

DIFFERENTIALLY PRIVATE MECHANISM DESIGN VIA QUANTILE ESTIMATION

Anonymous authors

Paper under double-blind review

ABSTRACT

We investigate the problem of designing differentially private (DP), revenue-maximizing single item auction. Specifically, we consider broadly applicable settings in mechanism design where agents’ valuation distributions are *independent, non-identical*, and can be either *bounded* or *unbounded*. Our goal is to design such auctions with *pure*, i.e., $(\epsilon, 0)$ privacy in polynomial time.

In this paper, we propose two computationally efficient auction learning framework that achieves *pure* privacy under bounded and unbounded distribution settings. These frameworks reduces the problem of privately releasing a revenue-maximizing auction to the private estimation of pre-specified quantiles. Our solutions increase the running time by polylog factors compared to the non-private version. As an application, we show how to extend our results to the multi-round online auction setting with non-myopic bidders. To our best knowledge, this paper is the first to efficiently deliver a Myerson auction with *pure* privacy and near-optimal revenue, and the first to provide such auctions for *unbounded* distributions.

1 INTRODUCTION

Though prior-dependent auctions, which adjust parameters based on samples of value distributions, often yield better revenue than prior-independent auctions, they risk leaking information about the bids they were trained upon. To address this issue, differential privacy (DP) offers a promising solution (Dwork, 2006; 2008; McSherry and Talwar, 2007; Pai and Roth, 2013), ensuring that a single data point minimally affects the algorithm’s output, thus preventing inference of a specific data point.

We study the problem of learning a single-item auction with near-optimal revenue from samples of independent and non-identical value distributions. In this context, the optimal auction (i.e., Myerson’s auction (Myerson, 1979)), which relies on value distributions (i.e., prior-dependent), achieves optimal revenue. However, releasing the learned Myerson’s auction raises privacy concerns, as the output mechanism may inadvertently reveal sensitive information about the distributions. To provably mitigate this risk, our goal is to integrate *pure* DP into the learning process of such auction.

Pure Differential Privacy. Given two datasets that differ in one data point, i.e., D, D' , we say an algorithm \mathcal{A} satisfies (ϵ, δ) -approximate DP if for any given output s : $\Pr[\mathcal{A}(D) = s] \leq e^\epsilon [\mathcal{A}(D') = s] + \delta$. We say \mathcal{A} satisfies *pure* DP if $\delta = 0$. Pure DP allows no slack in privacy protection, and hence is more challenging to achieve than approximate DP. Previous attempts (McSherry and Talwar, 2007; Nissim et al., 2012) to integrate DP with prior-dependent auctions have been computationally inefficient or guaranteed approximate rather than pure DP. To our knowledge, *no algorithm guarantees pure DP for Myerson’s auction in polynomial time*.

Efficiency. Incorporating DP into the mechanism often sacrifices efficiency, as achieving privacy guarantees typically incurs additional computational overhead (e.g., random noise addition or extra sampling procedure). This issue has been observed in similar contexts, such as online learning (Jain et al., 2012), federated learning (Zhang et al., 2023) and deep learning (Abadi et al., 2016). In our context, to achieve pure DP, implementing exponential mechanism (McSherry and Talwar, 2007) over all possible mechanisms would incur *exponential* time (See Appendix D). To obtain pure DP more efficiently, we apply recent advances (Durfee, 2023; Kaplan et al., 2022) in private quantile estimation. Our algorithm’s running time increases by only *polylog* factors compared to the non-private version.

Notations We use M_A to denote the optimal mechanism of distribution A , and we use $\text{Rev}(M, A)$ to denote the revenue of deploying mechanism M to distribution A . We restricted ourselves to single item auctions; hence, M_A denotes the Myerson auction fitted on distribution A , and we denote $\text{OPT}(A) := \text{Rev}(M_A, A)$ as the optimal revenue one could get from a distribution A . We use $\mathbf{1}_k$ to denote a k -dimensional vector with all entries equal to 1. We use \tilde{O} and $\tilde{\Theta}$ to hide polylog factors.

1.1 RESULTS

Formally, we define the problem of learning a near-optimal auction with a pure DP:

Problem 1.1 (Optimal Auction with $(\epsilon_p, 0)$ -DP). Given n samples of k -dimensional distribution \mathbf{D} , the goal is to learn a single item auction M with $(\epsilon_p, 0)$ -DP, whose expected revenue on \mathbf{D} is close to the optimal revenue, i.e., with prob. $1 - \delta^1$, $|\mathbb{E}[\text{Rev}(M, \mathbf{D}) - \text{OPT}(\mathbf{D})]| \leq \epsilon$ for some small ϵ .

Insight. To address this problem, we leverage the insight that, the expected optimal revenue from value distribution is *insensitive* to small statistical shifts and discretization in the quantile and value space. Additionally, we observe that the accuracy of the points returned by private quantile estimation (QE), assuming the data points follow a distribution, directly correlates with the statistical distance between the distribution formed by the returned points and the true distribution. Thus, we can reduce private Myerson fitting from samples to *private quantile estimation of pre-specified quantiles*.

Achieving pure DP while maintaining meaningful revenue guarantees is challenging. A crucial aspect is to ensure that the values (hence distribution) returned by DP Quantile Estimation (QE) possess meaningful and provable accuracy guarantees. To obtain such accuracy, our algorithm (Alg. 1) first additively discretize the empirical distribution in the value space to distribution \hat{D}^ϵ , then estimate the pre-specified quantiles with DPQE. We improved the accuracy bound of DPQE (DPQUANT, Kaplan et al. (2022)) to accommodate cases with duplicate values. This improved bound allows us to upper bound the statistical distance between the output distribution and \hat{D}^ϵ , thus upper bounding the revenue loss incurred from fitting a Myerson on the output distribution.

Theorem 1.2 briefly presents the near-optimal revenue of our proposed mechanism. The final privacy parameter has a dependency on k since the output of mechanism M is of dimension $2k$. We present complete details in Section 3 and the complete theorem statement in Theorems 3.2 and 3.3.

Theorem 1.2 (Revenue Guarantee of Private Myerson, Bounded). Given $n = \tilde{\Theta}(\epsilon^{-2})$ samples \hat{V} of the joint distribution $\mathbf{D} \in [0, h]^k$, there exist a mechanism M that is $2k\epsilon_p$ differentially private with running time $\tilde{\Theta}(kn)$ and takes $\tilde{\Theta}(1)$ pass of the distribution. With probability $1 - \delta$, this mechanism M satisfies: $|\mathbb{E}[\text{Rev}(M, \mathbf{D}) - \text{OPT}(\mathbf{D})]| \leq \tilde{O}((\epsilon + \epsilon^2/\epsilon_p)kh)$.

The prior algorithm does not work for *unbounded* distributions. Our second algorithm (Alg. 9) addresses the case for η -strongly regular value distributions by efficiently truncating them to bounded distributions with small expected revenue loss. This approach enables the application of our previous mechanism (Alg. 1) designed for the bounded distribution case. Since the truncation point is a function of the optimal revenue, we develop Alg. 7 to approximate this point by achieving a $\tilde{\Theta}(k)$ -approximation of the optimal revenue, where k denotes the dimension of the product distribution.

Theorem 1.3 outlines the accuracy of our proposed mechanism for certain parameter settings. Since this truncation point depends adaptively on the desired accuracy, the revenue gap exceeds that for the bounded case, and the tradeoff between privacy and revenue are more pronounced. We present more details in Section 4, and the complete theorem statement is in Theorems 4.1 and I.13.

Theorem 1.3 (Revenue Guarantee of Private Myerson, Unbounded). Given $n = \tilde{\Theta}(\epsilon^{-2})$ samples \hat{V} of η -strongly regular joint distribution $\mathbf{D} \in \mathbb{R}^k$, there exist a mechanism M for unbounded distribution that is $2k\epsilon_p$ differentially private with running time $\tilde{\Theta}(kn)$ and takes $O(n)$ passes. With probability $1 - \delta$, this mechanism M satisfies: $|\mathbb{E}[\text{Rev}(M, \mathbf{D}) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})]| \leq \tilde{O}(k^2\sqrt{\epsilon} + k^2\epsilon^{1.5}/\epsilon_p)$.

¹This failure probability δ is inevitable due to the inherent uncertainty in learning from a finite sample set, see Chapter 1 Kearns and Vazirani (1994)

Application: Online auction with nonmyopic bidders. We now describe how our mechanisms incentivize truthful bidding from nonmyopic bidders under practical online auction settings.² In the online setting, auctions are deployed iteratively and later auctions are informed by previous bids. Since future auctions can be affected by earlier bids, *nonmyopic* bidders may strategically bid in earlier rounds to increase winning chances and/or secure lower prices, increasing their utility.

To prevent from strategic bidding, we integrate our previous solutions (Alg. 1, Alg. 9) with a commitment mechanism. Our DP Myerson naturally upper bound the utility gain (of future rounds) by definition, in that the change of one bid affect the outcome’s probability by privacy parameter ϵ_p . Our algorithm operates in two stages. In the first stage, it employs a commitment mechanism that penalizes strategic bids. In the second stage, the algorithm fits a DP Myerson auction from the collected bids and generates revenue in the remaining rounds. This approach ensures that strategic bids only lies in a small neighbor of the true value; otherwise, the bidder’s utility becomes negative.

We present the *regret* (i.e., the time-averaged revenue of the proposed mechanism compared to the optimal one) of our proposed mechanism (Alg. 3) in Theorem 1.4, which shows the accuracy of our algorithm in terms of regret. We defer readers to Section 5 and Theorem 5.4 for further details.

Theorem 1.4 (Revenue Guarantee of Online Mechanism). *Given $\epsilon \in [0, 1/4]$, under the online auction setting described in Section 5.1), there exists an algorithm (Alg. 3) run with parameter $T = \tilde{\Theta}(\epsilon^{-2})$ that, with probability $1 - \delta$, achieves diminishing regret, i.e., $\text{REGRET} = \tilde{O}[(\epsilon + \sqrt{\eta\epsilon})kh]$, where η is a constant specific to bidders’ utility model.*

1.2 PRIOR WORK

DP Mechanism Design. Emerging from McSherry and Talwar (2007), there has been interest in delivering mechanisms with DP guarantees (Nissim et al., 2012; Huang et al., 2018a; Zhang and Zhong, 2022; Huh and Kandasamy, 2024). These mechanism are either *no longer optimal* in our setting, or doesn’t generalize to unbounded distribution setting.

Online Learning in Repeated Auction. Regarding the single item online auction setting, Kanoria and Nazerzadeh (2014); Huang et al. (2018a) established near-optimal solutions when bidders’ utility is discounted and valuations are i.i.d.. Deng et al. (2020); Abernethy et al. (2019) introduced specific incentive metrics to quantify bidders’ willingness to bid other than their true values and developed mechanisms that minimize incentives for strategic bidding under these metrics in large markets.

For a detailed, complete list of related work topics, please see Appendix C.

1.3 CONTRIBUTIONS

Revenue Maximizing Auctions with Pure Privacy Guarantee. Our work is the first to develop a mechanism with *pure* DP that obtains near optimal revenue for single item auction with independent and non-identical bidders, and for both *bounded* and *unbounded* η -strongly regular distributions. For bounded distributions, our mechanism achieves optimal time complexity within polylog factors.

Application to Online Auction Setting. We apply our mechanism into the online auction setting with nonmyopic, independent and non-identical bidders. Combined with our designed commitment strategy, the integrated solution restricts the bids to a small neighbor around the corresponding value. Consequently, these approximately truthful bids enables our solution to generate revenue guarantee that converges to the optimal revenue over time, for time-discounted, or large market bidders. We generalize the i.i.d bidder setting in Huang et al. (2018a) and solve the open problem they proposed.

Extended Analysis of Private Quantile Algorithm. We extend the analysis of the quantile estimation oracles employed in this paper. For quantile estimation on bounded datasets (Kaplan et al., 2022), the paper assumes that all data points are *distinct* and derive accuracy bounds dependent on the dataset’s range. We generalize their analysis to accommodate cases where multiple data points may share *identical* values. Additionally, for quantile estimation of unbounded distributions (Durfee, 2023), we provide theoretical accuracy guarantees, complementing the paper’s focus on empirical performance.

²In practice, recognizable non-i.i.d. value distributions are common, e.g., Meta Ad platform (met) requires that each advertiser selects one of six objectives, corresponding to different distributions based on the industry or advertisement topic.

2 PRELIMINARIES

In this section, we outline the preliminaries on mechanism design, differential privacy, and quantile estimation. Additional information can be found in Appendix E.

2.1 MECHANISM DESIGN BASICS

We now formally define the allocation rule and payment rule of a single item auction.

Definition 2.1 (Allocation Rule and Payment Rule). Given k bidders with bid $\mathbf{b} := (b_1, \dots, b_k)$, a single-item auction M consists of an allocation rule as $\mathbf{x}(\mathbf{b}) := (x_1(\mathbf{b}), \dots, x_k(\mathbf{b})) \in [0, 1]^k$ and a payment rule as $\mathbf{p}(\mathbf{b}) := (p_1(\mathbf{b}), \dots, p_k(\mathbf{b})) \in [0, 1]^k$, where x_j denotes the probability that the j -th bidder gets the item, and p_j denotes her payment.

Under truthful sample access, the Myerson’s auction maximizes the expected revenue.

Definition 2.2 (Myerson’s Single Item Auction (Myerson, 1981)). For a discrete product distribution $\mathbf{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_k$ (Elkind, 2007), the *virtual value* for \mathcal{D}_j at value v_i^j with support $\mathcal{V}_j = \{v_1^j, \dots, v_n^j\}$ is $\phi_j(v_i^j) = v_i^j - (v_{i+1}^j - v_i^j) \frac{1 - F_j(v_i^j)}{f_j(v_i^j)}$, where v_i^j s are ordered in increasing order of i , $f_j(v_i^j) = \mathbb{P}[v^j = v_i^j]$, and $F_j(v_i^j) = \sum_{k=1}^i f(v_k^j)$.

We say the product distribution \mathbf{D} is η -strongly regular if for all j , $\phi_j(v_i) - \phi_j(v_j) \geq \eta(v_i - v_j)$ for every $v_i > v_j \in \mathcal{V}$ and $\eta > 0$.

For these distributions \mathcal{D} with nondecreasing virtual value, Myerson’s allocation rule $x_i(v_i) = \mathbb{1}\{\phi_i(v_i) \geq \max(0, \max_{j \neq i} \phi_j(v_j))\}$, where $\mathbb{1}\{\cdot\}$ denotes the indicator function. The payment rule $p_i(v_i) = \mathbb{1}\{\phi_i(v_i) \geq \max(0, \max_{j \neq i} \phi_j(v_j))\} \phi_i^{-1}(\max(0, \max_{j \neq i} \phi_j(v_j)))$.³

2.2 DIFFERENTIAL PRIVACY BASICS

We present the definition of pure DP and approximate DP below.

Definition 2.3 (Differential privacy). An algorithm $\mathcal{A} : \mathbb{R}_+^n \rightarrow \mathbb{R}$ is (ϵ, δ) -approximate DP if for neighboring dataset $V, V' \in \mathbb{R}_+^n$ that differs in only one data point, and any possible output O , we have: $\Pr[\mathcal{A}(V) = O] \leq \exp(\epsilon) \Pr[\mathcal{A}(V') = O] + \delta$. We say it satisfies pure DP for $\delta = 0$.

A key property we leverage from differential privacy is its immunity to post-processing. Post-processing refers to any computation or transformation applied to the output of a DP algorithm after the data has been privatized. In our context, Myerson’s auction can be seen as a post-processing step. Therefore, applying Myerson’s auction to a differentially private release of the empirical distribution preserves the original privacy guarantees of the input distribution.

Lemma 2.4 (Immunity to Post-Processing). Let $\mathcal{A} : \mathbb{R}_+^n \rightarrow \mathbb{R}$ be an (ϵ, δ) -DP algorithm, and let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a random function. Then, $f \circ \mathcal{A} : \mathbb{R}_+^n \rightarrow \mathbb{R}$ is also (ϵ, δ) -DP.

2.3 QUANTILE ESTIMATION

Quantile estimation (QE) is used for estimating a value of specified quantiles from samples. Given samples from a distribution, an accurate QE from samples directly translates to an accurate CDF estimation of the underlying distribution. Below, we formally introduce the definition of QE.

Definition 2.5 (Quantile Estimation). Given a range of the data as H , a dataset $X \subseteq H^n$ containing n points from range H , and a set of m quantiles $0 \leq q_1, \dots, q_m < 1$, identify quantile estimations v_1, \dots, v_m such that for every $j \in [m]$, $|\{x \in X | x \leq v_j\}| \approx q_j \cdot n$.⁴

We now present the definition of *statistical dominance* and *KS-distance* below.

Definition 2.6 (Stochastic Dominance and KS-Distance). Given distribution \mathcal{D} and \mathcal{D}' , we denote the CDF of them as $F_{\mathcal{D}}, F_{\mathcal{D}'}$, respectively. Distribution \mathcal{D} stochastically dominates distribution \mathcal{D}' (denoted as $\mathcal{D} \succeq \mathcal{D}'$) if: (1) For any outcome x , $F_{\mathcal{D}(x)} \leq F_{\mathcal{D}'(x)}$. (2) For some x , $F_{\mathcal{D}(x)} < F_{\mathcal{D}'(x)}$. The KS distance between \mathcal{D} and \mathcal{D}' is $d_{\text{ks}}(\mathcal{D}, \mathcal{D}') = \sup_{x \in \mathbb{R}} |F_{\mathcal{D}(x)} - F_{\mathcal{D}'(x)}|$.

