
Learning Auctions with Robust Incentive Guarantees

Jacob Abernethy
Georgia Tech
prof@gatech.edu

Rachel Cummings
Georgia Tech
rachelc@gatech.edu

Bhuvesh Kumar
Georgia Tech
bhuvesh@gatech.edu

Jamie Morgenstern
Georgia Tech
jamiemt.cs@gatech.edu

Samuel Taggart
Oberlin College
sam.taggart@oberlin.edu

Abstract

We study the problem of learning Bayesian-optimal revenue-maximizing auctions. The classical approach to maximizing revenue requires a known prior distribution on the demand of the bidders, although recent work has shown how to replace the knowledge of a prior distribution with a polynomial sample. However, in an online setting, when buyers can participate in multiple rounds, standard learning techniques are susceptible to *strategic overfitting*: bidders can improve their long-term wellbeing by manipulating the trajectory of the learning algorithm through bidding. For example, they may be able to strategically adjust their behavior in earlier rounds to achieve lower, more favorable future prices. Such non-truthful behavior can hinder learning and harm revenue. In this paper, we combine tools from differential privacy, mechanism design, and sample complexity to give a repeated auction that (1) learns bidder demand from past data, (2) is approximately revenue-optimal, and (3) strategically robust, as it incentivizes bidders to behave truthfully.

1 Introduction

When we observe prices in market settings—stock exchanges, farmers’ markets, ad auctions—we understand that these prices were not chosen arbitrarily. Rather, the seller (auctioneer, market maker, etc.) selected these prices after observing a stream of previous transactions, which provide relevant information about the demands of buyers that are key to maximizing income as well as managing available inventory. The process of setting prices from a growing database of previous sales is fundamentally a learning problem, with all of the typical tradeoffs akin to bias versus variance, etc. In the case of repeated auctions, however, there is one additional challenge: market participants are often quite aware of the underlying learning procedures employed by the auctioneer and can seek to benefit using deceptive bidding strategies. Buyers, in other words, can aim to *induce* overfitting, introducing additional hurdles to learning problem at hand.

Under bayesian assumptions, and in a batch setting where agents only act once, auction pricing has been well-understood since the work of Myerson [39], who characterized the revenue-optimal scheme as a function of the prior distribution of values of the bidders. Frequentist alternatives to this model have been introduced in recent years [21, 13, 4, 18, 37, 5, 38, 25, 17, 23], with the goal of designing auctions with good revenue guarantees if one does not have a prior but instead is given only samples from the underlying distribution. These methods, however, still imagine only a one-shot mechanism and are not robust to multi-round strategic behavior of bidders.

This paper studies the design of multi-round auction-learning algorithms that exhibit theoretical guarantees that limit a buyer’s ability to manipulate the mechanism towards their own benefit. Our

results aim to nudge the development of optimal auctions closer to realistic environments where such mechanisms are deployed. We employ tools from *differential privacy* as our core technique to control the impact of any individual buyer’s strategy on her utility in future participation. A differentially private mechanism ensures that the output of a computation has only a small dependence on any one input data point. Privacy has been previously used as a tool to achieve truthfulness in a variety of game theoretic environments [14], including mechanism design [33, 40], mediated games [29, 16], and market design [15, 28, 44, 12]. Our seller’s learning algorithm is differentially private with respect to bid data, which limits the effect of each player’s bid on future choices of single-round auctions, thus disentangling incentives across rounds. In this sense we use differential privacy not as a tool for information security but instead for robustness; this in turn yields the desired incentive guarantees.

Our two main results are the first computationally tractable algorithms for learning nearly-optimal Bayesian revenue-maximizing auctions with an approximate truthfulness guarantee. We first give a learning algorithm for which there exists an approximate equilibrium in which all bidders report their true values. Along the way to this result, we provide several useful technical lemmas for comparing the revenue of mechanisms on two similar distributions, and comparing the revenue of two similar mechanisms on any fixed distribution. These may find future applications of independent interest beyond this work. Second, under an assumption limiting changes in buyer behavior from round to round, we construct a mechanism for which there is an exact equilibrium where every bidder bids within a small constant of their value, which also achieves a nearly optimal revenue guarantee. A mechanism with this guarantee is substantially more complex to achieve, requiring arguments about what bidders might learn about each other across rounds.

Related Work The classical Myerson auction [39] maximizes revenue when selling a single item to buyers drawn from a fixed, known distribution. More recent work investigated how to maximize revenue when that distribution is unknown but some small number of samples are available [21, 13, 4, 18, 37, 5, 38, 25, 17, 23] (or in an online model [9]). This line of work assumes every buyer in the sample reported their values *truthfully*, rather than manipulating their reported data. If each buyer appears in at most one auction, then they will not have incentive to misreport their bids. However, strategic buyers participating repeatedly in these auctions may have incentive to manipulate their behavior so as to guarantee themselves more utility in future rounds. We extend this line of work by assuming each buyer can participate in several auctions, and so we must analyze bidders’ incentives across rounds if we hope to learn a good auction from their past bidding behavior. Related line of work includes dynamic auctions [42, 3, 36, 31] where the buyers are strategic about how their behaviour today effects the seller’s behaviour in the future and [8] where the buyer uses a non regret learning algorithm to bid across the rounds.

Our results are related to work on iteratively learning prices [7, 35, 10, 41], although these results do not consider multiple appearances of bidders across rounds. Most closely related to our work is that of Liu et al. [30], which assumes buyers may appear more than once, and finds no-regret posted prices or anonymous reserves. We leverage several of their novel ideas, such as maintaining a differentially private internal state to guarantee approximate truthfulness. Further, our work optimizes over a substantially more complicated space of mechanisms.

With repeated appearances of each buyer, our auction learning problem comes to resemble dynamic mechanism design. We therefore review some of the relevant literature. A truthful mechanism is given in [6] that exactly maximizes social welfare in a dynamic environment, and [27, 43] extended this mechanism to maximize revenue. In contrast, our mechanism approximately maximizes revenue in a dynamic environment with much looser assumptions on buyers’ value distributions, but compares to the weaker per-round benchmark of the optimal single-shot revenue. Epasto et al. [22] takes a similar approach and proposes a dynamic mechanism for buyers who strategize across rounds under a *large market* assumption. In contrast to our proposed mechanism, theirs will not generally run in polynomial time. Additionally, we focus on mechanisms which *explicitly* limit an agent’s ability to manipulate them.

2 Model and Preliminaries

The model We consider a T -round auction, where a seller sells a supply of J identical items to n unit-demand bidders each round. For each round t and population i a value $v_{i,t}$ is sampled from

a fixed distribution D_i , representing the amount the bidder is willing to pay for the item. We let $\mathbf{D} = D_1 \times \dots \times D_n$ denote the product distribution of value distributions, and we use \mathbf{v} to denote a vector of values sampled from this distribution. Further, we let \mathbf{v}_{-i} denote \mathbf{v} with the i -th element removed, and use (v'_i, \mathbf{v}_{-i}) to denote the same vector with v'_i replacing the i -th element.

We consider a setting similar to Liu et al. [30] where a bidder from any population may appear several times over the course of the T rounds, drawing a fresh value each time.¹ In this setting, bidders may have an incentive to misreport their values in order to change the mechanism in future rounds, and their potential reward for doing so depends on the number of future rounds in which they expect to participate. Amin et al. [1] show that very little can be done when a bidder participates in every round, so we assume this cannot occur. Formally:

Assumption 1. *No bidder participates in more than k rounds of the T -round auction.*

Mechanism Design Basics One can view a *mechanism* (auction) $\mathcal{M} := (x, p)$ as having two components: (a) a possibly-randomized *allocation rule* $x : V^n \rightarrow \mathcal{X}$, which takes in a vector of values (bids) \mathbf{v} and returns a feasible allocation of the items, where $x_i(\mathbf{v})$ is 1 if i receives the item and 0 otherwise; and (b) a *payment rule* $p : V^n \rightarrow \mathbb{R}^n$, which takes \mathbf{v} and outputs a vector of payments demanded of each player. Assuming, for the moment, that bidders bid their true values \mathbf{v} , we can define the *expected revenue* of a mechanism \mathcal{M} as the expectation of the payments received,

$$\text{Rev}(\mathcal{M}; D) := \mathbb{E}_{\mathbf{v} \sim \mathbf{D}}[\sum_{i=1}^n p_i(\mathbf{v})].$$

We make the standard assumption that the participants have *quasi-linear utility*: for a vector of bids \mathbf{v}' (which may not necessarily match the values \mathbf{v}), a bidder's utility for allocation $x_i(\mathbf{v}')$ and payment $p_i(\mathbf{v}')$ is

$$u_i(x, p, \mathbf{v}') = v_i \cdot x_i(\mathbf{v}') - p_i(\mathbf{v}').$$

We may now introduce the notion of a *truthful* mechanism. A mechanism $\mathcal{M} := (x, p)$ is deemed truthful if, given a vector of true values \mathbf{v} and some arbitrary vector of bids \mathbf{v}' , each agent receives no less utility bidding v_i rather than v'_i ; that is, for every i , it must hold that $u_i(x, p, \mathbf{v}') \leq u_i(x, p, (v_i, \mathbf{v}'_{-i}))$.

Let us now recall a classical result in Bayesian-optimal mechanism design when the seller's goal is to maximize revenue. Myerson [39] essentially fully characterized the solution in this setting. The interested reader can learn more in Hartline [24]; we briefly review these results here in two pieces. The first piece states that payments in truthful mechanisms essentially depend solely on the allocation function.

Theorem 2 (Payment Identity, Myerson [39]). *A mechanism is truthful if and only if it has a monotone allocation rule and payments which for all valuation profiles \mathbf{v} satisfy*

$$p_i(\mathbf{v}) = v_i \cdot x_i(v_i, \mathbf{v}_{-i}) - \int_0^{v_i} x_i(u, \mathbf{v}_{-i}) du + p_i(0, \mathbf{v}_{-i}).$$

The second key result is that for truthful mechanisms, the expected revenue can be written in terms of welfare in a remapped *virtual value* space.

Theorem 3 (Myerson [39]). *For any truthful mechanism $\mathcal{M} = (x, p)$ with values distributed according to D , the expected revenue from player i can be written as $\mathbb{E}_{\mathbf{v} \sim D}[\phi_i(v_i)x_i(\mathbf{v})]$, where $\phi_i(v_i)$ is the virtual value, given by $\phi_i(v_i) = 1 - \frac{1-F_i(v_i)}{f_i(v_i)}$. So, $\text{Rev}(\mathcal{M}; D) = \mathbb{E}_{\mathbf{v} \sim D}[\sum_i \phi_i(v_i)x_i(\mathbf{v})]$.*

We will use the notation $\mathcal{M}_{\mathbf{D}}^*$ to denote the revenue-optimal mechanism for distribution \mathbf{D} —Myerson provides a precise construction of this auction.

Definition 4 (Myerson's Auction). *Fixing a prior distribution \mathbf{D} , given a value profile \mathbf{v} Myerson's revenue-optimal mechanism $\mathcal{M}_{\mathbf{D}}^*$ calculates virtual values $\phi_i(v_i) = v_i - \frac{1-F_{D_i}(v_i)}{f_{D_i}(v_i)}$ and (a) selects the feasible allocation which maximizes virtual welfare according to the virtual values and (b) charges payments according to the Payment Identity of Theorem 2.*

¹Our results can be shown to hold when values are not drawn fresh, in which case D_i is the distribution of drawn values, taking into account the process by which bidders are redrawn from population.

Truthfulness and Dynamic Equilibrium The mechanism design preliminaries discussed previously are for *one shot* games where players do not observe past actions of others and adjust their strategy accordingly. We now turn our attention to multi-round play, we need to expand our notion of player behavior and strategy. We will now assume that a strategy for each bidder maps the history of observed actions to contingency plans over actions in the current and future rounds: we assume agents observe their own outcomes x_i^t and p_i^t in rounds in which they participate, but not the full historical data used by the designer to produce the mechanism each round. A history h_i^t in round t for each agent i then consists of a bid, an allocation, and a payment for each round the agent has participated in. Given a history h_i^t , and value v_i^t in the current round t for agent i , we denote agent i 's strategy as $\sigma_i^t(v_i^t; h_i^t)$. We suppress the dependence on the history when clear from context.

Note that because agents do not observe the historical bids of others, they must form probabilistic beliefs about these bids, and we assume these beliefs are Bayesian in nature. Denote by $\mu_i^t(h_i^t)$ the beliefs of agent i in round t after observing history h_i^t . A profile of strategies is an *equilibrium* for this game if in every round t , every agent i , and every history h_i^t that agent i might have observed previously, agent i 's strategy maximizes their expected total utility over the current and future rounds; the expectation is taken over the randomness of agent i 's beliefs as well as the future values of all agents. Formally, let $U_i^t(\sigma, v_i^t, h_i^t)$ denote the total expected utility of agent i in rounds $t, t+1, \dots, T$ given a value of v_i^t in the current round and an observed history h_i^t , playing strategy σ .

Definition 5 (Perfect Bayesian Equilibrium (PBE)). *A profile of strategies σ is an η -approximate Perfect Bayesian equilibrium if for every agent i , round t , history h_i^t for agent i in round t , value v_i^t for i in round t , the strategy $\sigma_i^t(v_i^t; h_i^t)$ approximately maximizes agent i 's total expected utility from future rounds up to an additive η . That is, $U_i^t(\sigma_i^t(v_i^t; h_i^t), v_i^t, h_i^t) \geq U_i^t(\sigma', v_i^t, h_i^t) - \eta$ for every alternate strategy σ' . If $\eta = 0$, we say that σ is an exact Perfect Bayesian Equilibrium.*

In other words, a PBE is a strategy σ such that, for every bidder i and for every history of the game, if all other bidders besides i behave according to the strategy σ , then playing σ_i is approximately utility-maximizing behavior for bidder i .

