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 $\mathsf{Rev}(M, A)$ 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 $\mathsf{OPT}(A) := \mathsf{Rev}(M_A, A)$ 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 $\widetilde{O}$ and $\widetilde{\Theta}$ 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], $|\mathbb{E}[\mathsf{Rev}(M, \mathbf{D}) - \mathsf{OPT}(\mathbf{D})]| \leq \epsilon$ 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 $\widehat{D}^{\epsilon}$, 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 $\widehat{D}^{\epsilon}$, 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 $n = \widetilde{\Theta}(\epsilon^{-2})$ samples $\widehat{V}$ 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 $\widetilde{\Theta}(kn)$ and takes $\widetilde{\Theta}(1)$ pass of the distribution. With probability $1-\delta$, this mechanism M satisfies:* $|\mathbb{E}[\mathsf{Rev}(M, \mathbf{D}) - \mathsf{OPT}(\mathbf{D}))]| \leq \widetilde{O}((\epsilon + \epsilon^2/\epsilon_p)kh)$.

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 $\widetilde{\Theta}(k)$-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 $n = \widetilde{\Theta}(\epsilon^{-2})$ samples $\widehat{V}$ 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 $\widetilde{\Theta}(kn)$ and takes $O(n)$ passes. With probability $1-\delta$, this mechanism M satisfies:* $|\mathbb{E}[\mathsf{Rev}(M, \mathbf{D}) - \mathsf{Rev}(M_{\mathbf{D}}, \mathbf{D})]| \leq \widetilde{O}(k^2\sqrt{\epsilon} + k^2\epsilon^{1.5}/\epsilon_p)$.

[1]This 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 $T = \widetilde{\Theta}(\epsilon^{-2})$ that, with probability $1-\delta$, achieves diminishing regret, i.e.,* $\textsc{Regret} = \widetilde{O}[(\epsilon + \sqrt{\eta\epsilon})kh]$, *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 doen'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.

[2] In 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 $\mathbf{b} := (b_1, \ldots, b_k)$, a single-item auction M consists of an allocation rule as $\mathbf{x}(\mathbf{b}) := (x_1(\mathbf{b}), \ldots, x_k(\mathbf{b})) \in [0,1]^k$ and a payment rule as $\mathbf{p}(\mathbf{b}) := (p_1(\mathbf{b}), \ldots, p_k(\mathbf{b})) \in [0,1]^k$, 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 $\mathbf{D} = \mathcal{D}_1 \times \ldots \times \mathcal{D}_k$ (Elkind, 2007), the *virtual value* for $\mathcal{D}_j$ at value v_i^j with support $\mathcal{V}_j = \{v_1^j, \ldots, v_n^j\}$ 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^js 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 $\mathcal{A} : \mathbb{R}_+^n \to \mathbb{R}$ 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 $\mathcal{A} : \mathbb{R}_+^n \to \mathbb{R}$ be an (ϵ, δ)-DP algorithm, and let $f : \mathbb{R} \to \mathbb{R}$ be a random function. Then, $f \circ \mathcal{A} : \mathbb{R}_+^n \to \mathbb{R}$ 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 $0 \leq q_1, \ldots, q_m < 1$, identify quantile estimations $v_1, \ldots v_m$ 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 $d_{\text{ks}}(\mathcal{D}, \mathcal{D}') = \sup_{x\in\mathbb{R}} |F_{\mathcal{D}(x)} - F_{\mathcal{D}'}(x)|$.

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

[4] More formally, $v_j \in X$ is the minimum value that such 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 $[\epsilon_q, 2\epsilon_q, \ldots, 1]$ 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 $V \in R_+^{k\times n}$, 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 $\widehat{V}$.
2: Prepare the quantile to be estimated: $Q \leftarrow \{\epsilon_q, 2\epsilon_q, \ldots, \ldots, \lfloor(1/\epsilon_q)\rfloor \cdot \epsilon_q, 1\}$.
3: For each dimension $i \in [k]$, decide the prices $\widehat{S}_{[i,:]} \leftarrow$ QEstimate$(Q, V_{[i,:]}, \epsilon_p)$.
4: $\triangleright$ Estimate the quantiles by DPQuant (Alg. 4)
5: Construct distribution $\widetilde{D}$ based on $\widehat{S}$, treating the valuations in $\widehat{S}$ as if each has probability ϵ_q.
6: For each $i \in [k]$, shift the top ϵ_q quantile of $\widetilde{D}_i$ 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 reveue 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 $d_{ks}(\mathbf{D}_i, \mathbf{D}'_i) \le \alpha_i$ for any bidder i, the optimal revenue of these distribution satisfies:* $0 \le \mathbb{E}[\mathsf{Rev}(M_{\mathbf{D}}, \mathbf{D}) - \mathsf{Rev}(M_{\mathbf{D}'}, \mathbf{D}')] \le (\sum_{i\in[k]} \alpha_i)h$.

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 $\widetilde{\Theta}(\epsilon k h)$.

