
The Pseudo-Dimension of Near-Optimal Auctions

Jamie Morgenstern*

Computer and Information Science
University of Pennsylvania
Philadelphia, PA
jamiemor@cis.upenn.edu

Tim Roughgarden

Stanford University
Palo Alto, CA
tim@cs.stanford.edu

Abstract

This paper develops a general approach, rooted in statistical learning theory, to learning an approximately revenue-maximizing auction from data. We introduce *t-level auctions* to interpolate between simple auctions, such as welfare maximization with reserve prices, and optimal auctions, thereby balancing the competing demands of expressivity and simplicity. We prove that such auctions have small representation error, in the sense that for every product distribution F over bidders' valuations, there exists a t -level auction with small t and expected revenue close to optimal. We show that the set of t -level auctions has modest pseudo-dimension (for polynomial t) and therefore leads to small learning error. One consequence of our results is that, in arbitrary single-parameter settings, one can learn a mechanism with expected revenue arbitrarily close to optimal from a polynomial number of samples.

1 Introduction

In the traditional economic approach to identifying a revenue-maximizing auction, one first posits a prior distribution over all unknown information, and then solves for the auction that maximizes expected revenue with respect to this distribution. The first obstacle to making this approach operational is the difficulty of formulating an appropriate prior. The second obstacle is that, even if an appropriate prior distribution is available, the corresponding optimal auction can be far too complex and unintuitive for practical use. This motivates the goal of identifying auctions that are “simple” and yet nearly-optimal in terms of expected revenue.

In this paper, we apply tools from learning theory to address both of these challenges. In our model, we assume that bidders' valuations (i.e., “willingness to pay”) are drawn from an unknown distribution F . A learning algorithm is given i.i.d. samples from F . For example, these could represent the outcomes of comparable transactions that were observed in the past. The learning algorithm suggests an auction to use for future bidders, and its performance is measured by comparing the expected revenue of its output auction to that earned by the optimal auction for the distribution F .

The possible outputs of the learning algorithm correspond to some set \mathcal{C} of auctions. We view \mathcal{C} as a design parameter that can be selected by a seller, along with the learning algorithm. A central goal of this work is to identify classes \mathcal{C} that balance representation error (the amount of revenue sacrificed by restricting to auctions in \mathcal{C}) with learning error (the generalization error incurred by learning over \mathcal{C} from samples). That is, we seek a set \mathcal{C} that is rich enough to contain an auction that closely approximates an optimal auction (whatever F might be), yet simple enough that the best auction in \mathcal{C} can be learned from a small amount of data. Learning theory offers tools both for rigorously defining the “simplicity” of a set \mathcal{C} of auctions, through well-known complexity measures such as the

*Part of this work done while visiting Stanford University. Partially supported by a Simons Award for Graduate Students in Theoretical Computer Science, as well as NSF grant CCF-1415460.

pseudo-dimension, and for quantifying the amount of data necessary to identify the approximately best auction from \mathcal{C} . Our goal of learning a near-optimal auction also requires understanding the representation error of different classes \mathcal{C} ; this task is problem-specific, and we develop the necessary arguments in this paper.

1.1 Our Contributions

The primary contributions of this paper are the following. First, we show that well-known concepts from statistical learning theory can be directly applied to reason about learning from data an approximately revenue-maximizing auction. Precisely, for a set \mathcal{C} of auctions and an arbitrary unknown distribution F over valuations in $[1, H]$, $O(\frac{H^2}{\epsilon^2} d_{\mathcal{C}} \log \frac{H}{\epsilon})$ samples from F are enough to learn (up to a $1 - \epsilon$ factor) the best auction in \mathcal{C} , where $d_{\mathcal{C}}$ denotes the *pseudo-dimension* of the set \mathcal{C} (defined in Section 2). Second, we introduce the class of *t-level auctions*, to interpolate smoothly between simple auctions, such as welfare maximization subject to individualized reserve prices (when $t = 1$), and the complex auctions that can arise as optimal auctions (as $t \rightarrow \infty$). Third, we prove that in quite general auction settings with n bidders, the pseudo-dimension of the set of t -level auctions is $O(nt \log nt)$. Fourth, we quantify the number t of levels required for the set of t -level auctions to have low representation error, with respect to the optimal auctions that arise from arbitrary product distributions F . For example, for single-item auctions and several generalizations thereof, if $t = \Omega(\frac{H}{\epsilon})$, then for every product distribution F there exists a t -level auction with expected revenue at least $1 - \epsilon$ times that of the optimal auction for F .

In the above sense, the “ t ” in t -level auctions is a tunable “sweet spot”, allowing a designer to balance the competing demands of expressivity (to achieve near-optimality) and simplicity (to achieve learnability). For example, given a fixed amount of past data, our results indicate how much auction complexity (in the form of the number of levels t) one can employ without risking overfitting the auction to the data.

Alternatively, given a target approximation factor $1 - \epsilon$, our results give sufficient conditions on t and consequently on the number of samples needed to achieve this approximation factor. The resulting sample complexity upper bound has polynomial dependence on H , ϵ^{-1} , and the number n of bidders. Known results [1, 8] imply that any method of learning a $(1 - \epsilon)$ -approximate auction from samples must have sample complexity with polynomial dependence on all three of these parameters, even for single-item auctions.

1.2 Related Work

The present work shares much of its spirit and high-level goals with Balcan et al. [4], who proposed applying statistical learning theory to the design of near-optimal auctions. The first-order difference between the two works is that our work assumes bidders’ valuations are drawn from an unknown distribution, while Balcan et al. [4] study the more demanding “prior-free” setting. Since no auction can achieve near-optimal revenue ex-post, Balcan et al. [4] define their revenue benchmark with respect to a set \mathcal{G} of auctions on each input \mathbf{v} as the maximum revenue obtained by any auction of \mathcal{G} on \mathbf{v} . The idea of learning from samples enters the work of Balcan et al. [4] through the internal randomness of their partitioning of bidders, rather than through an exogenous distribution over inputs (as in this work). Both our work and theirs requires polynomial dependence on $H, \frac{1}{\epsilon}$: ours in terms of a necessary number of samples, and theirs in terms of a necessary number of bidders; as well as a measure of the complexity of the class \mathcal{G} (in our case, the pseudo-dimension, and in theirs, an analogous measure). The primary improvement of our work over of the results in Balcan et al. [4] is that our results apply for single item-auctions, matroid feasibility, and arbitrary single-parameter settings (see Section 2 for definitions); while their results apply only to single-parameter settings of unlimited supply.¹ We also view as a feature the fact that our sample complexity upper bounds can be deduced directly from well-known results in learning theory — we can focus instead on the non-trivial and problem-specific work of bounding the pseudo-dimension and representation error of well-chosen auction classes.

Elkind [12] also considers a similar model to ours, but only for the special case of single-item auctions. While her proposed auction format is similar to ours, our results cover the far more general

¹See Balcan et al. [3] for an extension to the case of a large finite supply.

case of arbitrary single-parameter settings and non-finite support distributions; our sample complexity bounds are also better even in the case of a single-item auction (linear rather than quadratic dependence on the number of bidders). On the other hand, the learning algorithm in [12] (for single-item auctions) is computationally efficient, while ours is not.

Cole and Roughgarden [8] study single-item auctions with n bidders with valuations drawn from independent (not necessarily identical) “regular” distributions (see Section 2), and prove upper and lower bounds (polynomial in n and ϵ^{-1}) on the sample complexity of learning a $(1 - \epsilon)$ -approximate auction. While the formalism in their work is inspired by learning theory, no formal connections are offered; in particular, both their upper and lower bounds were proved from scratch. Our positive results include single-item auctions as a very special case and, for bounded or MHR valuations, our sample complexity upper bounds are much better than those in Cole and Roughgarden [8].

