

Sample Complexity of Multi-Item Profit Maximization*

Maria-Florina Balcan
ninamf@cs.cmu.edu

Tuomas Sandholm
sandholm@cs.cmu.edu

Ellen Vitercik
vitercik@cs.cmu.edu

May 16, 2017

Abstract

We study the design of pricing mechanisms and auctions when the mechanism designer does not know the distribution of buyers' values. Instead, the mechanism designer receives a set of samples from this distribution and his goal is to use the sample to design a pricing mechanism or auction with high expected profit. We provide *generalization guarantees* which bound the difference between average profit on the sample and expected profit over the distribution. These bounds are directly proportional to the intrinsic complexity of the mechanism class the designer is optimizing over. We present a single, general theorem that uses *empirical Rademacher complexity* to measure the intrinsic complexity of a variety of widely-studied single- and multi-item auction classes, including affine maximizer auctions, mixed-bundling auctions, and second-price item auctions. This theorem also applies to multi- and single-item pricing mechanisms in both multi- and single-unit settings, such as linear and non-linear pricing mechanisms. Despite the extensive applicability of our main theorem, we match or improve over the best-known generalization guarantees for many mechanism classes. Finally, our central theorem allows us to easily derive generalization guarantees for every class in several finely grained hierarchies of auction and pricing mechanism classes. We demonstrate how to determine the precise level in a hierarchy with the optimal tradeoff between profit and generalization using *structural profit maximization*. The mechanism classes we study are significantly different from well-understood function classes typically found in machine learning, so bounding their complexity requires a sharp understanding of the interplay between mechanism parameters and buyer valuations.

1 Introduction

Machine learning is an indispensable tool for large-scale mechanism design given the vast quantity of consumer data companies have at their disposal. Its applicability to mechanism design is a natural consequence of a standard assumption made in economics: a buyer's value for a bundle of goods is defined by a probability distribution over all the possible valuations he might have for that bundle. In the model most applicable to machine learning, the mechanism designer receives a sample from this distribution and his goal is to derive a mechanism with high expected profit.

A recent line of work has augmented the sample-based mechanism design literature with provable guarantees via learning-theoretic analyses. For example, given a set of samples from the distribution over buyers' values, a natural way to determine a mechanism that will likely have strong expected performance is to choose one with high average profit over the sample. Implicit in this procedure is the assumption that a mechanism's performance on the sample will generalize to the distribution. For a fixed mechanism class \mathcal{A} , a bound on the difference between the average profit over the sample and the expected profit on the distribution for any mechanism in \mathcal{A} is known as a *generalization*

*This work was posted on arXiv on April 29, 2017 [Balcan et al., 2017].

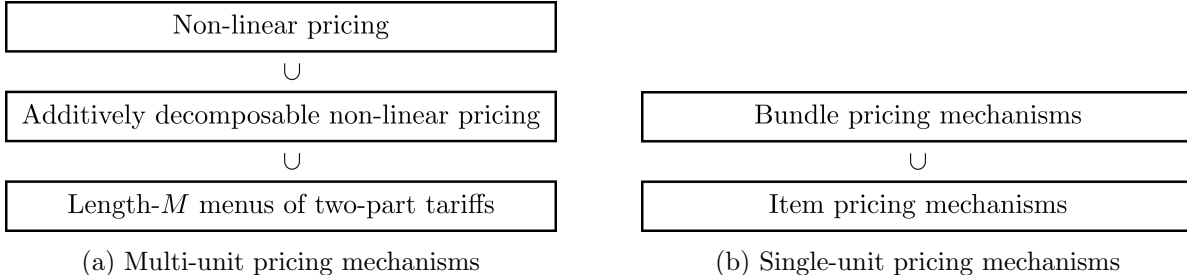


Figure 1: The hierarchies of pricing mechanism families studied in this paper. Generality increases upwards in the hierarchies.

guarantee. These guarantees are often highly dependent on the intrinsic complexity of the class \mathcal{A} and we use tools from learning theory such as *pseudo-dimension* and *Rademacher complexity* to bound this quantity. Bounding the complexity of a class is the bottleneck in deriving strong generalization guarantees.