³We define the virtual value inverse $\phi_i^{-1}(\phi)$ as $\arg \min_{v \in \mathcal{V}} \phi_i(v) \geq \phi$.

⁴More formally, $v_j \in X$ is the minimum value such that this quantity exceeds $q_j n$.

3 PRIVATE MYERSON’S AUCTION FOR BOUNDED DISTRIBUTIONS

In this section, we introduce the algorithm for fitting a Myerson’s auction with a pure privacy guarantee. To ensure pure privacy, since DP is immune to postprocessing, it is sufficient to input a private distribution estimated from samples to the Myerson. The challenge lies in finding such distributions that still yield near-optimal revenue.

Our approach leverages private quantile estimation (QE) over samples to achieve the desired guarantee. However, the standard guarantees of DPQE collapse when the dataset contains points that are extremely close. **This is a critical issue in our setting, as increasing the sample size n from continuous value distributions inherently causes the minimum distance between samples to approach zero.** To address this, we introduce additional discretization steps to prevent non-identical points from being too close together, and we develop new DPQE guarantees specifically tailored to handle samples with identical values.

3.1 PRIVATE MYERSON FOR BOUNDED DISTRIBUTIONS

Next, we present DPMYER algorithm (Alg. 1). The algorithm first value-discretize the samples of the distribution additively by ϵ_a , then quantile-discretize these samples by ϵ_q with pure privacy guarantee. Specifically, the quantile discretization estimates the values of the quantile set $[\epsilon_q, 2\epsilon_q, \dots, 1]$ with pure privacy. Next, DPMYER use the estimated quantile values and the quantile set to construct a distribution, then perturb it to a final distribution that is stochastically dominated by the ground truth. Finally, the final distribution is then used to implement Myerson’s mechanism.

Algorithm 1 DP Myerson, Bounded Distribution DPMYER($V, \epsilon_q, \epsilon_a, h, \epsilon_p$)

Input: n samples $V \in \mathbb{R}_+^{k \times n}$, discretization parameter ϵ_q, ϵ_a , upper bound h , privacy parameter ϵ_p

- 1: Discretize all values into multiples of ϵ_a ; let the resulting samples be \hat{V} .
 - 2: Prepare the quantile to be estimated: $Q \leftarrow \{\epsilon_q, 2\epsilon_q, \dots, \lfloor (1/\epsilon_q) \rfloor \cdot \epsilon_q, 1\}$.
 - 3: For each dimension $i \in [k]$, decide the prices $\hat{S}_{[i,:]} \leftarrow \text{QESTIMATE}(Q, V_{[i,:]}, \epsilon_p)$.
 - 4: ▷ Estimate the quantiles by DPQUANT (Alg. 4)
 - 5: Construct distribution \tilde{D} based on \hat{S} , treating the valuations in \hat{S} as if each has probability ϵ_q .
 - 6: For each $i \in [k]$, shift the top ϵ_q quantile of \tilde{D}_i to the bottom, fit Myerson on this distribution.
-

3.2 REVENUE OPTIMALITY AND RUNNING TIME

Next, we show the revenue optimality and the efficiency of our algorithm. To upper bound the revenue loss, we derive the revenue shift theorem, which upper bounds the revenue difference between two distributions by a linear function of their statistical distance.

Theorem 3.1 (Revenue Shift Theorem). *Given two product distribution $\mathbf{D} \succeq \mathbf{D}'$ whose valuations are bounded by h , with $d_{ks}(\mathbf{D}_i, \mathbf{D}'_i) \leq \alpha_i$ for any bidder i , the optimal revenue of these distribution satisfies: $0 \leq \mathbb{E}[\text{Rev}(M_{\mathbf{D}}, \mathbf{D}) - \text{Rev}(M_{\mathbf{D}'}, \mathbf{D}')] \leq (\sum_{i \in [k]} \alpha_i)h$.*

We apply this theorem to upper bound the revenue loss between 1) the quantile-discretized distribution and its pre-quantized counterpart, and 2) the distribution obtained from private quantile estimation and that from the groundtruth quantile estimation. The first one is evident, while the second arises from DPQUANTILE’s ability to control the KS-distance between the estimation and the ground truth.

We now present the accuracy guarantee of the private Myerson algorithm. Provided the privacy parameter is not too small (i.e., $\epsilon_p = \Omega(\epsilon^{-1})$), our guarantee implies that the optimal revenue of the distribution does not exceed the revenue of our algorithm on its samples by more than $\tilde{\Theta}(\epsilon kh)$.

Theorem 3.2 (Revenue Guarantee of Private Myerson (Alg. 1)). *Given n samples $\hat{V} \in [0, h]^{k \times n}$ of the joint distribution \mathbf{D} , DPMYER (Alg. 1) is $(2k\epsilon_p, 0)$ -DP, and the expected revenue of this mechanism is close to the optimal revenue of distribution \mathbf{D} , i.e., with probability $1 - \delta$:*

$$|\mathbb{E}[\text{Rev}(M_{\text{DPMYER}}, \mathbf{D}) - \text{OPT}(\mathbf{D})]| \leq \tilde{O}((\epsilon + \epsilon^2/\epsilon_p)kh).$$

under parameter $\epsilon_a = \epsilon_q = \epsilon$ and $n = \tilde{\Theta}(\epsilon^{-2})$, where we hide the polylog factors in $\tilde{\Theta}$ and \tilde{O} .

Proof Sketch. We begin by deriving the privacy guarantee of our algorithm. Next, we establish an upper bound on the distance between the private distribution \hat{D}^p and the additively discretized distribution \hat{D}^ϵ . This enables us to apply the revenue shift theorem (Thm.3.1) to upper bound the revenue loss from private quantile estimation. By aggregating this loss with the revenue loss due to value discretization, we arrive at the final result. In this proof sketch, we omit the polylog factors that depends on $k, n, \delta, \epsilon_a, \epsilon_p, \epsilon_q$ for a clear presentation. Further details are provided in Appendix H.2.

Privacy Guarantee. We know that the quantile estimates from DPQE is $(\epsilon_p, 0)$ private (Lem. H.2). Since DP is immune to post-processing (Lem. E.4), and that the output of allocation and payment combination is $2k$ dimensional, by composition theorem (Lem. E.5), our algorithm is $(2k\epsilon_p, 0)$ -DP.

Upper Bounding the Statistical Distance The distribution \hat{D}^p is obtained by changing from distribution \mathbf{D} through distribution \hat{D} , the distribution \hat{D}^ϵ and \hat{D}^q (Figure 1). We upper bound the statistical KS distance of these distributions: 1) By DKW inequality, we upper bound the KS-distance between \hat{D} and \mathbf{D} by $\tilde{\Theta}(1/\sqrt{n})$ for each coordinate i (with probability $1 - \delta/2$). 2) By definition, we upper bound the KS-distance between \hat{D}^ϵ and \hat{D}^q by $k\epsilon_q$. 3) By developing and converting the bound of the DP quantile algorithm (Lem. H.3) into a bound on the CDF, we upper bound the KS-distance between \hat{D}^q and \hat{D}^p by $k\hat{\epsilon}$ for $\hat{\epsilon} := \tilde{\Theta}(1/(\epsilon_p n))$ (with probability $1 - \delta/2$).

Upper Bounding the Revenue Loss. We then upper bound optimal revenue loss from \mathbf{D} to \hat{D}^p . This upper bound can be obtained by combining the revenue loss from the aforementioned distributions (by revenue shift theorem), with an additive ϵ_a revenue loss from discretization (by Lem. F.1). The revenue loss from statistical shift aggregates to $\tilde{\Theta}((1/\sqrt{n} + \epsilon_q + \hat{\epsilon})kh)$ with probability $1 - \delta$.

Putting it all together. Finally, condition on the DPQUANT proceeds successfully and the samples are close to the underlying distribution (with probability $1 - \delta$), we get that the expected revenue of DPQUANT on the underlying distribution is at least the optimal revenue from this distribution minus the revenue difference between \mathbf{D} and \hat{D}^p by the following inequality:

$$0 \geq \mathbb{E}[\text{Rev}(M_{\hat{D}^p}, \mathbf{D}) - \text{OPT}(\mathbf{D})] \geq \mathbb{E}[\text{Rev}(M_{\hat{D}^p}, \mathbf{D}) - \text{OPT}(\hat{D}^p)] - |\text{OPT}(\hat{D}^p) - \text{OPT}(\mathbf{D})|$$

where the first inequality follows from the optimality of $M_{\mathbf{D}}$ on \mathbf{D} and the second inequality follows from adding $\text{OPT}(\hat{D}^p)$. By our construction of \hat{D}^p , this distribution is stochastically dominated by \mathbf{D} , thus from the strong revenue monotonicity (Lem. F.3), we get that $\mathbb{E}[\text{Rev}(M_{\hat{D}^p}, \mathbf{D}) - \text{OPT}(\hat{D}^p)] \geq 0$. Thus, we concluded that the revenue gap is upper bounded by $\tilde{\Theta}((1/\sqrt{n} + \epsilon_q + \hat{\epsilon})kh + \epsilon_a)$. We set δ in the statement as $1/k$ of the δ we used in this proof to generate the final revenue guarantee. \square




Figure 1: **Distribution analyzed for DPMYER(Alg. 1).** We establish connections between the accuracy/revenue guarantee of the original distribution \mathbf{D} with the empirical distribution \hat{D} , the value-discretized \hat{D}^ϵ , the quantile-discretized \hat{D}^q and the distribution \hat{D}^p returned by DPQUANT(Alg. 4).

Next, we demonstrate the efficiency of our algorithm, which is achieved through a organized implementation of the DP Quantile algorithm. Intuitively, given m ordered quantiles, the algorithm iteratively identifies and estimates the median (the $m/2$ -th), followed by the $m/4$ and the $3m/4$ quantiles, and so on. This hierarchical structure ensures that each data point is used in at most $\log m$ quantile estimates (of a single quantile). For more details, we refer readers to Appendix H.1.

Theorem 3.3 (Time Complexity for Private Myerson, Bounded). *Given the same parameters as stated in Theorem 3.2, DPMYER (Alg.1) runs in $\tilde{\Theta}(kn)$ time and requires $\tilde{\Theta}(1)$ passes of the samples.*

Proof Sketch. The time dominant step is *quantile estimation*, which requires $\log(\lfloor 1/\epsilon_q \rfloor + 1)$ passes of the dataset. It takes $O(k \log(\lfloor 1/\epsilon_q \rfloor + 1)/(\epsilon_a \epsilon_q)) = \tilde{\Theta}(kn)$ time, since $n = \tilde{\Theta}(\epsilon^{-2})$. This step calculates the utility of $k \lfloor h/\epsilon_a \rfloor$ over $\lfloor 1/\epsilon_q \rfloor$ quantiles for at most $\tilde{\Theta}(1)$ time. For full version of this proof, please refer to Appendix H.3 \square

4 GENERALIZATION TO UNBOUNDED DISTRIBUTIONS

Generalizing the DP Myerson mechanism to unbounded distributions introduces new challenges. The revenue loss upper bound produced by previously introduced *quantile estimation* algorithm and *revenue shift* theorem both depends (positively) on the range of the distribution. Without a finite range, these upper bound becomes infinite and fail to effectively control the revenue loss.

We consider the widely accepted η -strongly regular distributions, which decays at least as fast as exponential distributions. A key element of our approach is appropriately truncating the distribution, which enables us to extend the discretize-then-DP-quantile method to the unbounded setting. Specifically, we apply the property of the regular distribution that (Devanur et al., 2016), truncating the distribution by $\frac{1}{\epsilon} \text{OPT}(\mathbf{D})$ costs at most 2ϵ fraction of the optimal revenue (Lem. I.1). Hence, for the truncation to work, it is essential to approximate the optimal revenue based on sample data. Meanwhile, incorporating the truncation with pure DP introduces additional complexities.

We are now ready to present our approach for a k -approximation of the optimal revenue with pure DP for η -strongly regular product distributions. Our DPKOPT (Alg. 2) algorithm approximates the optimal revenue by running a empirical reserve (ER) over *each* bidder’s distribution truncated at the top $\eta^{1/(1-\eta)}/4$ quantile.⁵ Summing up these estimates gives us a $\Theta(k)$ -approximation of the optimal revenue, by the fact that $k\text{OPT}(\mathbf{D}) \geq \sum_{i \in [k]} \text{OPT}(\mathcal{D}_k) \geq \text{OPT}(\mathbf{D})$.

Algorithm 2 DP Estimation for Optimal Revenue DPKOPT($V, \epsilon_q, \epsilon_a, \epsilon_p, \eta$)

Input: n samples $V = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$, quantile discretization ϵ_q , additive discretization ϵ_a , privacy parameter ϵ_p , regularity parameter η .

- 1: **for** $d = 1 \rightarrow k$ **do**
- 2: $\hat{q} \leftarrow 1/4 \cdot \eta^{1/(1-\eta)}$
- 3: Let $ub_d \leftarrow \text{DPQUANTU}(V_{[d,:]}, 1 - \hat{q})$. ▷ Estimate the truncation point of D_d .
- 4: Truncate distribution D_d at ub_d as \hat{D}_d , and discretize \hat{D}_d by additive ϵ_a in the value space.
- 5: Prepare the quantile to be estimated, $Q \leftarrow \{1 - \hat{q}, 1 - \hat{q} - \epsilon_q, \dots, 1 - \hat{q} - \lfloor \frac{1-\hat{q}}{\epsilon_q} \rfloor \cdot \epsilon_q, 0\}$.
- 6: $\hat{S}_{[d,:]} \leftarrow \text{QESTIMATE}(Q, V_{[d,:]}, \epsilon_p)$ ▷ Apply DP quantile estimate (Alg. 4).
- 7: Let \hat{F}_d be the distribution generated by value profile $\hat{S}_{[d,:]}$ and quantile set Q .
- 8: $\text{SREV}_d \leftarrow \max_{r \in \hat{S}} r(1 - \hat{F}_d(r))$. ▷ Estimate the optimal revenue from \hat{F}_d (Alg. 6).
- 9: **end for**
- 10: $\text{KREV} \leftarrow \sum_{d \in [k]} \text{SREV}_d$
- 11: **return** KREV

To guarantee pure privacy, our algorithm estimates the optimal revenue using a DP-estimated proxy $\hat{F}_{[k]}$ derived from the sample data. This proxy is obtained from truncating the distribution by DPQUANTU (Alg. 7) and quantile-discretizing the distribution by DPQUANT. During this process, the truncation by DPQUANTU cost at most a constant fraction of the optimal revenue, and DPQUANT cost at most an additional $\tilde{\Theta}(\frac{1}{\epsilon_p n} k + \epsilon_a)$. Aggregating these revenue loss concludes that the output is a $\Theta(k)$ -approximation of the optimal revenue. See Appendix I.4 for more details.

Our private Myerson algorithm for the unbounded distribution (Alg. 9) integrate DPKOPT and yields the following accuracy bound. See Appendix I.5 for formal statements and more details.

Theorem 4.1 (Revenue Guarantee of Private Myerson, Unbounded). *Given $\epsilon \in [0, 1/4]$, n samples \hat{V} of the joint distribution $\mathbf{D} \in [0, h]^k$, the output of Myerson fitted under DPMYERU (Alg. 9) is $(2k\epsilon_p, 0)$ -DP, and under $\epsilon_a = \epsilon_q = \epsilon$, $n = \tilde{\Theta}(\epsilon^2)$, $n_1 = \tilde{\Theta}(\epsilon^2)$, $\epsilon_t = \sqrt{\epsilon}$, with probability $1 - \delta$,*

$$\mathbb{E}[\text{Rev}(M_{\text{DPMYERU}}, \mathbf{D}) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})] \leq \tilde{O}(k^2 \sqrt{\epsilon} + k^2 \epsilon^{1.5} / \epsilon_p)$$

⁵Without privacy constraints, truncating at the top $\eta^{1/(1-\eta)}$ -suffices by Lem. I.4. Our algorithm adopt a looser truncation since the DPQUANTU algorithm only return the value of given quantiles *approximately*.

5 APPLICATION: ONLINE MECHANISM DESIGN FROM BIDS

We now study how to integrate our previous solutions into the online auction setting, such that, the algorithm produces time-averaged revenue guarantee that converges to the optimal. The auction now spans multiple rounds, where each auction is informed by the bids from previous rounds. We consider the setting where bidders are non-myopic bidders, and have incentives to bid strategically in the current round to increase their utilities over future auctions.

5.1 APPLICATION BACKGROUND

Before presenting our algorithm, we first provide the formal problem definition of the online auction setting. We study online mechanism design over a time horizon of T , where an identical item is sold at each iteration. Each bidder has a *publicly observable* attribute. Bidders with the same attribute have the same valuation distribution.

We are now ready to describe interactions between bidders and the auctioneer over time horizon T , as shown in Figure 2. We defer to Appendix J.2 for more details how bidder generates the samples.

For each time $t \in [T]$:

- The learner/auctioneer sells a fresh copy of the item.
- The learner collects the bids in the form of (b_j, a_j) , where b_j and a_j denote the bid and the attribute of the $j \in [d_t]$ -th bidder, respectively.
- The learner decides the allocation rule \mathbf{x}_t and payment \mathbf{p}_t accordingly.

Figure 2: Online Auction with k Attributes.

Each item the auctioneer sells is identical, and each bidder has an additive (discounted) utility of the items across rounds. We consider the bidders either have *discounted utility* or are in a *large market*.

Definition 5.1 (Bidder’s Utility). Each bidder j has a quasi-linear utility function at time t : $u_j^t = x_j^t(v_j^t - p_j^t)$, where x_j^t, v_j^t, p_j^t are the allocation, value, price for bidder j at time t , respectively. We consider two *nonmyopic* bidders’ utility models:

Discounted Utility: For discount factor $\gamma \in [0, 1]$, the bidders seek to maximize the sum of utilities discounted by γ . At the t -th iteration, the discounted utility is $\hat{u}_j^t = \sum_{r=t}^T u_j^r \gamma^{r-t}$.