Definition 6 (η -utility-approximate BIC). *A mechanism is η utility-approximately Bayesian incentive compatible if the strategy profile where every agent bids truthfully in every history is an η -approximate Perfect Bayesian equilibrium.*

We also consider a more robust notion of incentive compatibility, where there exists an exact equilibrium with each agent bidding η -close to her true value, a notion also used in Liu et al. [30].

Definition 7 (η -bid-approximate BIC). *A mechanism is η bid-approximate BIC if \exists an exact PBE where each bidder bids within η of their value in every history.*

Definition 6 guarantees that all bidders bidding truthfully in all rounds is an (approximate) Bayes-Nash equilibrium (BNE). In proving a mechanism utility-approximate BIC, one therefore may assume bidders report truthfully in future rounds. Consequently, the only impact an agent's bid has on their future utility is through their impact on future mechanisms. Definition 7 guarantees the existence of an exact equilibrium in which all bidders bid within η of their value. Therefore, mechanisms with this guarantee will need to ensure bidders do not change their behavior in later rounds to influence other bidders' behavior in earlier rounds.

Differential Privacy Background We now provide some basics on differential privacy, our main technique that helps guarantee approximate truthfulness in equilibrium. We refer to a *database* $Z \in \mathcal{Z}^n$ as a collection of data from n individuals, and we say that two databases are *neighboring* if they differ in at most one entry.

Definition 8 (Differential Privacy [19]). *An algorithm (mapping) $\mathcal{A} : \mathcal{Z}^n \rightarrow \mathcal{R}$ is (ϵ, δ) -differentially private if for neighboring databases $Z, Z' \in \mathcal{Z}^n$ and subsets of possible outputs $\mathcal{S} \subseteq \mathcal{R}$, we have $\mathbb{P}[\mathcal{A}(Z) \in \mathcal{S}] \leq \exp(\epsilon) \mathbb{P}[\mathcal{A}(Z') \in \mathcal{S}] + \delta$.*

The parameter ϵ quantifies the algorithm's privacy guarantee; smaller ϵ corresponds to stronger privacy. A key property of differential privacy is that it is robust to *post-processing*.

Lemma 9 (Post-processing [19]). *Let $\mathcal{A} : \mathcal{Z}^n \rightarrow \mathcal{R}$ be an (ϵ, δ) -differentially private algorithm and let $f : \mathcal{R} \rightarrow \mathcal{R}'$ be a random function. Then $f \circ \mathcal{A} : \mathcal{Z}^n \rightarrow \mathcal{R}'$ is also (ϵ, δ) -differentially private.*

We need a more precise notion of privacy when multiple agents are involved receiving different information. For notation, we say two databases are i -neighbors if they differ only in the i -th entry. Also, let $\mathcal{A}(Z)_{-i}$ denote the vector of outputs to all players except player i .

Definition 10 (Joint Differential Privacy [29]). *An algorithm $\mathcal{A} : \mathcal{Z}^n \rightarrow \mathcal{R}^n$ is (ϵ, δ) -jointly differentially private if for every $i \in [n]$, every pair of i -neighbors $Z, Z' \in \mathcal{Z}^n$, and every $\mathcal{S} \subseteq \mathcal{R}^{n-1}$,*

$$\mathbb{P}[\mathcal{A}(Z)_{-i} \in \mathcal{S}] \leq \exp(\epsilon) \mathbb{P}[\mathcal{A}(Z')_{-i} \in \mathcal{S}] + \delta.$$

Finally, we can reason about the joint differential privacy of mechanism decomposed into a public sanitized broadcast, i.e. as if on a “billboard,” and a private non-sanitized portion. The following lemma shows that privacy is still preserved under such a decomposition.

Lemma 11 (Billboard Lemma [26]). *Suppose $\mathcal{A} : \mathcal{Z}^n \rightarrow \mathcal{R}$ is (ϵ, δ) -differentially private. Consider any collection of functions $f_i : \mathcal{Z}_i \times \mathcal{R} \rightarrow \mathcal{R}'$, for $i \in [n]$, where \mathcal{Z}_i is the portion of the database containing i 's data. Then the composition $\{f_i(\Pi_i(Z), \mathcal{A}(Z))\}$ is (ϵ, δ) -jointly differentially private, where $\Pi_i : \mathcal{Z} \rightarrow \mathcal{Z}_i$ is the projection to i 's data.*

The final tools we borrow from differential privacy are the exponential mechanism [34], and the ability to maintain a histogram estimate of values which arrive in one at a time. The primary technique for the latter involves data structures known as tree-based aggregations [20, 11]. This protocol is a differentially private method for calculating the cumulative sum of elements from 1 to t for any $t \leq T$, for which at any round t the protocol can return an estimate of the number of elements prior to round t , for which the entire execution is differentially private. We provide more details about our instances of these algorithms in Appendices B.1.3 and A.1.1.

3 Utility-Approximate Bayesian Incentive Compatibility

In this section, we give an online algorithm (Algorithm 1) for learning the optimal auction which is utility-approximate BIC. The main idea is to use differential privacy to explicitly control the amount of information the auctioneer takes forward from round t to later rounds. We do so by maintaining a differentially private estimate $H'_{i,t}$ of each empirical bid distribution, and choosing future auctions based only on this differentially private estimate. Thus, from the perspective of any bidder, her behavior in round t has very little chance of affecting any of the auctions selected in later rounds. In round t , we run Myerson’s mechanism with prior $H'_{i,t}$ to compute allocations and payments. Thus the one-shot mechanism in round t is exactly incentive compatible with respect to the current round. The differentially private subroutine used to compute estimate $H'_{i,t}$ is described in full detail in Appendix A, with privacy, truthfulness, and revenue guarantees presented in Section 3.1.

Algorithm 1: Utility-Approximate BIC Online Auction

Parameters: discretization β , privacy ϵ , upper bound on support h , num. of rounds T

Initialize: $H'_{i,0} \leftarrow \text{Uniform}(0, h)$ for $i = 1, \dots, n$

for $t = 1, \dots, T$ **do**

 Receive bid profile $\mathbf{v}_t = (v_{1,t}, \dots, v_{n,t})$, rounded down to integer multiple of β

 Run Myerson (Def. 4) with \mathbf{H}'_{t-1} as prior and \mathbf{v}_t as bid for allocations/payments.

for $i = 1, \dots, n$ **do**

 | Update $H'_{i,t}$ via two-fold tree aggregation (Algorithm 3), giving as input $v_{i,t}$

end

end

3.1 Privacy, Truthfulness, and Revenue Guarantees

We now prove that the learning subroutine of Algorithm 1 satisfies differential privacy (Theorem 12), use this to prove the mechanism is utility-approximate BIC (Theorem 13), and show that Algorithm 1 achieves a $o(1)$ additive approximation to the optimal revenue (Theorem 14).

Theorem 12. *The stream of estimates $\{\mathbf{H}'_t\}_{t=1}^T$ maintained by Algorithm 1 is $(\epsilon, \epsilon/T)$ -differentially private with respect to the stream of input bids $\{v_t\}_{t=1}^T$.*

We emphasize that Theorem 12 does not claim that Algorithm 1 is itself differentially private, it only states that the procedure rests on a differentially private subroutine. This distinction is critical: our

algorithm is not differentially private in its selection of allocations and payments in round t . However, the information the mechanism carries forward (namely, the estimated empirical distribution) is maintained in a differentially private manner. This is sufficient for guaranteeing that bidders' behavior in round t does not significantly affect which auctions are selected in later rounds. This will allow us to prove a utility-approximate BIC guarantee, but will not be sufficiently strong to argue about bid-approximate BIC. For that, we will need to additionally ensure that the allocations (and payments) in round t are differentially private (jointly differentially private), see Section 4.

We now turn our attention to proving a guarantee on the truthfulness of Algorithm 1, which will lean heavily on the privacy guarantee given in Theorem 12. We note that if our mechanism were $(\epsilon, 0)$ -differentially private then a result of [33], stating that any $(\epsilon, 0)$ -DP mechanism is 2ϵ -dominant strategy incentive compatible.

Two issues arise if one were to try this approach in our setting. First, the entire mechanism is not differentially private as discussed above. A bidder i 's behavior might have significant impact on other bidder's allocations and payments, and those bidders may as a result choose to behave differently in later rounds based on that information. Thus we relax to the weaker incentive guarantee of utility-approximate BIC, avoiding the issue of other bidders behaving differently in response to activity from earlier rounds. Second, the stream of estimates maintained by our mechanism is (ϵ, δ) -differentially private for $\delta = \epsilon/T > 0$ and not $(\epsilon, 0)$ -differentially private which is necessary for the result of [33] to hold.

Theorem 13. *Algorithm 1 is $k h \epsilon (2 + \frac{1}{T})$ -utility approximate BIC when $\epsilon < 1$.*

Finally, we consider the revenue-optimality of our proposed mechanism. Our revenue guarantees will rely on tools developed and described in Appendix A.2. Recall that we use $\text{Rev}(\mathcal{M}; D)$ to denote the expected revenue generated by the mechanism \mathcal{M} on a value distribution D , and that \mathbf{D} and \mathbf{D}' respectively denote the joint distributions of true values and true values rounded down to the nearest multiple of β . Let $\mathcal{M}_{\mathbf{H}'_t}^*$, $\mathcal{M}_{\mathbf{D}'}^*$, and $\mathcal{M}_{\mathbf{D}}^*$ be the truly revenue-optimal mechanisms for the distributions \mathbf{H}'_t , \mathbf{D}' , and \mathbf{D} , respectively. In each round of our mechanism, we get a sample from \mathbf{D}' and run Myerson's auction with \mathbf{H}'_t as the prior; that is, we run $\mathcal{M}_{\mathbf{H}'_t}^*$. We use $\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{D}')$ to denote the expected revenue we achieve in round t . The bidders true values are drawn from \mathbf{D} and we use $\text{Rev}(\mathcal{M}_{\mathbf{D}}^*; \mathbf{D})$ to denote the optimal expected revenue in any round. In this section, we will compare the revenue $\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{D}')$ of our mechanism to the optimal revenue $\text{Rev}(\mathcal{M}_{\mathbf{D}}^*; \mathbf{D})$.

The main result of this section, Theorem 14, bounds the difference between our average expected revenue, $\frac{1}{T} \sum_{t=1}^T \text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{D}')$, and that of the optimum. We show that over T rounds, with high probability the average expected revenue is within an additive $o(1)$ error of the optimal.

Theorem 14. *With probability at least $1 - \alpha$, Algorithm 1 satisfies $\frac{1}{T} \sum_{t=1}^T \text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{D}') \geq \text{Rev}(\mathcal{M}_{\mathbf{D}}^*; \mathbf{D}) - \beta J - 4hn^2 \tilde{O}\left(\frac{1}{\sqrt{T}} + \frac{1}{T\epsilon}\right)$ for regular distributions \mathbf{D} and $\epsilon < 1$.*

Proof. We start by instantiating Lemma 34 for every round t instantiated with failure probability α/T . Then taking a union bound over all T rounds and summing over t , ensures that with probability $1 - \alpha$, $\frac{1}{T} \sum_{t=1}^T \text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{D}') \geq \text{Rev}(\mathcal{M}_{\mathbf{D}}^*; \mathbf{D}) - \beta J - \frac{4hn^2}{T} \sum_{t=1}^T \gamma_t$ for

$$\gamma_t = \sqrt{\frac{\log \frac{n}{\alpha}}{2t}} + \frac{\sigma}{t} \sqrt{\log \frac{h}{\beta} \log T} \sqrt{2 \log \left(\frac{2hn}{\beta\alpha} \right)} \text{ and } \sigma = \frac{8 \log T \log \frac{h}{\beta}}{\epsilon} \sqrt{\ln \frac{T \log T \log \frac{h}{\beta}}{\epsilon}}$$

In the remainder of the proof, we bound $\frac{1}{T} \sum_{t=1}^T \gamma_t$. (Recall that the α in Lemma 34 is α/T here.)

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \gamma_t &= \frac{1}{T} \sum_{t=1}^T \left(\sqrt{\frac{\log \frac{nT}{\alpha}}{2t}} + \frac{\sigma}{t} \sqrt{\log \frac{h}{\beta} \log T} \sqrt{2 \log \left(\frac{2hnT}{\beta\alpha} \right)} \right) \\ &\leq \sqrt{\frac{2 \log \frac{nT}{\alpha}}{T}} + \frac{2\sigma \log T}{T} \sqrt{\log \frac{h}{\beta} \log T} \sqrt{2 \log \left(\frac{2hnT}{\beta\alpha} \right)} = \tilde{O}\left(\frac{1}{\sqrt{T}} + \frac{1}{T\epsilon}\right) \end{aligned}$$

The first inequality comes from the facts that $\sum_{t=1}^T \frac{1}{\sqrt{t}} \leq 2\sqrt{T}$ and $\sum_{t=1}^T \frac{1}{t} = H_T \leq \log T + 1 \leq 2 \log T$. The following equality come from plugging in the expression of σ and combining terms.

Hence, w.p. $\geq 1 - \alpha$, $\frac{1}{T} \sum_{t=1}^T \text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{D}') \geq \text{Rev}(\mathcal{M}_{\mathbf{D}}^*; \mathbf{D}) - \beta J - 4hn^2 \tilde{O}\left(\frac{1}{\sqrt{T}} + \frac{1}{T\epsilon}\right)$. \square

Notice that this statement says that if we set β , the discretization parameter to be $o(1)$ in terms of T then the revenue one can earn is a $(1 - o(1))$ -approximation to the optimal revenue.

4 Bid-Approximate Bayesian Incentive Compatibility

We now describe an algorithm for “training” a mechanism to achieve nearly optimal revenue in an exact equilibrium, where each bidder bids close to their true value. This contrasts with the result in the prior section, which shows that bidders bidding their exact true values is an approximate equilibrium. While the equilibrium we describe in this section is not quite a truthful one, we can compare its revenue with the revenue of the optimal mechanism facing truthful bids.

