in Figure 5 clearly demonstrate that Mechanism 1 outperforms Mechanism 2. As observed, to attain a confidence rate of 0.9, Mechanism 1 requires $51\% = 1 - (m^* + n^*)/mN = 1 - (126 + 21)/300$ of the total queries, while Mechanism 2 still requires $93\% = 1 - n^*/mN = 1 - 21/300$ of the total queries. Additionally, both mechanisms have a practical regret of no larger than 2, which is relatively small when compared to the revenue of the original VCG mechanism, which is 543.

Overall, the combination of Figures 4 and 5 reveal that Mechanism 1 has both the lowest practical and theoretical regret, requiring a mere $58\% = 1 - m^*/mN = 1 - 126/300$ of total queries. Furthermore, when a confidence rate of 0.9 is desired, Mechanism 1 only needs 51% of the total queries, with a minimal practical revenue reduction of 2. These results highlight the effectiveness and efficiency of the proposed algorithm, making it an ideal choice in practical applications.

D.2 LARGE-SCALE DATA

D.2.1 ADVANTAGE WHEN THE NUMBER OF BIDDERS IS LARGE




Figure 6: Comparison of the revenue and regret when $m = 50, N = 30$.

In our experiments utilizing large-scale data, we first set the number of bidders to $m = 50$ and the number of items to $N = 30$, creating a scenario in which the bidders must compete for the items as $m > N$.

The notations used in Figure 6 are consistent with those in Figure 4. The revenue of the original VCG mechanism in this case is 4635.7. The results in Figure 6 reveal that the revenue generated by Mechanism 3 is consistently lower than that of Mechanism 1 and Mechanism 2 in most cases. By comparing the results in Figure 4 and Figure 6, we know that the use of kernel density estimation yields superior performance when the data scale is increased. In this scenario, the maximum theoretical regret kd is 300, which is relatively small in comparison to the revenue of the original VCG

mechanism, which is 4635.7. These results demonstrate the significant advantage of our mechanism in scenarios where the number of items is limited but the number of bidders is large, as previously analyzed in Section 3.4.

Additionally, we also evaluate our method under a scenario where the number of bidders is less than the number of items. Specifically, we set $m = 50$ and $N = 100$. The results, shown in Figure 7, exhibit a similarity to those presented in Figure 6, resulting in a revenue of 50457.2 for the original VCG mechanism. However, as the number of items m increases, a substantial increase in both practical and theoretical regret is observed, with the maximum regret reaching approximately 175 and 1000, respectively. In our scenario, the basic value of items increases as the number of items increases, making a maximum regret of 1000 acceptable due to the high true revenue. Nonetheless, if the basic value of items remains constant, this outcome may not be optimal. Hence our method is most appropriate in scenarios where the number of items is limited, or the values of items are high but exhibit limited variance.




Figure 7: Comparison of the revenue and regret when $m = 50, N = 100$.

D.2.2 PERFORMANCE OF MECHANISM 1

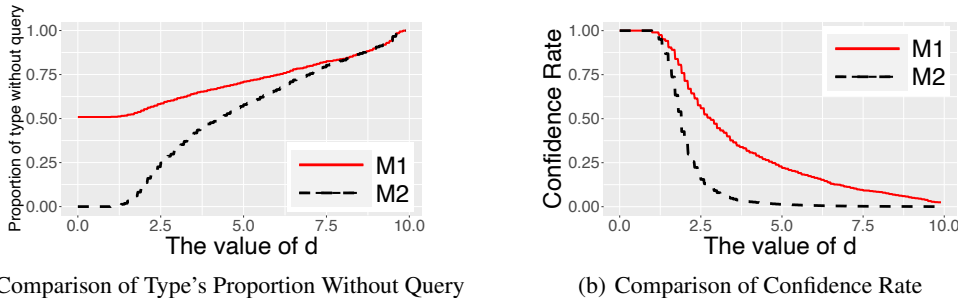


Figure 8: Comparison of the proportion of bidder types without queries and comparison of the confidence rate when $m = 50, N = 30$.