Huang et al. [15] consider learning the optimal price from samples when there is a single buyer and a single seller; this problem was also studied implicitly in [10]. Our general positive results obviously cover the bounded-valuation and MHR settings in [15], though the specialized analysis in [15] yields better (indeed, almost optimal) sample complexity bounds, as a function of ϵ^{-1} and/or H .

Medina and Mohri [17] show how to use a combination of the pseudo-dimension and Rademacher complexity to measure the sample complexity of selecting a single reserve price for the VCG mechanism to optimize revenue. In our notation, this corresponds to analyzing a single set \mathcal{C} of auctions (VCG with a reserve). Medina and Mohri [17] do not address the expressivity vs. simplicity trade-off that is central to this paper.

Dughmi et al. [11] also study the sample complexity of learning good auctions, but their main results are negative (exponential sample complexity), for the difficult scenario of multi-parameter settings. (All settings in this paper are single-parameter.)

Our work on t -level auctions also contributes to the literature on simple approximately revenue-maximizing auctions (e.g., [6, 14, 7, 9, 21, 23, 2]). Here, one takes the perspective of a seller who knows the valuation distribution F but is bound by a “simplicity constraint” on the auction deployed, thereby ruling out the optimal auction. Our results that bound the representation error of t -level auctions (Theorems 3.4, 4.1, 5.4, and 6.2) can be interpreted as a principled way to trade off the simplicity of an auction with its approximation guarantee. While previous work in this literature generally left the term “simple” safely undefined, this paper effectively proposes the pseudo-dimension of an auction class as a rigorous and quantifiable simplicity measure.

2 Preliminaries

This section reviews useful terminology and notation standard in Bayesian auction design and learning theory.

Bayesian Auction Design We consider *single-parameter settings* with n bidders. This means that each bidder has a single unknown parameter, its *valuation* or willingness to pay for “winning.” (Every bidder has value 0 for losing.) A setting is specified by a collection \mathcal{X} of subsets of $\{1, 2, \dots, n\}$; each such subset represent a collection of bidders that can simultaneously “win.” For example, in a setting with k copies of an item, where no bidder wants more than one copy, \mathcal{X} would be all subsets of $\{1, 2, \dots, n\}$ of cardinality at most k .

A generalization of this case, studied in the supplementary materials (Section 5), is *matroid* settings. These satisfy: (i) whenever $X \in \mathcal{X}$ and $Y \subseteq X$, $Y \in \mathcal{X}$; and (ii) for two sets $|I_1| < |I_2|$, $I_1, I_2 \in \mathcal{X}$, there is always an augmenting element $i_2 \in I_2 \setminus I_1$ such that $I_1 \cup \{i_2\} \in \mathcal{X}$. The supplementary materials (Section 6) also consider arbitrary single-parameter settings, where the only assumption is that $\emptyset \in \mathcal{X}$. To ease comprehension, we often illustrate our main ideas using single-item auctions (where \mathcal{X} is the singletons and the empty set).

We assume bidders’ valuations are drawn from the continuous joint cumulative distribution F . Except in the extension in Section 4, we assume that the support of F is limited to $[1, H]^n$. As in most of optimal auction theory [18], we usually assume that F is a product distribution, with $F = F_1 \times F_2 \times \dots \times F_n$ and each $v_i \sim F_i$ drawn independently but not identically. The *virtual*

value of bidder i is denoted by $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$. A distribution satisfies the *monotone-hazard rate (MHR)* condition if $f_i(v_i)/(1 - F_i(v_i))$ is nondecreasing; intuitively, if its tails are no heavier than those of an exponential distribution. In a fundamental paper, [18] proved that when every virtual valuation function is nondecreasing (the “regular” case), the auction that maximizes expected revenue for n Bayesian bidders chooses winners in a way which maximizes the sum of the virtual values of the winners. This auction is known as Myerson’s auction, which we refer to as \mathcal{M} . The result can be extended to the general, “non-regular” case by replacing the virtual valuation functions by “ironed virtual valuation functions.” The details are well-understood but technical; see Myerson [18] and Hartline [13] for details.

Sample Complexity, VC Dimension, and the Pseudo-Dimension This section reviews several well-known definitions from learning theory. Suppose there is some domain \mathcal{Q} , and let c be some unknown target function $c : \mathcal{Q} \rightarrow \{0, 1\}$. Let \mathcal{D} be an unknown distribution over \mathcal{Q} . We wish to understand how many labeled samples $(x, c(x))$, $x \sim \mathcal{D}$, are necessary and sufficient to be able to output a \hat{c} which agrees with c almost everywhere with respect to \mathcal{D} . The distribution-independent sample complexity of learning c depends fundamentally on the “complexity” of the set of binary functions \mathcal{C} from which we are choosing \hat{c} . We define the relevant complexity measure next.

Let S be a set of m samples from \mathcal{Q} . The set S is said to be *shattered* by \mathcal{C} if, for every subset $T \subseteq S$, there is some $c_T \in \mathcal{C}$ such that $c_T(x) = 1$ if $x \in T$ and $c_T(y) = 0$ if $y \notin T$. That is, ranging over all $c \in \mathcal{C}$ induces all $2^{|S|}$ possible projections onto S . The *VC dimension* of \mathcal{C} , denoted $\mathcal{VC}(\mathcal{C})$, is the size of the largest set S that can be shattered by \mathcal{C} .

Let $\text{err}_S(\hat{c}) = (\sum_{x \in S} |c(x) - \hat{c}(x)|)/|S|$ denote the empirical error of \hat{c} on S , and let $\text{err}(\hat{c}) = \mathbb{E}_{x \sim \mathcal{D}}[|c(x) - \hat{c}(x)|]$ denote the true expected error of \hat{c} with respect to \mathcal{D} . A key result from learning theory [22] is: for every distribution \mathcal{D} , a sample S of size $\Omega(\epsilon^{-2}(\mathcal{VC}(\mathcal{C}) + \ln \frac{1}{\delta}))$ is sufficient to guarantee that $\text{err}_S(\hat{c}) \in [\text{err}(\hat{c}) - \epsilon, \text{err}(\hat{c}) + \epsilon]$ for every $\hat{c} \in \mathcal{C}$ with probability $1 - \delta$. In this case, the error on the sample is close to the true error, simultaneously for every hypothesis in \mathcal{C} . In particular, choosing the hypothesis with the minimum sample error minimizes the true error, up to 2ϵ . We say \mathcal{C} is (ϵ, δ) -uniformly learnable with sample complexity m if, given a sample S of size m , with probability $1 - \delta$, for all $c \in \mathcal{C}$, $|\text{err}_S(c) - \text{err}(c)| < \epsilon$: thus, any class \mathcal{C} is (ϵ, δ) -uniformly learnable with $m = \Theta(\frac{1}{\epsilon^2}(\mathcal{VC}(\mathcal{C}) + \ln \frac{1}{\delta}))$ samples. Conversely, for every learning algorithm \mathcal{A} that uses fewer than $\frac{\mathcal{VC}(\mathcal{C})}{\epsilon}$ samples, there exists a distribution \mathcal{D}' and a constant q such that, with probability at least q , \mathcal{A} outputs a hypothesis $\hat{c}' \in \mathcal{C}$ with $\text{err}(\hat{c}') > \text{err}(\hat{c}) + \frac{\epsilon}{2}$ for some $\hat{c} \in \mathcal{C}$. That is, the true error of the output hypothesis is more than $\frac{\epsilon}{2}$ larger the best hypothesis in the class.