Despite the growing literature on sample-based mechanism design, there is not yet a unifying framework guiding the derivation of complexity bounds for both simple mechanisms such as take-it-or-leave-it pricing mechanisms and more combinatorially challenging VCG-based mechanisms such as affine maximizer auctions (AMAs). Existing frameworks apply to “simple” mechanisms that can be reduced to the single-buyer setting [Morgenstern and Roughgarden, 2016] or only to specific learning algorithms that return the mechanism with empirically maximal profit [Syrkkanis, 2017]. Determining the empirically optimal mechanism is typically not feasible for highly complex mechanism classes such as affine maximizer auctions.

In this work, we extract the common features linking mechanisms that have been studied in the sample-based mechanism design literature as well as many more. Taking advantage of this structure,

...we present one general theorem that bounds the intrinsic complexity of a wide swath of mechanism classes, from simple take-it-or-leave-it pricing mechanisms to complex VCG-based auctions such as affine maximizer auctions.

We measure intrinsic complexity using *empirical Rademacher complexity*, a learning theoretic notion that we define in Section 3. Our central theorem immediately implies strong generalization guarantees for every mechanism class it applies to. See Figures 1 and 2 for hierarchical depictions of the mechanism classes we study. Surprisingly, despite the theorem’s generality and widespread applicability, it matches or improves upon the best-known generalization guarantees for many of the auction classes already studied in the literature.

Next, we apply this machinery to a learning framework known as *structural profit maximization* (*SPM*). Oftentimes, the mechanism classes we study exhibit a finely grained hierarchical structure and the intrinsic complexity of the subclasses decreases as we traverse down the hierarchy. Of course, the mechanism designer should not choose the simplest class available simply to guarantee good generalization because the more complex a mechanism class is, the more likely it is to contain a nearly optimal mechanism. Inevitably, there is a tradeoff between profit maximization and generalization. The mechanism designer can apply our main theorem to each level in the hierarchy to derive subclass-specific generalization guarantees and use structural profit maximization to determine the precise level in the hierarchy that will assure him the optimal tradeoff between profit maximization and generalization. For example, we present several hierarchies of affine maximizer auctions, a family that was first introduced by Roberts [Roberts, 1979]. At a high level, an AMA

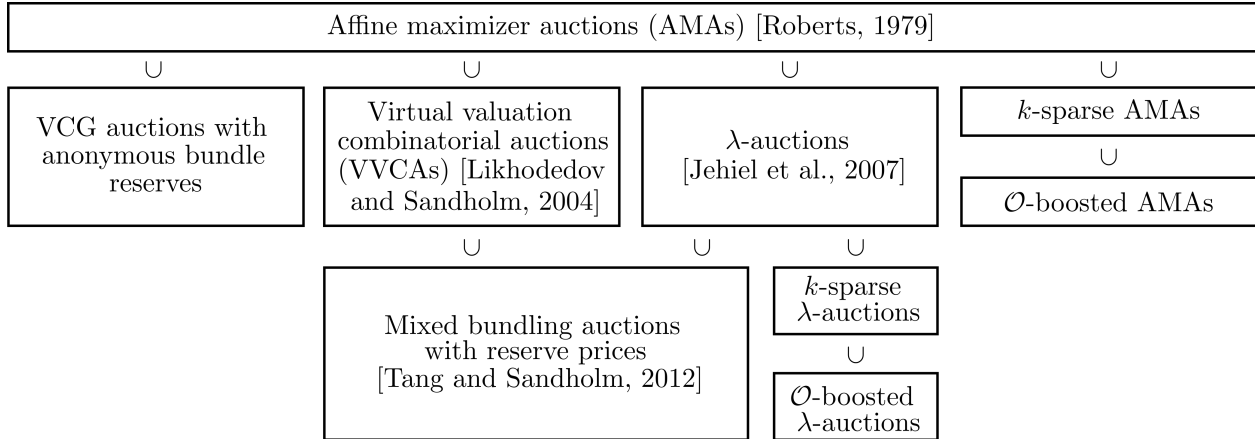


Figure 2: The hierarchy of auction families studied in this paper. Generality increases upwards in the hierarchy.