Large Market: (Anari et al., 2014; Jalaly Khalilabadi and Tardos, 2018; Chen et al., 2016): The bidder only participates in a subset S_j of auctions, i.e., for each $u^{1:T} = \sum_{t \in S_j} u^t$, with subset $|S_j| < l$.

Ideally, the learner’s objective is maximize time-averaged revenue with high probability. Our regret compare this revenue against the optimal revenue of the (unobservable) value history.

Definition 5.2 (Learner’s Objective). Given δ , the learner’s objective is to decide an allocation $\mathbf{x}_{1:T}$ and a payment $\mathbf{p}_{1:T}$ that achieves sublinear regret, i.e., with probability $1 - \delta$,

$$\text{REGRET} := \frac{1}{T} \sum_{t \in [T]} \mathbb{E}[\text{Rev}(\mathbf{x}_t, \mathbf{p}_t, \mathbf{b}_t) - \mathbb{E}[\text{OPT}(\mathbf{v}_t)]] = o(1),$$

with the expectation taken over the value distribution.

5.2 TWO-STAGE MECHANISM FOR BOUNDED DISTRIBUTION

This two-stage algorithm (Alg. 3) consists of repeated auctions over T rounds, and the participating bidders’ values in each round are upper bounded by a *known* constant h . The algorithm first collects the samples for the first T_1 rounds, by running a commitment algorithm (Alg. 10) that punishes nontruthful bids. Then, the algorithm deploys our previously developed DP Myerson’s Algorithm (Alg. 1, Alg. 9) for the remaining rounds to obtain near optimal revenue. In addition to these two steps, our algorithm includes a step where all samples are reduced by ν (line 4 of Alg. 3) and projected onto nonnegative value spaces. This step is designed to offset the impact of strategic bidding.

Algorithm 3 Two-Stage Algorithm $\mathcal{A}_{\text{BOUNDED}}$

Input: Rounds T , learning rounds T_1 , parameter $\epsilon_a, \epsilon_q, \epsilon_p, \nu$, upper bound h .

- 1: **for** $t \leftarrow 1, \dots, T_1$ **do** ▷ Collection Stage
- 2: Receive bids \mathbf{b}^t , and attributes \mathbf{a}^t .
- 3: Return $(\mathbf{x}^t, \mathbf{p}^t) \leftarrow \text{COMMIT}(\mathbf{b}^t)$. ▷ Commitment Algorithm(Alg. 10)
- 4: $\tilde{\mathbf{b}}^t \leftarrow \mathbb{P}_{[0, h]}[\mathbf{b}^t - \nu \mathbf{1}_k]$
- 5: **end for**
- 6: $(\tilde{\mathbf{x}}(\cdot), \tilde{\mathbf{p}}(\cdot)) \leftarrow \text{DPMYER}(\hat{\mathbf{b}}^{1:T_1}, \mathbf{a}^{1:T_1}, \epsilon_q, \epsilon_a, h, \epsilon_p)$ ▷ Fit Myerson's auction (Alg. 1, or Alg. 9)
- 7: **for** $t \leftarrow T_1 + 1, \dots, T$ **do** ▷ Revenue Stage
- 8: Receive bids \mathbf{b}^t , and attributes \mathbf{a}^t .
- 9: $(\mathbf{x}^t, \mathbf{p}^t) \leftarrow \text{MYERSON}(\tilde{\mathbf{x}}(\cdot), \tilde{\mathbf{p}}(\cdot))$;
- 10: **end for**

Specifically, the parameter ν is carefully calibrated to ensure that the bid distribution fed into the private Myerson mechanism is stochastically dominated by the empirical distribution. Our algorithm provides an incentive guarantee that bids lie within a small, controllable neighborhood of the true values. The range of this neighborhood is determined by the privacy parameter ϵ_p (hence is controlled by our algorithm), and the bidders' utility functions. By setting ν to match the range of this neighborhood, the resulting distribution is dominated by the empirical distribution.

5.3 REVENUE GUARANTEE OF THE ALGORITHM

Before presenting the revenue guarantee of our main algorithm, we first introduce a lemma that upper bounds how a bidder's bid deviates from its true value during the collection stage. Intuitively, by the design of our commitment algorithm the bidder will incur a loss that scales (positively) with the bid deviation, compared to truthful bidding. Furthermore, our private Myerson ensures that the bidder's future utility gain is upper bounded (Lem. J.5). Thus, bidders are incentivized to report bids within a certain range of their true values to optimize their overall utility. More details in Appendix J.4.

Lemma 5.3 (Bid Deviation). *For any $t \in [0, T_1]$, the bidder will bid only b_t such that $|b_t - v_t| \leq 2\alpha$, where $\alpha = \sqrt{2(l-1)\epsilon_p}hk$ for bidders in a large market; and $\alpha = \sqrt{\frac{2\gamma\epsilon_p}{1-\gamma}}kh$ for discounting bidder.*

From this lemma, we get that selecting a small ϵ_p would incentivize bid distributions that are close to the ground-truth. Let $\nu = 2\alpha$ in our algorithm (line 4, Alg. 3) would yield a distribution that is stochastically dominated by, yet close in revenue guarantee to, the true distribution. Run our DP Myerson algorithm on this distribution would give us sublinear regret, as stated below.

Theorem 5.4 (Accuracy Guarantee of Two-stage Mechanism). *Given $\epsilon \in [0, 1/4]$, n samples of the joint distribution $\mathbf{D} \in [0, h]^k$, and $T_1 = \Theta(\epsilon^{-2} \log(k/\delta))$, $T = \Omega(T_1)$, $\epsilon_a = \epsilon_q = \epsilon_p = \epsilon$, with probability $1 - \delta$, Alg. 3 generates sublinear regret, i.e.,*

Under a large market, the regret is upper bounded by $\tilde{O}[(\epsilon + \sqrt{l\epsilon})kh]$, for $\nu = 2\sqrt{2(l-1)\epsilon_p}hk$.

Under discounting bidder, the regret is upper bounded by $\tilde{O}[(\epsilon + \sqrt{\frac{\gamma\epsilon}{1-\gamma}})kh]$, for $\nu = 2\sqrt{\frac{2\gamma\epsilon_p}{1-\gamma}}kh$.

Proof Sketch. We denote the empirical distribution as $\hat{\mathbf{D}}$, the distribution after subtraction in line 4 of Alg. 3 as $\tilde{\mathbf{D}}$, and the (final) output distribution as $\hat{\mathbf{D}}^p$. Then these distribution satisfies $\hat{\mathbf{D}} \succeq \tilde{\mathbf{D}} \succeq \hat{\mathbf{D}}^p$. By strong monotonicity(Lem. F.3), we know that $\mathbb{E}[\text{Rev}(M_{\hat{\mathbf{D}}^p}, \mathbf{D})] \geq \mathbb{E}[\text{OPT}(\hat{\mathbf{D}}^p)]$. Since $M_{\hat{\mathbf{D}}^p}$ need not be optimal over \mathbf{D} , we have that:

$$\begin{aligned}
0 &\geq \mathbb{E}[\text{Rev}(M_{\hat{\mathbf{D}}^p}, \mathbf{D}) - \text{OPT}(\mathbf{D})] \\
&\geq \mathbb{E}[\text{Rev}(M_{\hat{\mathbf{D}}^p}, \mathbf{D}) - \text{OPT}(\hat{\mathbf{D}}^p)] + \mathbb{E}[\text{OPT}(\hat{\mathbf{D}}^p) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})] \\
&\geq \mathbb{E}[\text{OPT}(\hat{\mathbf{D}}^p) - \text{OPT}(\hat{\mathbf{D}})] - |\mathbb{E}[\text{OPT}(\hat{\mathbf{D}}) - \text{OPT}(\mathbf{D})]| \geq -\tilde{\Theta}((\epsilon + \epsilon^2/\epsilon_p)kh + \nu).
\end{aligned}$$

where in the last inequality we apply revenue shift theorem (Thm. 3.1) to upper bound the first term and apply Lemma J.9 to upper bound the second term. Please refer to Appendix J.3 for more details.

□

6 EXPERIMENTS

In this section, we present the experimental results for the Differentially Private (DP) Myerson mechanism, comparing its performance against two standard mechanism design baselines: the *Myerson* (optimal) auction and the *Vickrey* (second-price) auction. The former is designed to achieve near-optimal revenue for a given value distribution, whereas the latter, while strategy-proof, offers no revenue guarantees in settings with independent and non identical value distributions.

Our experiments are conducted on normal and lognormal distributions truncated to positive domains. The lognormal distribution is widely considered a representative or “groundtruth” model in many auction settings, thanks to its capacity to capture a broad spectrum of value distributions commonly observed in economic and market contexts (Gorbenko and Malenko, 2014). A random variable V is said to be lognormal distributed with parameter (μ, σ) , if $\ln(V)$ follows normal distribution $\mathcal{N}(\mu, \sigma)$.

For each value profile, we test various hyperparameters—additive discretization (ϵ_a), quantile discretization (ϵ_q), and the privacy parameter (ϵ_p)—and select the configuration with the *best* performance. For details on DP Myerson’s sensitivity to hyperparameters, see Appendix A.

Bidder Profile	DP Myerson	Second Price	Myerson	Ref.
Normal $\mathcal{N}(0.3, 0.5)$	0.25272	0.15154 (66.7 %)	0.32598	Table 2
Lognormal $(\mu, \sigma) = (-1.87, 1.15)$				
Normal $\mathcal{N}(0.3, 0.5)$	0.37691	0.33741 (11.7 %)	0.50204	Table 3
Normal $\mathcal{N}(0.5, 0.7)$				
Lognormal $(\mu, \sigma) = (-1.87, 1.15)$	0.13912	0.11578 (20.2 %)	0.21292	Table 4
Lognormal $(\mu, \sigma) = (-1.24, 1.04)$				

Table 1: Empirical Revenue of DPMYerson (Alg. 1) under 2-dimensional non-identical value distributions. Each DPMYerson configuration is averaged over 50 draws, with revenue evaluated on 10,000 samples. Percentages in parentheses represent the improvement over the second-price mechanism.

In Table 1, under non i.i.d distribution settings where there is a significant revenue gap between the Vickrey auction and the Myerson auction, DPMYerson achieves a notable revenue increase (at least 11%) over the second-price mechanism.

7 CONCLUSION

We investigate the problem of learning a single-item auction (i.e., Myerson) from samples with *pure* DP. We consider the broader setting where the agents’ valuations are *independent, non-identical*, and can either be *bounded* or *unbounded*. By recognizing that the optimal auction mechanism exhibits robustness to small statistical perturbations in the underlying distribution, we reduce the challenge of privately learning an optimal auction from sample data to the task of privately approximating pre-specified quantiles. Specifically, our approach ensures pure privacy while generating a distribution that is closely aligned with the underlying distribution in terms of expected revenue.

We then extend this framework to the online auction setting, where later auctions are fitted on bids from previous auctions. In this setting, non-myopic bidders reason about their utility across rounds, and can bid strategically under (one-shot) truthful auctions. By leveraging our private Myerson mechanisms with an extra commitment mechanism, we achieve near-optimal revenue outcomes over the bidders’ (unobservable) value samples, despite the strategic complexity introduced by non-myopic behavior (i.e., time discounting bidder and/or non-discounting bidders in a large market). This result highlights the robustness of our approach in both protecting privacy and maintaining near optimal expected revenue in dynamic, strategic environments.

REFERENCES

- The ad auction explained. URL <https://www.facebook.com/business/ads/ad-auction>. (page 3)
- Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318, 2016. (page 1)
- Jacob Abernethy, Chansoo Lee, Abhinav Sinha, and Ambuj Tewari. Online linear optimization via smoothing. In *Conference on learning theory*, pages 807–823. PMLR, 2014. (page 18)
- Jacob D Abernethy, Rachel Cummings, Bhuvish Kumar, Sam Taggart, and Jamie Morgenstern. Learning auctions with robust incentive guarantees. In *NeurIPS*, pages 11587–11597, 2019. (page 3, 18)
- Kareem Amin, Afshin Rostamizadeh, and Umar Syed. Learning prices for repeated auctions with strategic buyers. *Advances in Neural Information Processing Systems*, 26, 2013. (page 19)
- Nima Anari, Gagan Goel, and Afshin Nikzad. Mechanism design for crowdsourcing: An optimal $1-1/e$ competitive budget-feasible mechanism for large markets. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 266–275. IEEE, 2014. (page 8, 36)
- M.-F. Balcan, A. Blum, J.D. Hartline, and Y. Mansour. Mechanism design via machine learning. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS’05)*, pages 605–614, 2005. doi: 10.1109/SFCS.2005.50. (page 18)
- Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Item pricing for revenue maximization. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 50–59, 2008. (page 18)
- Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik. Estimating approximate incentive compatibility. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 867–867, 2019. (page 18)
- Maria-Florina F Balcan, Tuomas Sandholm, and Ellen Vitercik. Sample complexity of automated mechanism design. In *Advances in Neural Information Processing Systems*, pages 2083–2091, 2016. (page 18)
- Santiago R Balseiro, Omar Besbes, and Francisco Castro. Mechanism design under approximate incentive compatibility. *Operations Research*, 72(1):355–372, 2024. (page 18)
- Sébastien Bubeck, Nikhil Devanur, Zhiyi Huang, and Rad Niazadeh. Multi-scale online learning: Theory and applications to online auctions and pricing. *Journal of Machine Learning Research*, 2019. (page 18)
- Yang Cai and Constantinos Daskalakis. Extreme-value theorems for optimal multidimensional pricing. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 522–531. IEEE, 2011. (page 18)
- Ning Chen, Xiaotie Deng, Bo Tang, and Hongyang Zhang. Incentives for strategic behavior in fisher market games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016. (page 8, 36)
- Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 243–252, 2014. (page 18, 30)
- Xiaotie Deng, Ron Lavi, Tao Lin, Qi Qi, Wenwei Wang, and Xiang Yan. A game-theoretic analysis of the empirical revenue maximization algorithm with endogenous sampling. *Advances in Neural Information Processing Systems*, 33:5215–5226, 2020. (page 3, 18)

- Yuan Deng, Sébastien Lahaie, Vahab Mirrokni, and Song Zuo. Revenue-incentive tradeoffs in dynamic reserve pricing. In *International Conference on Machine Learning*, pages 2601–2610. PMLR, 2021. (page 18)
- Nikhil R Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. The sample complexity of auctions with side information. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 426–439, 2016. (page 7, 18, 24, 30)
- David Durfee. Unbounded differentially private quantile and maximum estimation. *Advances in Neural Information Processing Systems*, 36, 2023. (page 1, 3, 31, 35)
- Paul Dütting, Zhe Feng, Harikrishna Narasimhan, David C Parkes, and Sai Srivatsa Ravindranath. Optimal auctions through deep learning: Advances in differentiable economics. *Journal of the ACM*, 71(1):1–53, 2024. (page 18)
- Aryeh Dvoretzky, Jack Kiefer, and Jacob Wolfowitz. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *The Annals of Mathematical Statistics*, pages 642–669, 1956. (page 22)
- Cynthia Dwork. Differential privacy. In *International colloquium on automata, languages, and programming*, pages 1–12. Springer, 2006. (page 1)
- Cynthia Dwork. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*, pages 1–19. Springer, 2008. (page 1)
- Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology-EUROCRYPT 2006: 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28-June 1, 2006. Proceedings 25*, pages 486–503. Springer, 2006. (page 22)
- Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential privacy under continual observation. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 715–724, 2010. (page 18)
- Edith Elkind. Designing and learning optimal finite support auctions. 2007. (page 4, 18, 21)
- European Commission. Regulation of the european parliament and of the council on the transparency and targeting of political advertising, 2024. URL <https://data.consilium.europa.eu/doc/document/PE-90-2023-INIT/en/pdf>. (page 17)
- Alexander S Gorbenko and Andrey Malenko. Strategic and financial bidders in takeover auctions. *The Journal of Finance*, 69(6):2513–2555, 2014. (page 10)
- Chenghao Guo, Zhiyi Huang, and Xinzhi Zhang. Settling the sample complexity of single-parameter revenue maximization. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 662–673, 2019. (page 18)
- Wenshuo Guo, Michael Jordan, and Emmanouil Zampetakis. Robust learning of optimal auctions. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021. (page 25)
- Jason Hartline, Vahab Mirrokni, and Mukund Sundararajan. Optimal marketing strategies over social networks. In *Proceedings of the 17th international conference on World Wide Web*, pages 189–198, 2008. (page 30)
- Jason D Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 225–234, 2009. (page 18, 22)
- Zhiyi Huang, Jinyan Liu, and Xiangning Wang. Learning optimal reserve price against non-myopic bidders. In *Advances in Neural Information Processing Systems*, 2018a. (page 3, 18)
- Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. Making the most of your samples. *SIAM Journal on Computing*, 47(3):651–674, 2018b. (page 18, 30)