To learn real-valued functions, we need a generalization of VC dimension (which concerns binary functions). The *pseudo-dimension* [19] does exactly this. Formally, let $c : \mathcal{Q} \rightarrow [0, H]$ be a real-valued function over \mathcal{Q} , and \mathcal{C} the class we are learning over. Let S be a sample drawn from \mathcal{D} , $|S| = m$, labeled according to c . Both the empirical and true error of a hypothesis \hat{c} are defined as before, though $|\hat{c}(x) - c(x)|$ can now take on values in $[0, H]$ rather than in $\{0, 1\}$. Let $(r^1, \dots, r^m) \in [0, H]^m$ be a set of *targets* for S . We say (r^1, \dots, r^m) *witnesses* the shattering of S by \mathcal{C} if, for each $T \subseteq S$, there exists some $c_T \in \mathcal{C}$ such that $f_T(x^i) \geq r^i$ for all $x^i \in T$ and $c_T(x^i) < r^i$ for all $x^i \notin T$. If there exists some \vec{r} witnessing the shattering of S , we say S is *shatterable* by \mathcal{C} . The *pseudo-dimension* of \mathcal{C} , denoted $d_{\mathcal{C}}$, is the size of the largest set S which is shatterable by \mathcal{C} . The sample complexity upper bounds of this paper are derived from the following theorem, which states that the distribution-independent sample complexity of learning over a class of real-valued functions \mathcal{C} is governed by the class’s pseudo-dimension.

Theorem 2.1 [E.g. [1]] *Suppose \mathcal{C} is a class of real-valued functions with range in $[0, H]$ and pseudo-dimension $d_{\mathcal{C}}$. For every $\epsilon > 0$, $\delta \in [0, 1]$, the sample complexity of (ϵ, δ) -uniformly learning f with respect to \mathcal{C} is $m = O\left(\left(\frac{H}{\epsilon}\right)^2 \left(d_{\mathcal{C}} \ln\left(\frac{H}{\epsilon}\right) + \ln\left(\frac{1}{\delta}\right)\right)\right)$.*

Moreover, the guarantee in Theorem 2.1 is realized by the learning algorithm that simply outputs the function $c \in \mathcal{C}$ with the smallest empirical error on the sample.

Applying Pseudo-Dimension to Auction Classes For the remainder of this paper, we consider classes of truthful auctions \mathcal{C} .² When we discuss some auction $c \in \mathcal{C}$, we treat $c : [0, H]^n \rightarrow \mathbb{R}$ as the function that maps (truthful) bid tuples to the revenue achieved on them by the auction c . Then, rather than minimizing error, we aim to maximize revenue. In our setting, the guarantee of Theorem 2.1 directly implies that, with probability at least $1 - \delta$ (over the m samples), the output of the *empirical revenue maximization* learning algorithm — which returns the auction $c \in \mathcal{C}$ with the highest average revenue on the samples — chooses an auction with expected revenue (over the true underlying distribution F) that is within an additive ϵ of the maximum possible.

3 Single-Item Auctions

We focus on single-item auctions, a setting in which our key definitions and proof techniques are easily understood; the results of this section are generalized significantly in the full version of the paper. Section 3.1 defines the class of t -level single-item auctions, gives an example, and interprets the auctions as approximations to virtual welfare maximizers. Section 3.2 proves that the pseudo-dimension of the set of such auctions is $O(nt \log nt)$, which by Theorem 2.1 implies a sample-complexity upper bound. Section 3.3 proves that taking $t = \Omega(\frac{H}{\epsilon})$ yields low representation error.

3.1 t -Level Auctions: The Single-Item Case

We now introduce t -level auctions, or \mathcal{C}_t for short. Intuitively, one can think of each bidder as facing one of t possible prices; the price they face depends upon the values of the other bidders. Consider, for each bidder i , t numbers $0 \leq \ell_{i,0} \leq \ell_{i,1} \leq \dots \leq \ell_{i,t-1}$. We refer to these t numbers as *thresholds*. This set of tn numbers defines a t -level auction with the following allocation rule. Consider a valuation tuple \mathbf{v} :

1. For each bidder i , let $t_i(v_i)$ denote the index τ of the largest threshold $\ell_{i,\tau}$ that lower bounds v_i (or -1 if $v_i < \ell_{i,0}$). We call $t_i(v_i)$ the *level* of bidder i .
2. Sort the bidders from highest level to lowest level and, within a level, use a fixed lexicographical tie-breaking ordering \succ to pick the winner.³
3. Award the item to the first bidder in this sorted order (unless $t_i = -1$ for every bidder i , in which case there is no sale).

The payment rule is the unique one that renders truthful bidding a dominant strategy and charges 0 to losing bidders — the winning bidder pays the lowest bid at which she would continue to win. It is important for us to understand this payment rule in detail; there are three interesting cases. Suppose bidder i is the winner. In the first case, i is the only bidder who might be allocated the item (other bidders have level -1), in which case her bid must be at least her lowest threshold. In the second case, there are multiple bidders at her level, so she must bid high enough to be at her level (and, since ties are broken lexicographically, this is her threshold to win). In the final case, she need not compete at her level: she can choose to either pay one level above her competition (in which case her position in the tie-breaking ordering does not matter) or she can bid at the same level as her highest-level competitors (in which case she only wins if she dominates all of those bidders at the next-highest level according to \succ). Formally, the payment p of the winner i (if any) is as follows. Let $\bar{\tau}$ denote the highest level τ such that there are at least two bidders at or above level τ , and I be the set of bidders other than i whose level is at least $\bar{\tau}$.

Monop If $\bar{\tau} = -1$, then $p_i = \ell_{i,0}$ (she is the only potential winner, but must have level ≥ 0 to win).

Mult If $t_i(v_i) = \bar{\tau}$ then $p_i = \ell_{i,\bar{\tau}}$ (she needs to be at level $\bar{\tau}$).

²An auction is *truthful* if truthful bidding is a dominant strategy for every bidder. That is: for every bidder i , and all possible bids by the other bidders, i maximizes its expected utility (value minus price paid) by bidding its true value. In the single-parameter settings that we study, the expected revenue of the optimal non-truthful auction (measured at a Bayes-Nash equilibrium with respect to the prior distribution) is no larger than that of the optimal truthful auction.

³When bidder's valuation distributions are regular, this tie-breaking can be done by value, or randomly; when it is done by value, this equates to a generalization of VCG with nonanonymous reserves (and is IC and has identical representation error as this analysis when bidders are regular).

Unique If $t_i(v_i) > \bar{\tau}$, if $i \succ i'$ for all $i' \in I$, she pays $p_i = \ell_{i,\bar{\tau}}$, otherwise she pays $p_i = \ell_{i,\bar{\tau}+1}$ (she either needs to be at level $\bar{\tau} + 1$, in which case her position in \succ does not matter, or at level $\bar{\tau}$, in which case she would need to be the highest according to \succ).

We now describe a particular t -level auction, and demonstrate each case of the payment rule.

Example 3.1 Consider the following 4-level auction for bidders a, b, c . Let $\ell_{a,\cdot} = [2, 4, 6, 8]$, $\ell_{b,\cdot} = [1.5, 5, 9, 10]$, and $\ell_{c,\cdot} = [1.7, 3.9, 6, 7]$. For example, if bidder a bids less than 2 she is at level -1 , a bid in $[2, 4)$ puts her at level 0, a bid in $[4, 6)$ at level 1, a bid in $[6, 8)$ at level 2, and a bid of at least 8 at level 3. Let $a \succ b \succ c$.

Monop If $v_a = 3, v_b < 1.5, v_c < 1.7$, then b, c are at level -1 (to which the item is never allocated). So, a wins and pays 2, the minimum she needs to bid to be at level 0.

Mult If $v_a \geq 8, v_b \geq 10, v_c < 7$, then a and b are both at level 3, and $a \succ b$, so a will win and pays 8 (the minimum she needs to bid to be at level 3).