is defined by a set of bidder weights w_i , where all of Bidder i 's bids are increased multiplicatively by a factor of w_i , and outcome boosts, where the social welfare of any allocation Q is increased additively by a factor of $\lambda(Q)$. The Vickrey-Clarke-Groves mechanism is then run on this transformed bid space. For any set of allocations \mathcal{O} , we introduce the class of \mathcal{O} -boosted AMAs, where only allocations in \mathcal{O} are boosted. More generally, the class of k -sparse AMAs consists of those auctions where at most k outcomes are boosted. By varying the value of k or the set \mathcal{O} , the mechanism designer can carefully control the intrinsic complexity of the auction class. Structural profit maximization allows the mechanism designer to determine the class \mathcal{O} or the sparsity k that guarantees the optimal tradeoff between profit and generalization.

We also apply our main theorem and the SPM learning framework to pricing mechanisms. The pricing mechanisms we study fall into two categories: single-unit and multi-unit. In the multi-unit setting, we study both linear and non-linear unit prices. Non-linear pricing allows the mechanism designer to set differing costs per unit. For example, he may offer a discount for each additional unit that the buyer buys. Non-linear pricing is ubiquitous throughout many sectors of the economy, such as the shipping industry, electricity market, cellular service market, air travel industry, and advertising industry [Wilson, 1993]. A *two-part tariff* is one of the simplest examples of non-linear pricing. It consists of a fixed, up-front fee charged to each buyer that buys at least one unit of the good and a price per unit bought. The fixed fee is often described as an installation, access, or subscription charge. Oftentimes, producers will offer a menu of two-part tariffs from which the buyer chooses a payment plan. As Wilson describes, offering a menu of two-part tariffs is often equivalent to offering a *multipart tariff*. A multipart tariff consists of a fixed fee and b different marginal prices that apply in different volume bands or intervals. When the marginal prices are successively decreasing, a multipart tariff is equivalent to offering a menu of b two-part tariffs [Wilson, 1993].

In the single-unit setting, the simplest mechanism we study is an item-pricing mechanism, where each item has a price. Buyers arrive one at a time in a fixed but arbitrary order to buy the bundle maximizing their utility among the remaining items. We also study generalizations such as \mathcal{B} -pricing mechanisms, where \mathcal{B} is a set of bundles and every bundle in \mathcal{B} has a price. Buyers arrive one at a time to buy the bundle maximizing their utility among the remaining items. The class of $(\mathcal{B}_1, \dots, \mathcal{B}_n)$ -pricing mechanisms generalizes the class of \mathcal{B} -pricing mechanisms to include bidder-specific prices, as we describe in Section 3.

Reference theorem and mechanism class \mathcal{A}	Price category	$\epsilon_{\mathcal{A}}(\delta, N)$ upper bound
Theorem 4.5: Non-linear pricing	Anonymous	$O\left(U\sqrt{\frac{1}{N}\prod_{i=1}^m\kappa_i\left(\log n + \sum_{i=1}^m\log\kappa_i\right)} + U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
	Non-anonymous	$O\left(U\sqrt{\frac{n}{N}\prod_{i=1}^m\kappa_i\left(\log n + \sum_{i=1}^m\log\kappa_i\right)} + U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
Theorem 4.6: Additively decomposable non-linear pricing	Anonymous	$O\left(U\sqrt{\frac{1}{N}\sum_{i=1}^m\kappa_i\left(\log n + \sum_{i=1}^m\log\kappa_i\right)} + U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
	Non-anonymous	$O\left(U\sqrt{\frac{n}{N}\sum_{i=1}^m\kappa_i\left(\log n + \sum_{i=1}^m\log\kappa_i\right)} + U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
Theorem 4.7: Length- M menus of two-part tariffs	Anonymous	$O\left(U\sqrt{\frac{M}{N}\log(n\kappa M)} + U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
	Non-anonymous	$O\left(U\sqrt{\frac{nM}{N}\log(n\kappa M)} + U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$

Table 1: Generalization guarantees for multi-unit pricing mechanisms. The variables κ and κ_i for all $i \in [m]$ are upper bounds on the number of units available, as described in Section 3.