In Figure 8, we present the results of our method under the scenario of large-scale data when $m = 50, N = 30$. Specifically, we set $m^* = 763$ and $mN = 1500$, which results in an initial query rate of less than 49.2%, as approximated by $(mN - m^*)/mN$. With a confidence rate η of 0.9, Mechanism 1 further reduces the query rate to less than $(mN - m^* - n^*)/mN = 47.74\%$, resulting in just 2.5 practical regrets and 24 theoretical regrets. These results demonstrate that the proposed Mechanism 1 works well, even when dealing with large datasets.

In Figure 9, we also obtain results for an even larger dataset with $m = 50$ and $N = 100$. In this scenario, we set $m^* = 2456$ and $mN = 5000$, which results in an initial query rate of less than 50.88%, as approximated by $(mN - m^*)/mN$. With a confidence rate of 0.9, Mechanism 1 further reduces the query rate to 50.46% as approximated by $(mN - m^* - n^*)/mN$, resulting in just 1.6 practical regrets and 72 theoretical regrets. We note that as the value of mN increases, the confidence rate decreases at an accelerated rate. In this scenario, the seller must make a trade-off between classifying the intervals and making all queries to the winnowed-down data ($d = 0$), or reducing the number of queries by choosing a larger value of d , at the risk of sacrificing some

revenue. Our results show that when $d = 3.75$, although the confidence rate is low (less than 0.06), we can still reduce the number of queries to less than $37.5\% = 1 - 62.5\%$. This reduction results in a reduction of more than $mN * (50 - 37.5)\% = 625$ in the number of queries, with only a negligible reduction in the revenue of the original VCG mechanism (50 out of a total of 50457.2). This analysis supports the discussion in Section C.3.1 and highlights the importance of carefully considering the choice of d in practical applications.

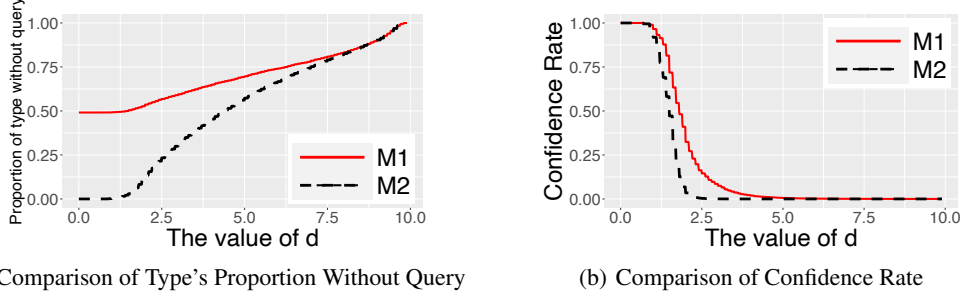


Figure 9: Comparison of the proportion of bidder types without queries and comparison of the confidence rate when $m = 50$, $N = 100$.

D.3 ANALYSIS OF ERROR BARS

In the bidding process, we make the assumption that the distributions of the bidder types remain fixed, as mentioned in the initial paragraph of the main results. Consequently, once the distributions are fixed, their credible intervals should also be fixed. In other words, $t_{i,j}^L$ and $t_{i,j}^U$ represent fixed constants for all $i \in [m]$ and $j \in [N]$. Therefore, the intervals obtained from the estimation based on simulated data may contain some errors when compared to the true intervals. In this section, we will analyze these errors.

First, we determine the true intervals by giving the following proposition.