Unique If $v_a \geq 8, v_b \in [5, 9], v_c \in [3.9, 6]$, then a is at level 3, and b and c are at level 1. Since $a \succ b$ and $a \succ c$, a need only pay 4 (enough to be at level 1). If, on the other hand, $v_a \in [4, 6], v_b = [5, 9]$ and $v_c \geq 6$, c has level at least 2 (while a, b have level 1), but c needs to pay 6 since $a, b \succ c$.

Remark 3.2 (Connection to virtual valuation functions) t -level auctions are naturally interpreted as discrete approximations to virtual welfare maximizers, and our representation error bound in Theorem 3.4 makes this precise. Each level corresponds to a constraint of the form ‘‘If any bidder has level at least τ , do not sell to any bidder with level less than τ .’’ We can interpret the $\ell_{i,\tau}$ ’s (with fixed τ , ranging over bidders i) as the bidder values that map to some common virtual value. For example, 1-level auctions treat all values below the single threshold as having negative virtual value, and above the threshold uses values as proxies for virtual values. 2-level auctions use the second threshold to refine virtual value estimates, and so on. With this interpretation, it is intuitively clear that as $t \rightarrow \infty$, it is possible to estimate bidders’ virtual valuation functions and thus approximate Myerson’s optimal auction to arbitrary accuracy.

3.2 The Pseudo-Dimension of t -Level Auctions

This section shows that the pseudo-dimension of the class of t -level single-item auctions with n bidders is $O(nt \log nt)$. Combining this with Theorem 2.1 immediately yields sample complexity bounds (parameterized by t) for learning the best such auction from samples.

Theorem 3.3 *The pseudo-dimension of the set of n -bidder single-item t -level auctions is $O(nt \log(nt))$, fixing \succ , the tie-breaking order.*

Proof: Recall from (Section 2) that we need to upper bound the size of every set that is shatterable using t -level auctions. Fix a set of samples $S = (\mathbf{v}^1, \dots, \mathbf{v}^m)$ of size m and a potential witness $R = (r^1, \dots, r^m)$. Each auction c induces a binary labeling of the samples \mathbf{v}^j of S (whether c ’s revenue on \mathbf{v}^j is at least r^j or strictly less than r^j). The set S is shattered with witness R if and only if the number of distinct labelings of S given by any t -level auction is 2^m .

We upper-bound the number of distinct labelings of S given by t -level auctions (for some fixed potential witness R), counting the labelings in two stages. Note that S involves nm numbers — one value v_i^j for each bidder for each sample, and a t -level auction involves nt numbers — t thresholds $\ell_{i,\tau}$ for each bidder. Call two t -level auctions with thresholds $\{\ell_{i,\tau}\}$ and $\{\hat{\ell}_{i,\tau}\}$ *equivalent* if merging the sorted list of the v_i^j ’s with the sorted list of the $\ell_{i,\tau}$ ’s yields the same partition of the v_i^j ’s as does merging it with the sorted list of the $\hat{\ell}_{i,\tau}$ ’s.

Note that this is an equivalence relation. If two t -level auctions are equivalent, every comparison between a valuation and a threshold or two valuations is resolved identically by those auctions. Using the defining properties of equivalence, a crude upper bound on the number of equivalence classes is

$$(nt)! \cdot \binom{nm + nt}{nt} \leq (nm + nt)^{2nt}. \quad (1)$$

We now upper-bound the number of distinct labelings of S that can be generated by any auction in a single equivalence class C . First, as all comparisons between two numbers (valuations or thresholds) are resolved identically for all auctions in C , each bidder i in each sample \mathbf{v}^j of S is assigned the same level (across auctions in C), and the winner (if any) in each sample \mathbf{v}^j is constant across all of C . By the same reasoning, the identity of the parameter that gives the winner’s payment (some $\ell_{i,\tau}$) is uniquely determined by pairwise comparisons (recall Section 3.1) and hence is common across all auctions in C . The payments $\ell_{i,\tau}$, however, can vary across auctions in the equivalence class.

For a bidder i and level $\tau \in \{0, 1, 2, \dots, t-1\}$, let $S_{i,\tau} \subseteq S$ be the subset of samples in which bidder i wins and pays $\ell_{i,\tau}$. The revenue obtained by each auction in C on a sample of $S_{i,\tau}$ is simply $\ell_{i,\tau}$ (and independent of all other parameters of the auction). Thus, ranging over all t -level auctions in C generates at most $|S_{i,\tau}|$ distinct binary labelings of $S_{i,\tau}$ — the possible subsets of $S_{i,\tau}$ for which an auction meets the corresponding target r^j form a nested collection.

Summarizing, within the equivalence class C of t -level auctions, varying a parameter $\ell_{i,\tau}$ generates at most $|S_{i,\tau}|$ different labelings of the samples $S_{i,\tau}$ and has no effect on the other samples. Since the subsets $\{S_{i,\tau}\}_{i,\tau}$ are disjoint, varying all of the $\ell_{i,\tau}$ ’s (i.e., ranging over C) generates at most

$$\prod_{i=1}^n \prod_{\tau=0}^{t-1} |S_{i,\tau}| \leq m^{nt} \quad (2)$$

distinct labelings of S .

Combining (1) and (2), the class of all t -level auctions produces at most $(nm + nt)^{3nt}$ distinct labelings of S . Since shattering S requires 2^m distinct labelings, we conclude that $2^m \leq (nm + nt)^{3nt}$, implying $m = O(nt \log nt)$ as claimed. ■

3.3 The Representation Error of Single-Item t -Level Auctions

In this section, we show that for every bounded product distribution, there exists a t -level auction with expected revenue close to that of the optimal single-item auction when bidders are independent and bounded. The analysis “rounds” an optimal auction to a t -level auction without losing much expected revenue. This is done using thresholds to approximate each bidder’s virtual value: the lowest threshold at the bidder’s monopoly reserve price, the next $\frac{1}{\epsilon}$ thresholds at the values at which bidder i ’s virtual value surpasses multiples of ϵ , and the remaining thresholds at those values where bidder i ’s virtual value reaches powers of $1 + \epsilon$. Theorem 3.4 formalizes this intuition.

Theorem 3.4 *Suppose F is production over $[1, H]^n$. If $t = \Omega\left(\frac{1}{\epsilon} + \log_{1+\epsilon} H\right)$, C_t contains a single-item auction with expected revenue at least $1 - \epsilon$ times the optimal expected revenue.*

Theorem 3.4 follows immediately from the following lemma, with $\alpha = \gamma = 1$.

Lemma 3.5 *Consider n bidders with valuations in $[0, H]$ and with $\mathbb{P}[\max_i v_i > \alpha] \geq \gamma$. Then, C_t contains a single-item auction with expected revenue at least a $1 - \epsilon$ times that of an optimal auction, for $t = \Theta\left(\frac{1}{\gamma\epsilon} + \log_{1+\epsilon} \frac{H}{\alpha}\right)$.*

Proof: Consider a fixed bidder i . We define t thresholds for i , bucketing i by her virtual value, and prove that the t -level auction \mathcal{A} using these thresholds for each bidder closely approximates the expected revenue of the optimal auction \mathcal{M} . Let ϵ' be a parameter defined later.

Set $\ell_{i,0} = \phi_i^{-1}(0)$, bidder i ’s monopoly reserve.⁴ For $\tau \in [1, \lceil \frac{1}{\gamma\epsilon'} \rceil]$, let $\ell_{i,\tau} = \phi_i^{-1}(\tau \cdot \alpha\gamma\epsilon')$ ($\phi_i \in [0, \alpha]$). For $\tau \in [\lceil \frac{1}{\gamma\epsilon'} \rceil, \lceil \frac{1}{\gamma\epsilon'} \rceil + \lceil \log_{1+\frac{\epsilon}{2}} \frac{H}{\alpha} \rceil]$, let $\ell_{i,\tau} = \phi_i^{-1}(\alpha(1 + \frac{\epsilon}{2})^{\tau - \lceil \frac{1}{\gamma\epsilon'} \rceil})$ ($\phi_i > \alpha$).