The primary technical meat of this section has two parts. First, we describe how to modify Algorithm 1 to get the stronger incentive guarantee of bid-approximate Bayesian incentive compatibility. Theorem 13 gave utility-approximate BIC, which means all bidders behaving truthfully in all rounds is an approximate Bayes-Nash equilibrium (BNE). Bid-approximate BIC means there is an *exact* Bayes-Nash equilibrium where all bidders bid within η of their values. The main challenge here comes from whether bidders change their behavior in later rounds based on the non-truthful behavior of other bidders in earlier rounds.² To make this jump, we guarantee that round t 's allocations and payments leak very little information about a bidder's behavior to other bidders. So, bidders will have very little ability to condition on one another's behavior in round t when selecting a strategy in later rounds. To this end, we ensure that the allocations in round t are differentially private and the payments in the round t are jointly differentially private, informally stated below.

Theorem 15. *Algorithm 2 runs in polynomial time, is η_t -bid approximate BIC in round t for $\eta_t = c + \tilde{O}(t^{-1/4})$ and achieves expected revenue $Rev = OPT - c' - \tilde{O}(t^{-1/4})$ for small constants c, c' .*

The more precise result, given in Theorem 19, shows that under mild assumptions about bidders' behavior, every equilibrium bid under our mechanism lies close to the true value in the bid space. That is, in round t , we show that no bidder in equilibrium will underbid by more than a small factor η_t . The main challenge in getting such a result requires us to bound the extent to which bidders' behavior in the current round can affect their future utilities. If we can control this quantity, we can control the amount by which people can game the system in the current round and also ensures that bidders are unable to learn much about the value distributions of other bidders. The formal revenue guarantees for our mechanism are presented in Theorem 20 where we show that with high probability, the average expected revenue obtained by our mechanism is close to the optimal expected revenue of the mechanism with complete information about the value distribution of the bidders (i.e., Myerson's auction on \mathbf{D}).

4.1 A mechanism with private payments and allocations

We now describe Algorithm 2 which is a Bid-approximate BIC Online Auction algorithm. The primary challenge is to ensure that the allocation and payments in round t do not leak much information about bidders' behavior to one another. The mechanism ensures this by making choices which are all jointly differentially private, ensuring that bidders do not substantially affect either the mechanism's state or what other bidders learn about them. The mechanism maintains a private estimate of the empirical distribution (as before), and also computes prices and allocations in round t using jointly differentially private algorithms. We use the exponential mechanism (Algorithm 6) for picking round t 's allocation, black box payments (Algorithm 5) to pick payments in round t (which in expectation yields payments close to the truthful payments). This additional step of ensuring round t decisions are differentially private is crucial when bidders might condition their behavior in later rounds based upon what they learn from round t (which might be the case in bid-approximate equilibria).

To describe exact equilibria, we need to argue about each bidder's utility for modifying her bid, and argue that shading by more than η_t hurts her utility. We do this using a *punishing* mechanism, which penalizes bidders who shade their bid by using a strictly truthful mechanism with some probability ρ_t in each round. Since we are going to show that there exists an equilibrium where in round t every bidder bids within η_t of their true value, it will be convenient to define the round t equilibrium bid distributions $\mathbf{F}_t = \mathbf{F}_{1,t} \times \cdots \times \mathbf{F}_{n,t}$.

²This is not a challenge when we assume all bidders behave truthfully, since bidders won't have non-truthful behavior in earlier rounds to condition their behavior upon.

Algorithm 2: Bid-approximate BIC Online Auction

Parameters: discretization β , privacy ϵ , upper bound on support h , num. of rounds T

Initialize: $H'_{i,0} \leftarrow \text{Uniform}(0, h)$ for $i = 1, \dots, n$

```
for  $t = 1, \dots, T$  do
  Receive vector of bids  $\mathbf{b}_t = (b_{1,t}, \dots, b_{n,t})$ , rounded down to multiple of  $\beta$ 
  With probability  $\rho_t$  run mechanism  $\text{StrictlyTruthful}(\mathbf{b}_t)$  (Algorithm 4)
  else
    for  $i = 1, \dots, n$  do
      Use  $H'_{i,t-1}$  to calculate  $\phi_{i,t}(b_{i,t})$  (Theorem 3)
    end
    Use exponential mechanism (Algorithm 6) to select allocation  $x_t(\phi_t(\mathbf{b}))$ 
    Use Black box payments (Algorithm 5) to calculate payments  $p_t(\mathbf{b}_t)$ .
  end
  for  $i = 1, \dots, n$  do
    Update  $H'_{i,t}$  via two-fold tree aggregation (Algorithm 3), giving as input  $b_{i,t}$ 
  end
end
```

4.2 Privacy, Truthfulness, and Revenue Guarantees

In this subsection, we provide the differential privacy and incentive guarantees of Algorithm 2. We refer to Appendix B for the and additional technical details and proofs.

Lemma 16. *Algorithm 2 is $(3\epsilon, 3\epsilon/T)$ -jointly differentially private in the bids of users.*

Lemma 17. *Fix a round t , population i , and the bidder from population i in round t . This bidder's total utility in rounds $t' > t$ for misreporting in round t is $\leq \epsilon kh(2 + \frac{1}{T})$ more than truthful reporting.*

Finally, we introduce an assumption we make about the equilibrium behavior of bidders facing this mechanism. Informally, the assumption states that a mechanism run over many rounds should have equilibrium bid distributions which have similar probability of any bid between adjacent rounds.

Assumption 18 (λ -Stable Bid Distribution). *We say the mechanism \mathcal{M} supports a λ -stable bid distribution if it has some BNE with distribution of equilibrium bids \mathbf{F}_t in round t such that for all populations i and rounds t , there exist λ_t s.t. $\|\mathbf{F}_t - \mathbf{F}_{t-1}\|_\infty \leq \lambda_t$.*

Remark. Consider a mechanism with very similar behavior in round t versus $t + 1$: both the mechanism's distribution over allocation and payment rules in round t and $t + 1$ are very close. If all other bidders strategies from round t to $t + 1$ are also very close, then the utility for any particular bid b from a bidder with valuation v from population i will be quite close, but that bidder's utility-optimal bid may or may not be identical or even particularly close in the two rounds.

This suggests that analyzing exact equilibria in iterated settings is quite complex, in that the distribution over utility-optimal bids might shift quite substantially from round to round. So, mechanisms and equilibria without this property may have highly erratic behavior, and such equilibria may not support a learning procedure which competes with the (truthful) optimal revenue. Hence we assume that the mechanism in Algorithm 2 supports a λ -Stable Bid Distribution with the condition that the quantity $\Sigma_T := \sum_{t=1}^{T-1} \frac{t\lambda_t}{T} = o(1)$. We now present the truthfulness guarantee of Algorithm 2.

Theorem 19. *Algorithm 2 in round t is η_t -bid approximate BIC, i.e. in round t , any bidder i with value $v_{i,t}$ reports bid $b_{i,t}$ which satisfies $v_{i,t} - \eta_t \leq b_{i,t} \leq v_{i,t}$ where $\eta_t = h\sqrt{\frac{8n^2\gamma_t + 6k\epsilon}{\rho_t J}}$,*

$$\gamma_t = \sqrt{\frac{\log(2hn/\beta\alpha)}{2t}} + \Sigma_t + \frac{\sigma}{t} \sqrt{\log \frac{h}{\beta} \log T} \sqrt{\log \left(\frac{hn}{\beta\alpha} \right)}, \sigma = \frac{8 \log T \log \frac{h}{\beta}}{\epsilon} \sqrt{\ln \frac{\log T \log \frac{h}{\beta}}{\delta}} \text{ and } \delta = \frac{\epsilon}{T}.$$

Let \mathbf{F}'_t be distribution of round t bids (\mathbf{F}_t) rounded down to the nearest multiple of β . Let the mechanism we run in round t on the bids we receive from \mathbf{F}'_t be $\mathcal{M}_{\mathbf{H}'_t}$. The expected revenue we achieve in round t is $\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}; \mathbf{F}'_t)$ which we will compare with the optimal revenue $\text{Rev}(\mathcal{M}_{\mathbf{D}}^*; \mathbf{D})$. We now present the main revenue guarantee for Algorithm 2, i.e we compare the average expected revenue of Algorithm 2 with respect to optimal revenue. We refer to the Appendix B.4 for the proof.

Theorem 20. Using Algorithm 2 for T rounds, with probability at least $1 - \alpha$ the average expected revenue obtained by the mechanism over the T rounds satisfies,

$$\frac{1}{T} \sum_{t=1}^T \text{Rev}(\mathcal{M}_{\mathbf{H}'_t}; \mathbf{F}'_t) \geq \text{Rev}(\mathcal{M}_{\mathbf{D}}^*; \mathbf{D}) - \left(hnJ^{2/3} \tilde{O}\left(\frac{1}{T^{1/4}}\right) + \frac{J \ln n}{\epsilon} + \frac{h}{3\epsilon} \ln \frac{J}{k\epsilon} + hJ^{2/3} (12k\epsilon)^{1/3} + \beta J \right)$$

References

- [1] Kareem Amin, Afshin Rostamizadeh, and Umar Syed. 2013. Learning prices for repeated auctions with strategic buyers. In *Advances in Neural Information Processing Systems*. 1169–1177.
- [2] Aaron Archer, Christos Papadimitriou, Kunal Talwar, and Éva Tardos. 2004. An approximate truthful mechanism for combinatorial auctions with single parameter agents. *Internet Mathematics* 1, 2 (2004), 129–150.
- [3] Itai Ashlagi, Constantinos Daskalakis, and Nima Haghpanah. 2016. Sequential mechanisms with ex-post participation guarantees. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, 213–214.
- [4] Maria-Florina Balcan, Avrim Blum, Jason D Hartline, and Yishay Mansour. 2005. Mechanism design via machine learning. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*. IEEE, 605–614.
- [5] Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik. 2016. Sample complexity of automated mechanism design. In *Advances in Neural Information Processing Systems*. 2083–2091.
- [6] Dirk Bergemann and Juuso Valimaki. 2010. The Dynamic Pivot Mechanism. *Econometrica* 78, 2 (2010), 771–789.
- [7] Avrim Blum and Jason D Hartline. 2005. Near-optimal online auctions. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 1156–1163.
- [8] Mark Braverman, Jieming Mao, Jon Schneider, and Matt Weinberg. 2018. Selling to a no-regret buyer. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. ACM, 523–538.
- [9] Sebastien Bubeck, Nikhil R Devanur, Zhiyi Huang, and Rad Niazadeh. 2017. Online auctions and multi-scale online learning. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. ACM, 497–514.
- [10] Nicolo Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. 2015. Regret minimization for reserve prices in second-price auctions. *IEEE Transactions on Information Theory* 61, 1 (2015), 549–564.
- [11] T-H Hubert Chan, Elaine Shi, and Dawn Song. 2011. Private and continual release of statistics. *ACM Transactions on Information and System Security (TISSEC)* 14, 3 (2011), 26.
- [12] Yiling Chen, Stephen Chong, Ian A Kash, Tal Moran, and Salil Vadhan. 2016. Truthful mechanisms for agents that value privacy. *ACM Transactions on Economics and Computation (TEAC)* 4, 3 (2016), 13.
- [13] Richard Cole and Tim Roughgarden. 2014. The sample complexity of revenue maximization. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*. ACM, 243–252.
- [14] Rachel Cummings. 2017. Differential Privacy As a Tool for Truthfulness in Games. *XRDS* 24, 1 (Sept. 2017), 34–37.

- [15] Rachel Cummings, Stratis Ioannidis, and Katrina Ligett. 2015. Truthful Linear Regression. In *Proceedings of The 28th Conference on Learning Theory (COLT '15)*. 448–483.
- [16] Rachel Cummings, Michael Kearns, Aaron Roth, and Zhiwei Steven Wu. 2015. Privacy and Truthful Equilibrium Selection for Aggregative Games. In *Proceedings of the 11th International Conference on Web and Internet Economics (WINE '15)*. 286–299.
- [17] Nikhil R Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. 2016. The sample complexity of auctions with side information. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. ACM, 426–439.
- [18] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. 2015. Revenue maximization with a single sample. *Games and Economic Behavior* 91 (2015), 318–333.
- [19] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*. Springer, 265–284.
- [20] Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. 2010. Differential privacy under continual observation. In *Proceedings of the forty-second ACM symposium on Theory of computing*. ACM, 715–724.
- [21] Edith Elkind. 2007. Designing and Learning Optimal Finite Support Auctions. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '07)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 736–745. <http://dl.acm.org/citation.cfm?id=1283383.1283462>
- [22] Alessandro Epasto, Mohammad Mahdian, Vahab Mirrokni, and Song Zuo. 2018. Incentive-Aware Learning for Large Markets. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 1369–1378.
- [23] Yannai A. Gonczarowski and Noam Nisan. 2017. Efficient Empirical Revenue Maximization in Single-parameter Auction Environments. In *Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC '17)*. 856–868.
- [24] Jason Hartline. 2013. Mechanism design and approximation. *Book draft*. October 122 (2013).
- [25] Jason Hartline and Samuel Taggart. 2016. Non-Revelation Mechanism Design. arXiv preprint 1608.01875. (2016).
- [26] Justin Hsu, Zhiyi Huang, Aaron Roth, Tim Roughgarden, and Zhiwei Steven Wu. 2016. Private matchings and allocations. *SIAM J. Comput.* 45, 6 (2016), 1953–1984.
- [27] Sham Kakade, Ilan Lobel, and Hamid Nazerzadeh. 2013. Optimal Dynamic Mechanism Design and the Virtual-Pivot Mechanism. *Operations Research* 64, 4 (2013), 837–854.
- [28] Sampath Kannan, Jamie Morgenstern, Ryan Rogers, and Aaron Roth. 2015. Private Pareto Optimal Exchange. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation (EC '15)*. 261–278.
- [29] Michael Kearns, Mallesh Pai, Aaron Roth, and Jonathan Ullman. 2014. Mechanism Design in Large Games: Incentives and Privacy. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science (ITCS '14)*. ACM, 403–410.
- [30] Jinyan Liu, Zhiyi Huang, and Xiangning Wang. 2018. Learning Optimal Reserve Price against Non-myopic Bidders. In *Advances in Neural Information Processing Systems*. 2038–2048.
- [31] Siqi Liu and Christos-Alexandros Psomas. 2018. On the competition complexity of dynamic mechanism design. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, 2008–2025.
- [32] Pascal Massart. 1990. The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. *The annals of Probability* (1990), 1269–1283.