- Joon Suk Huh and Kirthevasan Kandasamy. Nash incentive-compatible online mechanism learning via weakly differentially private online learning. *ICML*, 2024. (page 3)
- Prateek Jain, Pravesh Kothari, and Abhradeep Thakurta. Differentially private online learning. In *Conference on Learning Theory*, pages 24–1. JMLR Workshop and Conference Proceedings, 2012. (page 1)
- Pooya Jalaly Khalilabadi and Éva Tardos. Simple and efficient budget feasible mechanisms for monotone submodular valuations. In *International Conference on Web and Internet Economics*, pages 246–263. Springer, 2018. (page 8, 36)
- Yash Kanoria and Hamid Nazerzadeh. Dynamic reserve prices for repeated auctions: Learning from bids. In *Web and Internet Economics: 10th International Conference, WINE 2014, Beijing, China, December 14-17, 2014, Proceedings*, volume 8877, page 232. Springer, 2014. (page 3)
- Yash Kanoria and Hamid Nazerzadeh. Incentive-compatible learning of reserve prices for repeated auctions. *Operations Research*, 69(2):509–524, 2021. (page 18)
- Haim Kaplan, Shachar Schnapp, and Uri Stemmer. Differentially private approximate quantiles. In *International Conference on Machine Learning*, pages 10751–10761. PMLR, 2022. (page 1, 2, 3, 26, 27)
- Michael Kearns, Mallesh Pai, Aaron Roth, and Jonathan Ullman. Mechanism design in large games: Incentives and privacy. In *Proceedings of the 5th conference on Innovations in theoretical computer science*, pages 403–410, 2014. (page 17)
- Michael J Kearns and Umesh Vazirani. *An introduction to computational learning theory*. MIT press, 1994. (page 2)
- Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 94–103. IEEE, 2007. (page 1, 3, 18, 37)
- Jamie H Morgenstern and Tim Roughgarden. On the pseudo-dimension of nearly optimal auctions. *Advances in Neural Information Processing Systems*, 28, 2015. (page 18)
- Roger B Myerson. Incentive compatibility and the bargaining problem. *Econometrica: journal of the Econometric Society*, pages 61–73, 1979. (page 1, 18)
- Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981. (page 4, 21)
- Kobbi Nissim, Rann Smorodinsky, and Moshe Tennenholtz. Approximately optimal mechanism design via differential privacy. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 203–213, 2012. (page 1, 3, 18)
- Mallesh M Pai and Aaron Roth. Privacy and mechanism design. *ACM SIGecom Exchanges*, 12(1): 8–29, 2013. (page 1)
- Tim Roughgarden and Okke Schrijvers. Ironing in the dark. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 1–18, 2016. (page 18)
- Adam Smith. Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 813–822, 2011. (page 26)
- Junhua Zhang and Caiming Zhong. Differential privacy-based double auction for data market in blockchain-enhanced internet of things. *Wireless Communications and Mobile Computing*, 2022 (1):8038846, 2022. (page 3)
- Xiaojin Zhang, Yan Kang, Kai Chen, Lixin Fan, and Qiang Yang. Trading off privacy, utility, and efficiency in federated learning. *ACM Transactions on Intelligent Systems and Technology*, 14(6): 1–32, 2023. (page 1)

A MORE DETAILS FROM EXPERIMENTS

In this section, we present the experimental results for the bounded Differentially Private (DP) Myerson mechanism, comparing its performance against two standard mechanism design baselines: the *Myerson* (optimal) auction and the *Vickrey* (second-price) auction. The *Myerson* auction is designed to achieve near-optimal revenue for a given value distribution, whereas the *Vickrey* auction, while strategy-proof, offers no revenue guarantees in settings with non-i.i.d. value distributions. All reported results are based on experiments conducted with at least 10,000 samples. In the following subsections, we detail the various experimental settings and configurations considered in this study.

A.1 DISTRIBUTION

Definition A.1 (Log Normal Distribution). A random variable V is said to be lognormal distributed with parameter (μ, σ) , if $\ln(V)$ follows normal distribution $\mathcal{N}(\mu, \sigma)$.

Given mean μ_V and standard deviation σ_V , the parameter of the underlying normal distribution is as follows:

$$\mu = \ln\left(\frac{\mu_V^2}{\sqrt{\mu_V^2 + \sigma_V^2}}\right), \sigma = \sqrt{\ln\left(1 + \frac{\sigma_V^2}{\mu_V^2}\right)}$$

A.2 EMPIRICAL REVENUE ANALYSIS OF DP MYERSON ACROSS HYPERPARAMETERS

In this subsection, we evaluate the performance of DP Myerson across various hyperparameter configurations, including additive discretization (ϵ_a), privacy parameter (ϵ_p), and quantile discretization (ϵ_q), using 2-dimensional non-i.i.d. distributions. We compare the revenue guarantees of DP Myerson against those of the second-price auction and the Myerson auction. It is important to note that we do not expect DP Myerson to outperform the Myerson auction, as the latter converges to the optimal auction as the sample size increases.

Bidder 1	Normal $\mathcal{N}(0.3, 0.5)$			
Bidder 2	Lognormal $(\mu, \sigma) = (-1.87, 1.15)$			
Privacy Parameter ϵ_p	0.1	0.2	0.4	0.7
<hr/>				
<i>w/ additive $\epsilon_a = 0.05$, Upper Bound $h = 1$.</i>				
DPMY w/ $\epsilon_q = 0.05$	0.14974	0.15777	0.17104	0.14598
DPMY w/ $\epsilon_q = 0.1$	0.15564	0.17322	0.17218	0.18670
DPMY w/ $\epsilon_q = 0.14$	0.18674	0.19108	0.18960	0.18429
Second Price	0.17755			
Myerson	0.33291			
<hr/>				
<i>w/ additive $\epsilon_a = 0.1$, Upper Bound $h = 1$.</i>				
DPMY w/ $\epsilon_q = 0.26$	0.2500	0.25272	0.248056	0.24658
DPMY w/ $\epsilon_q = 0.31$	0.24780	0.24387	0.24540	0.24689
DPMY w/ $\epsilon_q = 0.36$	0.24534	0.24899	0.24687	0.24723
Second Price	0.15154			
Myerson	0.32598			

Table 2: Average Empirical Revenue of DP Myerson under non-i.i.d. Value Distributions with Varying Discretization Parameters (additive ϵ_a , quantile ϵ_q) and Privacy Parameter ϵ_p . The performance of each DP Myerson is averaged over 50 draws, fitted on 100,000 samples, with empirical average revenue evaluated over another 10,000 samples. Best performing for each ϵ_p are marked **bold**.

In Table 2, we evaluate the performance of DPMYerson in a setting where one bidder’s value distribution follows a normal distribution, and the other follows a lognormal distribution with parameters $(\mu, \sigma) = (-1.87, 1.15)$ (i.e., mean 0.3 and standard deviation 0.5). The revenue achieved by the best DPMYerson configuration significantly exceeds that of the second-price auction.

Bidder 1	Normal $\mathcal{N}(0.3, 0.5)$			
Bidder 2	Normal $\mathcal{N}(0.5, 0.7)$			
Privacy Parameter ϵ_p	0.1	0.2	0.4	0.8
w/ additive $\epsilon_a = 0.1$, Upper Bound $h = 1.5$, w/ 100,000 samples.				
DPMY w/ $\epsilon_q = 0.05$	0.33918	0.32878	0.34927	0.33501
DPMY w/ $\epsilon_q = 0.2$	0.36784	0.36229	0.36408	0.37160
DPMY w/ $\epsilon_q = 0.3$	0.37521	0.37691	0.37436	0.35579
Second Price	0.33741			
Myerson	0.50204			

Table 3: Average Empirical Revenue of DP Myerson under non-i.i.d. Value Distributions with Varying Discretization Parameters (additive ϵ_a , quantile ϵ_q) and Privacy Parameter ϵ_p . The performance of each DP Myerson is averaged over 50 draws, with empirical average revenue evaluated over another 10,000 samples. Best performing for each ϵ_p are marked **bold**.

In Table 3, we evaluate the performance of DPMYerson under non-i.i.d. normal bid distributions. The best DPMYerson configuration achieves a significant revenue improvement over the second-price auction, with an increase of 12%.

Bidder 1	Lognormal $(\mu, \sigma) = (-1.8685, 1.1528)$			
Bidder 2	Lognormal $(\mu, \sigma) = (-1.2357, 1.0417)$			
Privacy Parameter ϵ_p	0.1	0.2	0.4	0.8
w/ additive $\epsilon_a = 0.1$, Upper Bound $h = 1.0$.				
DPMY w/ $\epsilon_q = 0.1$	0.13536	0.13156	0.11881	0.11976
DPMY w/ $\epsilon_q = 0.2$	0.13912	0.13448	0.12737	0.13947
DPMY w/ $\epsilon_q = 0.3$	0.035531	0.03761	0.03568	0.03952
Second Price	0.11578			
Myerson	0.21292			

Table 4: Average Empirical Revenue of DP Myerson under non-i.i.d. Value Distributions with Varying Discretization Parameters (additive ϵ_a , quantile ϵ_q) and Privacy Parameter ϵ_p . The performance of each DP Myerson is averaged over 50 draws, fitted on 100,000 samples, with empirical average revenue evaluated over another 10,000 samples. Best performing for each ϵ_p are marked **bold**.

In Table 4, we evaluate the performance of DPMYerson under non-i.i.d. lognormal bid distributions. The first bidder’s value distribution follows a lognormal distribution with parameters $(\mu, \sigma) = (-1.8685, 1.1528)$ (mean 0.3, std 0.5), while the second bidder’s value distribution follows a lognormal distribution with parameters $(\mu, \sigma) = (-1.2357, 1.0417)$ (mean 0.5, std 0.7).

Remark A.2. The tables above compare the revenue of all mechanisms on distributions discretized and truncated using the same hyperparameters (i.e., ϵ_a and h) as DP Myerson. As a result, the empirical revenue of the baseline mechanisms varies with these hyperparameters.

Remark A.3. The revenue loss incurred in the above tables are expected and consistent with our theoretical analysis. More specifically, the additional revenue loss compared to the optimal Myerson mechanism on the given distribution consists of two components: 1) Quantile Discretization Cost, and 2) Private Quantile Estimation Cost. For the tables above, the revenue loss of DP Myerson under the best hyperparameters is upper-bounded by $h\epsilon_q$, which is the upper bound of the former component.




Figure 3: The impact of different privacy parameter ϵ_p w/ other discretization parameters fixed.

In Figure 3, we illustrate the revenue performance of DP Myerson across different values of the privacy parameter ϵ_p . As shown, ϵ_p acts as a hyperparameter, and the revenue does not vary monotonically with changes in ϵ_p . Interestingly, while one might intuitively expect higher values of ϵ_p (indicating weaker privacy constraints) to result in consistently improved revenue, this is not observed uniformly. Instead, the relationship between ϵ_p and revenue appears non-linear, suggesting a complex interplay between privacy guarantees and auction performance.

Notably, the performance of DP Myerson demonstrates significant robustness to the choice of ϵ_p under *large* sample sizes. This robustness implies that careful tuning of ϵ_p may not always be critical for achieving competitive revenue, particularly in data-rich settings. However, for smaller sample sizes, the choice of ϵ_p could have a more pronounced impact, potentially requiring more nuanced calibration to balance privacy and revenue effectively.

B REMARKS

B.1 GENERALIATION TO JOINT DIFFERENTIAL PRIVACY

A related but weaker privacy notion for *multi-player* setting, i.e, jointly differential privacy (JDP) (Kearns et al., 2014) also applies to our setting. Standard Differential Privacy (DP) requires that changing one entry in the dataset affects the probability of every possible output *vector* by at most ϵ_p . In contrast, joint differential privacy only requires that changes in one player’s input (multi-party collision) do not significantly affect other player’s *scalar* outcome, without imposing restrictions on how those changes impact the outcomes of the players whose inputs were altered. Thus, DP implies JDP, and our algorithm provides a JDP guarantee that is $1/k$ of its DP guarantee.

Notice that JDP doesn’t implies DP in general. JDP guarantees privacy only in an incomplete information setting where each bidder sees only their own outcome, which is often impractical in auction settings due to (1) Sybil attacks, where bidders may create multiple identities/attributes within the same auction, and (2) transparency requirements, such as EU regulations mandating public disclosure of political ad payments and allocations (European Commission, 2024).

B.2 GENERALIZATION OF DP MYERSON

Our DP Myerson for unbounded distributions (Alg. 9) can be generalized to unbounded and irregular distribution settings with *light tail*. The effectiveness of the truncation depends solely on whether the tail of the distribution decays faster than that of the equal revenue distribution, which has support over $[1, +\infty]$ and has a cdf $F(v) = 1 - 1/v$. When the condition is met, the truncated distribution (hence the integrated algorithm), still approximately maintains the revenue guarantee.

B.3 GENERALIZATION TO OTHER ONLINE AUCTION SETTINGS

Varying Bidder Counts. Our mechanism generalizes to settings where, in certain iterations, there are multiple bidders or occurrences for a single attribute or an absence of a single attribute, as long as there are sufficient samples for each attribute over the collection stage.

Generalization to Bounded Rationality. If we adopt the weaker assumption that the bidders has bounded rational, i.e., they will bid truthfully if strategic bidding only gives them at most a small fraction of the extra utility. Then running our DP Myerson *alone* with an appropriate privacy parameter ϵ_p would incentivize truthful bidding. In this context, the commitment mechanism can be replaced by any truthful and prior-independent auctions, for example, second price auction.

C PRIOR WORK DISCUSSION

DP Mechanism Design. Emerging from McSherry and Talwar (2007), there has been interest in delivering mechanisms with DP guarantees Nissim et al. (2012). However, their designs focused more on approximating optimal utility and less on running time efficiency. Consequently, these algorithms incur exponential time in our setting, even when the distribution is of finite support(See Appendix D).

The most relevant recent work is Huang et al. (2018a), which developed an (ϵ, δ) -approximate private empirical reserve mechanism by applying the Gaussian mechanism via two-fold aggregation (Dwork et al., 2010). The added noise follows a mean-0 normal distribution and hence coincides with the smoothed analysis framework, allowing this work to apply solutions from there (Abernethy et al., 2014). However, getting a stronger pure DP mechanism requires the added noise to be *non-normal*. Thus, existing technical solutions from smoothed analysis do not apply to our setting.

Sample Complexity of Auctions One line of related research problem is to show provably sample complexity guarantees for learning in auctions from *truthful* samples, i.e., how many samples are needed to learn auctions that approximately maximize revenue. This problem was first introduced by (Balcan et al., 2005), and led to an explosion of work on the topic (Bubeck et al., 2019; Huang et al., 2018b; Morgenstern and Roughgarden, 2015; Elkind, 2007; Balcan et al., 2005; 2008; Roughgarden and Schrijvers, 2016; Cai and Daskalakis, 2011; Hartline and Roughgarden, 2009; Devanur et al., 2016; Balcan et al., 2016; Cole and Roughgarden, 2014; Guo et al., 2019). Typically, they upper bound the sample complexity of certain auctions by proposing a mechanism that achieves the proposed complexity, and their lower bound for independent, non-i.i.d single-item auctions apply to our setting.

However, it is non-trivial to extend their mechanism to our setting with non-myopic bidders, in that this line of work assumes the learner/ auctioneer has access to *truthful samples*. In contrast, in our setting, if the bidders participate in multiple rounds of the auction, they can bid strategically to maximize their own total utility over all their rounds, and hence the samples are no longer truthful.

Online learning in repeated auction Reserve-price style strategies achieve near-optimal revenue for the i.i.d setting and can be learned within given incentive guarantees (Deng et al., 2020; Kanoria and Nazerzadeh, 2021). However, these methods capture only a constant fraction of the optimal revenue in our setting where the bidders are from *different* distribution.⁶

Huang et al. (2018a); Abernethy et al. (2019) have applied differential privacy as a solution to achieve incentive compatibility. However, their methods rely heavily on the existence of an upper bound of the value distribution and are thus not applicable to the unbounded setting.

Incentive measurements Another relevant topic is to measure and guarantee the incentive compatibility (IC) Myerson (1979) in a mechanism. A mechanism is IC if truthful bidding outperforms other strategies. In the absence of truthful samples from value distribution, strict-IC and the optimality of revenue guarantee cannot be achieved simultaneously.

Hence, previous works have designed several approximate IC metrics and methods to evaluate them from samples (Balcan et al., 2019). This includes approximate Bayesian Incentive Compatibility Bal-seiro et al. (2024), approximate Dominant Strategy Incentive Compatibility Dütting et al. (2024), Stage Incentive Compatibility Deng et al. (2021), and etc.

⁶See example 3.11 in Hartline and Roughgarden (2009).

D FAILED ATTEMPTS

D.1 FAILED ATTEMPTS FOR DP MYERSON

Failed Attempt 1: Deploying the Exponential Algorithm. The exponential mechanism is one typical solution to integrate pure DP with Myerson, but applying it directly to a continuous distribution is computationally inefficient. This is because it requires fitting an (ironed) virtual value curve for each dimension—a continuous function with an unknown exact form, aside from the fact that it is monotonically increasing with the value. A plausible fix is to discretize the distribution into a finite number of values, and then deploy exponential over possible Myerson’s over the discretized distribution, assuming the distribution is bounded. This doesn’t fully resolve the computational challenges. Specifically, when each value distribution is of finite support l and there are k different attributes/distributions, then Myerson’s auction corresponds to ranking all kl values in increasing order of their virtual values. According to this ranking, the mechanism picks the bidder with the highest-ranked value as the winner and charges them the minimum value that maintains a higher rank than the value of the second-highest-ranked bid.

Thus, the exponential mechanism corresponds to sampling an ordering over all $(kl)!/(l!)^k$ possible rankings, with each requiring $O(kn)$ time to evaluate revenue, resulting in *exponential* running time. Suppose we additive discretize each (bounded) distribution by ϵ_a , the number of distinct ordering will be $\Theta(k^{k/\epsilon_a})$, which blow up exponentially with $1/\epsilon_a$. This approach is inefficient since the number of possible rankings grows exponentially with l and does not generalize to unbounded distributions.