Consider a fixed valuation profile \mathbf{v} . Let i^* denote the winner according to \mathcal{A} , and i' the winner according to the optimal auction \mathcal{M} . If there is no winner, we interpret $\phi_{i^*}(v_{i^*})$ and $\phi_{i'}(v_{i'})$ as 0.

⁴Recall from Section 2 that ϕ_i denotes the virtual valuation function of bidder i . (From here on, we always mean the ironed version of virtual values.) It is convenient to assume that these functions are strictly increasing (not just nondecreasing); this can be enforced at the cost of losing an arbitrarily small amount of revenue.

Recall that \mathcal{M} always awards the item to a bidder with the highest positive virtual value (or no one, if no such bidders exist). The definition of the thresholds immediately implies the following.

1. \mathcal{A} only allocates to non-negative ironed virtual-valued bidders.
2. If there is no tie (that is, there is a unique bidder at the highest level), then $i' = i^*$.
3. When there is a tie at level τ , the virtual value of the winner of \mathcal{A} is close to that of \mathcal{M} :
 If $\tau \in [0, \lceil \frac{1}{\gamma\epsilon'} \rceil]$ then $\phi_{i'}(v_{i'}) - \phi_{i^*}(v_{i^*}) \leq \alpha\gamma\epsilon'$;
 if $\tau \in [\lceil \frac{1}{\gamma\epsilon'} \rceil, \lceil \frac{1}{\gamma\epsilon'} \rceil + \lceil \log_{1+\frac{\epsilon}{2}} \frac{H}{\alpha} \rceil]$, $\frac{\phi_{i^*}(v_{i^*})}{\phi_{i'}(v_{i'})} \geq 1 - \frac{\epsilon}{2}$.

These facts imply that

$$\mathbb{E}_{\mathbf{v}}[\text{Rev}(\mathcal{A})] = \mathbb{E}_{\mathbf{v}}[\phi_{i^*}(v_{i^*})] \geq (1 - \frac{\epsilon}{2}) \cdot \mathbb{E}_{\mathbf{v}}[\phi_{i'}(v_{i'})] - \alpha\gamma\epsilon' = (1 - \frac{\epsilon}{2}) \cdot \mathbb{E}_{\mathbf{v}}[\text{Rev}(\mathcal{M})] - \alpha\gamma\epsilon'. \quad (3)$$

where the first and final equality follow from \mathcal{A} and \mathcal{M} 's allocations depending on ironed virtual values, not on the values themselves, thus, the ironed virtual values are equal in expectation to the unironed virtual values, and thus the revenue of the mechanisms (see [13], Chapter 3.5 for discussion).

The assumption that $\mathbb{P}[\max_i v_i > \alpha] \geq \gamma$ implies (since a feasible auction prices the good at α and awards it to any bidder with value at least α , if any) that $\mathbb{E}[\text{Rev}(\mathcal{M})] \geq \alpha\gamma$. Combining this with (3), and setting $\epsilon' = \frac{\epsilon}{2}$ implies $\mathbb{E}_{\mathbf{v}}[\text{Rev}(\mathcal{A})] \geq (1 - \epsilon) \mathbb{E}_{\mathbf{v}}[\text{Rev}(\mathcal{M})]$. ■

Combining Theorems 2.1 and 3.4 yields the following Corollary 3.6.

Corollary 3.6 *Let F be a product distribution with all bidders' valuations in $[1, H]$. Assume that $t = \Theta(\frac{1}{\epsilon} + \log_{1+\epsilon} H)$ and $m = O\left(\left(\frac{H}{\epsilon}\right)^2 (nt \log(nt) \log \frac{H}{\epsilon} + \log \frac{1}{\delta})\right) = \tilde{O}\left(\frac{H^2 n}{\epsilon^3}\right)$. Then with probability at least $1 - \delta$, the single-item empirical revenue maximizer of \mathcal{C}_t on a set of m samples from F has expected revenue at least $1 - \epsilon$ times that of the optimal auction.*

4 Unbounded MHR Distributions

This section shows how to replace the assumption of bounded valuations by the assumption that each valuation distribution satisfies the monotone hazard rate (MHR) condition, meaning that $\frac{f_i(v_i)}{1-F_i(v_i)}$ is nondecreasing. Our resulting sample complexity bounds depend on the number of bidders n and the error parameter ϵ only. This extension is based on previous work [5] that effectively reduces the case of MHR valuations to the case of valuations lying in the interval $[\beta\epsilon, 2\beta \log \frac{1}{\epsilon}]$ for a suitable choice of β , allowing us to reduce the problem to the bounded case. Our analysis works with η -truncated t -level auctions, where each t -level auction f is replaced with $f_\eta = \min(f, \eta)$.

Theorem 4.1 *Suppose F is a product distribution and each bidder's valuation distribution satisfies the MHR condition. Then, for each $\epsilon > 0$, and each $\hat{\beta} \geq \beta$ such that $\mathbb{P}\left[\max_i v_i \geq \frac{\hat{\beta}}{2}\right] \geq 1 - \frac{1}{\sqrt{\epsilon}} - \epsilon'$, there is a t -level $\left(\hat{\beta} \log \frac{1}{\epsilon'}\right)$ -truncated auction with expected revenue at least $1 - \epsilon$ times that of an optimal auction, where $t = \Theta\left(\frac{1}{\epsilon'} + \log_{1+\epsilon'}\left(\log \frac{1}{\epsilon'}\right)\right)$ and $\epsilon' = O\left(\frac{\epsilon}{\log \frac{1}{\epsilon}}\right)$.*

Before proving Theorem 4.1, we quote a key fact about MHR distributions Cai and Daskalakis [5].

Theorem 4.2 (Theorem 19 and Lemma 38 of [5]) *Let X_1, \dots, X_n be a collection of independent random variables whose distributions satisfy the MHR condition. Then there exists an anchoring point β such that*

$$\mathbb{P}\left[\max_i X_i \geq \frac{\beta}{2}\right] \geq 1 - \frac{1}{\sqrt{\epsilon}},$$

and for all $\epsilon > 0$,

$$\int_{2\beta \log \frac{1}{\epsilon}}^{\infty} z f_{\max_i X_i}(z) dz \leq 36\beta\epsilon \log \frac{1}{\epsilon}.$$

Now, we proceed to prove Theorem 4.1.