Proposition 3 *Suppose the truncated Gaussian distribution D_i^j has the mean $\mu_{i,j} \in [10(j-1), 10j]$ and the variance $\sigma_{i,j}^2 > 0$, and it is re-normalized within the interval $[10(j-1), 10j]$. Let $\phi(x; \mu_{i,j}, \sigma_{i,j}^2)$ and $\Phi(x; \mu_{i,j}, \sigma_{i,j}^2)$ represent the probability density function and cumulative distribution function of $\mathcal{N}(\mu_{i,j}, \sigma_{i,j}^2)$, respectively. Then, the $\alpha/2$ th quantile $t_{i,j}^L$ and the $(1 - \alpha/2)$ th quantile $t_{i,j}^U$ of the D_i^j are calculated as follows:*

$$t_{i,j}^L = \Phi^{-1} \left(\Phi(10j; \mu_{i,j}, \sigma_{i,j}^2) - (1 - \alpha/2) \Phi(10(j-1); \mu_{i,j}, \sigma_{i,j}^2); \mu_{i,j}, \sigma_{i,j}^2 \right),$$

$$t_{i,j}^U = \Phi^{-1} \left(\Phi(10j; \mu_{i,j}, \sigma_{i,j}^2) - (\alpha/2) \Phi(10(j-1); \mu_{i,j}, \sigma_{i,j}^2); \mu_{i,j}, \sigma_{i,j}^2 \right).$$

for any $\alpha \in [0, 1]$.

Proof D.1

$$\begin{aligned} \phi(t_{i,j}; \mu_{i,j}, \sigma_{i,j}^2 | 10(j-1) < t_{i,j} < 10j) &= \frac{\phi(t_{i,j}; \mu_{i,j}, \sigma_{i,j}^2)}{\mathbb{P}(10(j-1) < t_{i,j} < 10j)} \\ &= \frac{\phi(t_{i,j}; \mu_{i,j}, \sigma_{i,j}^2)}{\Phi(10j; \mu_{i,j}, \sigma_{i,j}^2) - \Phi(10(j-1); \mu_{i,j}, \sigma_{i,j}^2)} \\ &\equiv \frac{\phi(t_{i,j}; \mu_{i,j}, \sigma_{i,j}^2)}{C}, \end{aligned}$$

where we define $C = \Phi(10j; \mu_{i,j}, \sigma_{i,j}^2) - \Phi(10(j-1); \mu_{i,j}, \sigma_{i,j}^2)$.

With reference to the definition of $t_{i,j}^L$, we can derive the following expression:

$$\begin{aligned}
& \int_{10(j-1)}^{t_{i,j}^L} \frac{\phi(t_{i,j}; \mu_{i,j}, \sigma_{i,j}^2)}{C} dt_{i,j} = (\alpha/2) \\
\iff & \int_{10(j-1)}^{t_{i,j}^L} \phi(t_{i,j}; \mu_{i,j}, \sigma_{i,j}^2) dt_{i,j} = (\alpha/2)C \\
\iff & \Phi(t_{i,j}^L; \mu_{i,j}, \sigma_{i,j}^2) - \Phi(10(j-1); \mu_{i,j}, \sigma_{i,j}^2) = (\alpha/2)C \\
\iff & \Phi(t_{i,j}^L; \mu_{i,j}, \sigma_{i,j}^2) = (\alpha/2)C + \Phi(10(j-1); \mu_{i,j}, \sigma_{i,j}^2) \\
\iff & t_{i,j}^L = \Phi^{-1}\left((\alpha/2)C + \Phi(10(j-1); \mu_{i,j}, \sigma_{i,j}^2); \mu_{i,j}, \sigma_{i,j}^2\right) \\
\iff & t_{i,j}^L = \Phi^{-1}\left(\Phi(10j; \mu_{i,j}, \sigma_{i,j}^2) - (1 - \alpha/2)\Phi(10(j-1); \mu_{i,j}, \sigma_{i,j}^2); \mu_{i,j}, \sigma_{i,j}^2\right).
\end{aligned}$$

Likewise, by replacing $t_{i,j}^L$ with $t_{i,j}^U$ and $\alpha/2$ with $1 - \alpha/2$, we obtain:

$$t_{i,j}^U = \Phi^{-1}\left(\Phi(10j; \mu_{i,j}, \sigma_{i,j}^2) - (\alpha/2)\Phi(10(j-1); \mu_{i,j}, \sigma_{i,j}^2); \mu_{i,j}, \sigma_{i,j}^2\right).$$

Based on the true intervals, we can introduce two additional mechanisms to visualize the error of the random simulation points.