Failed Attempt 2: Pre-processing Via Tree Aggregation. One might wonder whether we could use tree aggregation with differentially private (DP) noise (e.g., Laplacian noise) on its cumulative density function after discretizing the value space into intervals when value distribution is bounded. While this method could maintain DP guarantee, the noise added only maintains “close” approximations to the zeroth order information of the revenue curve, and could lead to *negative probability mass* on certain intervals. This solution indeed is feasible for mechanisms that only require the accuracy of the zeroth order information of the revenue curve, e.g., empirical reserve. However, for the Myerson auction in the non i.i.d case, approximating this mechanism requires both the zeroth order and the first order information of the revenue curve, hence tree aggregation is not feasible.

Failed Attempt 3: Postprocessing. Another common solution to differentially private release mechanisms is to add noise to the output. Although this method witnessed its success in robust and differentially private mean estimation, it’s unclear how to handle noise in the mechanism design setting even when the value distribution is bounded by, say H . The reason is due to sensitivity: For mean estimation, the sensitivity (hence the necessary level of added noise) grows smaller with a larger number of samples. On the contrary, for the mechanism design setting, it’s not clear how the efficiency guarantee (e.g. revenue, social welfare) scales with the inverse of the number of bidders.⁷ Unfortunately, if we add DP noise according to sensitivity H , the noise level is too large and fails to guarantee a near-optimal target.

D.2 FAILED ATTEMPTS AND LOWER BOUNDS FOR ONLINE MECHANISM DESIGN

If the bidder is non-discounting and participates in every round of the auction, then it is known that it is not possible to obtain sublinear regret against bid history.

Lemma D.1 (Regret Lower Bound for Additive Bidder (Theorem 3, Amin et al. (2013))). *Let \mathcal{A} be any seller algorithm for the repeated setting; then, there exists a valuation \mathcal{D} such that:*

$$\text{Regret} \geq 1/12.$$

For application to our online setting, bounding the regret against *bid* history is not enough to guarantee a near-optimal revenue for the value distribution.

Failed Attempt: Bounding the Regret Only Another way of thinking about the repeated auction problem is to reduce it to the online setting, and the benchmark is the revenue produced by best-fixed mechanism over the *bid history*. At every iteration, however, the adversary would produce bids as the

⁷In fact, this efficiency guarantee in the worst case will be exactly H .

1026 best (or better) response to the mechanism, hence this bid could deviate *a lot* from the value sequence.
1027 Thus, the revenue from the benchmark could be arbitrarily worse than the revenue from the best-fixed
1028 mechanism over the *value history*.
1029

1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

E MORE DETAILS FROM PRELIMINARIES

Notation. For a mechanism M and a k dimensional product distribution $\mathbf{D} = D_1 \times \dots \times D_k$, denote $\text{Rev}(M, \mathbf{D})$ as the expected revenue by running auction M on \mathbf{D} . Let $M_{\mathbf{D}}$ be the revenue-maximizing auction on \mathbf{D} and its expected revenue be $\text{OPT}_{\mathbf{D}} = \text{Rev}(M_{\mathbf{D}}, \mathbf{D})$. Denote $x_j(v) \in [0, 1]$, $p_j(v)$ as the allocation probability and the payment for the bidder with value v and distribution \mathcal{D}_j . We overload j to denote both the index of bidder and the index of distribution, and we denote $\mathbf{x} = (x_1, \dots, x_k)$ and $\mathbf{p} = (p_1, \dots, p_k)$ as the allocation vector and payment vector, respectively. We use $\text{Rev}(M, D)$ to denote the expected revenue for mechanism M on distribution D , where $M = (\mathbf{x}, \mathbf{p})$ denotes the allocation and payment as a function of bids.

For bidder j at round $t \in [T]$, denote the bidder value as v_j^t , bid as b_j^t , the allocation rule as x_j^t , and the set of indices of rounds the bidder participates in the mechanism as S_j . WLOG, for $t \notin S_j$, we let $b_j^t = v_j^t = 0$. We denote $v_j^{[T]} = (v_j^1, \dots, v_j^T)^\top$ and $b_j^{[T]} = (b_j^1, \dots, b_j^T)^\top$ as the batched value and bid vectors, respectively. Assume the utility of bidder j at time t with bid b_j^t is $u_j^t(b_j^t) = x_j^t(b_j^t) \cdot (v_j^t - p_j^t)$ and over all T rounds is $u_j^{[T]}(b_j^{[T]}, v_j^{[T]}, h_j^t) = \sum_{t \in S_j} x_j^t(b_j^t) \cdot (v_j^t - p_j^t)$. Denote $U_j^t(b_j^{[T]}) = \mathbb{E}[u_j^{[T]}(b_j^{[T]}, v_j^{[T]}, h_j^t) - u_j^{[t]}(b_j^{[T]}, v_j^{[T]}, h_j^t)]$ as the expected utility of bidder j from the t -th round to final round T if the bidder's bid vector is $b_j^{[T]}$, where h_j^t is the history at the t -th round, including the price history up to the t -th round.

E.1 MECHANISM DESIGN BASICS

In this section, we present a detailed definitions on the machineries we use in this paper. We begin with the formal definition of Myerson's auction, which maximizes revenue in Bayesian environments.

Definition E.1 (Myerson's auction, formal version of Definition 2.2). Myerson (1981) Myerson's auction maximizes the expected revenue of a single-item single round auction on product distribution $\mathbf{D} = D_1 \times \dots \times D_k$. Consider the single round auction where there k bidders, where bidder i is from distribution \mathcal{D}_i , and let F_i and f_i denote the cdf and pdf of her value distribution.

For continuous product distribution \mathbf{D} , the virtual value $\phi_i(v_i)$ of the bidder i with value v_i is $\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$. For the case where the product distribution \mathbf{D} is discrete, the virtual value function from distribution D_j at value v_i^j with support $\mathcal{V}_j = \{v_1^j, \dots, v_n^j\}$, is defined as (Elkind (2007)):

$$\phi_j(v_i^j, v_{i+1}^j) = v_i^j - (v_{i+1}^j - v_i^j) \frac{1 - F_j(v_i^j)}{f_j(v_i^j)}$$

where v_i^j s are ordered in increasing order of i , and $f_j(v_i^j) = \mathbb{P}[v^j = v_i^j]$, and $F_j(v_i^j) = \sum_{k=1}^i f(v_k^j)$.

We say a distribution \mathcal{D}_j is η -strongly regular if for every distribution j , $\phi_j(v_i) - \phi_j(v_j) \geq \eta(v_i - v_j)$, for every $v_i > v_j \in \mathcal{V}$. When $\eta = 1$, we say the distribution is monotone hazard rate (MHR) distribution. For any distribution \mathcal{D} with the above property, Myerson's allocation rule is

$$x_i(v_i) = \begin{cases} 1 & \text{if } \phi_i(v_i) \geq \max(0, \max_{j \neq i} \phi_j(v_j)) \\ 0 & \text{otherwise} \end{cases}$$

and payment⁸ rule is

$$p_i(v_i) = \begin{cases} \phi_i^{-1}(\max(0, \max_{j \neq i} \phi_j(v_j))) & \text{if } \phi_i(v_i) \geq \max(0, \max_{j \neq i} \phi_j(v_j)) \\ 0 & \text{otherwise} \end{cases}$$

⁸The virtual value inverse $\phi_i^{-1}(v)$ for *discrete* distribution is defined as $\arg \min_{v \in \mathcal{V}} \phi_i(v) \geq v$, where \mathcal{V} is the support for distribution \mathcal{D}_i .

For regular distributions, Myerson’s auction allocates the good to the bidder with highest non-negative virtual value⁹, and the winner pays the threshold value¹⁰. If there are no bidders with non-negative virtual value, no one wins the item.

For irregular distributions, Myerson’s auction requires an extra “ironing” procedure. The “ironed” virtual value $\tilde{\phi}$ is a monotonic increasing function of the value, and the above payment/allocation rule are defined based on the “ironed” virtual value.

Next, we introduce the definition of Vickrey auction.

Definition E.2 (Vickrey Auction). For a single item auction with multiple bidders, the Vickrey auction allocates the item to the highest bidder and charges them the second highest bid.

When all bidders’ values are i.i.d distributed, the Vickrey auction with Myerson Reserve ($r = \phi^{-1}(0)$) gets optimal revenue in expectation Hartline and Roughgarden (2009).

E.2 DIFFERENTIAL PRIVACY BASICS

We present the definition of pure DP and approximate DP below.

Definition E.3 (Differential privacy). An algorithm $\mathcal{A} : \mathbb{R}_+^n \rightarrow \mathbb{R}$ is (ϵ, δ) -approximate DP if for neighboring dataset $V, V' \in \mathbb{R}_+^n$ that differs in only one data point, and any possible output O , we have: $\Pr[\mathcal{A}(V) = O] \leq \exp(\epsilon) \Pr[\mathcal{A}(V') = O] + \delta$. We say it satisfies *pure* DP for $\delta = 0$.

A key property we leverage from differential privacy is its immunity to post-processing. Post-processing refers to any computation or transformation applied to the output of a DP algorithm after the data has been privatized. In our context, Myerson’s auction can be seen as a post-processing step. Therefore, applying Myerson’s auction to a differentially private release of the empirical distribution preserves the original privacy guarantees of the input distribution.

Lemma E.4 (Immunity to Post-Processing). Let $\mathcal{A} : \mathbb{R}_+^n \rightarrow \mathbb{R}$ be an (ϵ, δ) -DP algorithm, and let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a random function. Then, $f \circ \mathcal{A} : \mathbb{R}_+^n \rightarrow \mathbb{R}$ is also (ϵ, δ) -DP.

Lemma E.5 (Basic Composition Theorem (Dwork et al., 2006)). Let $\mathcal{M}_1 : \mathcal{D} \rightarrow \mathcal{R}_1$ and $\mathcal{M}_2 : \mathcal{D} \rightarrow \mathcal{R}_2$ be two mechanisms that are (ϵ_1, δ_1) -differentially private and (ϵ_2, δ_2) -differentially private, respectively. Then, the composition of \mathcal{M}_1 and \mathcal{M}_2 , denoted as $(\mathcal{M}_1, \mathcal{M}_2) : \mathcal{D} \rightarrow (\mathcal{R}_1 \times \mathcal{R}_2)$, satisfies $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$ -differential privacy.

E.3 PROBABILITY INEQUALITIES

Next, we present the Dvoretzky–Kiefer–Wolfowitz(DKW) inequality that, upper bound the probability that the empirical CDF differs from the CDF of the true distribution.

Lemma E.6 (DKW Inequality (Dvoretzky et al. (1956))). Let X_1, X_2, \dots, X_m be i.i.d random variables with cumulative distribution function $F(\cdot)$. Let F_m denote the associated empirical distribution function defined as $F_m(x) = \frac{1}{m} \sum_{i \in [m]} \mathbb{1}_{\{X_i \leq x\}}$. Then, we have the following probability bound:

$$\Pr[\max_{x \in \mathbb{R}} |F_m(x) - F(x)| > \epsilon] \leq 2 \exp(-2m\epsilon^2)$$

E.4 STOCHASTIC DOMINANCE BASICS

Next, we introduce some technical preliminaries on the distance and dominance between distributions.

Definition E.7 (First Order Stochastic Dominance). For distribution \mathcal{D} and \mathcal{D}' , we denote the cdf of them as $F_{\mathcal{D}}, F_{\mathcal{D}'}$, respectively. Distribution \mathcal{D} first order stochastically dominates distribution \mathcal{D}' if:

- For any outcome x , $F_{\mathcal{D}(x)} \leq F_{\mathcal{D}'}(x)$.
- For some x , $F_{\mathcal{D}(x)} < F_{\mathcal{D}'}(x)$

⁹If there are multiple bidders with highest virtual value, break ties arbitrarily, e.g., in lexicographical order

¹⁰the max of the value at which her virtual value is zero and the value at which her virtual value becomes largest

We denote $\mathcal{D} \succeq \mathcal{D}'$ for \mathcal{D} first order stochastically dominates \mathcal{D}' . For product distribution \mathbf{D} and \mathbf{D}' , if for every i , $\mathcal{D}_i \succeq \mathcal{D}'_i$, we say that $\mathbf{D} \succeq \mathbf{D}'$.

Definition E.8 (Kolmogorov-Smirnov distance). For probability distributions $\mathcal{D}_1, \mathcal{D}_2$ on \mathbb{R} , and let F_1, F_2 denote the cumulative function of $\mathcal{D}_1, \mathcal{D}_2$. Then, the Kolmogorov-Smirnov distance of \mathcal{D}_1 and \mathcal{D}_2 is defined as follows:

$$d_{ks}(\mathcal{D}_1, \mathcal{D}_2) = \sup_{x \in \mathbb{R}} |F_1(x) - F_2(x)|.$$

Moreover, we call \mathcal{D}_1 and \mathcal{D}_2 t -close if $d_{ks}(\mathcal{D}_1, \mathcal{D}_2) \leq t$.

F RESULTS FROM BAYESIAN MECHANISM DESIGN

F.1 REVENUE LOSS

We first state a lemma that guarantee the expected revenue loss by additive discretization by ϵ is upper bounded by ϵ :

Lemma F.1 (Additive Discretization of Value Space (Lemma 6.3 in arXiv version of Devanur et al. (2016))). *Given any product distribution \mathbf{D} and \mathbf{D}' , where \mathbf{D}' is obtained by rounding down the values from \mathbf{D} to the closest multiples of ϵ , we have:*

$$\text{OPT}(\mathbf{D}') \geq \text{OPT}(\mathbf{D}) - \epsilon$$

Lemma F.2 (Weak Revenue Monotonicity (Devanur et al., 2016)). *Suppose \mathbf{D}, \mathbf{D}' be two product distribution such that \mathbf{D}' is first order stochastic dominated by \mathbf{D} , then the optimal revenue for these distributions satisfies the following:*

$$\text{OPT}(\mathbf{D}) \geq \text{OPT}(\mathbf{D}')$$

Lemma F.3 (Strong Revenue Monotonicity (Devanur et al., 2016)). *Let \mathbf{D}' be a product distribution with finite support. There exists a mechanism M_0 such that M_0 is an optimal auction for \mathbf{D} , and for all finite support distributions $\mathbf{D} \succeq \mathbf{D}'$:*

$$\text{Rev}(M_0, \mathbf{D}) \leq \text{Rev}(M_0, \mathbf{D}')$$

G REVENUE SHIFT THEOREM

In this subsection, we introduce technicals to show our revenue shift theorem. We present the definition of an increasing function w.r.t vector input below.

Definition G.1 (Increasing Functions). Let $u : \mathbb{R}^n \rightarrow \mathbb{R}$, we say that u is increasing if for every $\mathbf{v} = (v_1, \dots, v_k)$, $\mathbf{v}' = (v'_1, \dots, v'_k)$ such that $v'_i \geq v_i$, it holds that $u(\mathbf{v}') \geq u(\mathbf{v})$.

We now extend this definition to the scenario where the input vectors follow *distributions*. In this case, the difference in the expected output on these vectors can be bounded by a function of the statistical distance between their underlying distributions.

Lemma G.2 (Utility Difference for Bounded Distribution (Guo et al., 2021)). *Let $\mathbf{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_k$, $\mathbf{D}' = \mathcal{D}'_1 \times \dots \times \mathcal{D}'_k$ be product k -dimensional distributions with $d_{ks}(\mathcal{D}_i, \mathcal{D}'_i) \leq \alpha_i$. Then for every increasing function $u : \mathbb{R}^k \rightarrow [0, \bar{u}]$, it holds that:*

$$|\mathbb{E}_{\mathbf{v} \sim \mathbf{D}}[u(\mathbf{v})] - \mathbb{E}_{\mathbf{v}' \sim \mathbf{D}'}[u(\mathbf{v}')]| \leq \bar{u} \cdot \left(\sum_{j=1}^n \alpha_j \right)$$

Our proof of the revenue shift theorem relies on the property that the optimal revenue (as characterized by Myerson in our setting) equals the maximum payment achievable from a given value profile and is an increasing function of the observed bids. The formal proof is provided below:

Theorem G.3 (Revenue Shift). *Given two product distribution $\mathbf{D} \succeq \mathbf{D}'$ whose valuations are bounded by $[0, h]$, with $d_{ks}(\mathbf{D}_i, \mathbf{D}'_i) \leq \alpha_i$ for any bidder/entry i , the optimal revenue of these distribution satisfies:*

$$0 \leq \mathbb{E}[\text{Rev}(M_{\mathbf{D}}, \mathbf{D}) - \text{Rev}(M_{\mathbf{D}'}, \mathbf{D}')] \leq \left(\sum_{i \in [k]} \alpha_i \right) h$$

Proof. According to weak revenue monotonicity F.2, we get that the optimal revenue of a distribution \mathcal{D} over support $[0, h]$ is an increasing function defined in Def. G.1, and the optimal revenue is upper bounded by h . Hence, by lemma G.2, we get that for any distribution $\mathbf{D} \succeq \mathbf{D}'$.

$$0 \leq \mathbb{E}[\text{Rev}(M_{\mathbf{D}}, \mathbf{D}) - \text{Rev}(M_{\mathbf{D}'}, \mathbf{D}')] \leq \left(\sum_{i \in [k]} \alpha_i \right) h$$

□

H MORE DETAILS FOR BOUNDED DISTRIBUTIONS

H.1 PRIVATE QUANTILE ESTIMATION

It’s worthwhile to present and state the quantile estimation oracle below. In our paper, we apply a similar quantile estimation algorithm as in Kaplan et al. (2022), which applies the exponential mechanism (Smith, 2011) efficiently on the dataset/samples.

Algorithm 4 DP Quantile, Bounded Distribution DPQUANT (Kaplan et al., 2022)

Input: n samples $V = \{v_1, \dots, v_n\}$, range $[lb, ub]$, set of quantiles $Q := \{q_1, \dots, q_m\}$, privacy parameter ϵ_p , DP oracle \mathcal{A} that estimate a single quantile with privacy $\epsilon_p/(\log_2 m + 1)$ (Alg. 5).