Proof of Theorem 4.1: Fix ϵ' , to be defined later. Conditioning on all bids being at most $2\hat{\beta} \log \frac{1}{\epsilon'}$ allows us apply Lemma 3.5 as though the valuations are bounded. In particular, since $\mathbb{P}[\max_i v_i \geq \frac{\hat{\beta}}{2}] \geq 1 - \frac{1}{\sqrt{e}} - \epsilon'$, Lemma 3.5 implies for $\alpha = \frac{\hat{\beta}}{2}$, $\gamma = 1 - \frac{1}{\sqrt{e}} - \epsilon'$ and $H = 2\hat{\beta} \log \frac{1}{\epsilon'}$, implies the existence of a $t = O\left(\frac{1}{\epsilon'} + \log_{1+\epsilon'}\left(\log \frac{1}{\epsilon'}\right)\right)$ -level truncated⁵ auction \mathcal{A} such that:

$$\mathbb{E} \left[\text{Rev}(\mathcal{A}) \mid \max_i \mathbf{v}_i \leq 2\hat{\beta} \log \frac{1}{\epsilon'} \right] \geq (1 - \epsilon') \mathbb{E} \left[\text{Rev}(\mathcal{M}) \mid \max_i \mathbf{v}_i \leq 2\hat{\beta} \log \frac{1}{\epsilon'} \right] \quad (4)$$

Thus, we have

$$\begin{aligned} \mathbb{E} [\text{Rev}(\mathcal{A})] &\geq \mathbb{E} \left[\text{Rev}(\mathcal{A}) \mid \max_i \mathbf{v}_i \leq 2\hat{\beta} \log \frac{1}{\epsilon'} \right] \mathbb{P} \left[\max_i \mathbf{v}_i \leq 2\hat{\beta} \log \frac{1}{\epsilon'} \right] \\ &\geq (1 - \epsilon') \mathbb{E} \left[\text{Rev}(\mathcal{M}) \mid \max_i \mathbf{v}_i \leq 2\hat{\beta} \log \frac{1}{\epsilon'} \right] \mathbb{P} \left[\max_i \mathbf{v}_i \leq 2\hat{\beta} \log \frac{1}{\epsilon'} \right] \\ &\geq (1 - \epsilon') \mathbb{E} [\text{Rev}(\mathcal{M})] - 36\hat{\beta}\epsilon' \log \frac{1}{\epsilon'} \\ &\geq \left(1 - O\left(\epsilon' \log \frac{1}{\epsilon'}\right) \right) \mathbb{E} [\text{Rev}(\mathcal{M})] \end{aligned}$$

where the first inequality comes from the fact that \mathcal{A} only sells to agents with non-negative virtual value (so the revenue on a smaller region of bids is only less), the second from Equation 4 and probabilities being at most 1, the penultimate from Theorem 4.2, and the final from the fact that $\text{Rev}(\mathcal{M}) \geq \frac{1 - \frac{1}{\sqrt{e}}}{2} \hat{\beta}$. Setting $\epsilon = \epsilon' \log \frac{1}{\epsilon'}$ and noticing that this implies $\epsilon' \leq \frac{\epsilon}{\log \frac{1}{\epsilon}}$ yields the desired result. ■

Corollary 4.3 follows as a corollary of Theorems 2.1 and 4.1.

Corollary 4.3 *With probability $1 - \delta$, the empirical revenue maximizer for a sample of size m of the class of t -level η -truncated single-item auctions is a $1 - O(\epsilon)$ -approximation of the optimal auction for n MHR bidders, for $t = O\left(\frac{1}{\epsilon'} + \log_{1+\epsilon'}\left(\log \frac{1}{\epsilon'}\right)\right)$, $\epsilon' = \frac{\epsilon}{\log \frac{1}{\epsilon}}$ and*

$$m = O\left(\left(\frac{1}{\epsilon'}\right)^2 \left(nt \log(nt) \ln \frac{1}{\epsilon'} + \ln \frac{1}{\delta}\right)\right) = \tilde{O}\left(\frac{n}{\epsilon^3}\right)$$

where η can be learned from the set of m samples.

Proof: We first argue that one can learn some η from the sample. Let $Q(S, g) = \frac{\sum_{\mathbf{v} \in S_{\epsilon'}} \mathbb{I}[\max_i \mathbf{v}_i \geq g]}{|S|}$ and $q(g) = \mathbb{P}[\max_i \mathbf{v}_i \geq g]$ (the empirical and true probability that the maximum bid is at least g , respectively). Given ϵ' , consider a set of samples $S_{\epsilon'}$ of profiles, and compute the largest $\hat{\beta}$ such that $Q(S_{\epsilon'}, \frac{\hat{\beta}}{2}) \geq 1 - \frac{1}{\sqrt{e}} - \epsilon'$. Standard VC-bounds imply that $|q(\rho) - Q(S_{\epsilon'}, \rho)| \leq \epsilon'$ for all ρ with probability at least $1 - \delta$, provided $|S_{\epsilon'}| \geq \Theta\left(\left(\frac{1}{\epsilon'}\right)^2 \ln \frac{1}{\delta}\right)$.⁶ In particular, with probability $1 - \delta$, we will have $q(S_{\epsilon'}, \frac{\hat{\beta}}{2}) \geq 1 - \frac{1}{\sqrt{e}} - \epsilon'$, so it will be the case that $\hat{\beta} \geq \beta$, and also $q(\hat{\beta}) \geq 1 - \frac{1}{\sqrt{e}} - 2\epsilon'$. Then, let $\eta = 2\hat{\beta} \log \frac{1}{\epsilon'}$.

Now, Theorem 4.1 implies the existence of a η -truncated t -level auction which $(1 - \epsilon')$ -approximates the optimal auction. The argument is completed using the fact that, if the auctions' values are upper-bounded by η , one can equivalently think of the *values* being upper-bounded by η , so Theorem 2.1

⁵Lemma 3.5 only implies the existence of such a t -level auction. However, when the bids are all below some η , one can always find an η -truncated auction which is equivalent to each untruncated auction.

⁶This can be thought of as a class of binary classifiers with VC-dimension one.

implies the sample complexity bound allowing additive error $\epsilon' \eta = \frac{\epsilon' \hat{\beta}}{2} \log \frac{1}{\epsilon'}$. This error is multiplicatively at most $\epsilon = \epsilon' \log \frac{1}{\epsilon'}$, since $\text{Rev}(\mathcal{M}) = \Omega\left(\frac{\hat{\beta}}{2}\right)$. ■

Remark 4.4 (Near-optimality of sample complexity) Can we do better than Theorem 3.3? Can we learn an approximately optimal auction from a simpler class — with pseudo-dimension $\text{poly}(\log H, \log n, \frac{1}{\epsilon})$, say — allowing for much smaller sample complexity than achieved here? The answer is negative: lower bounds in Cole and Roughgarden [8] imply that approximate revenue maximization requires sample complexity at least linear in the number of bidder n , even when bidders' valuations independently drawn from MHR distributions. Thus, every class of auctions that is sufficiently expressive to guarantee expected revenue at least $1 - \epsilon$ times optimal must have pseudo-dimension that grows polynomially with n .

5 t -Level Matroid Auctions

This section extends the ideas and techniques from Section 3.2 to matroid environments. The straightforward generalization of t -level auctions to matroid environments suffices: we order the bidders by level, breaking ties within a level by some fixed linear ordering over agents \succ , and greedily choose winners according to this ordering (subject to feasibility and to bids exceeding the lowest threshold). The next theorem bounds the pseudo-dimension of this more general class of auctions.

Theorem 5.1 *The pseudo-dimension of t -level matroid auctions with n bidders is $O(nt \log(nt))$.*

The proof is conceptually similar to that of Theorem 3.3, though we require a more general argument. Our proof uses a couple of standard results from learning theory (see e.g. [16] for details). The first, also known as Sauer's Lemma, states that the number of distinct projections of a set S induced by a set system with bounded VC dimension grows only polynomially in $|S|$.

Lemma 5.2 *Let \mathcal{C} be a set of functions from \mathcal{Q} to $\{0, 1\}$ with VC dimension d , and $S \subseteq \mathcal{Q}$. Then*

$$|\{S \cap \{x \in \mathcal{Q} : c(x) = 1\} : c \in \mathcal{C}\}| \leq |S|^d.$$

Recall that a *linear separator* in \mathbb{R}^d is defined by coefficients a_1, \dots, a_d , and assigns $x \in \mathbb{R}^d$ the value 1 if $\sum_{i=1}^d a_i x_i \geq 0$ and the value 0 otherwise.

Lemma 5.3 *The set of linear separators in \mathbb{R}^d has VC dimension $d + 1$.*

Proof of Theorem 5.1: Consider a set of samples S of size m which can be shattered by t -level matroid auctions with revenue targets (r^1, \dots, r^m) . We upper-bound the number of labelings of S possible using t -level auctions, which again yields an upper bound on m .