- [33] Frank McSherry and Kunal Talwar. 2007. Mechanism design via differential privacy. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*. IEEE, 94–103.
- [34] Frank McSherry and Kunal Talwar. 2007. Mechanism Design via Differential Privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*. 94–103.
- [35] Andres M Medina and Mehryar Mohri. 2014. Learning theory and algorithms for revenue optimization in second price auctions with reserve. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*. 262–270.
- [36] Vahab S Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. 2016. Dynamic Auctions with Bank Accounts.. In *IJCAI*. 387–393.
- [37] Jamie Morgenstern and Tim Roughgarden. 2016. Learning simple auctions. In *Conference on Learning Theory*. 1298–1318.
- [38] Jamie H Morgenstern and Tim Roughgarden. 2015. On the pseudo-dimension of nearly optimal auctions. In *Advances in Neural Information Processing Systems*. 136–144.
- [39] Roger B Myerson. 1981. Optimal auction design. *Mathematics of operations research* 6, 1 (1981), 58–73.
- [40] Kobbi Nissim, Claudio Orlandi, and Rann Smorodinsky. 2012. Privacy-aware Mechanism Design. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC '12)*. ACM, 774–789.
- [41] Renato Paes Leme, Martin Pal, and Sergei Vassilvitskii. 2016. A field guide to personalized reserve prices. In *Proceedings of the 25th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 1093–1102.
- [42] Christos Papadimitriou, George Pierrakos, Christos-Alexandros Psomas, and Aviad Rubinfeld. 2016. On the complexity of dynamic mechanism design. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 1458–1475.
- [43] Alessandro Pavan, Ilya Segal, and Juuso Toikka. 2014. Dynamic Mechanism Design: A Myersonian Approach. *Econometrica* 82, 2 (2014), 601–653.
- [44] David Xiao. 2013. Is privacy compatible with truthfulness?. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*. ACM, 67–86.

A Omitted Proofs and Subroutines for Utility-Approximate BIC Online Auction (Algorithm 1)

A.1 Differentially Private Distribution Estimation

We now describe a differentially private procedure for estimating value distributions $H'_{i,t}$ for each bidder population. This corresponds to the final operation in each round of our mechanism (Algorithm 1).

When we receive value $v_{i,t}$ we round it down to nearest integer multiple of β . Recall that the value distribution for bidder population i rounded down to the nearest multiple of β is D'_i , and that D'_i has finite support $\{0, \beta, 2\beta, \dots, h\}$ of size $h/\beta + 1$, and thus the estimates $H'_{-i,t}$ we maintain will have the same finite support. We work under the assumption that bids are truthfully reported values, which we later validate with the incentive guarantees shortly.

Let $H_{i,t}$ be the empirical (non-private) estimate of D'_i at round t . The following lemma establishes that the empirical distribution of rounded values provides a good estimate of the true distribution of rounded values, with respect to the supremum norm $\|\cdot\|_\infty$.

Lemma 21. *Let $H_{i,t}$ be the empirical distribution of t i.i.d. samples from D'_i . Then, with probability*

$$\text{at least } 1 - \alpha, \|D'_i - H_{i,t}\|_\infty \leq \sqrt{\frac{\log \frac{2}{\alpha}}{2t}}.$$

Proof. This is a direct result of the Dvoretzky-Kiefer-Wolfowitz inequality [32], which establishes the concentration of the empirical CDF for any distribution. \square

A.1.1 Two-fold tree aggregation

Standard tree-based aggregation [20, 11] maintains online counts of a single quantity (e.g., the number of 1s in a stream of bits); we will use an extension of these known as *two-fold tree aggregation* [30] to maintain an online estimate of the CDF of a probability distribution \mathbf{D}' in a differentially private manner. Informally, one achieves this by maintaining a private counter for each bin in a histogram over possible values which might update the counters. Since \mathbf{D}' has a discrete support of $\{0, \beta, 2\beta, \dots, h\}$, the non-private empirical CDF $H_{i,t}$ can be described by a simple (increasing) step function, with steps occurring at integer multiples of β . To compute $H_{i,t}(u)$, the empirical CDF at value u , we need only count the number of samples from $v_{i,1}, \dots, v_{i,t}$ which are less than u , i.e., $H_{i,t}(u) = (\sum_{\tau=1}^t \mathbf{1}\{v_{i,\tau} \leq u\})/t$. Two-fold tree aggregation allows us to privately maintain these cumulative sums for all points $u \in \{\beta, 2\beta, \dots, h\}$ in the support of our distributions.

The algorithm maintains n instances of the two-fold tree aggregation procedure, one for each bidder population, where each instance has its own distinct internal state. The i th instance maintains $H'_{i,t}$ in round t . In each round t , the mechanism receives a value profile \mathbf{v}_t . For each population i , $v_{i,t}$ is used to update the internal state of population i 's tree aggregation instance.

The algorithm is given formally in Algorithm 3, and require the following additional notation. Consider any $t \in [T]$ with binary representation $(t_{\lceil \log T \rceil}, \dots, t_1, t_0)$. That is, $t = \sum_0^{\lceil \log T \rceil} t_j 2^j$. Let j_t be the lowest nonzero bit, and let

$$\Lambda_t = \{t - 2^{j_t} + 1, t - 2^{j_t} + 2, \dots, t - 1, t\}.$$

We also define the set,

$$\Gamma_t = \{t' : t' = t - \sum_{j=0}^{h-1} t_j 2^h, h = 1, \dots, \lceil \log T \rceil\}.$$

We note that Γ_t has size at most $\lceil \log T \rceil$, and the set $[t]$ can be described as the union of $\lceil \log T \rceil$ sets, i.e., $[t] = \cup_{j \in \Gamma_t} \Lambda_j$.

In two-fold tree aggregation, we have these aggregations over two axes: time t and value u . Thus we maintain $\frac{h}{\beta} \cdot T$ partial sums, denoted in Algorithm 3 as internal states $A_{tq} = H_{i,t}(u)$ for $u = q\beta$. One sample $v_{i,t}$ contributes to at most $\log \frac{h}{\beta} \log T$ partial sums, and each $H'_{i,t}(u)$ can be written as a sum of at most $\log \frac{h}{\beta} \log T$ partial sums.

Algorithm 3: Two-fold tree aggregation for population i [30]

Input: discretization parameter β , privacy parameter ϵ , upper bound on support h , number of rounds T

Internal State: Noisy partial sums A_{tq} for all $t \in [T]$ and $q \in \left[\frac{h}{\beta}\right]$

Initialize: Set $\sigma = \frac{8 \log T \log \frac{h}{\beta}}{\epsilon} \sqrt{\ln \frac{T \log T \log \frac{h}{\beta}}{\epsilon}}$ and sample $A_{tq} \sim_{i.i.d.} \mathcal{N}(0, \sigma^2)$ for all t and q
for $t = 1, \dots, T$ **do**

 Receive $v_{i,t} = p\beta$ for some $p \in \left[\frac{h}{\beta}\right]$

for j, k satisfying $t \in \Lambda_j$ and $p \in \Lambda_k$ **do**

$A_{jk} = A_{jk} + 1$

end

for $q \in \left[\frac{h}{\beta}\right]$ **do**

 Sample $\nu_{tq} \sim \mathcal{N}(0, ((\log \frac{h}{\beta} + 1)(\log T + 1) - |\Gamma_t| |\Gamma_q|) \sigma^2)$

end

 Output $H'_{i,t}$, the estimated CDF: $H'_{i,t}(x) := \sum_{j \in \Gamma_t} \sum_{k \in \Gamma_q} \frac{A_{jk} + \nu_{tq}}{t}$, where $q = \lfloor x/\beta \rfloor$.

end

The following lemma shows that Algorithm 3 is differentially private, and guarantees that privacy is maintained throughout the entire run of the algorithm.

Lemma 22 (Liu et al. [30]). *The entire stream of estimates $\{H'_{i,t}\}_{t=1}^T$ output by Algorithm 3 is $(\epsilon, \frac{\epsilon}{T})$ -differentially private with respect to the input stream of bids $\{v_{i,t}\}_{t=1}^T$.*

The construction of Algorithm 3 ensures that every value $H'_{i,t}(u)$ is obtained by perturbing $H_{i,t}(u)$ with the t -normalized sum of Gaussian variables, each with variance σ^2 . The total number of Gaussian noise terms added to obtain $H'_{i,t}(u)$ is no more than $\log \frac{h}{\beta} \log T$ because the sets Γ_t and Γ_q used in the final output of Algorithm 3 have size at most $\lceil \log T \rceil$ and $\lceil \log h/\beta \rceil$, respectively. That is, for each fixed u , we have $H'_{i,t}(u) - H_{i,t}(u) \sim N\left(0, \frac{\sigma^2}{t^2} \log \frac{h}{\beta} \log T\right)$. Lemma 23 uses this fact to bound the distance between $H_{i,t}$ and $H'_{i,t}$.

Lemma 23. *After t rounds, for a fixed population i , with probability at least $1 - \alpha$ the empirical distribution $H_{i,t}$ and the differentially private estimate $H'_{i,t}$ produced by Algorithm 3 will satisfy*

$$\|H_{i,t} - H'_{i,t}\|_{\infty} \leq \frac{\sigma}{t} \sqrt{\log \frac{h}{\beta} \log T} \sqrt{2 \log \left(\frac{2h}{\beta\alpha}\right)},$$

$$\text{for } \sigma = \frac{8 \log T \log \frac{h}{\beta}}{\epsilon} \sqrt{\ln \frac{T \log T \log \frac{h}{\beta}}{\epsilon}}.$$

Proof. For a Gaussian random variable $Z \sim \mathcal{N}(\mu, \rho^2)$ it holds that $\mathbb{P}[|Z - \mu| > x\rho] \leq 2 \exp(-x^2/2)$, by a standard tail bound. We apply this inequality to each point $u \in \{0, \beta, 2\beta, \dots, h\}$ to see that, with probability at least $1 - \frac{\beta\alpha}{h}$, we have $|H_{i,t}(u) - H'_{i,t}(u)| \leq \frac{\sigma}{t} \sqrt{\log \frac{h}{\beta} \log T} \sqrt{2 \log \left(\frac{2h}{\beta\alpha}\right)}$. Applying a union bound over all of the $\frac{h}{\beta}$ values of u completes the proof. \square

We can now combine the previous lemmas to relate the distributions $H'_{i,t}$ and D'_i .

Theorem 24. *After t rounds Algorithm 1, it holds with probability at least $1 - \alpha$ that*

$$\|H'_{i,t} - D'_i\|_{\infty} \leq \gamma_t \quad \text{for every } i \in [n],$$

$$\text{where } \gamma_t = \sqrt{\frac{\log \frac{n}{2t}}{2t}} + \frac{\sigma}{t} \sqrt{\log \frac{h}{\beta} \log T} \sqrt{2 \log \left(\frac{2hn}{\beta\alpha}\right)} \text{ and } \sigma = \frac{8 \log T \log \frac{h}{\beta}}{\epsilon} \sqrt{\ln \frac{T \log T \log \frac{h}{\beta}}{\epsilon}}.$$

Proof. Applying triangle inequality on Lemma 23 and Lemma 21 gives the bound for a single bidder population i , and a union bound over all n bidder populations proves the lemma. \square

In the remainder of this section, the same definition of γ_t will be used.

A.2 Revenue Maximization on Similar Distributions

In the previous section we showed that the differentially private estimate of the value distribution we maintain is close to the true distribution from where the values are being drawn. We now introduce tools we need to argue that the revenue of our mechanism, facing truthful bidding, is approximately optimal. In particular, describe how one can compare the revenue of a fixed, well-behaved mechanism on two similar product distributions \mathbf{D} and $\tilde{\mathbf{D}}$. This will imply that the problem of optimizing with respect to $\tilde{\mathbf{D}}$ will yield approximately optimal revenue with respect to \mathbf{D} . These results are broadly stated and should be of independent interest.

More formally, we will consider distributions that are close in ℓ_∞ distance, and mechanisms which are well-behaved in the sense that allocating one agent leads to the exclusion of others. The relevant definitions follow. Recall that we overload the notation and D and \tilde{D} refer to the corresponding CDFs when used as functions.

Definition 25 (τ -closeness). *We call two distributions D and \tilde{D} τ -close if $\|D - \tilde{D}\|_\infty \leq \tau$.*

Definition 26 (Competitiveness). *A truthful mechanism is competitive if for any valuation profile \mathbf{v} and any pair of bidders i and j , the allocation probability $x_i(\mathbf{v})$ for bidder i is a non-increasing function of v_j .*

In multi-unit settings, mechanisms that exactly maximize virtual surplus for any monotone virtual value function (e.g. ironed Myerson virtual value) satisfy this property, as does the exponential mechanism with score for each allocation equal to its virtual surplus. Given these definitions, we may state the main result of this section:

Theorem 27. *Let \mathcal{M} be a competitive mechanism, and let \mathbf{D} and $\tilde{\mathbf{D}}$ be two product distributions of values such that for every bidder i , D_i and \tilde{D}_i are τ -close. Then the expected revenue of \mathcal{M} on \mathbf{D} is within an additive $2n^2h\tau$ of the revenue from \mathcal{M} on $\tilde{\mathbf{D}}$. That is, $|\text{Rev}(\mathcal{M}; \mathbf{D}) - \text{Rev}(\mathcal{M}; \tilde{\mathbf{D}})| \leq 2n^2h\tau$.*

In fact, we prove a stronger statement than Theorem 27: we show that the revenue from each bidder is within $2nh\tau$ in each mechanism. To prove this stronger statement, we argue from the perspective of an individual bidder, and consider three steps. First, we show that switching the distributions of all other bidders from \mathbf{D}_{-i} to $\tilde{\mathbf{D}}_{-i}$ does not significantly change i 's allocation probability. We then show that because of the relationship between allocation and payments in truthful mechanisms an insignificant change in allocation probability implies an insignificant change in revenue. We then show that switching bidder i 's distribution from D_i to \tilde{D}_i does not significantly impact the revenue of any mechanism. The result will follow from the triangle inequality. We begin with the first of our three steps by defining new notation for the allocation probability secured by each bidder in expectation over the other bidders.