Theorem 3.2 (Revenue Guarantee of Private Myerson (Alg. 1))**.** *Given n samples $\widehat{V} \in [0,h]^{k\times n}$ 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 $1-\delta$:*

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

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

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 $\widehat{D}^p$ and the additively discretized distribution $\widehat{D}^\epsilon$. 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, δ, ϵ_a, ϵ_p, ϵ_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 $\widehat{D}^p$ is obtained by changing from distribution $\mathbf{D}$ through distribution $\widehat{D}$, the distribution $\widehat{D}^\epsilon$ and $\widehat{D}^q$ (Figure 1). We upper bound the statistical KS distance of these distributions: 1) By DKW inequality, we upper bound the KS-distance between $\widehat{D}$ and $\mathbf{D}$ by $\widetilde{\Theta}(1/\sqrt{n})$ for each coordinate i (with probability $1-\delta/2$). 2) By definition, we upper bound the KS-distance between $\widehat{D}^\epsilon$ and $\widehat{D}^q$ 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 $\widehat{D}^q$ and $\widehat{D}^p$ by $k\widehat{\epsilon}$ for $\widehat{\epsilon} := \widetilde{\Theta}(1/(\epsilon_p n))$ (with probability $1-\delta/2$).

Upper Bounding the Revenue Loss. We then upper bound optimal revenue loss from $\mathbf{D}$ to $\widehat{D}^p$. 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 $\widetilde{\Theta}((1/\sqrt{n}+\epsilon_q+\widehat{\epsilon})kh)$ 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 $\widehat{D}^p$ by the following inequality:

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

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

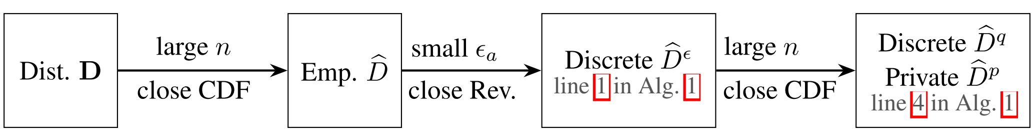

Figure 1: **Distribution analyzed for DPMYER(Alg. 1)**. We establish connections between the accuracy/revenue guarantee of the original distribution $\mathbf{D}$ with the empirical distribution $\widehat{D}$, the value-discretized $\widehat{D}^\epsilon$, the quantile-discretized $\widehat{D}^q$ and the distribution $\widehat{D}^p$ 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 $\widetilde{\Theta}(kn)$ time and requires $\widetilde{\Theta}(1)$ 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 $O(k\log(\lfloor 1/\epsilon_q \rfloor + 1)/(\epsilon_a\epsilon_q)) = \widetilde{\Theta}(kn)$ time, since $n = \widetilde{\Theta}(\epsilon^{-2})$. This step calculates the utility of $k\lfloor h/\epsilon_a \rfloor$ over $\lfloor 1/\epsilon_q \rfloor$ quantiles for at most $\widetilde{\Theta}(1)$ 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 $\frac{1}{\epsilon}\mathsf{OPT}(\mathbf{D})$ 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 $k\mathsf{OPT}(\mathbf{D}) \geq \sum_{i\in[k]} \mathsf{OPT}(\mathcal{D}_k) \geq \mathsf{OPT}(\mathbf{D})$.

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 discetization ϵ_a, privacy parameter ϵ_p, regularity parameter η.

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

To guarantee pure privacy, our algorithm estimates the optimal revenue using a DP-estimated proxy $\widehat{F}_{[k]}$ 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 $\widetilde{\Theta}(\frac{1}{\epsilon_p n}k + \epsilon_a)$. 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 $\widehat{V}$ 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 $\epsilon_a = \epsilon_q = \epsilon, n = \widetilde{\Theta}(\epsilon^2)$, $n_1 = \widetilde{\Theta}(\epsilon^2)$, $\epsilon_t = \sqrt{\epsilon}$, with probability $1-\delta$,*

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

[5]Without 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 *nonmypoic* 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 $\widehat{u}_j^t = \sum_{r=t}^{T} u_j^r \gamma^{r-t}$.

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 $1 - \delta$,

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

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 $(\mathbf{x}^t, \mathbf{p}^t) \leftarrow \textsc{Commit}(\mathbf{b}^t)$. ▷ Commitment Algorithm(Alg. 10)
4: $\widehat{\mathbf{b}}^t \leftarrow \mathbb{P}_{[0,h]}[\mathbf{b}^t - \nu \mathbf{1}_k]$
5: **end for**
6: $(\widetilde{\mathbf{x}}(\cdot), \widetilde{\mathbf{p}}(\cdot)) \leftarrow \text{DPM}\textsc{yer}(\widehat{\mathbf{b}}^{1:T_1}, \mathbf{a}^{1:T_1}, \epsilon_q, \epsilon_a, h, \epsilon_p)$ ▷ 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: $(\mathbf{x}^t, \mathbf{p}^t) \leftarrow \textsc{Myerson}(\widetilde{\mathbf{x}}(\cdot), \widetilde{\mathbf{p}}(\cdot))$;
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 incurs 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 $|b_t - v_t| \le 2\alpha$, 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 stochatically 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 $\widetilde{O}[(\epsilon + \sqrt{l\epsilon})kh]$, for $\nu = 2\sqrt{2(l-1)\epsilon_p} hk$.*

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

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

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

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

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 reasons about their utility accross 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, Bhuvesh 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)