- 1: Rank the quantiles Q in increasing order.
- 2: **if** $m == 1$ **then return** $\{\mathcal{A}((lb, ub), V, q_1)\}$
- 3: **end if**
- 4: $s \leftarrow \mathcal{A}((lb, ub), V, q_{\lfloor m/2 \rfloor})$.
- 5: Separate the samples V by s into left and right quantiles, i.e., $V_l := \{v < s | v \in V\}$, $V_r := \{v > s | v \in V\}$.
- 6: Update the candidate quantile into Q_l and Q_r , where $Q_l := \{q < q_{\lfloor m/2 \rfloor} | q \in Q\}$, $Q_r = \{q > q_{\lfloor m/2 \rfloor} | q \in Q\}$
- 7: **return** $\{\text{DPQUANT}(V_l, (lb, s), Q_l, \epsilon_p, \mathcal{A})\} \cup \{s\} \cup \{\text{DPQUANT}(V_r, (s, ub), Q_r, \epsilon_p, \mathcal{A})\}$

We define the utility function measuring the accuracy of a given quantile estimation I of quantile q . This function measures the number of points between the true value of a given quantile q and its estimation I .

Definition H.1 (Utility Function of Quantile Estimation (Smith, 2011)). Given a dataset $V \in [lb, ub]^n$, a quantile q , and an estimation I , the utility function of quantile estimation is defined as:

$$u(V, I, q) := -|\{v \in V | v < I\} - \lfloor q \cdot n \rfloor|$$

For multiple quantiles $Q = (q_1, \dots, q_m)$ and estimation $I = (I_1, \dots, I_m)$, the utility is defined as the worst utility of these quantile estimations, i.e.,

$$u(V, I, Q) := \min_{r \in [m]} u(V, I_r, q_r)$$

Here we present the DP Single Quantile estimation below. WE apply this algorithm in thhe DP Quantile with privacy $\epsilon_p/(\log_2 m + 1)$, for number of quantiles m and required (pure) privacy ϵ_p . The DP Single Quantile guarantees pure privacy by efficiently implementing the exponential algorithm with the utility we previously described. For DP Single Quantile, we use the convention that for $v_k = v_{k-1}$, the sample probability is 0, under this convention, only interval with *positive* length will have the probability to be sampled.

Algorithm 5 DP Single Quantile

Input: n samples $V = \{v_1, \dots, v_n\}$, range $[lb, ub]$, quantile q , privacy parameter ϵ_p

- 1: Rank samples in V in increasing order.
- 2: **for** $i = 1 \rightarrow n + 1$ **do**
- 3: Let interval $I_i = [v_{i-1}, v_i]$, where $x_0 = lb$, $x_{n+1} = ub$.
- 4: **end for**
- 5: Sample an interval I_k from this set of intervals, with probability $\exp(\epsilon_p u(V, I_k, q)/2) \cdot (v_k - v_{k-1})$.
- 6: **return** a uniformly random point from interval I_k .

Now we present the utility for single quantile for duplicates value below. This proof is similar to the one adopted in the Appendix A.1 of Kaplan et al. (2022).

Lemma H.2 (Utility for Single Quantile, Discrete Distribution). *With probability $1 - \delta$, given samples $V \in [0, h]^n$ and quantile $q \in [0, 1]$, where samples may have the same value, but those with different*

values will differ by at least ϵ_a . Then, the exponential mechanism described in Alg. 5 would output s with $(\epsilon_p, 0)$ -DP and:

$$|u(V, I, q)| \leq 2 \cdot \frac{\log \phi - \log \delta}{\epsilon_p}.$$

where $\phi := h/\epsilon_a$.

Proof. Let I_t be an interval that $u(V, I_t, q) \leq -\gamma$. Then the probability that we sample a point from I_t is at most:

$$\begin{aligned} \Pr[\mathcal{A}(V) = I_t] &\leq \frac{\exp(-\epsilon_p \gamma / 2)(v_i - v_{i-1})}{\sum_{i \in [n], x_i \neq x_{i-1}} \exp(\epsilon_p u(V, I_i, q)) \epsilon_a} \\ &\leq \frac{\exp(-\epsilon_p \gamma / 2) h}{\exp(\epsilon_p u(V, I_o, q)) \epsilon_a} \\ &\leq \phi \exp(-\epsilon_p \gamma / 2) \end{aligned}$$

where \mathcal{A} denotes the output by DP Single Quantile, and o denote the optimal interval with for quantile q that has zero utility. Next, it follows that with probability less than δ , the returned interval will have a utility at most $-\gamma$ for $\gamma = \Theta(\log \phi + \log(1/\delta))/\epsilon_p$, which completes the proof. \square

We have proven that the single quantile algorithm still holds similar accuracy guarantee under our assumption on the dataset which the data points can have same value, but if they are different, then their values will differ by at least ϵ_a . Following the same logic as in Kaplan et al. (2022), we can also demonstrate the accuracy guarantee of the quantile estimation algorithm in our distribution setting.

Lemma H.3 (Utility of DP Quantile (Thm. 3.3 in Kaplan et al. (2022))). *With probability $1 - \delta$, given samples $V \in [0, h]^n$ and quantile $Q = (q_1, \dots, q_m)$ and privacy parameter ϵ_p , where samples may have the same value, but those with different values will differ by at least ϵ_a . Then, the Alg 4 will output $S = (s_1, \dots, s_m)$ with $(\epsilon_p, 0)$ -DP, such that:*

$$\text{ERR}(V, S, Q) := -u(V, S, Q) \leq 2(\log m + 1) \cdot \frac{\log \phi + \log m - \log \delta}{\epsilon_p}$$

where $\phi := h/\epsilon_a$.

H.2 REVENUE GUARANTEE OF PRIVATE MYERSON

We now present the complete proof of the accuracy guarantee for the private Myerson mechanism under the bounded distribution setting.

Theorem H.4 (Revenue Guarantee of Private Myerson (Alg. 1), formal version of Theorem 3.2). *Given n samples $\hat{V} \in [0, h]^{k \times n}$ of the joint distribution \mathbf{D} , DPMYER (Alg. 1) is $(2k\epsilon_p, 0)$ -DP, and the expected revenue of this mechanism is close to the optimal revenue of distribution \mathbf{D} , i.e., with probability $1 - \delta$:*

$$|\mathbb{E}[\text{Rev}(M_{\text{DPMYER}}, \mathbf{D}) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})]| \leq \tilde{\Theta}\left(\left(\epsilon_q + \frac{1}{\epsilon_p \cdot n} + \frac{1}{\sqrt{n}}\right) \cdot kh + \epsilon_a\right).$$

for $n \geq \max\{2 \log(4k/\delta)/\epsilon_q^2, 8(\log(1/\epsilon_q) + 1)(\log(hk \log(1/\epsilon_q)/(\epsilon_a \delta)))/(\epsilon_q \epsilon_a)\}$, which can be further simplifies to $n = \tilde{\Omega}(\max\{1/(\epsilon_p \epsilon_q), 1/\epsilon_q^2\})$. Furthermore, under $\epsilon_a = \epsilon_q = \epsilon$, and that $n = \Theta(\max\{\epsilon^{-2} \log(k/\delta), \epsilon^{-2} \log(1/\epsilon) \log(\frac{hk \log(2/\epsilon)}{\epsilon \delta}), \epsilon^{-2} \log(k/\delta)\})$, which can be further simplifies to $n = \tilde{\Theta}(\epsilon^{-2})$, we have:

$$|\mathbb{E}[\text{Rev}(M_{\text{DPMYER}}, \mathbf{D}) - \text{OPT}(\mathbf{D})]| \leq \tilde{\Theta}((\epsilon + \epsilon^2/\epsilon_p)kh).$$

Proof. Privacy We know that the quantile estimates from DPQE is $(\epsilon_p, 0)$ private (Lem. H.2). Since DP is immune to post-processing (Lem. E.4), and that the output of allocation and payment combination is $2k$ dimensional, by composition theorem (Lem. E.5), our algorithm is $(2k\epsilon_p, 0)$ -DP.

We include all distributions considered in this proof in Figure 1 below.




Figure 4: **Distribution analyzed for DPMYER(Alg. 1).** We establish connections between the accuracy/revenue guarantee of the original distribution \mathbf{D} with the empirical distribution \hat{D} , the value-discretized \hat{D}^ϵ , the quantile-discretized \hat{D}^q and the distribution \hat{D}^p returned by DPQUANT(Alg. 4).

From the DKW inequality E.6, we know that with probability $1 - \delta_1$, the cumulative density function of the empirical distribution is close to the true distribution, i.e., for all attribute $i \in [k]$,

$$d_{\text{ks}}(\mathbf{D}_i, \hat{D}_i) := \max_v |(F_{\mathbf{D}}(v) - F_{\hat{D}}(v))| \leq \sqrt{\log(2k/\delta_1)/2n}.$$

We condition on this event holds since after discretization by ϵ_a , each value at most decreases by ϵ_a in the new distribution \hat{D}^ϵ . By Lemma F.1, we note that the optimal revenue at most decreases by ϵ_a :

$$\mathbb{E}[\text{Rev}(M_{\hat{D}}, \hat{D}) - \text{Rev}(M_{\hat{D}^\epsilon}, \hat{D}^\epsilon)] \leq \epsilon_a.$$

Next, we discretize again on distribution \hat{D}^ϵ , which additively discretizes this distribution in the quantile space. We denote this distribution as \hat{D}^q . Consequently, for any attribute $i \in [k]$, we have:

$$d_{\text{ks}}(\hat{D}^\epsilon, \hat{D}^q) \leq \epsilon_q.$$

Next, from Lemma H.3, we get that with probability $1 - \delta_2$, for each attribute $i \in [k]$, the following holds:

$$\begin{aligned} d_{\text{ks}}(\hat{D}_i^q, \hat{D}_i^p) &= |-\text{ERR}(V, S, Q)|/n \\ &\leq 2(\log m + 1) \cdot \frac{\log h - \log \epsilon_a + \log m + \log k - \log \delta_2}{\epsilon_p \cdot n} := \hat{\epsilon}. \end{aligned}$$

Next, we show by $n \geq \max\{2 \log(4k/\delta)/\epsilon_q^2, 4(\log m + 1)(\log(hmk/(2\epsilon_a\delta)))/(\epsilon_q\epsilon_q)\}^{11}$, we have $\mathbf{D} \succeq \hat{D}^q$: 1) For $n \geq 2 \log(4k/\delta)/\epsilon_q^2$, we have that $d_{\text{ks}}(\mathbf{D}_i, \hat{D}_i) \leq \epsilon_q/2$. 2) For $n \geq 4(\log m + 1)(\log(hmk/(2\epsilon_a\delta)))/(\epsilon_q\epsilon_q)$, we have that $\hat{\epsilon} \leq \epsilon_q/2$. Since quantile discretization shift the distribution down by $[\epsilon_q, 2\epsilon_q]$, and that additive discretization only shift the distribution to a one that is dominated by it, we have that \hat{D}^q is still dominated by \mathbf{D} .

We condition on both events holding and denote the KS-distance upper bound between the privacy estimation vs the ground truth of the discretized distribution of one attribute as $\hat{\epsilon}$. Thus, we get that with probability $1 - \delta_1 - \delta_2$, we have the following revenue bound of the final mechanism:

$$\begin{aligned} 0 &\geq \mathbb{E}[\text{Rev}(M_{\hat{D}^p}, \mathbf{D}) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})] \\ &\geq \mathbb{E}[\text{Rev}(M_{\hat{D}^p}, \mathbf{D}) - \text{Rev}(M_{\hat{D}^p}, \hat{D}^p)] - |\text{Rev}(M_{\hat{D}^p}, \hat{D}^p) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})| \\ &\geq -|\text{Rev}(M_{\hat{D}^p}, \hat{D}^p) - \text{Rev}(M_{\hat{D}^q}, \hat{D}^q)| - |\text{Rev}(M_{\hat{D}^q}, \hat{D}^q) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})| \\ &\geq -k\hat{\epsilon}h - |\text{Rev}(M_{\hat{D}^q}, \hat{D}^q) - \text{Rev}(M_{\hat{D}^\epsilon}, \hat{D}^\epsilon)| - |\text{Rev}(M_{\hat{D}^\epsilon}, \hat{D}^\epsilon) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})| \\ &\geq -k(\hat{\epsilon} + \epsilon_q)h - |\text{Rev}(M_{\hat{D}^\epsilon}, \hat{D}^\epsilon) - \text{Rev}(M_{\hat{D}}, \hat{D})| - |\text{Rev}(M_{\hat{D}}, \hat{D}) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})| \\ &\geq -k(\hat{\epsilon} + \epsilon_q)h - \epsilon_a - \sqrt{\log(2k/\delta_1)/2n}kh, \end{aligned}$$

where the first inequality follows from the optimality of mechanism $M_{\mathbf{D}}$ on distribution \mathbf{D} , and that $\mathbf{D} \succeq \hat{D}^q$ by our choice of n . The second inequality follows from rearranging the term. The

¹¹The quantity in the theorem statement is $8(\log(1/\epsilon_q) + 1)(\log(hk \log(1/\epsilon_q)/(\epsilon_a\delta)))/(\epsilon_q\epsilon_q)$ and is greater than the second term in the max here.

third inequality follows from strong revenue monotonicity F.3 and that $\mathbf{D} \succeq \hat{D}^p$, we get that the term $\mathbb{E}[\text{Rev}(M_{\hat{D}^p}, \mathbf{D}) - \text{Rev}(M_{\hat{D}^p}, \hat{D}^p)] \geq 0$. The next few inequalities follows from applying the revenue shift theorem (Thm G.3) iteratively for: 1) \hat{D}^p and \hat{D}^q with distance $\hat{\epsilon}$, 2) \hat{D}^ϵ and \hat{D}^q with distance ϵ_q , and 3) \hat{D} and \mathbf{D} with distance $\sqrt{\log(2k/\delta_1)/2n}$. We also apply Lemma F.1 to upper bound the revenue loss from additive discretization.

Next, we plug in the value of $\hat{\epsilon}$ and $m = \lfloor 1 + 1/\epsilon_q \rfloor$ to upper bound the value of $\hat{\epsilon}$, i.e.,

$$\hat{\epsilon} \leq \frac{4}{\epsilon_p \cdot n} \cdot \log\left(\frac{1}{\epsilon_q}\right) \cdot \left(\frac{m h k}{\epsilon_a \cdot \delta_2}\right).$$

Finally, letting $\delta_1 = \delta_2 = \delta/2$, we get the final result: with probability $1 - \delta$:

$$\begin{aligned} & |\mathbb{E}[\text{Rev}(M_{\hat{D}^p}, \mathbf{D}) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})]| \\ & \leq (\epsilon_q + \frac{10}{\epsilon_p \cdot n} \log\left(\frac{1}{\epsilon_q}\right) \cdot \log\left(\frac{h k}{\epsilon_q \epsilon_a \cdot \delta}\right) + \sqrt{\frac{\log(4k/\delta)}{2n}}) \cdot k h + \epsilon_a \\ & \leq \tilde{\Theta}\left((\epsilon_q + \frac{1}{\epsilon_p \cdot n} + \frac{1}{\sqrt{n}}) \cdot k h + \epsilon_a\right), \end{aligned}$$

where the $\tilde{\Theta}$ hide the polylog factors. Furthermore, let $\epsilon_q = \epsilon_a$ and let $n \geq \epsilon^{-2} \log(2k/\delta_1)/2$ give us the following bound:

$$|\mathbb{E}[\text{Rev}(M_{\hat{D}^p}, \mathbf{D}) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})]| \leq \tilde{\Theta}((\epsilon + \epsilon^2/\epsilon_p) k h).$$

Next, we let δ in the statement to be $1/k$ of the δ we used in this proof. This δ would only affect the revenue by poly log factors hence are hidden in the $\tilde{\Theta}$. \square

H.3 RUNNING TIME FOR PRIVATE MYERSON

Theorem H.5 (Running time, DP Myerson for Bounded Distribution). *Given the same parameters as stated in Theorem 3.2, the running time of DPMYER (Alg.1) is $\Theta(\log(1/\epsilon_q)n + kn) = \tilde{\Theta}(kn)$ and requires $\Theta(\log(1/\epsilon_q)) = \tilde{\Theta}(1)$ passes of the samples.*

Proof. We describe below how to implement DPMYER efficiently.

- For the additive discretization step, we rounded down each value to the closest multiples of ϵ_a . This step runs in $O(kn)$ time and requires 1 pass of the dataset.
- For quantile preparation, this step takes $\lfloor 1/\epsilon_q \rfloor + 1$ time.
- For the DP quantile estimation step, we know that it requires $\log(\lfloor 1/\epsilon_q \rfloor + 1)$ passes of the dataset. The running time of each of these pass depends not on n , but on the number of distinct value we have after additive discretization (i.e., $O(h/\epsilon_a)$), and the number of quantiles we want to calculate our utility (Def. H.3) on (i.e., $O(1/\epsilon_q)$). Each pass of the distribution will take $O(k/(\epsilon_a \epsilon_q))$ time. Since $\epsilon_a = \epsilon_q = \epsilon$, we have that the total running time is $\tilde{\Theta}(k\epsilon^{-2}) = \tilde{\Theta}(kn)$.

Summing them up provides the final guarantee. \square

I MORE DETAILS FOR UNBOUNDED DISTRIBUTIONS

I.1 TRUNCATION POINT LEMMA

Here we present a lemma on how to truncate the regular distribution. Notice that this truncation point depends on the optimal revenue produced by the distribution itself. In order to estimate this truncation point up by approximation, an approximation on the revenue is needed.