Mechanism 4 (True intervals & Winnow down data): This mechanism does not utilize simulated historical data but directly employs the true intervals. To implement this, we apply Algorithms 2 and 3, considering the true values for $t_{i,j}^L$ and $t_{i,j}^U$ as described in Proposition 3.

Mechanism 5 (True intervals & Full data): In comparison to Mechanism 4, this mechanism solely employs Algorithm 3 and sets $\mathcal{B}_j = [m], \forall j \in [N]$, and $m^* = 0$ as inputs to Algorithm 3.

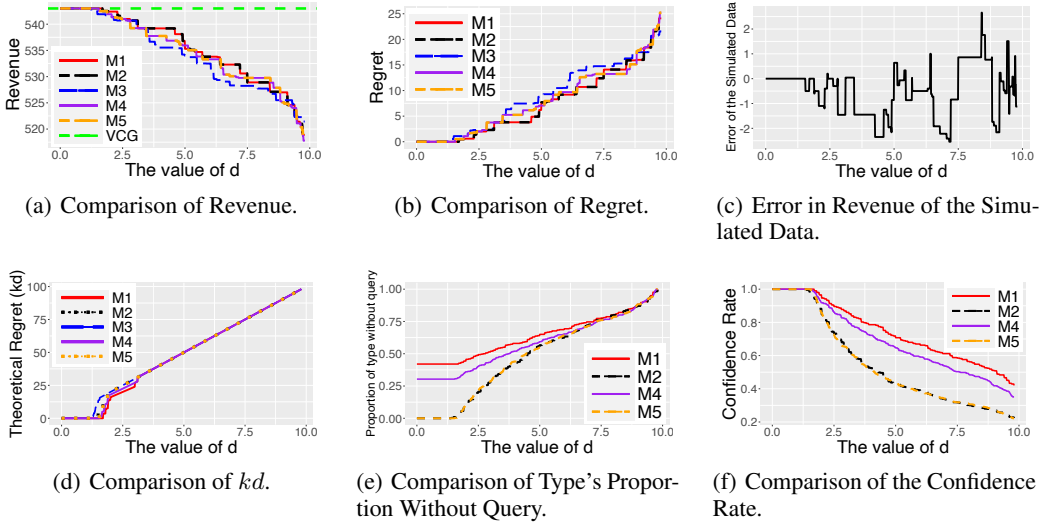


Figure 10: The visualization of error when $m = 30$ and $N = 10$.

Figure 10, 11, and 12 depict the visualization of errors in each auction size. Mechanism 4 (M4) is represented by the solid purple line, while Mechanism 5 (M5) is indicated by the dashed or dotted orange line. Mechanism 4 and Mechanism 5 can be considered as benchmarks for Mechanism 1 and Mechanism 2, respectively, with fixed true intervals. We can apply the same analysis discussed earlier by replacing Mechanism 1 with Mechanism 4 and Mechanism 2 with Mechanism 5, yielding similar results.