We partition auctions into equivalence classes, identically to the proof of Theorem 3.3. Recall that, across all auctions in an equivalence class, all comparisons between two thresholds or a threshold and a bid are resolved identically. Recall also that the number of equivalence classes is at most $(nm + nt)^{2nt}$. We now upper bound the number of distinct labelings any fixed equivalence class C of auctions can generate.

Consider a class C of equivalent auctions. The allocation and payment rules are more complicated than in the single-item case but still relatively simple. In particular, whether or not a bidder wins depends only on the ordering of the bidders (by level) and the fixed tie-breaking rule \succ , and thus is a function only of comparisons between bids and thresholds. This implies that, for every sample in S , all auctions in the class C declare the same set of winners. It also implies that the payment of each winning bidder is a fixed threshold $\ell_{i^*, \tau}$, and the identity of this parameter is the same across all auctions in C .

Now, encode each auction $\mathcal{A} \in C$ and sample \mathbf{v}^j as an $nt + 1$ -dimensional vector as follows. Let $x_{i, \tau}^{\mathcal{A}}$ equal the value of $\ell_{i, \tau}$ in the auction \mathcal{A} . Define $x_{nt+1}^{\mathcal{A}} = 1$ for every $\mathcal{A} \in C$. Define $y_{i, \tau}^j = 1$

if bidder i is a winner paying her threshold $\ell_{i,\tau}$ for auctions in C and 0 otherwise. Finally, define $y_{nt+1}^j = -r^j$. The point is that, for every auction \mathcal{A} in the class C and sample \mathbf{v}^j ,

$$x^{\mathcal{A}} \cdot y^j \geq 0$$

if and only if $\text{Rev}(\mathcal{A}) \geq r^j$. Thus, the number of distinct labelings of the samples generated by auctions in C is bounded above by the number of distinct sign patterns on m points in \mathbb{R}^{nt+1} generated by all linear separators. (The y^j -vectors are constant across C and can be viewed as m fixed points in \mathbb{R}^{nt+1} ; each auction $\mathcal{A} \in C$ corresponds to the vector $x^{\mathcal{A}}$ of coefficients.) Applying Lemmas 5.2 and 5.3, t -level matroid auctions can generate at most m^{nt+2} labelings per equivalence class, and hence at most $(nm + nt)^{3nt+2}$ distinct labelings in total. This imposes the restriction

$$2^m \leq (nm + nt)^{3nt+2};$$

solving for m yields the desired bound. ■

We now extend our representation error bound for t -level single-item auctions to matroids.

Theorem 5.4 *Consider an arbitrary matroid environment. Suppose F is a production distribution with valuations in $[1, H]$. Provided $t = \Omega\left(\frac{1}{\epsilon} + \log_{1+\epsilon} H\right)$, there exists a t -level matroid auction with expected revenue at least a $1 - \epsilon$ fraction of the optimal expected revenue.*

The key new idea in the proof is to exhibit a bijection between the feasible sets \mathcal{I}^* (our winning set) and \mathcal{I}' (the optimal winning set) such that each bidder from \mathcal{I}^* has a level at least as high as their bijective partner in \mathcal{I}' . To implement this, we use the following property of matroids (e.g. [14, 20]).

Proposition 5.5 *Let OPT denote the largest-weight set of a matroid, and let B be any other feasible set such that $|B| = |\text{OPT}|$, and OPT_i, B_i denote the i -th largest element of OPT and B , respectively. Then $w(\text{OPT}_i) \geq w(B_i)$ for all i .*

Proof of Theorem 5.4: Define bidders' thresholds exactly as in the proof of Theorem 3.4 and let \mathcal{A} denote the corresponding t -level auction. Fix an arbitrary valuation profile \mathbf{v} . Let \mathcal{I}^* denote the set of winning bidders in \mathcal{A} and \mathcal{I}' the set of winning bidders in \mathcal{M} . Recall that the latter is the feasible set that maximizes the sum of virtual valuations. Both sets are maximally independent amongst those bidders with non-negative virtual value (\mathcal{M} , by virtual of being welfare-maximal, and \mathcal{A} , by definition). Then, we claim $|\mathcal{I}^*| = |\mathcal{I}'|$ (if not, by the augmentation property of matroids, the smaller set could be extended to include an element of the larger while maintaining independence, violating their maximality).

Notice that \mathcal{I}^* is lexicographically optimal with respect to the *levels*, rather than the exact weights. Proposition 5.5 implies that \mathcal{I}' is also lexicographically optimal with respect to the levels; thus, the level of the i th largest bidder in \mathcal{I}' has the same level as the i th largest bidder in \mathcal{I}^* . Then, by an accounting argument identical to the one for the single-item case (comparing virtual values for the i th bidder in \mathcal{I}^* to the i th bidder in \mathcal{I}') summing up over all bidders completes the proof. ■

Thus, we have the following corollary about the sample complexity of $1 - \epsilon$ -approximating Myerson in matroid environments with t -level auctions, noting that the maximum revenue is now nH rather than H .

Corollary 5.6 *With probability $1 - \delta$, the empirical revenue maximizer for a sample of size m of the class of t -level single-item auctions is a $1 - O(\epsilon)$ -approximation to Myerson for n bidders whose valuations are in $[1, H]$, for $t = O\left(\frac{1}{\epsilon} + \log_{1+\epsilon} H\right)$ and*

$$m = O\left(\left(\frac{Hn}{\epsilon}\right)^2 \left(nt \log(nt) \ln \frac{Hn}{\epsilon} + \ln \frac{1}{\delta}\right)\right) = \tilde{O}\left(\frac{H^2 n^3}{\epsilon^3}\right).$$

6 Single-parameter t -level auctions

In this section, we show how to extend the ideas and techniques from Section 3.2 to any single-parameter environment which has the empty set as a feasible outcome. With this mild assumption, the results in this section do not require the environment to be a matroid or even downwards-closed. Before we state this result, we need a slight generalization of the t -level auction to this setting. Previously, no t -level auction would allocate to any bidder whose value was below their lowest threshold, and this will not be a possibility in environments which are not downwards-closed. Instead, in this setting, if any bidder fails to pass her lowest threshold, we will assume \mathcal{A} will choose the empty outcome. Moreover, a t -level auction will now need more about what the various levels represent: previously, we implicitly used the k th threshold to correspond to a value where each bidder's virtual value would pass some quantity q_k .

In this general setting, we make that connection explicit. There will still be nt numbers which define a particular t -level auction, the t threshold locations per bidder. In addition, we will consider a fixed vector $\Phi \in \mathbb{R}^t$ (not parameterizing the auction class) which, for all τ , intuitively assigns an estimate of $\phi_i^{-1}(\ell_{i,\tau})$, which is the same for all bidders i . Formally, Φ_τ will be used to assign a real value to a feasible set $X \in \mathcal{X}$ with a valuation profile \mathbf{v} as follows. Let $e_X = \sum_{i \in X} \Phi_{t_i(v_i)}$, where $t_i(v_i)$ as before is the level agent i 's bid according to \mathbf{v}_i . Then, a particular t -level auction will choose the winning set $X \in \mathcal{X}$ which maximizes e_X (breaking ties in some fixed way which does not depend upon the bids). If, for each bidder i and level τ , the threshold $\ell_{i,\tau}$ is placed exactly at the value at which i 's virtual valuation surpasses Φ_τ , then this auction is approximately optimizing the virtual surplus of the winning set.