Lemma I.1 (Truncation of Regular Distribution (Devanur et al., 2016)). *For any product regular distribution $\mathbf{D} = (D_1, \dots, D_k)$, given any $\epsilon \in (0, 1/4]$, let $\bar{v} \geq \frac{1}{\epsilon} \text{OPT}(\mathbf{D})$ be the truncation point, and let $\bar{D}_1, \dots, \bar{D}_k$ be the distribution after truncating \mathbf{D} by point \bar{v} . Then, we have*

$$\text{OPT}(\bar{\mathbf{D}}) \geq (1 - 2\epsilon) \text{OPT}(\mathbf{D}).$$

I.2 THE EMPIRICAL RESERVE ALGORITHM

In this section, we formally introduce the details of *Empirical Reserve* algorithm, and how it approximates the optimal revenue when there is only one bidder. We run a $\delta/2$ -guarded reserve mechanism to collect an estimation on the optimal revenue of product distribution \mathbf{D} . Then, we analyze the approximation guarantee, the incentive robustness and the convergence of this algorithm. In this subsection, the quantile q is defined as the value corresponds to the top q quantile as opposed to in out context, the quantile is defined as the bottom q quantile. First, we describe the β -guarded empirical reserve algorithm (Alg. 6).

Algorithm 6 Empirical Reserve ER (Huang et al., 2018b)

Parameters: distribution \mathcal{D} , failure probability δ , guarded parameter β , accuracy parameter ϵ_{ER} .

Input: $m = \Theta(\beta^{-1} \epsilon^{-2} \log(\beta^{-1} \epsilon^{-1}) \log(1/\delta_{\text{ER}}))$ samples from distribution \mathcal{D}

1: Sort m samples in the decreasing order, i.e., $v_1 \geq v_2 \geq \dots \geq v_m$.

2: Find the smallest index $j \in [\beta \cdot m, m]$ that maximizes the empirical revenue, i.e.,

$$j = \arg \max_{\beta m \leq i \leq m} i \cdot v_i$$

3: $r_{\mathcal{D}} \leftarrow v_j$

\triangleright β -guarded empirical reserve

4: $R_{\mathcal{D}} \leftarrow j \times v_j / m$

\triangleright Empirical revenue

Output: $r_{\mathcal{D}}, R_{\mathcal{D}}$

Definition I.2 (β -guarded reserve). Given m samples $v_1 \geq v_2 \geq \dots \geq v_m$, the *empirical reserve* is

$$\arg \max_{i \geq 1} i \cdot v_i$$

If we only consider $i \geq \beta m$ for some parameter β , it is called the β -guarded empirical reserve.

Lemma I.3 (Empirical Reserve, γ -strongly-Regular, Thm 3.3 in Huang et al. (2018b)). *The empirical reserve with $m = \Theta(\epsilon^{-3/2} \log(\epsilon^{-1}))$ samples is $(1 - \epsilon)$ -approximate for all γ -strongly regular distributions, for a constant $\alpha > 0$.*

Lemma I.4 (Optimal Quantile). *Let q denote the quantile, $v(q)$ denote the value of that quantile, and let $R(q) = qv(q)$ be the revenue as a function over the quantile space. Let q^* and $v^* = v(q^*)$ be the revenue-optimal quantile and reserve price respectively. Then,*

- (Hartline et al. (2008)) For every MHR distribution, $q^* \geq \frac{1}{e}$.

- (Cole and Roughgarden (2014)) For any γ -strongly regular distribution, $q^* \geq \gamma^{\frac{1}{1-\gamma}}$.

Lemma I.5 (Empirical Reserve, Bounded, Thm 3.6 in Huang et al. (2018b)). *The empirical reserve with $m = \Theta(H \epsilon^{-2} \log(H \epsilon^{-1}))$ samples is $(1 - \epsilon)$ -approximate for all distributions with support $[1, H]$.*

Lemma I.6 ($\beta_0/2$ -guarded empirical reserve, Thm. 3.5 in Arxiv Version of Huang et al. (2018b)). *The $\frac{\beta_0}{2}$ -guarded empirical reserve with $m = \Theta(\beta_0^{-1} \epsilon^{-2} \log(\beta_0^{-1} \epsilon^{-1}))$ gives revenue at least $(1 - \epsilon) R_{\beta_0}^*$ for all distributions, where $R_{\beta_0}^*$ is the optimal revenue by prices with sale probability at least β_0 , in expectation.*

Algorithm 7 DP Quantile Estimation for Unbounded Distribution, DPQUANTU(V, Q) (Durfee, 2023)

Lemma I.7 (DP Guarantee of DP Quantile, Unbounded (Durfee, 2023)). *Alg. 7 is ϵ_p -DP.*

We derive the accuracy guarantee of DP Quantile for unbounded distribution below.

- **Unbiased:** $\mathbb{E}[f_{i-1}(V)] \leq qn \leq \mathbb{E}[f_i(V)]$, for i is the iteration of halt.

$$F((\hat{S} + 1)/\beta - 1) - \epsilon_p \cdot \log(1/\delta)/2n \leq F(S) \leq F(\hat{S}) + \epsilon_p \cdot \log(1/\delta)/2n$$

Proof. We prove the unbiasedness and approximation guarantee below:

- **Asymptotic β -approximation:** This follows from the tail bound of exponential distribution:

Reorganizing we have, with probability $1 - \delta$:

Since $\epsilon_i, \epsilon \geq 0$, we have:

I.4 ANALYSIS FOR DPKOPT

In this subsection, we formally present our algorithm for estimate an $\Theta(k)$ -approximation of the optimal revenue. We private estimate the maximum revenue (of a single item single bidder setting) from each bidder's distribution. Aggregating these private estimation of the empirical revenue gives as a $\Theta(k)$ -approximation of the optimal revenue for the product distribution.

Algorithm 8 DP Estimation for Optimal Revenue DPKOPT($V, \epsilon_q, \epsilon_a, \epsilon_p, \eta$)

Input: n samples $V = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$, quantile discretization ϵ_q , additive discretization ϵ_a , privacy parameter ϵ_p , regularity parameter η .

- 1: **for** $d = 1 \rightarrow k$ **do**
- 2: $\hat{q} \leftarrow 1/4 \cdot \eta^{1/(1-\eta)}$
- 3: Let $ub_d \leftarrow \text{DPQUANTU}(V_{[d,:]}, 1 - \hat{q}, \epsilon_p)$. ▷ Estimate the truncation point of D_d .
- 4: Truncate distribution D_d at ub_d as \hat{D}_d , and discretize \hat{D}_d by additive ϵ_a in the value space.
- 5: Prepare the quantile to be estimated, $Q \leftarrow \{1 - \hat{q}, 1 - \hat{q} - \epsilon_q, \dots, 1 - \hat{q} - \lfloor \frac{1-\hat{q}}{\epsilon_q} \rfloor \cdot \epsilon_q, 0\}$.
- 6: $\hat{S}_{[d,:]} \leftarrow \text{QESTIMATE}(Q, V_{[d,:]}, \epsilon_p)$ ▷ Apply DP quantile estimate(Alg. 4).
- 7: Let \hat{F}_d be the distribution generated by value profile $\hat{S}_{[d,:]}$ and quantile set Q .
- 8: $\text{SREV}_d \leftarrow \max_{r \in \hat{S}} r(1 - \hat{F}_d(r))$. ▷ Estimate the optimal revenue from \hat{F}_d (Alg. 6).
- 9: **end for**
- 10: $\text{KREV} \leftarrow \sum_{d \in [k]} \text{SREV}_d$
- 11: **return** KREV

Lemma I.10 (Expected Revenue Guarantee of DPKOPT). *Given $\epsilon \in (1, 1/4)$ and $n = \Theta(\epsilon^{-2} \log(\epsilon^{-1}))$ samples \hat{V} of the joint distribution $\mathbf{D} \in [0, \mathbf{h}]^k$, the expected revenue of Myerson fitted under DPKOPT (Alg. 2) over distribution \mathbf{D} is close to the optimal revenue of distribution \mathbf{D} , i.e., with probability $1 - k\delta$, for every $i \in [k]$, $|E[\text{Rev}(M_{ER(\mathcal{D}_p)}, \mathcal{D}_p) - \text{OPT}(\mathcal{D}_i)]| \leq \tilde{O}(k(\frac{\epsilon^2}{\epsilon_p} + \epsilon + \epsilon_q + \epsilon_a))$.*

Proof. Since bidders' distributions are independent, we analyze the the revenue guarantee of each SREV_i for $i \in [k]$ separately, and we omitted the subscription on i in the proofs.

We denote $q_0 = \eta^{1/(1-\eta)} = 4\hat{q}$ as the (top) optimal quantile for the empirical revenue (Lem. I.4). We denote the output value from line 3 as \hat{v}_{\max} .

From Lemma I.9, with probability $1 - \delta_2$, for $n \geq 4\epsilon_p \log(1/\delta_2)/\hat{q} = \Theta(\epsilon_p \log(1/\delta_2))$ samples:

$$F(\hat{v}_{\max}) \geq 1 - \hat{q} - \epsilon_p \log(1/\delta_2)/2n \geq 1 - \frac{3}{8}q_0.$$

This is saying that with high probability, the returned value is greater than that of the top $3q_0/8$ quantile. Thus, conditioned on this event, the revenue from applying the empirical reserve on the truncated distribution (which equals the optimal revenue from the truncated distribution) is equivalent to that of applying the empirical reserve on the original distribution. The revenue generated by empirical reserve equals the expected optimal revenue from the same distribution. Hence, we concluded that this truncation won't affect the optimal revenue.

Next, we privately estimate the pre-specified quantiles of the truncated distribution, and output the revenue generated from it. We denote the output distribution as \mathcal{D}_p , and the truncated distribution as \mathcal{D}_{TR} . By similar arguments as in Theorem 3.2, we know that, with probability $1 - \delta_2$, for $n \geq \tilde{\Theta}(1/\epsilon_p)$:

$$\begin{aligned} d_{\text{ks}}(\mathcal{D}_p, \mathcal{D}_{\text{TR}}) &\leq 2(\log m + 1) \frac{\log \hat{v}_{\max} - \log \epsilon_a + \log m - \log \delta_2}{\epsilon_p \cdot n} + \epsilon_q \\ &:= \epsilon_{\text{PTR}} = \tilde{\Theta}\left(\frac{1}{\epsilon_p \cdot n}\right) (\leq 1/16q_0) \end{aligned}$$

where $m = \lfloor \frac{1-\hat{q}}{\epsilon_q} \rfloor + 1$. Then, from Thm G.3, the optimal revenue from these distributions differs by at most $\hat{v}_{\max} \epsilon_{\text{PTR}}$. Notice again from Lemma I.9, with probability $1 - \delta_1$, $\hat{v}_{\max} \epsilon_{\text{PTR}}$ also satisfies:

$$F(\hat{v}_{\max}/\beta) \leq 1 - q_0/8$$

This results in $\hat{v}_{\max} \leq \beta \cdot C_0$, where C_0 is the true value of quantile $1 - 1/8q_0$, and can be treated as a constant since q_0 is a constant. Aggregating these together gives us the optimal revenue loss from the second private algorithm is upper bounded by $C_0 \beta \epsilon_{\text{PTR}}$, where β is the parameter used by the DP quantile for unbounded distribution.

The final ingredient is how the empirical reserve algorithm works for distribution \mathcal{D}_p ; from Lemma I.6, we know that when $n = \theta(\epsilon^{-2} \log(\epsilon^{-1}))$, this revenue is at least $(1 - \epsilon) \text{OPT}(\mathcal{D}_p)$, i.e.,

$$|\text{Rev}(M_{\text{ER}}, \mathcal{D}_p) - \text{OPT}(\mathcal{D}_p)| \leq \epsilon \cdot \text{OPT}(\mathcal{D}_p).$$

Hence, assuming $\delta_1 = \delta/4, \delta_2 = \delta/2$ and noticing that \hat{v}_{\max} still exceeds the optimal quantile for distribution \mathcal{D}_p , we can upper bound the expected revenue gap between empirical reserves on the private distribution. The optimal revenue of the original distribution becomes:

$$\begin{aligned} |\mathbb{E}[\text{Rev}(M_{\text{ER}(\mathcal{D}_p)}, \mathcal{D}_p) - \text{OPT}(\mathcal{D})]| &\leq |\text{Rev}(M_{\text{ER}(\mathcal{D}_p)}, \mathcal{D}_p) - \text{OPT}(\mathcal{D}_p)| + |\text{OPT}(\mathcal{D}_p) - \text{OPT}(\mathcal{D})| \\ &\leq \epsilon \cdot \text{OPT}(\mathcal{D}_p) + \epsilon_a + \Theta(\beta \epsilon_{\text{PTR}}) \leq \Theta(\epsilon + \epsilon_a + \epsilon_{\text{PTR}}) \\ &\leq \tilde{O}\left(\frac{\epsilon^2}{\epsilon_p} + \epsilon + \epsilon_q + \epsilon_a\right), \end{aligned}$$

where the second to last inequality following from C is a constant, and the last inequality follows from hide the log factors. Since each of the k distribution would contribute this amount to the revenue loss, our statement as an revenue error bound as $\tilde{O}(k(\frac{\epsilon^2}{\epsilon_p} + \epsilon + \epsilon_q + \epsilon_a))$. \square

Notice that this lemma only guarantees that the expectation of SREV_i is close to the expected optimal revenue of \mathcal{D}_i . We still needs to prove that the SREV_i converges to its expectation quickly, hence is close to the underlying expected optimal revenue.

Theorem I.11 (Accuracy Guarantee of DPKOPT). *Let all parameters be the same as stated in Lemma I.10, we have that: with probability $1 - \delta$,*

$$|\text{KREV} - \sum_{i \in [k]} \mathbb{E}[\text{OPT}(\mathcal{D}_i)]| \leq \tilde{\Theta}(k(\epsilon + \epsilon^2/\epsilon_p + \epsilon_q + \epsilon_a))$$

Proof. From the proofs of previous lemma, we notice that the distribution \hat{S} is upper bounded by \hat{v}_{\max} , hence upper bounded by a βC_0 . We denote $C_1 := \beta C_0$ here, and applies the Chernoff to upper bound how SREV_i might deviates from its expectation. With probability $1 - \delta_3$,

$$|\text{SREV}_i - \mathbb{E}[\text{Rev}(M_{\text{ER}(\mathcal{D}_p)}^i, \mathcal{D}_p^i)]| \leq \Theta(\sqrt{1/n \cdot \log(1/\delta_3)})$$

Plugging in $n = \epsilon^{-2} \log(\epsilon^{-1})$ gives us that, with probability $1 - \delta_3$,

$$\text{SREV}_i - \mathbb{E}[\text{Rev}(M_{\text{ER}(\mathcal{D}_p)}^i, \mathcal{D}_p^i)] \leq \tilde{\Theta}(\epsilon)$$

Thus, we have that, with probability $1 - k\delta_3$, we have:

$$\begin{aligned} \text{KREV} &= \sum_{i \in [k]} \text{SREV}_i \\ &\leq \sum_{i \in [k]} \mathbb{E}[\text{Rev}(M_{\text{ER}(\mathcal{D}_p)}^i, \mathcal{D}_p^i)] + \tilde{\Theta}(k\epsilon) \\ &\leq \sum_{i \in [k]} \mathbb{E}[\text{OPT}(\mathcal{D}_i)] + \tilde{\Theta}(k(\epsilon + \epsilon^2/\epsilon_p + \epsilon_q + \epsilon_a)) \end{aligned}$$

At the same time, $\text{KREV} \geq \sum_{i \in [k]} \mathbb{E}[\text{OPT}(\mathcal{D}_i)] - \tilde{\Theta}(k(\epsilon + \epsilon^2/\epsilon_p + \epsilon_q + \epsilon_a))$. Now, let $\delta_3 = \delta$, thus we have that with probability $1 - 2k\delta$,

$$|\text{KREV} - \sum_{i \in [k]} \mathbb{E}[\text{OPT}(\mathcal{D}_i)]| \leq \tilde{\Theta}(k(\epsilon + \epsilon^2/\epsilon_p + \epsilon_q + \epsilon_a))$$

Now let δ in the statement be $1/2k$ of the δ applied in the proof gives us the desired results. \square

I.5 DP MYERSON FOR UNBOUNDED DISTRIBUTION

Algorithm 9 DP Myerson, Unbounded Distribution DPMYERU($V, \epsilon_q, \epsilon_a, h, \epsilon_p$)

Input: n samples $V = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$, parameter n_1 , quantile discretization ϵ_q , additive discretization ϵ_a , regularity parameter η , privacy parameter ϵ_p , truncation parameter ϵ_t

- 1: $\text{KREV} \leftarrow \text{DPKOPT}(\{\mathbf{v}_1, \dots, \mathbf{v}_{n_1}\}, \epsilon_q, \epsilon_a, \epsilon_p, \eta)$.
 \triangleright Use n_1 samples to get a k -approximation of the optimal revenue.
- 2: Truncation all remaining samples by $1/\epsilon_t \cdot \text{KREV}$.
- 3: Discretize all the values into multiples of ϵ_a , let the resulting samples be \hat{V} .
- 4: Prepare the quantile to be estimated: $Q \leftarrow \{\epsilon_q, 2\epsilon_q, \dots, \lfloor (1/\epsilon_q) \rfloor \cdot \epsilon_q, 1\}$
- 5: For each dimension d , decide the prices on remaining samples:
 $\hat{S}_{[d,:]} \leftarrow \text{QESTIMATE}(Q, V_{[d,n_1:]}, \epsilon_p)$
 \triangleright Apply DP quantile estimate on the discretized value space(Alg. 4).
- 6: Fit Myerson's mechanism as if the valuations is in \hat{S} , each associated with probability ϵ_q .