Definition 28 (Interim allocation rule). *Let $\mathcal{M} = (\mathbf{x}, \mathbf{p})$ be a mechanism for the single-round game. For any value v_i , the interim allocation rule for i at v_i is given by $x_i(v_i) = \mathbb{E}_{\mathbf{v}_{-i}}[x_i(\mathbf{v})]$.*

Our first step is to show that each bidder's interim allocation rule under competitive mechanisms is robust to small changes in other bidders' value distributions. Formally:

Lemma 29. *Let \mathbf{D}_{-i} and $\tilde{\mathbf{D}}_{-i}$ be value distributions for bidders other than i , with D_j and \tilde{D}_j τ -close for all $j \neq i$. Consider any truthful competitive mechanism, and let $x_i(\cdot)$ and $\tilde{x}_i(\cdot)$ denote the interim allocation rules of bidder i under \mathbf{D}_{-i} and $\tilde{\mathbf{D}}_{-i}$, respectively. Then for any value v_i , $|x_i(v_i) - \tilde{x}_i(v_i)| \leq (n-1)\tau$.*

Proof. We will consider changing the value of just one bidder, j , and observing the impact on the interim allocation rule of bidder i . The lemma will follow from repeating this argument once for

each bidder other than i . To show that slightly changing bidder j 's distribution has a minimal effect, we write the following sequence of equalities and inequalities, which we justify afterward.

$$\begin{aligned}
x_i(v_i) &= \int_0^h \mathbb{E}_{\mathbf{v}_{-i,j}}[x_i(\mathbf{v})] D_j(v_j) dv_j \\
&= \mathbb{E}_{\mathbf{v}_{-i,j}}[x_i(\mathbf{v})] D_j(v_j) \Big|_0^h - \int_0^h \frac{d}{dv_j} \mathbb{E}_{\mathbf{v}_{-i,j}}[x_i(\mathbf{v})] D_j(v_j) dv_j \\
&= \mathbb{E}_{\mathbf{v}_{-i,j}}[x_i(v_i, h, \mathbf{v}_{-i,j})] - \int_0^h \mathbb{E}_{\mathbf{v}_{-i,j}}[\frac{d}{dv_j} x_i(\mathbf{v})] D_j(v_j) dv_j \\
&\geq \mathbb{E}_{\mathbf{v}_{-i,j}}[x_i(v_i, h, \mathbf{v}_{-i,j})] - \int_0^h \mathbb{E}_{\mathbf{v}_{-i,j}}[\frac{d}{dv_j} x_i(\mathbf{v})] (\tilde{D}_j(v_j) - \tau) dv_j \\
&= \mathbb{E}_{\mathbf{v}_{-i,j}}[x_i(v_i, h, \mathbf{v}_{-i,j})] - \int_0^h \mathbb{E}_{\mathbf{v}_{-i,j}}[\frac{d}{dv_j} x_i(\mathbf{v})] \tilde{D}_j(v_j) dv_j - \tau \\
&= \mathbb{E}_{\mathbf{v}_{-i,j}}[x_i(\mathbf{v})] \tilde{F}_j(v_j) \Big|_0^h - \int_0^h \mathbb{E}_{\mathbf{v}_{-i,j}}[\frac{d}{dv_j} x_i(\mathbf{v})] \tilde{D}_j(v_j) dv_j - \tau \\
&= \int_0^h \mathbb{E}_{\mathbf{v}_{-i,j}}[x_i(\mathbf{v})] \tilde{D}_j'(v_j) dv_j - \tau
\end{aligned}$$

The reasoning is as follows. The first equality is from the definition of the interim allocation rule $x_i(v_i)$. The second and third equalities follow by integration by parts and interchanging the derivative and integral, respectively. The third inequality follows from the τ -closeness of D_j and \tilde{D}_j , and from the fact that $\frac{d}{dv_j} x_i(\mathbf{v}) \leq 0$ by the competitiveness of the mechanism. The remaining equalities follow from the same reasoning as the first three. Hence, changing bidder j 's value distribution from D_j to \tilde{D}_j can decrease bidder i 's allocation probability by at most τ . A symmetric argument bounds the increase. Further applying this same argument to all bidders other than i , one at a time, implies the lemma. \square

The payment identity (Theorem 2) characterizes the payments of an individual bidder with a realized type in a truthful mechanism, and shows that this payment is completely determined by the allocation rule the agent faces. Taking expectations over the values of other agents yields a characterization of an agent's expected payments in terms of their interim allocation rule. This characterization will allow us to show that the revenue from any bidder under two similar interim allocation rules is similar.

Corollary 29.1 (of Theorem 2). *In any truthful mechanism, for any bidder i with value v_i , the expected revenue of bidder i satisfies:*

$$\mathbb{E}_{\mathbf{v}_{-i}}[p_i(\mathbf{v})] = v_i x_i(v_i) - \int_0^{v_i} x_i(z) dz + \mathbb{E}_{\mathbf{v}_{-i}}[p_i(0, \mathbf{v}_{-i})] \quad (1)$$

Lemma 30. *Let x_i and \tilde{x}_i be interim allocation rules for bidder i such that $|x_i(v_i) - \tilde{x}_i(v_i)| \leq \tau$ for all $v_i \in [0, h]$. If $\mathbb{E}_{\mathbf{v}_{-i}}[p_i(0, \mathbf{v}_{-i})] = 0$ under both allocation rules, then for any value v_i , the expected payments made by a bidder with that value differ by at most $2v_i\tau$ under the two allocation rules.*

Proof. By equation (1), the difference in revenue between the two mechanisms is given by

$$v_i(x_i(v_i) - \tilde{x}_i(v_i)) - \int_0^{v_i} (x_i(z) - \tilde{x}_i(z)) dz.$$

The first term is at most $v_i\tau$. Moreover the second term is at most $\int_0^{v_i} \tau dz$, which is equal to $v_i\tau$. \square

Combining Lemmas 29 and 30 yields:

Corollary 30.1. *Let \mathbf{D}_{-i} and $\tilde{\mathbf{D}}_{-i}$ be value distributions for bidders other than i , with D_j and \tilde{D}_j τ -close for all $j \neq i$. Consider any truthful competitive mechanism \mathcal{M} where the bidders with value 0 make no payments. Then the expected revenue of \mathcal{M} from bidder i differs by at most $2(n-1)h\tau$.*

We finally show that holding other bidders' value distributions fixed and switching bidder i from value distribution D_i to a τ -close distribution \tilde{D}_i yields similar revenue. Formally:

Lemma 31. *Let D_i and \tilde{D}_i be τ -close value distributions for bidder i . For any truthful mechanism \mathcal{M} and any value distributions \mathbf{D}_{-i} for other bidders, the expected revenue from bidder i under $D_i \times \mathbf{D}_{-i}$ and $\tilde{D}_i \times \mathbf{D}_{-i}$ differ by at most $h\tau$.*

To prove Lemma 31, we use a standard characterization of a bidder's expected payments in a truthful mechanism, which can be obtained by integrating equation (1) over all values v_i and integrating by parts.

Corollary 31.1 (of Theorem 2). *In any truthful mechanism where bidders with value 0 make no payments, for any bidder i with value distribution D_i , the expected revenue from bidder i is given by*

$$\mathbb{E}_{\mathbf{v}}[p_i(\mathbf{v})] = \int_0^h x'_i(v_i) R_i(v_i) dv_i \quad (2)$$

where $R_i(v_i)$ is bidder i 's price posting revenue function, given by $R_i(v_i) = v_i(1 - D_i(v_i))$

Proof of Lemma 31. By equation (2), the difference in expected revenue between the two distributions is given by

$$\int_0^h x'_i(v_i) v_i (\tilde{D}_i - D_i) dv_i \leq h\tau \int_0^h x'_i(v_i) dv_i$$

Since $\int_0^h x'_i(v_i) dv_i \leq 1$, the result follows. \square

Proof of Theorem 27. Combining Corollary 30.1 with Lemma 31 and using the triangle inequality implies that the revenue of any individual bidder i differs by at most $2(n-1)h\tau + h\tau \leq 2nh\tau$ under \mathbf{D} rather than $\tilde{\mathbf{D}}$. Summing over all bidders yields the desired bound. \square

A.3 Omitted Proofs in Section 3

Proof of Theorem 12. Algorithm 1's only record of bids which persists across rounds is its distribution estimate \mathbf{H}'_t . In each round, it chooses an auction as a post-processing step over those estimated distributions. The two-fold tree aggregation step is $(\epsilon, \epsilon/T)$ -differentially private by Lemma 22. Thus, the algorithm's post-processing to estimate the virtual value distribution and select future auctions is $(\epsilon, \epsilon/T)$ -differentially private by Lemma 9. \square

Proof of Theorem 13. Consider a bidder deciding how to bid in round t . She has three considerations: how her behavior will affect (1) the learning algorithm in future rounds, (2) the behavior of other bidders in future rounds, and (3) her utility in round t .

Since we seek to show that the mechanism is utility-approximate BIC, we can assume all other bidders behave truthfully in every history (by Definition 6). Therefore, other bidders will not change their behavior in future rounds, and the value of (2) is 0. The value of (3) is also 0 because the empirical Myerson auction run in each round t is chosen to be exactly truthful as a one-shot (static) mechanism, so no payer can gain in expected utility by misreporting her bid. Thus the only utility a player can gain by lying about her value is from (1).

Next we analyze (1). Since our mechanism's differential privacy guarantee limits the extent to which a player i 's report in round t affects $H'_{i,s}$ for each $s > t$, and hence limits how it affects future choices of the mechanism, this allows us to control the amount of future utility she can gain from misreporting at t . Consider the change in this player's utility in all rounds $s > t$ that results from changing her bid from her true value $v_{i,t}$ to any other misreport $v'_{i,t}$. Let Y be the event that she bids truthfully in round t , and let \bar{Y} be the event that she misreports. Let $u(Y)$ and $u(\bar{Y})$ respectfully be the total utility the bidder achieves in all future rounds conditioned on events Y and \bar{Y} . Let \mathcal{S} be the set of all possible outcomes (i.e., allocations and payments) from all future rounds of the auction that this bidder may participate in, and let $w(s)$ be the utility the bidder from outcome $s \in \mathcal{S}$. We

can now bound the player's gain in expected utility from lying by bounding the expected value of $u(\bar{Y})$.

$$\begin{aligned}
\mathbf{E}[u(\bar{Y})] &= \int_{\mathcal{S}} w(s) \mathbf{P}[s|\bar{Y}] ds \\
&\leq \int_{\mathcal{S}} w(s) (e^\epsilon \mathbf{P}[s|Y] + \frac{\epsilon}{T}) ds \\
&\leq \int_{\mathcal{S}} w(s) ((1 + 2\epsilon) \mathbf{P}[s|Y] + \frac{\epsilon}{T}) ds \\
&= (1 + 2\epsilon) \mathbf{E}[u(Y)] + \int_{\mathcal{S}} w(s) \frac{\epsilon}{T} ds \\
&\leq \mathbf{E}[u(Y)] + 2\epsilon kh + kh \frac{\epsilon}{T}.
\end{aligned}$$

The first inequality follows from the (ϵ, δ) -DP guarantee of Theorem 12, the second from the fact that $e^\epsilon \leq (1 + 2\epsilon)$ for $\epsilon < 1$, and the final inequality from the fact that each bidder participates in at most k rounds and her maximum utility is any round is h , so both $\mathbf{E}[u(Y)]$ and $\int_{\mathcal{S}} w(s) ds$ are upper bounded by kh . Thus the maximum change in utility over all future rounds between any two behaviors in the current round is therefore $2\epsilon kh + kh \frac{\epsilon}{T} = kh\epsilon(2 + \frac{1}{T})$.

Thus, the overall utility the bidder might gain from misreporting in round t is $kh\epsilon(2 + \frac{1}{T})$, which converges to $2kh\epsilon$ as $T \rightarrow \infty$. □

A.4 Revenue Analysis for Algorithm 1

To prove the main revenue guarantee of Utility-Approximate BIC Online Auction presented in Algorithm 1, i.e Theorem 14 of our, we present a few lemmas to bound the difference in revenue obtained in each round and then sum it over the T rounds to bound the average expected revenue.

In each round, Algorithm 1 runs the optimal mechanism for \mathbf{H}'_t , but discretized value profiles are sampled from \mathbf{D}' . The following lemma uses tools from the previous subsection to show that the expected revenue of running $\mathcal{M}_{\mathbf{H}'_t}^*$ on samples from \mathbf{H}'_t is not much worse than running $\mathcal{M}_{\mathbf{D}'}^*$ on samples from \mathbf{D}' .

Lemma 32. *With probability at least $1 - \alpha$, $|\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{H}'_t) - \text{Rev}(\mathcal{M}_{\mathbf{D}'}^*; \mathbf{D}')| \leq 2hn^2\gamma_t$.*

Proof. We start by re-writing the revenue difference we wish to bound, $|\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{H}'_t) - \text{Rev}(\mathcal{M}_{\mathbf{D}'}^*; \mathbf{D}')|$ as follows,

$$|(\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{H}'_t) - \text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{D}')) + (\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{D}') - \text{Rev}(\mathcal{M}_{\mathbf{D}'}^*; \mathbf{D}'))|.$$

For the first term inside the absolute value, Theorem 24 says that with probability at least $1 - \alpha$, $\|D'_j - H'_{j,t}\|_\infty \leq \gamma_t$ for all j , and therefore D'_j and $H'_{j,t}$ are γ_t -close for all j . Applying Theorem 27 gives that $\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{H}'_t) - \text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{D}') \leq 3hn^2\gamma_t$ with the same probability. The second term is 0 because $\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{D}') \leq \text{Rev}(\mathcal{M}_{\mathbf{D}'}^*; \mathbf{D}')$, since $\mathcal{M}_{\mathbf{D}'}^*$ is the revenue-optimal mechanism for \mathbf{D}' . Therefore, we get $\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{H}'_t) - \text{Rev}(\mathcal{M}_{\mathbf{D}'}^*; \mathbf{D}') \leq 2hn^2\gamma_t$.