Theorem 6.1 *The pseudo-dimension of t -level single-parameter auctions with n bidders is $O(nt \log(nt))$.*

These auctions have slightly more complicated payment rules than those for matroids, where each agent i was intuitively competing with (at most) one other bidder for inclusion in the winning set. Now, a bidder i who is in the winning set X will have a payment of the following form. For a fixed assignment of levels to bidders, sort the alternatives according to their values e_Y for all $Y \in \mathcal{X}$. Let Y be the highest-ranked alternative set which does not contain i . Then, i 's payment will be the threshold corresponding to the minimal τ such that $\sum_{i' \in X, i' \neq i} \Phi_{t_{i'}(v_{i'})} + \Phi_\tau \geq \sum_{i' \in Y} \Phi_{t_{i'}(v_{i'})}$ ⁷ (namely, the minimal bid which keeps X preferred to Y in terms of the estimated virtual values)⁷. While this rule is more complicated, it is *still* the case that, once each bidder is assigned to some level, each of the bidders in the winning set's payment is just one of their thresholds. Thus, the proof of Theorem 6.1 is identical to the one of Theorem 5.1 without ties.

When considering non-downwards-closed environments, the optimal revenue may be arbitrarily close to 0, making it difficult to argue about multiplicative approximations to the optimal revenue. Instead, we will give a weaker guarantee, namely, that the empirical revenue maximizer will have expected revenue which is *additively* close to the optimal expected revenue. If one is willing to restrict the environment to be downwards-closed, it is possible to achieve a multiplicative guarantee, since in that case the optimal revenue is at least 1. We now state the Theorem which bounds the representation error of t -level auctions for single-parameter environments.

Theorem 6.2 *There is a t -level auction whose expected revenue is within an additive ϵ of optimal, in any single-parameter setting \mathcal{X} such that $\emptyset \in \mathcal{X}$, for $t = O\left(\frac{Hn^2}{\epsilon}\right)$, for n bidders with valuations distributions which are product and bounded in $[0, H]$.*

Proof: Let $t = \frac{Hn^2}{\epsilon} + \frac{Hn}{\epsilon}$. We will begin by defining Φ , the t -dimensional vector corresponding to the estimated virtual values. Let $\Phi_0 = -Hn$ (if any bidder has virtual value $< -Hn$, the virtual value of any set containing her is negative, since virtual values are upper-bounded by values and the value of the remaining set may be at most $H(n-1)$, so in this case one should allocate to \emptyset). Then, let $\Phi_\tau = \Phi_{\tau-1} + \frac{\epsilon}{n}$. Thus, we partition the space of virtual values into additive sections of width $\frac{\epsilon}{n}$.

⁷Since we assume ties are broken in a way which does not depend on the bids, we can ignore ties in the payment rule, and agents will only ever pay thresholds.

Then, for each bidder i and τ , let $\ell_{i,\tau} = \phi_i^{-1}(\Phi_\tau)$, the value at which bidder i 's virtual value surpasses Φ_τ . Then, consider a valuation profile \mathbf{v} on which \mathcal{M} and this particular \mathcal{A} disagree on the winning sets $\mathcal{I}^*, \mathcal{I}' \in \mathcal{X}$. Notice that each bidder i 's virtual value is estimated correctly within an additive $\frac{\epsilon}{n}$ by $\Phi_{t_i(v_i)}$ (assuming no bidder has highly negative virtual value, in which case \mathcal{M} and \mathcal{A} both choose outcome \emptyset), and are never overestimated. Thus, it is the case that

$$\sum_{i^* \in \mathcal{I}^*} \phi_{i^*}(\mathbf{v}_{i^*}) \geq \sum_{i^* \in \mathcal{I}^*} \Phi_{t_{i^*}(v_{i^*})} \geq \sum_{i' \in \mathcal{I}'} \Phi_{t_{i'}(v_{i'})} \geq \sum_{i' \in \mathcal{I}'} \phi_{i'}(\mathbf{v}_{i'}) - \epsilon$$

and the claim follows. ■

Thus, we have the following sample complexity result for general single-parameter settings.

Corollary 6.3 *With probability $1 - \delta$, the empirical revenue maximizer on m samples S from the class of t -level auctions has true expected revenue within an additive ϵ of Myerson's expected revenue, for the single-parameter environment \mathcal{X} when bidders have valuations in $[0, H]$, for $t = O\left(\frac{Hn^2}{\epsilon}\right)$, and*

$$m = O\left(\left(\frac{Hn}{\epsilon}\right)^2 \left(nt \log(nt) \log \frac{Hn}{\epsilon} + \log \frac{1}{\delta}\right)\right) = \tilde{O}\left(\frac{H^3 n^5}{\epsilon^3}\right).$$

References

- [1] Martin Anthony and Peter L. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, NY, NY, USA, 1999.
- [2] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. *SIGecom Exch.*, 13(2):31–35, January 2015.
- [3] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Single price mechanisms for revenue maximization in unlimited supply combinatorial auctions. Technical report, Carnegie Mellon University, 2007.
- [4] Maria-Florina Balcan, Avrim Blum, Jason D Hartline, and Yishay Mansour. Reducing mechanism design to algorithm design via machine learning. *Jour. of Comp. and System Sciences*, 74(8):1245–1270, 2008.
- [5] Yang Cai and Constantinos Daskalakis. Extreme-value theorems for optimal multidimensional pricing. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 522–531, Palm Springs, CA, USA., Oct 2011. IEEE.
- [6] Shuchi Chawla, Jason Hartline, and Robert Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conf. on Electronic Commerce*, pages 243–251, NY, NY, USA, 2007. ACM.
- [7] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the Forty-second ACM Symposium on Theory of Computing*, pages 311–320, NY, NY, USA, 2010. ACM.
- [8] Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 243–252, NY, NY, USA, 2014. SIAM.
- [9] Nikhil Devanur, Jason Hartline, Anna Karlin, and Thach Nguyen. Prior-independent multi-parameter mechanism design. In *Internet and Network Economics*, pages 122–133. Springer, Singapore, 2011.
- [10] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. In *Proceedings of the 11th ACM Conf. on Electronic Commerce*, pages 129–138, NY, NY, USA, 2010. ACM.
- [11] Shaddin Dughmi, Li Han, and Noam Nisan. Sampling and representation complexity of revenue maximization. In *Web and Internet Economics*, volume 8877 of *Lecture Notes in Computer Science*, pages 277–291. Springer Intl. Publishing, Beijing, China, 2014.
- [12] Edith Elkind. Designing and learning optimal finite support auctions. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 736–745. SIAM, 2007.
- [13] Jason Hartline. *Mechanism design and approximation*. Jason Hartline, Chicago, Illinois, 2015.
- [14] Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *ACM Conf. on Electronic Commerce*, Stanford, CA, USA., 2009. ACM.

- [15] Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. Making the most of your samples. abs/1407.2479: 1–3, 2014. URL <http://arxiv.org/abs/1407.2479>.
- [16] Michael J. Kearns and Umesh V. Vazirani. *An Introduction to Computational Learning Theory*. MIT Press, Cambridge, MA, 1994.
- [17] Andres Munoz Medina and Mehryar Mohri. Learning theory and algorithms for revenue optimization in second price auctions with reserve. In *Proceedings of The 31st Intl. Conf. on Machine Learning*, pages 262–270, 2014.
- [18] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- [19] David Pollard. *Convergence of stochastic processes*. David Pollard, New Haven, Connecticut, 1984.
- [20] T. Roughgarden and O. Schrijvers. Ironing in the dark. Submitted, 2015.
- [21] Tim Roughgarden, Inbal Talgam-Cohen, and Qiqi Yan. Supply-limiting mechanisms. In *Proceedings of the 13th ACM Conf. on Electronic Commerce*, pages 844–861, NY, NY, USA, 2012. ACM.
- [22] Vladimir N Vapnik and A Ya Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16(2):264–280, 1971.
- [23] Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 92–109, San Diego, CA, USA., 2015. ACM.