Intergrating the bound of the DPKOPT(Alg. 2) into Alg. 9 gives us the following accuracy bound:

Theorem I.12 (Revenue Guarantee of Private Myerson, Unbounded (Alg. 9)). *Given $\epsilon \in [0, 1/4]$, n samples \hat{V} of the joint distribution $\mathbf{D} \in [0, h]^k$, the output of Myerson fitted under DPMYERU (Alg. 9) is $(2k\epsilon_p, 0)$ -DP, and the expected revenue of this mechanism is close to the optimal revenue of distribution \mathbf{D} , i.e., for $\epsilon_a = \epsilon_q = \epsilon$, $n = \Theta(\epsilon^2 \log(k/\delta))$ and $n_1 = \epsilon^{-2} \log(\epsilon^{-1})$, with probability $1 - \delta$,*

$$|\mathbb{E}[\text{Rev}(M_{\text{DPMYERU}}, \mathbf{D}) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})]| \leq \tilde{O}(\epsilon_t + k^2\epsilon/\epsilon_t + k^2\epsilon^2/\epsilon_p\epsilon_t)$$

Furthermore, when $\epsilon_t = \sqrt{\epsilon}$, the bounds simplifies to:

$$|\mathbb{E}[\text{Rev}(M_{\text{DPMYERU}}, \mathbf{D}) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})]| \leq \tilde{O}(k^2\sqrt{\epsilon} + k^2\epsilon^{1.5}/\epsilon_p)$$

Proof. From Thm H.4, we get that from $n = \tilde{\Theta}(\epsilon^{-2})$ samples, with probability $1 - \delta$ and $\epsilon_a = \epsilon_q = \epsilon$ we get that:

$$|\mathbb{E}[\text{Rev}(M_{\text{DPMYER}}, \mathbf{D}) - \text{OPT}(\mathbf{D})]| \leq \tilde{\Theta}((\epsilon + \epsilon^2/\epsilon_p)kh).$$

Next, we upper bound the value of the truncation point, from Thm I.11, with probability $1 - \delta$, we get:

$$|\text{KREV} - \sum_{i \in [k]} \mathbb{E}[\text{OPT}(\mathcal{D}_i)]| \leq \tilde{\Theta}(k(\epsilon + \epsilon^2/\epsilon_p))$$

Then, we have:

$$\begin{aligned} \text{KREV} &\leq \sum_{i \in [k]} \mathbb{E}[\text{OPT}(\mathcal{D}_i)] + \tilde{\Theta}(k(\epsilon + \epsilon^2/\epsilon_p + \epsilon_q + \epsilon_a)) \\ &\leq k\text{OPT}(\mathbf{D}) + \tilde{\Theta}(k(\epsilon + \epsilon^2/\epsilon_p + \epsilon_q + \epsilon_a)) \\ &= \tilde{\Theta}(k(1 + \epsilon + \epsilon^2/\epsilon_p)) \end{aligned}$$

Now we plug in $h = \frac{1}{\epsilon_t} \text{KREV}$, after simplification, this will gives us that, with probability $1 - 2\delta$:

$$|\mathbb{E}[\text{Rev}(M_{\text{DPMYERU}}, \mathbf{D}) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})]| \leq \tilde{O}(\epsilon_t + k^2\epsilon/\epsilon_t + k^2\epsilon^2/\epsilon_p\epsilon_t)$$

Now let the δ in the statement be the $1/2$ of the δ applied in this proof gives us the desired results.

□

Since our DP Myerson for unbounded distribution integrates the DP Myerson for bounded distribution, and the estimation for optimal revenue at most take $O(n)$ pass of the whole distribution, we have the running time guarantee below. Notice that here the ϵ_t doesn't affects the running time

Theorem I.13 (Running time, DP Myerson for Unbounded Distribution). *Given n samples, and quantile discretization parameter ϵ_q , the running time of DPMYER (Alg.9) is $\Theta(k \log(1/\epsilon_q)n) = \tilde{\Theta}(kn)$ and requires $\tilde{O}(n)$ pass of the distribution.*

Proof. From previous running time analysis of the DP Myerson for bounded distribution(Theorem H.5), we get that this part will take $\tilde{\Theta}(1)$ pass of the distribution and has running time $\tilde{\Theta}(kn)$.

The major component that contributes to the time complexity is the DP Quantile Estimation for Unbounded distribution, which, in the worst case, will take $O(n)$ passes over the dataset Durfee (2023), with each subsequent query takes $O(1)$ time for each distribution. Thus, the running time of DPQE for unbounded distribution is $O(kn)$.

Summing these up gives us a total running time of $\tilde{\Theta}(kn)$.

□

J MORE DETAILS FOR APPLICATION TO ONLINE MECHANISM DESIGN

J.1 COMMITMENT MECHANISM

Algorithm 10 Commitment Mechanism for bounded distribution

Input: Bids $\mathbf{b} \in \mathbb{R}_+^n$, distribution upper bound h .
 1: Sample a price $p \in [0, h]$ uniformly at random.
 2: Sample a bidder $i \in [n]$ uniformly at random.
 3: **if** $b_i > p$ **then** allocate the item to bidder i with price p .
 4: **else** No bidder gets the item.
 5: **end if**

Our commitment algorithm (Alg. 10) selects each bidder with equal probability, with a price drawn uniformly from $[0, h]$.

J.2 ONLINE MECHANISM DESIGN PRELIMINARIES

Assumption J.1 (Bidders' Distributions). We assume there are k publicly available bidder attributes, corresponding to k different distributions, i.e., each bidder with the attribute $a \in [k]$ will sample their valuations from \mathcal{D}_a . These distributions are unknown to the learner (i.e., prior independent). At every iteration t , one bidder from each attribute participates in the auction, sees the item, and decides their valuations.¹² In addition, these valuations are independent across different bidders and rounds.

Definition J.2 (Bidder's Utility). Each bidder j has a quasi-linear utility function at time t : $u_j^t = x_j^t(v_j^t - p_j^t)$. In our paper, we consider the following bidder models:

- **Discounted Utility.** For some discount factor $\gamma \in [0, 1]$, all bidders discount future utility by γ and seek to maximize the sum of discounted utilities. At the t -th iteration, the discounted utility is $\hat{u}_j^t = \sum_{r=t}^T u_j^r \gamma^{r-t}$.
- **Large Market** (Anari et al., 2014; Jalaly Khalilabadi and Tardos, 2018; Chen et al., 2016): $u_j^{1:T} = \sum_{t \in S_j} u_j^t$, with $|S_j| < l$.

where S_j is the set of iterations that the bidder participates in for the auction and x_j^t, v_j^t, p_j^t is the allocation, value, price for bidder j at time t , respectively.

This assumption is essential to optimize a near-optimal revenue since it is impossible to obtain more than a constant fraction of the revenue in a single bidder setting if each bidder participates in every round of the mechanism (Lem D.1). Ideally, the learner's objective is to learn a revenue-maximizing auction with a small failure probability. This regret is comparable to the revenue of the best fixed mechanism against the (nonobservable) value history; hence, it is stronger than the traditional regret, which is comparable to the revenue of the best-fixed mechanism against the bid history.

Definition J.3 (Learner's Objective). Given δ , the goal of the learner is to decide an allocation $\mathbf{x}_{1:T}$ and $\mathbf{p}_{1:T}$ such that the cumulative revenue is near optimal and with sublinear regret, i.e., with probability $1 - \delta$:

$$\widehat{\text{REGRET}} := \frac{1}{T} \sum_{t \in [T]} \mathbb{E}[\text{Rev}(\mathbf{x}_t, \mathbf{p}_t, \mathbf{b}_t) - \mathbb{E}[\text{OPT}(\mathbf{v}_t)]] = o(1),$$

where the expectation is taken over the bidders' distribution.

J.3 REVENUE GUARANTEE

Theorem J.4 (Accuracy Guarantee of Two-stage Mechanism). Given $\epsilon \in [0, 1/4]$, n samples of the joint distribution $\mathbf{D} \in [0, h]^k$, let Alg. 3 run with parameter $T_1 = \Theta(\epsilon^{-2} \log(k/\delta))$, $T = \Omega(T_1)$, $\epsilon_a = \epsilon_q = \epsilon_p = \epsilon$, $\nu = 2\alpha$ as calculated by Lemma J.7. Then, with probability $1 - \delta$, the regret is upper bounded, i.e.,

¹²We assume the bidder cannot see her valuation of all items at the start of the process.

- **Large Market Bidder:** $\text{REGRET} = \tilde{\Theta}[(\epsilon + \sqrt{l\epsilon})kh]$
- **Discounting Bidder:** $\text{REGRET} = \tilde{\Theta}[(\epsilon + \sqrt{\frac{\gamma\epsilon}{1-\gamma}})kh]$.

Proof. First note that each mechanism in the rounds T_1 to T is $(2k\epsilon_p, 0)$ -DP. From Lemma J.7, we get that by our choice of ν , during the first stage of our mechanism, each bid will at most deviate by ν . Then, by Lemma J.9, the expected utility of the Myerson deployed in the second stage (without DP) satisfies that

$$|\mathbb{E}[\text{Rev}(M_{\hat{\mathbf{D}}}, \mathbf{D}) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})]| \leq 2\nu,$$

where $\hat{\mathbf{D}}$ is the distribution after subtracting in line 4 of Alg. 3. Then, from the proofs of the DP mechanism for bounded distribution (Thm. 3.2), we have that with probability $1 - \delta$, the expected utility of DPMYER on distribution $\hat{\mathbf{D}}$ will have the following guarantee under $n = \Theta(\epsilon^{-2} \log(k/\delta))$:

$$|\mathbb{E}[\text{Rev}(M_{\hat{p}_p}, \hat{\mathbf{D}}) - \text{Rev}(M_{\hat{\mathbf{D}}}, \hat{\mathbf{D}})]| \leq \tilde{\Theta}[(\epsilon + \epsilon^2/\epsilon_p)kh],$$

where $M_{\hat{p}_p}$ denotes the algorithm applied in the second stage of our mechanism. Aggregating this guarantee gives us:

$$\begin{aligned} & |\mathbb{E}[\text{Rev}(M_{\hat{p}_p}, \mathbf{D}) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})]| \\ & \leq |\mathbb{E}[\text{Rev}(M_{\hat{p}_p}, \mathbf{D}) - \text{OPT}(\hat{\mathbf{D}})]| + |\mathbb{E}[\text{OPT}(\hat{\mathbf{D}}) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})]| \\ & \leq \tilde{\Theta}[(\epsilon + \epsilon^2/\epsilon_p)kh + \nu]. \end{aligned}$$

We can now analyze the entire regret of our 2-stage algorithm. With probability $1 - \delta$,

$$\text{REGRET} = \mathbb{E}\left[\frac{T - T_1}{T}(\text{OPT}(\mathbf{D}) - \tilde{\Theta}[(\epsilon + \epsilon^2/\epsilon_p)kh + \nu]) - \text{OPT}(\mathbf{D})\right].$$

Plugging in the ν for the large market and the discounting bidder gives us the desired result. \square

J.4 PROOF OF LEMMAS

Lemma J.5 (Approximate Truthfulness via DP (Lemma 3 in McSherry and Talwar (2007))). *Given a mechanism M with $(\epsilon_p, 0)$ -DP and database V, \hat{V} that differs only in one entry, for any nonnegative function g , we have that*

$$\mathbb{E}[g(M(V))] \leq \exp(\epsilon_p) \cdot \mathbb{E}[g(M(\hat{V}))].$$

Moreover, this implies that $\mathbb{E}[g(M(V))] \in [(1 \pm 2\epsilon_p) \mathbb{E}[g(M(\hat{V}))]]$.

We now give a lemma controlling the distance between a bid and its value.

Lemma J.6 (Bid Utility). *Given $\alpha \in [0, 1]$ and $t \in [0, T_1]$, the current utility at the t -th round of truthful bidding v_t exceeds that of strategic bidding b_t such that $|b_t - v_t| > 2\alpha$ by at least $\frac{2\alpha^2}{kh}$.*

Proof. WLOG, we assume $b_t < v_t - 2\alpha$. Then, the utility loss of strategic bidding is calculated as follows:

$$\int_{p=v_t-2\alpha}^{v_t} (v_t - p) \frac{1}{kh} = \frac{\alpha v_t}{kh} + \frac{2\alpha^2}{kh} \geq \frac{2\alpha^2}{kh}.$$

\square

Lemma J.7 (Bid Deviation). *Given any $t \in [0, T_1]$, the bidder will bid only b_t such that $|b_t - v_t| \leq 2\alpha$ for:*

- **Large Market:** $\sqrt{2(l-1)\epsilon_p}hk$.
- **Discounting Bidder:** $\sqrt{\frac{2\gamma\epsilon_p}{1-\gamma}}kh$.

Proof. The future utility of a bidder is upper bounded by: (1) $(l-1)h$ for the large market and (2) $h[1 + \frac{1}{\gamma} + \frac{1}{\gamma^2} + \dots] = \frac{\gamma}{1-\gamma}h$ for the discounting bidder. Notice that for each round in $[T_1, T]$, the mechanism is $2k\epsilon_p$ -DP; then, from Lemma J.5, we get that the future utility of strategic bidding is upper bounded by: (1) $4(l-1)hk\epsilon_p$ for the large market and (2) $\frac{\gamma}{1-\gamma}4kh\epsilon_p$ for the discounting bidder. Letting the current utility loss (i.e., $\frac{2\alpha^2}{kh}$) exceed the future rounds gives us a lower bound on α : (1) $\sqrt{2(l-1)\epsilon_p}hk$ for the large market and (2) $\sqrt{\frac{2\gamma\epsilon_p}{1-\gamma}}kh$ for the discounting bidder. \square

We now present an auxiliary lemma that bounds the revenue gap between the optimal mechanism under \mathcal{D} versus $\mathcal{D} - 2\nu$.

Lemma J.8 (Loss under additive ν). *Given a product distribution $\mathbf{D} \in \mathbb{R}_+^k$, let $\hat{\mathbf{D}} := \mathbb{P}_{\mathbb{R}_+}[\mathbf{D} - \nu]$, which results from subtracting ν for each value in \mathbf{D} then projected onto positive value domain. Then, we have that*

$$\mathbb{E}[\text{OPT}(\mathbf{D}) - \text{OPT}(\hat{\mathbf{D}})] \leq \nu.$$

Proof. Let $\tilde{\mathbf{D}} := \mathbf{D} - \nu$, then we couple distribution \mathbf{D} and $\tilde{\mathbf{D}}$ such that each $v \in \mathbb{R}_+^k$ in \mathbf{D} corresponds to $v - \gamma$ in $\tilde{\mathbf{D}}$, where $\mathbf{1}_k$ denotes the all one vector in k -dimensional space.

Then, we construct mechanism M according to the optimal mechanism $M_{\mathbf{D}}$ as follows: For each value profile v , $x_M(v) = x_{M_{\mathbf{D}}}(v + \nu\mathbf{1}_k)$ and $p_M(v) = p_{M_{\mathbf{D}}}(v + \nu\mathbf{1}_k) - \mathbf{1}_k\nu$. Since $M_{\tilde{\mathbf{D}}}$ is monotonic (hence, truthful), M is also truthful. Since all virtual value shifts equally, the allocation of M corresponds to the allocation of the optimal mechanism for $\tilde{\mathbf{D}}$. Hence, we have

$$\begin{aligned} \mathbb{E}[\text{OPT}(\mathbf{D}) - \text{OPT}(\hat{\mathbf{D}})] &\geq \mathbb{E}[\text{OPT}(\mathbf{D}) - \text{OPT}(\tilde{\mathbf{D}})] + \mathbb{E}[\text{OPT}(\tilde{\mathbf{D}}) - \text{OPT}(\hat{\mathbf{D}})] \\ &\geq \mathbb{E}[\text{Rev}(M_{\mathbf{D}}, \mathbf{D}) - \text{Rev}(M_{\tilde{\mathbf{D}}}, \hat{\mathbf{D}})] \\ &\geq \mathbb{E}[\text{Rev}(M_{\mathbf{D}}, \mathbf{D}) - \text{Rev}(M, \hat{\mathbf{D}})] = \nu, \end{aligned}$$

where the first line follows from the definition of the optimal mechanism. \square

We can now give our bound on the revenue deviation due to untruthful bidding.

Lemma J.9 (Revenue Deviation). *Given distribution \mathbf{D} , for any (bid) distribution $\tilde{\mathbf{D}}$ resulting from perturbing at most $\pm\nu$ for each value of \mathbf{D} and any distribution $\hat{\mathbf{D}}$ resulting from subtracting γ for each value of $\tilde{\mathbf{D}}$, we have that*

$$0 \geq \mathbb{E}[\text{Rev}(M_{\tilde{\mathbf{D}}}, \mathbf{D}) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})] \geq -2\nu.$$

Proof. We clearly see that $\hat{\mathbf{D}}$ satisfies $\mathbf{D} \succeq \hat{\mathbf{D}}$. Since $M_{\mathbf{D}}$ is optimal for \mathbf{D} , we have $0 \geq \mathbb{E}[\text{Rev}(M_{\tilde{\mathbf{D}}}, \mathbf{D}) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})]$. It remains to give a lower bound of this revenue difference. By strong monotonicity (Lemma F.3), we have

$$\mathbb{E}[\text{Rev}(M_{\tilde{\mathbf{D}}}, \mathbf{D}) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D})] \geq \text{Rev}(M_{\tilde{\mathbf{D}}}, \hat{\mathbf{D}}) - \text{Rev}(M_{\mathbf{D}}, \mathbf{D}) \geq -2\nu,$$

where the last inequality follows from Lemma J.8. \square