A symmetric argument using $\mathcal{M}_{\mathbf{D}'}^*$, gives that $\text{Rev}(\mathcal{M}_{\mathbf{D}'}^*; \mathbf{D}') - \text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{H}'_t) \leq 2hn^2\gamma_t$, which completes the proof. □

Now we present a result from Devanur et al. [17] (generalized to multiple item auction) which states that discretization of the value space by rounding down to nearest multiple of β only reduces the optimal revenue by an additive factor of $J\beta$ for a J -item auction. Intuitively, since bids are always rounded down, this can result in a loss of at most β revenue from each of the J items.

Lemma 33 (Devanur et al. [17]). *$\text{Rev}(\mathcal{M}_{\mathbf{D}'}^*; \mathbf{D}') \geq \text{Rev}(\mathcal{M}_{\mathbf{D}}^*; \mathbf{D}) - \beta J$.*

Combining these results, we can now bound the expected revenue of our mechanism for a fixed round.

Lemma 34. *With probability at least $1 - \alpha$, the expected revenue obtained by Algorithm 1 in the t^{th} round, $\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{D}')$, satisfies,*

$$\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{D}') \geq \text{Rev}(\mathcal{M}_{\mathbf{D}}^*; \mathbf{D}) - \beta J - 4hn^2\gamma_t,$$

for $\gamma_t = \sqrt{\frac{\log \frac{n}{\alpha}}{2t}} + \frac{\sigma}{t} \sqrt{\log \frac{h}{\beta} \log T} \sqrt{2 \log \left(\frac{2hn}{\beta\alpha} \right)}$ and $\sigma = \frac{8 \log T \log \frac{h}{\beta}}{\epsilon} \sqrt{\ln \frac{T \log T \log \frac{h}{\beta}}{\epsilon}}$.

Proof. Using Lemma 32 gives,

$$\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{D}') \geq \text{Rev}(\mathcal{M}_{\mathbf{D}'}^*; \mathbf{D}') - 4hn^2\gamma_t,$$

and applying Lemma 33 gives,

$$\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{D}') \geq \text{Rev}(\mathcal{M}_{\mathbf{D}}^*; \mathbf{D}) - \beta J - 4hn^2\gamma_t.$$

□

Now that we have bounded the expected revenue in a fixed round t , we can bound the average revenue over T rounds with a union bound over all the rounds to obtain a guarantee for the average expected revenue of our Algorithm 1 in Theorem 14.

B Omitted Proofs and Subroutines for Bid-approximate BIC online auctions (Algorithm 2)

B.1 Subroutines used in Algorithm 2

B.1.1 Strictly Truthful mechanism

To ensure that underbidding is costly for the bidders in the current round t , with some probability ρ_t , we use a naïve mechanism which is strictly truthful. We simply select a random subset S of J bidders and charge them a uniformly random price p between 0 and h . If bidder i is selected in the set S and has bid $b_{i,t}$ is above p , they get the item and pay p . Otherwise, they pay 0 and are not allocated an item.

Algorithm 4: Strictly Truthful

Input: Bid profile \mathbf{b}_t
 Select a subset $S \subseteq [n]$ of size J uniformly at random
 Select a price $p \in [0, h]$ uniformly at random
for Each $s \in S$ **do**
 | **if** $b_{s,t} \geq p$ **then** allocate item to s and charge payment p ;
end

As the maximum possible revenue in a single round is hJ , the expected loss in revenue from Algorithm 4 in round t of the mechanism is at most $\rho_t hJ$.

B.1.2 Private Payments from Black Box Payments

We now describe how to privately compute payments to charge buyers in round t . Our goals for this payment computation are first, to ensure our round t mechanism charges payments according to Theorem 2 (ensuring incentive compatibility for the one-round game), and second, for player i 's round t payment to be differentially private in all other round t bidders' bids.

In order to achieve the latter goal of private payments, we select the round t payment for each winning bidder i as a function of $b_{i,t}$, their bid, the private estimated distributions of other bidder types \mathbf{H}'_t , and the allocation algorithm \mathcal{A} , but *not* as a function of the other bidders' round t bids. This ensures the payment of bidder i is 0 jointly differentially private in round t bids, and (ϵ, δ) -differentially private in all previous rounds' bids. Concretely, we use black box payments (Algorithm 5) introduced in Archer et al. [2].

Critically, bidder i 's payment does not depend on the bid of any other player, save for the extent to which the allocation is affected by other the other players' bids. For a differentially private allocation rule, this implies that player i 's payment is differentially private in all other players' bids. i therefore learns almost nothing about other players' behavior from their payment in round t .

Algorithm 5: Black box payments

Input: Value distribution \mathbf{H}'_t , Allocation algorithm \mathcal{A} , agent i 's bid $b_{i,t}$
if algorithm \mathcal{A} does not allocate to bidder i **then** agent i 's payment $p_i \leftarrow 0$;
else
 Choose b'_i uniformly from $[0, b_{i,t}]$
 Draw $\mathbf{b}'_{-i} \sim H'_{-i}$ and run $\mathcal{A}(b'_i, \mathbf{b}'_{-i})$
 if algorithm \mathcal{A} allocates to bidder i in the previous step **then** $X \leftarrow b_{i,t}$;
 else $X \leftarrow 0$;
 if $X \neq 0$ **then**
 repeatedly draw values $\mathbf{b}'_{-i} \sim \mathbf{H}'_{-i,t}$ and run $\mathcal{A}(b'_i, \mathbf{b}'_{-i})$ until the algorithm allocates to
 player i , and let L be the number of iterations required
 end
 Agent i 's payment $p_{i,t} \leftarrow b_{i,t} - L \cdot X$
end

We make one remark about this variant of black-box payments giving the formal guarantees we will use to analyze our mechanism. These black-box payments are computed from the (strategically) shaded, differentially private estimate of the empirical bid distribution, rather than the true, unaltered distribution, since (a) the empirical, shaded distribution is the best estimate the mechanism has, and (b) maintaining only a differentially private estimate of this distribution ensures that no single bidder's prior rounds' behavior significantly change these payments.

We now present the lemmas about Algorithm 5 which help us analyze the revenue and (approximate) incentive compatibility of the overall mechanism.

Fixing the distribution \mathbf{H}'_t , the expected payment for bidder i given her value $b_{i,t}$ will be $p_{i,t}(b_{i,t}) = \mathbf{E}_{\mathbf{b}_{-i}}[p_{i,t}(b_{i,t}, \mathbf{b}_{-i})]$, and the probability she is allocated an item using the allocation mechanism \mathcal{A} is $x_{i,t}(b_{i,t}) = \mathbf{E}_{\mathbf{b}_{-i}}[x_{i,t}(b_{i,t}, \mathbf{b}_{-i})]$. We have the following lemma.

Lemma 35 (Archer et al. [2]). *Fixing the distribution \mathbf{H}'_t and the allocation mechanism \mathcal{A} , the expected payment $p_{i,t}(b_{i,t})$ returned by Black Box payments (Algorithm 5) satisfies*

$$p_{i,t}(b_{i,t}) = b_{i,t} \cdot x_{i,t}(b_{i,t}) - \int_0^{b_{i,t}} x_{i,t}(u) du$$

This property about the payments calculated by Black Box payments directly helps us to translate the expected virtual surplus from the allocation rule i.e the Exponential Mechanism to the expected revenue.

Lemma 36. *Fix a distribution \mathbf{H}'_t . Then, allocating according to the exponential mechanism \mathcal{M}_E and charging black-box payments yields expected revenue on \mathbf{H}'_t equal to the expected virtual surplus of the allocation selected by \mathcal{M}_E .*

Proof. The exponential mechanism is monotone, by Lemma 40. Black-box payments are designed to satisfy the payment identity, and so the two together have expected revenue equal to their expected virtual surplus, by Theorem 2. \square

B.1.3 Private, efficient allocation via the Exponential mechanism

We now describe the exponential mechanism, and how we use it to select an allocation in round t . The primary difference between this choice and many other uses of the exponential mechanism is that we find a way to select allocations with it in a computationally tractable way (details to follow). Let \mathcal{X} be the set of possible allocations. For any mechanism \mathcal{M} , we need an allocation rule \mathcal{A} , a function that takes as input the bid profile \mathbf{b} , the value distribution H' , and returns an allocation $\mathbf{x}(\mathbf{b}) \in \mathcal{X}$. \mathcal{A} can be a randomized algorithm, in which case $x_i(\mathbf{b})$ is a random variable indicating the probability of allocation.

In general, the exponential mechanism runs in time polynomial in the number of elements it selects amongst. If one is using it to select an allocation of J items to some set of bidders, this will often result in exponential runtime.

We use the exponential mechanism to select an allocation in a slightly nonstandard way, in that the quality score we use to measure the utility of outcomes is only an approximation of what we would ideally like to maximize.

We use the private, learned distributions to define estimated virtual values for each bid, and use those estimated virtual values as the arguments to the exponential mechanism. Let $\hat{\phi}(\mathbf{b})$ be the virtual value profile given as input to the exponential mechanism. The exponential mechanism is parameterized by a *quality score* $Q(\hat{\phi}(\mathbf{b}), \mathbf{x})$, a function mapping inputs and outputs of the desired optimization task to a measurement of the outcome's approximation to the optimization task on the input. In our mechanism, we select the quality score for the Exponential Mechanism \mathcal{M}_E to be the estimated virtual welfare of the allocation on the reported bids:

$$Q(\hat{\phi}(\mathbf{b}), \mathbf{x}) = \sum_{i=0}^n \hat{\phi}_i(b_i)x_i.$$

If \mathbf{b} and \mathbf{b}' are two neighboring bid profiles, i.e they differ in at most one bidder's bid, we define the sensitivity Δ of the quality score as maximum change in quality score over all pairs of neighbouring bid profiles \mathbf{b} and \mathbf{b}' and all possible feasible allocations:

$$\Delta = \max_{\mathbf{x} \in \mathcal{X}, \mathbf{b}, \mathbf{b}'} \left| Q(\hat{\phi}(\mathbf{b}), \mathbf{x}) - Q(\hat{\phi}(\mathbf{b}'), \mathbf{x}) \right|$$

Given a virtual value profile $\phi(\mathbf{b})$ the Exponential Mechanism $\mathcal{M}_E(\mathcal{X}, Q, \phi(\mathbf{b}), \epsilon)$ chooses an allocation $\mathbf{x} \in \mathcal{X}$ with probability

$$\mathbb{P} \left[\mathcal{M}_E(\mathcal{X}, Q, \hat{\phi}(\mathbf{b}), \epsilon) = \mathbf{x} \right] \propto \exp \frac{\epsilon Q(\hat{\phi}(\mathbf{b}), \mathbf{x})}{\Delta} \leq h$$

where the upper bound follows from the fact that as one bidder changing her bid can affect the revenue of any allocation by at most h . Since expected revenue is equal to the expected virtual welfare, then the exponential mechanism designed to maximize expected estimated virtual welfare on the reported bids will also approximately maximize expected revenue. We now state the main theorem about the exponential mechanism we will use in what follows.

Theorem 37 (McSherry and Talwar [33]). *The Exponential Mechanism $\mathcal{M}_E(\mathcal{X}, Q, \hat{\phi}(\mathbf{b}), \epsilon)$ is ϵ -differentially private, and for any $\alpha > 0$,*

$$\Pr[Q(\hat{\phi}(\mathbf{b}), \mathcal{M}_E(\mathcal{X}, Q, \hat{\phi}(\mathbf{b}), \epsilon)) \geq \max_{\mathbf{x} \in \mathcal{X}} Q(\hat{\phi}(\mathbf{b}), \mathbf{x}) - \frac{\ln|\mathcal{X}|}{\epsilon} - \frac{h}{\epsilon} \ln \left(\frac{1}{\rho} \right)] \geq 1 - \rho$$

This result holds with high probability for a value bid profile \mathbf{b} . Below we state a common corollary of this statement, which allows us to quantify the expected revenue loss over the mechanism's randomness:

Lemma 38. *For any virtual bid profile $\hat{\phi}(\mathbf{b})$ and $\rho \in (0, 1)$*

$$\max_{\mathbf{x} \in \mathcal{X}} Q(\hat{\phi}(\mathbf{b}), \mathbf{x}) - \mathbf{E}_{\mathcal{M}_E} \left[Q(\hat{\phi}(\mathbf{b}), \mathcal{M}_E(\mathcal{X}, Q, \hat{\phi}(\mathbf{b}), \epsilon)) \right] \leq \frac{J \ln n}{\epsilon} + \frac{h}{\epsilon} \ln \frac{1}{\rho} + \rho J h.$$

Proof. Using Theorem 37, we have

$$\begin{aligned} & \max_{\mathbf{x} \in \mathcal{X}} Q(\hat{\phi}(\mathbf{b}), \mathbf{x}) - \mathbf{E}_{\mathcal{M}_E} \left[Q(\hat{\phi}(\mathbf{b}), \mathcal{M}_E(\mathcal{X}, Q, \hat{\phi}(\mathbf{b}), \epsilon)) \right] \\ & \leq (1 - \rho) \left(\frac{\ln|\mathcal{X}|}{\epsilon} + \frac{h}{\epsilon} \ln \frac{1}{\rho} \right) + \rho J h \\ & \leq \frac{J \ln n}{\epsilon} + \frac{h}{\epsilon} \ln \frac{1}{\rho} + \rho J h \end{aligned}$$

□

The trivial way to implement $\mathcal{M}_E(\mathcal{X}, Q, \hat{\phi}(\mathbf{b}), \epsilon)$ would be to iterate over all $x \in \mathcal{X}$ to calculate $Q(\hat{\phi}(\mathbf{b}), \mathbf{x})$.