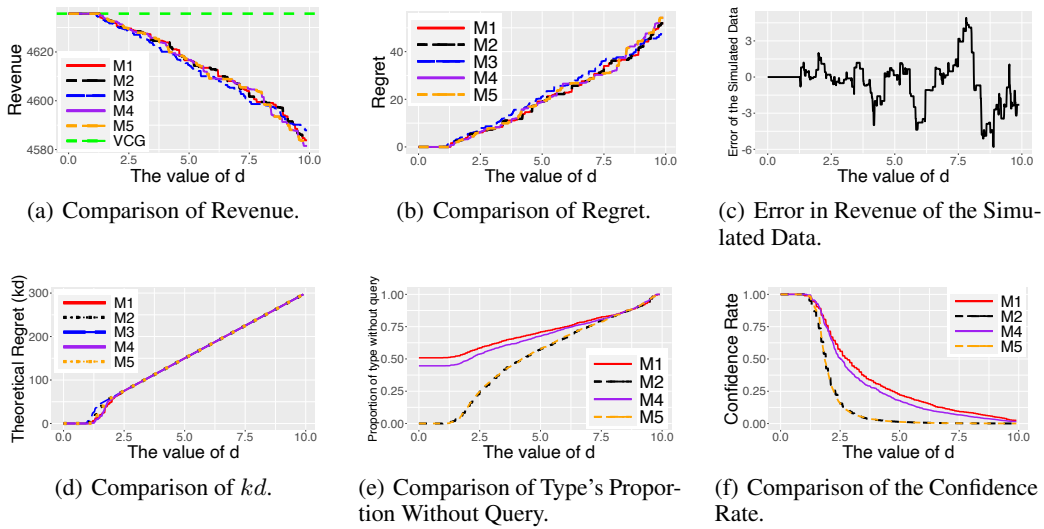


Figure 11: The visualization of error when $m = 50$ and $N = 30$.

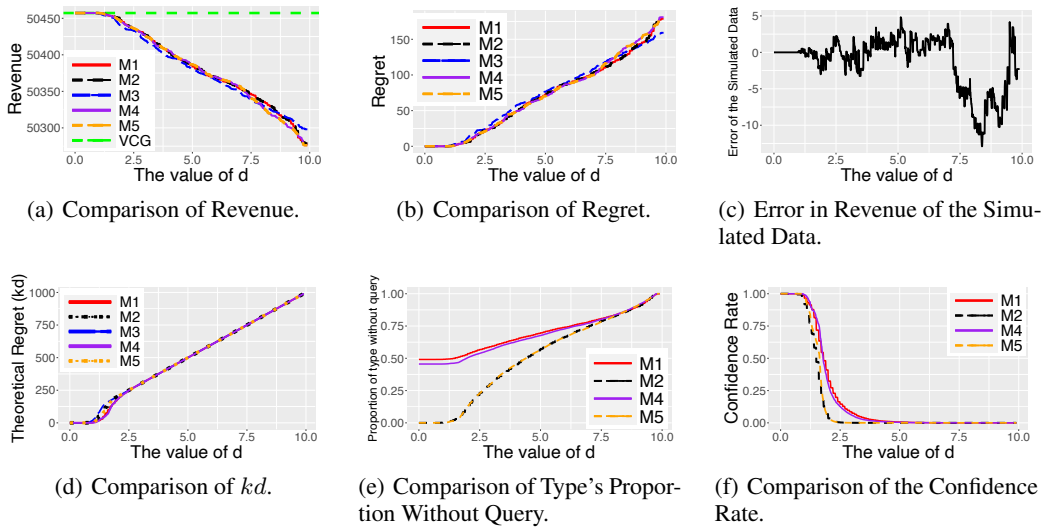


Figure 12: The visualization of error when $m = 50$ and $N = 100$.

The relative error decreases as the number of bidders and items increases, highlighting the advantages of our strategies in large-size auctions. Moreover, in Figure 10(c), 11(c), and 12(c), the errors in revenue are depicted, which are calculated by subtracting the revenue of Mechanism 1 (Mechanism 2) from the revenue of Mechanism 4 (Mechanism 5) for each scenario. It is evident that the error in revenue, obtained through our strategies utilizing credible intervals derived from the simulated data, is remarkably small compared to the revenue generated by the original VCG mechanism across all scenarios. This demonstrates that our proposed methods are not sensitive to changes in the credible intervals, indicating the robustness of our multi-item auction design.