Since we are running a J -item auction in each round t , our feasible allocations are the set

$$\mathcal{X} = \{ \mathbf{x} : \mathbf{x} \in \{0, 1\}^n, \|\mathbf{x}\|_1 = J \},$$

which has size $|\mathcal{X}| = \binom{n}{J} \leq n^J$. Thus, explicitly iterating over \mathcal{X} would result in an $\mathcal{O}(n^J)$ algorithm.

Algorithm 6: Poly-time Exponential Mechanism for selecting an allocation of J goods ($\mathcal{M}_{polyE}(\phi(\mathbf{b}), \epsilon)$)

Parameters: Privacy parameter ϵ , Sensitivity Δ , Number of bidders to select J

Input: Virtual bid profile $\hat{\phi}(\mathbf{b})$

$S_0 \leftarrow \{\}$

for $i = 1, \dots, n$ **and** $|S| < J$ **do**

$$\alpha_i = \frac{\exp\left(\frac{\epsilon \hat{\phi}_i(b_i)}{\Delta}\right) \sum_{A \in \binom{\{i+1, \dots, n\}}{J-|S_{i-1}|-1}} \left(\exp\left(\frac{\epsilon}{\Delta} \sum_{a \in A} \hat{\phi}_a(b_a)\right)\right)}{\sum_{B \in \binom{\{i+1, \dots, n\}}{J-|S_{i-1}|}} \left(\exp\left(\frac{\epsilon}{\Delta} \sum_{b \in B} \hat{\phi}_b(b_b)\right)\right)}$$

With probability α_i **do** $S_i \leftarrow S_{i-1} \cup \{i\}$ **else** $S_i \leftarrow S_{i-1}$

end

Return binary allocation \mathbf{x} such $x_i = 1$ for all $i \in S_n$

However, since our feasible set \mathcal{X} has a very simple structure (a set of n length binary vector with $\leq J$ ones), we can exploit this structure to select a subset of size J in polynomial time by carefully calculating the marginal probabilities of selecting individual bidders such that the resulting joint probabilities of the subsets is same as in $\mathcal{M}_E(\mathcal{X}, Q, \phi(\mathbf{b}), \epsilon)$.

To take advantage of this structure, we calculate the marginal probability of selecting bidder i conditioned on the bidders we have already sampled or rejected. Calculating these marginal probability can be done using a simple dynamic programming solution in $\mathcal{O}(nJ)$. We select or reject each bidder in turn, leading to a total running time of for Algorithm 6 of $\mathcal{O}(n^2J)$. We now prove that the allocation selected by Algorithm 6 has the same distribution as $\mathcal{M}_E(\mathcal{X}, Q, \phi(\mathbf{b}), \epsilon)$. Let the mechanism defined in Algorithm 6 be $\mathcal{M}_{polyE}(\phi(\mathbf{b}), \epsilon)$.

Lemma 39. $\mathbb{P}\left[\mathcal{M}_{polyE}(\hat{\phi}(\mathbf{b}), \epsilon) = \mathbf{x}\right] = \mathbb{P}\left[\mathcal{M}_E(\mathcal{X}, Q, \hat{\phi}(\mathbf{b}), \epsilon) = \mathbf{x}\right] = \frac{\exp\left(\frac{\epsilon Q(\hat{\phi}(\mathbf{b}), \mathbf{x})}{\Delta}\right)}{\sum_{\mathbf{x}' \in \mathcal{X}} \exp\left(\frac{\epsilon Q(\hat{\phi}(\mathbf{b}), \mathbf{x}')}{\Delta}\right)}$

Proof. Let

$$\mathbb{P}[\mathbf{x}] = \frac{\exp\left(\frac{\epsilon \sum_{i=0}^n \hat{\phi}_i(b_i) x_i}{\Delta}\right)}{\sum_{\mathbf{x}' \in \mathcal{X}} \exp\left(\frac{\epsilon \sum_{i=0}^n \hat{\phi}_i(b_i) x'_i}{\Delta}\right)}.$$

Note that α_i in Algorithm 6 is the marginal conditional probability $\mathbb{P}[x_i = 1 | x_1, \dots, x_{i-1}]$. Thus we set $x_i = 1$ with the corresponding conditional marginal. We conclude the proof using the chain rule of joint probability distributions, which states that

$$\mathbb{P}[x_1, x_2, \dots, x_n] = \mathbb{P}[x_1] \mathbb{P}[x_2 | x_1] \cdots \mathbb{P}[x_n | x_1, \dots, x_{n-1}].$$

□

The following is a well-understood fact, but we include a short proof for completeness.

Fact 40. For downward-closed environments, the Exponential Mechanism instantiated with any monotone virtual value function as its quality score has a monotone allocation rule.

Proof of 40. Fix the behavior \mathbf{v}_{-i} of players other than i . The probability the Exponential Mechanism selects allocation $a \in \mathcal{X}$ is

$$\frac{e^{(\epsilon \sum_{i \in a} \phi_i(v_i)) / \Delta}}{\sum_{a' \in \mathcal{X}} e^{(\epsilon \sum_{i \in a'} \phi_i(v_i)) / \Delta}}$$

We can view allocation $a \in \mathcal{X}$ as a set of bidders who will receive the item. Then, the probability that player i is in the allocated set is,

$$\sum_{a \ni i} \frac{e^{(\epsilon \sum_{i' \in a} \phi_{i'}(v_{i'})) / \Delta}}{\sum_{a' \in \mathcal{X}} e^{(\epsilon \sum_{i' \in a'} \phi_{i'}(v_{i'})) / \Delta}}$$

thus, the numerator increases for each allocation a containing i because ϕ_i is monotone, and because $\frac{a+x}{a+x+y} > \frac{x}{x+y}$, this probability increases in v_i . □

Now recall that we defined $\text{Rev}(\mathcal{M}; D)$ to be the expected revenue generated by the mechanism \mathcal{M} on a value distribution D . Let the mechanism we use in round t for allocation and pricing be $\mathcal{M}_{\mathbf{H}'_t}$ and let $\mathcal{M}_{\mathbf{H}'_t}^*$ be the optimal mechanism on \mathbf{H}'_t . Using Lemma 36, we can see that $\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{H}'_t) = \mathbf{E}_{\mathbf{b} \sim \mathbf{H}'_t} [\max_{\mathbf{x} \in \mathcal{X}} Q(\phi_t(\mathbf{b}), \mathbf{x})]$ where $\phi_t(\mathbf{b})$ calculates the virtual bid using \mathbf{H}'_t as prior. Combining this guarantees with the guarantee over the virtual surplus given by exponential mechanism, we get the following lemma.

Lemma 41. $\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}; \mathbf{H}'_t) \geq \text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{H}'_t) - (\frac{J \ln n}{\epsilon} + \frac{h}{\epsilon} \ln \frac{1}{\rho_t} + \rho_t J h) - \rho_t h J$.

Proof. The lemma is directly a consequence of taking expectation of over \mathbf{b} in Lemma 38 with respect to \mathbf{H}'_t with $\rho = \rho_t$ and then using Lemma 36. The last term $\rho_t h J$ comes from the fact in round t , we use the strictly truthful mechanism (Algorithm 4) with probability ρ_t and lose up to $h J$. \square

B.2 Differentially Private Bid Distribution Estimation

We now present a lemma which compares how close our differentially private empirical bid distribution estimates are to the true bid distributions, under this assumption that the bid distributions do not change too quickly.

Lemma 42. *Assume we have a sequence of distributions $F'_{i,t}$ with support on $[0, h]$, and assume that for every t we have $\|F'_{i,t} - F'_{i,t+1}\|_\infty \leq \lambda_t$ for some sequence $\lambda_1, \lambda_2, \dots$ satisfying $\Sigma_T := \sum_{t=1}^{T-1} \frac{t \lambda_t}{T} = o(1)$. Let $b_{i,t} \sim F'_{i,t}$ and let $H_{i,T}$ be the empirical distribution of $b_{i,1}, \dots, b_{i,T}$. Then with probability at least $1 - \alpha$ we have*

$$\|H_{i,T} - F'_{i,T}\|_\infty \leq \sqrt{\frac{\log(2h/(\beta\alpha))}{2T}} + \Sigma_T.$$

Proof. Let \bar{F} be the distribution constructed from the uniform mixture of the distributions $F'_{i,1}, \dots, F'_{i,T}$, that is $\bar{F} := \frac{1}{T} \sum_{t=1}^T F'_{i,t}$. Observe that, via repeated application of the triangle inequality,

$$\begin{aligned} \|H_{i,T} - F'_{i,T}\|_\infty &\leq \|H_{i,T} - \bar{F}\|_\infty + \|\bar{F} - F'_{i,T}\|_\infty \\ &\leq \|H_{i,T} - \bar{F}\|_\infty + \frac{1}{T} \sum_{t=1}^T \|F'_{i,t} - F'_{i,T}\|_\infty \\ &\leq \|H_{i,T} - \bar{F}\|_\infty + \frac{1}{T} \sum_{t=1}^{T-1} \sum_{j=t}^{T-1} \|F'_{i,j} - F'_{i,j+1}\|_\infty \\ &\leq \|H_{i,T} - \bar{F}\|_\infty + \frac{1}{T} \sum_{t=1}^{T-1} \sum_{j=t}^{T-1} \lambda_j = \|H_{i,T} - \bar{F}\|_\infty + \Sigma_T. \end{aligned}$$

Thus, to complete the proof, we must simply show that with probability at least $1 - \alpha$ that $\|H_{i,T} - \bar{F}\|_\infty \leq \sqrt{\frac{\log(2h/(\beta\alpha))}{2T}}$. To do this, first select some $u \in [0, h]$ and notice that

$$\mathbf{E}[H_{i,T}(u)] = \mathbf{E}\left[\frac{1}{T} \sum_{t=1}^T \mathbf{1}[b_{i,t} \leq u]\right] = \frac{1}{T} \sum_{t=1}^T \mathbf{P}[b_{i,t} \leq u] = \frac{1}{T} \sum_{t=1}^T F'_{i,t}(u) = \bar{F}(u);$$

in other words, $H_{i,T}(u)$ is an unbiased estimator of $\bar{F}(u)$. Since each indicator variable $\mathbf{1}[b_{i,t} \leq u]$ is bounded in $[0, 1]$, we may apply Hoeffding's Inequality to obtain

$$\mathbf{P}\left[\left|\frac{1}{T} \sum_{t=1}^T \mathbf{1}[b_{i,t} \leq u] - \bar{F}(u)\right| > \xi\right] \leq 2 \exp(-2T\xi^2). \quad (3)$$

Observe that, since $H_{i,T}$ and \bar{F} are step functions which only change at points $U = \{\beta, 2\beta, \dots, \frac{h}{\beta}\beta\}$, it holds for any $\xi > 0$ that

$$\|H_{i,T} - \bar{F}\|_\infty \geq \xi \iff \exists u \in [0, h] : |H_{i,T}(u) - \bar{F}(u)| \geq \xi \iff \exists u \in U : |H_{i,T}(u) - \bar{F}(u)| \geq \xi.$$

Thus we have that

$$\begin{aligned}
\mathbb{P} \left[\|H_{i,T} - \bar{F}\|_\infty \geq \xi \right] &= \mathbb{P} \left[\exists u \in U : \left| \frac{1}{T} \sum_{t=1}^T \mathbf{1}[b_{i,t} \leq u] - \bar{F}(u) \right| > \xi \right] \\
&\leq \sum_{u \in U} \mathbb{P} \left[\left| \frac{1}{T} \sum_{t=1}^T \mathbf{1}[b_{i,t} \leq u] - \bar{F}(u) \right| > \xi \right] \\
\text{(using (3))} \quad &\leq \frac{2h}{\beta} \exp(-2T\xi^2).
\end{aligned}$$

The proof is completed by setting α equal to the final expression and solving for the value ξ . \square

The next theorem states that our differentially private estimates of $H_{i,t}$ are fairly close to the true distributions.

Theorem 43. *After t rounds Algorithm 2, it holds with probability at least $1 - \alpha$ that*

$$\|H'_{i,t} - F'_{i,t}\|_\infty \leq \gamma_t \quad \text{for every } i \in [n],$$

$$\text{where } \gamma_t = \sqrt{\frac{\log(2hn/\beta\alpha)}{2t}} + \Sigma_t + \frac{\sigma}{t} \sqrt{\log \frac{h}{\beta} \log T} \sqrt{\log \left(\frac{hn}{\beta\alpha} \right)} \text{ and } \sigma = \frac{8 \log T \log \frac{h}{\beta}}{\epsilon} \sqrt{\ln \frac{\log T \log \frac{h}{\beta}}{\delta}}.$$

Proof. We assume the same Σ_t for each bidder type. Next, apply a union bound on Lemma 42 for all players and use the triangle inequality with Lemma 23. The result is analogous to Theorem 24. \square

The next lemma states that the distribution over critical prices induced by our learned distribution is close to that induced by the true bid distribution.

Lemma 44. *For a bidder i , let $G_{i,t}$ be the distribution of critical prices offered to i when the distribution of other bidders is $\mathbf{F}'_{-i,t}$ and let $G'_{i,t}$ be the distribution of critical prices when the distribution of other bidders is $\mathbf{H}'_{-i,t}$, then $\|G_{i,t} - G'_{i,t}\|_\infty \leq (n-1)\gamma_t$.*

Proof. From Theorem 43, we have that $\mathbf{F}'_{-i,t}$ and $\mathbf{H}'_{-i,t}$ are γ_t close. The exponential mechanism is a monotone allocation rule (Fact 40). The lemma then follows from Lemma 29. \square

B.3 Truthfulness Guarantees of Algorithm 2

Lemma 45. *Let $r_{i,t}(v)$ be the true BIC payment (from Theorem 2) for bidder i in round t when value is v and let $p_{i,t}(v)$ be the expected payment charged by Algorithm 1 by using Black Box payments (Algorithm 5), then $|r_{i,t}(v) - p_{i,t}(v)| \leq 2hn\gamma_t$.*

Proof. The proof follows from the fact the Black Box payments in round t calculate the BIC in expectation with respect to $H'_{i,t}$ and from using Lemma 30. \square

We now use these results to prove that our Online DP auction satisfies differential privacy.

Lemma 46. *Fix a round t and bidder population i . Then, consider the case when the mechanism runs the exponential allocation and black-box payments. Then, truthful reporting for the bidder from population i present in round t earns within $6hn\gamma_t$ of the optimal utility they could achieve.*

Recall that $r_{i,t}(v)$ is the true BIC payment and $p_{i,t}(v)$ be the expected payment charged by Algorithm 1. Let $v_{i,t}$ be the true value of the bidder i in round t , and let v' be some misreport. Since $r_{i,t}(v)$ is the correct BIC payments, the maximum utility bidder i could have received if the payments charged were BIC be $u(v_{i,t}) = v_{i,t}x_{i,t}(v_{i,t}) - r_{i,t}(v_{i,t}) \geq u(v')x_{i,t}(v') - r_{i,t}(v')$, but the actual utility achieved by the bidder in the current as a function of their bid is $u'(v) = vx_{i,t}(v) - p_{i,t}(v)$.

From Lemma 45, we have that $|r_{i,t}(v) - p_{i,t}(v)| \leq 2hn\gamma_t$. So if the bidder bids v' instead of $v_{i,t}$, then the utility achieved by i in round t is

$$\begin{aligned} u'(v') &= v'x_{i,t}(v') - p_{i,t}(v') \\ &\leq v'x_{i,t}(v') - r_{i,t}(v') + 2hn\gamma_t \\ &\leq v_{i,t}x_{i,t}(v_{i,t}) - r_{i,t}(v_{i,t}) + 2hn\gamma_t \\ &\leq v_{i,t}x_{i,t}(v_{i,t}) - p_{i,t}(v_{i,t}) + 4hn\gamma_t \\ &= u'(v_{i,t}) + 4hn\gamma_t. \end{aligned}$$

The next lemma shows that a bidder cannot gain a lot in future rounds by manipulating her bid in the current round.

Lemma 47. *In any round t , for bidder i with value $v_{i,t}$, let $u_i(w)$ be their expected utility in the current round as a function of their bid w . If they report bid $b_{i,t}$ such that $b_{i,t} \leq v_{i,t} - \eta$, then*

$$u_i(b_{i,t}) \leq u_i(v_{i,t}) - \frac{\rho_t J \eta^2}{2hn} + 4hn\gamma_t.$$

Proof. With probability ρ_t we use algorithm 4, and if we use algorithm 4 then with probability $\frac{J}{n}$, the bidder i is selected for allocation is offered a random price $p \in [0, h]$. If $b_{i,t} = v_{i,t} - y$ where $y \geq \eta$, then with probability at least $\frac{\eta}{h}$, $p \in [b_{i,t}, v_{i,t}]$, i.e bidder i loses utility in this case, which they could have gained if they report their value truthfully. If $p \in [b_{i,t}, v_{i,t}]$, and bidder i is selected then the utility lost by bidder i by not bidding truthfully is $v_{i,t} - p$. With probability $(1 - \rho_t)$, we use exponential mechanism and black box payments from lemma 46, the gain in utility from underbidding in this case be at most $4hn\gamma_t$. Combining the two cases, we can get a lower bound on the loss of utility in the current round

$$\begin{aligned} u_i(v_{i,t}) - u_i(b_{i,t}) &\geq \frac{\rho_t J}{n} \left(\int_{b_{i,t}}^{v_{i,t}} (v_{i,t} - p) \frac{1}{h} dp \right) - (1 - \rho_t) 4hn\gamma_t \\ &\geq \frac{\rho_t J (v_{i,t} - b_{i,t})^2}{2hn} - 4hn\gamma_t \\ &\geq \frac{\rho_t J \eta^2}{2hn} - 4hn\gamma_t. \end{aligned}$$

□

Lemma 47 shows that that the bidder can lose utility in the current round if they underbid because of the penalty imposed by the strictly truthful mechanism (Algorithm 4).

Now we have bounded all three sources of change in utility for shading one's bid: how much they can gain from future rounds if they underbid today, how much they are penalized by Algorithm 4 for underbidding, and how much they can gain from current round by underbidding. Combing these three results, we can now prove that the mechanism in round t is η_t -bid approximate BIC (Definition 7).

Proof for Theorem 19. If the $v_{i,t} - b_{i,t} > \eta_t = h\sqrt{\frac{8n^2\gamma_t + 6k\epsilon}{\rho_t J}}$, then from Lemma 47 the loss in utility in current round is

$$\geq \frac{\rho_t J \eta_t^2}{2hn} - 4hn\gamma_t = 4hn\gamma_t + 3k\epsilon - 4hn\gamma_t \geq k\epsilon(2 + \frac{1}{T})$$

and from Lemma 17, $k\epsilon(2 + \frac{1}{T})$ is the maximum gain in utility the player can achieve in future rounds by underbidding. Hence any bid $b_{i,t} < v_{i,t} - \eta_t$ cannot be an equilibrium bid because the total gain in utility in future rounds from underbidding in the current round is less than the loss in utility in current round because of underbidding by more than η_t . □

B.4 Revenue Analysis for Algorithm 2

In this section, we bound the expected revenue achieved by our algorithm in some exact equilibrium compared to the optimal expected revenue facing truthful bids, $\text{Rev}(\mathcal{M}_{\mathbf{D}}^*; \mathbf{D})$. In Algorithm 2, with probability ρ_t , we execute the strictly truthful mechanism (Algorithm 4) which is also has monotone allocation rule. As stated earlier, let the mechanism we run in round t be $\mathcal{M}_{\mathbf{H}'_t}$, it is clear that the allocation rule used in $\mathcal{M}_{\mathbf{H}'_t}$ is monotone. In round t , the bids received by Algorithm 2 come from \mathbf{F}'_t . Let $\mathcal{M}_{\mathbf{F}'_t}^*$ be the optimal mechanism on \mathbf{F}'_t . As we have shown that our Differentially private estimate is close to \mathbf{F}'_t , we can use the theory developed in subsection A.2 to bound the revenues. Particularly, we have the following lemma.

Lemma 48. *With probability at least $1 - \alpha$, $|\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}; \mathbf{H}'_t) - \text{Rev}(\mathcal{M}_{\mathbf{H}'_t}; \mathbf{F}'_t)| \leq 2hn^2\gamma_t$.*

Proof. Theorem 43 implies that for all j , $\|F'_{j,t} - H'_{j,t}\|_\infty \leq \gamma_t$; As $\mathcal{M}_{\mathbf{H}'_t}$ has a monotone allocation rule, applying Theorem 27 completes the proof. \square

We showed that $\mathcal{M}_{\mathbf{H}'_t}$ has similar revenue on both \mathbf{F}'_t and \mathbf{H}'_t , to prove our required bound, we also need to show that the optimal revenue on these distributions is also.

Lemma 49. *With probability at least $1 - \alpha$, $|\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{H}'_t) - \text{Rev}(\mathcal{M}_{\mathbf{F}'_t}^*; \mathbf{F}'_t)| \leq 2hn^2\gamma_t$.*

Proof. Using similar arguments as 48, $\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{H}'_t) - \text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{F}'_t) \leq 2hn^2\gamma_t$, and since $\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{F}'_t) \leq \text{Rev}(\mathcal{M}_{\mathbf{F}'_t}^*; \mathbf{F}'_t)$, we get $\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{H}'_t) - \text{Rev}(\mathcal{M}_{\mathbf{F}'_t}^*; \mathbf{F}'_t) \leq 2hn^2\gamma_t$. Using same argument on $\mathcal{M}_{\mathbf{F}'_t}^*$, we get $\text{Rev}(\mathcal{M}_{\mathbf{D}'_t}^*; \mathbf{F}'_t) - \text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{H}'_t) \leq 2hn^2\gamma_t$. \square

We recall that \mathbf{F}'_t is the rounded down distribution of the actual bid distribution \mathbf{F}_t in round t . Using the same result from [17], we have

Lemma 50 ([17]). $\text{Rev}(\mathcal{M}_{\mathbf{F}'_t}^*; \mathbf{F}'_t) \geq \text{Rev}(\mathcal{M}_{\mathbf{F}_t}^*; \mathbf{F}_t) - \beta J$.

In Theorem 19, we show that in round t , if the the value for a bidder i is $v_{i,t} \sim D_i$, then the bid $b_{i,t}$ is at most η_t less than it which means that the distribution of the bids and the values has a nice coupling which mentioned in the introduction of this section. Recall that $\mathbf{F}_{i,t}$ is the distribution of bids of player i in round t . With very similar arguments to Lemma 50 we can show the following.

Lemma 51. $\text{Rev}(\mathcal{M}_{\mathbf{F}_t}^*; \mathbf{F}_t) \geq \text{Rev}(\mathcal{M}_{\mathbf{D}}^*; \mathbf{D}) - \eta_t J$.

Proof. The proof follows identically to the proof of Lemma 50 in [17]. \square

Now we have all the pieces to bound the expected revenue achieved by Algorithm 2 in a fixed round t .

Theorem 52. *Using Algorithm 2, in round t , with probability at least $1 - \alpha$ the expected revenue $\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}; \mathbf{F}'_t)$ achieved by the mechanism satisfies:*

$$\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}; \mathbf{F}'_t) \geq \text{Rev}(\mathcal{M}_{\mathbf{D}}^*; \mathbf{D}) - 4hn^2\gamma_t - \frac{J \ln n}{\epsilon} - \frac{h}{\epsilon} \ln \frac{1}{\rho_t} - (\eta_t + \beta + 2\rho_t h)J$$

where $\gamma_t = \sqrt{\frac{\log(2hn/\beta\alpha)}{2t}} + \Sigma_t + \frac{\sigma}{t} \sqrt{\log \frac{h}{\beta} \log T} \sqrt{\log \left(\frac{hn}{\beta\alpha} \right)}$ and $\eta_t = h \sqrt{\frac{8n^2\gamma_t + 6k\epsilon}{\rho_t J}}$.

Proof. Using Lemma 41 and 48 and setting we have

$$\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}; \mathbf{F}'_t) \geq \text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{H}'_t) - 4hn^2\gamma_t - \frac{J \ln n}{\epsilon} - \frac{h}{\epsilon} \ln \frac{1}{\rho_t} - 2\rho_t hJ,$$

We apply Lemma 49 to get

$$\text{Rev}(\mathcal{M}_{\mathbf{H}'_t}; \mathbf{F}'_t) \geq \text{Rev}(\mathcal{M}_{\mathbf{H}'_t}^*; \mathbf{H}'_t) - 4hn^2\gamma_t - \frac{J \ln n}{\epsilon} - \frac{h}{\epsilon} \ln \frac{1}{\rho_t} - 2\rho_t hJ.$$

Now, applying Lemmas 51 and 50 completes the proof. \square

To get a bound on the average expected revenue achieved by Algorithm 2, we sum over all rounds and take union bound.

Proof of Theorem 20. Similar to Theorem 14, when we sum terms from Theorem 52, the leading term in $\frac{1}{T} \sum_{t=1}^T \gamma_t$, turns out to be $\tilde{O}(\frac{1}{\sqrt{T}})$. Keeping ρ_t as constant ρ , $\frac{1}{T} \sum_{t=1}^T \eta_t$, can be bounded by $hnJ^{-1/3}(\tilde{O}(\frac{1}{T^{1/4}})) + h\sqrt{\frac{6k\epsilon}{\rho J}}$. Adding the rest of the terms and keeping only leading terms in T , we get

$$\frac{1}{T} \sum_{t=1}^T \text{Rev}(\mathcal{M}_{\mathbf{H}'_t}; \mathbf{F}'_t) \geq \text{Rev}(\mathcal{M}_{\mathbf{D}}^*; \mathbf{D}) - hnJ^{2/3}\tilde{O}(\frac{1}{T^{1/4}}) - \frac{J \ln n}{\epsilon} - \frac{h}{\epsilon} \ln \frac{1}{\rho} - hJ(\sqrt{\frac{6k\epsilon}{\rho J}} + 2\rho) - \beta J$$

Optimizing over ρ and setting it to $(\frac{3k\epsilon}{2J})^{1/3}$ gives the theorem. \square

B.5 Omitted Proofs

Proof of Lemma 16. The mechanism makes 3 choices in each round: how to allocate items, how to charge bidders, and how to update the estimated distributions. The distribution estimates are maintained in a $(\epsilon, \epsilon/T)$ -private manner, by Theorem 22.

Allocations and payments are chosen in one of two ways: either according to the exponential mechanism, followed by black-box mechanism, or by using the strictly truthful mechanism Algorithm 4. In the former case, the allocation rule is ϵ differentially private by Theorem 37, and the payment for bidder i is a function of her bid in round t and postprocessing of the privately maintained distribution estimates, and are therefore jointly ϵ jointly differentially private for bidder i (by Lemma 9). In the latter case, the initial allocation is chosen uniformly at random and therefore perfectly private; buyer i 's ultimate allocation and payment is a postprocessing of this perfectly private selection along with her bid, and is therefore perfectly jointly differentially private.

Thus, since privacy guarantees sum under composition, the entire mechanism is $(3\epsilon, 3\epsilon/T)$ -jointly differentially private. \square

Proof of Lemma 17. This follows from the privacy guarantee of Lemma 16 and an argument identical to Theorem 13 upper-bounding the benefit of misreporting. Thus, the benefit of misreporting in round t to all future rounds in which this bidder participates is at most $\epsilon kh(2 + \frac{1}{T})$. \square