We show that when the buyers have simple valuations, such as additive or unit-demand valuations, “simple” mechanisms become even simpler, in that our upper bounds on their empirical Rademacher complexity shrink. Empirical Rademacher complexity is a modern complexity measure used for obtaining data-dependent, beyond-worst-case generalization guarantees in learning theory [Bartlett and Mendelson, 2002, Koltchinskii, 2001, Shalev-Shwartz and Ben-David, 2014, Mohri et al., 2012]. Empirical Rademacher complexity can be measured on the sample and implies generalization guarantees that improve based on structure exhibited by the sample. Therefore, the mechanism designer need not know *a priori* whether the buyers are additive or unit-demand in order to benefit from these improved guarantees. Rather, if their valuations are simple, this will be reflected in the empirical Rademacher complexity of the simple mechanism classes we study, and as a result, the mechanism designer can derive the strong guarantees afforded by simple buyers.

Tables 1, 2, and 3 display each mechanism class \mathcal{A} together with the generalization guarantee upper bound we derive in this work.¹ Formally, for a class of mechanisms \mathcal{A} , we denote our generalization guarantee using the notation $\epsilon_{\mathcal{A}}(N, \delta)$, where with probability at least $1 - \delta$ over the draw of a sample of size N , for all mechanisms in the class \mathcal{A} , average profit over the sample is $\epsilon_{\mathcal{A}}(N, \delta)$ -close to expected profit over the distribution. A generalization guarantee can easily be converted to a sample complexity guarantee $N_{\mathcal{A}}(\epsilon, \delta)$, where with probability at least $1 - \delta$ over the draw of a sample of size $N_{\mathcal{A}}(\epsilon, \delta)$, for every mechanism in \mathcal{A} , average profit is ϵ -close to expected profit.

For several of the mechanism classes we study, we improve over the previously best-known generalization guarantees. First of all, we prove generalization bounds for several “simple” mechanism classes when the buyers have valuation profiles not studied by Morgenstern and Roughgarden

¹For item pricing mechanisms, our bound for general valuations and anonymous prices matches that by Morgenstern and Roughgarden whenever $n \leq O(2^m)$ [Morgenstern and Roughgarden, 2016]. The same holds for additive buyers with anonymous prices whenever $n \leq O(m)$ and for non-anonymous prices for all values of n and m . Finally our bound for general valuations with non-anonymous buyers improves upon that by Morgenstern and Roughgarden.

Mechanism class \mathcal{A}	Reference theorem and valuation category	Price category	$\epsilon_{\mathcal{A}}(\delta, N)$ upper bound
Item pricing	Theorem 4.11: General	Anonymous	$O\left(U\sqrt{\frac{m(m+\log n)}{N}} + U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
		Non-anonymous	$O\left(U\sqrt{\frac{nm(m+\log n)}{N}} + U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
	Theorem 4.8: Additive	Anonymous	$O\left(U\sqrt{\frac{m\log(nm)}{N}} + U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
		Non-anonymous	$O\left(U\sqrt{\frac{nm\log(nm)}{N}} + U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
	Theorem 4.9: Unit-demand	Anonymous	$O\left(U\sqrt{\frac{m\log(nm)}{N}} + U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
		Non-anonymous	$O\left(U\sqrt{\frac{nm\log(nm)}{N}} + U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
\mathcal{B} -pricing	Theorem 4.10: General	Anonymous	$O\left(U\sqrt{\frac{ \mathcal{B} \log(m+\log(n \mathcal{B}))}{N}} + U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
		Non-anonymous	$O\left(U\sqrt{\frac{1}{N}\sum_{i=1}^n \mathcal{B}_i \left(m+\log\left(n\sum_{i=1}^n \mathcal{B}_i \right)\right)} + U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$

Table 2: Generalization guarantees for single-unit pricing mechanisms.

[Morgenstern and Roughgarden, 2016]. For example, Morgenstern and Roughgarden studied item pricing mechanisms when the buyers have general and additive valuations, and we prove improved bounds when the buyers have unit-demand valuations. Similarly, we provide generalization bounds for VCG auctions with item reserves when the buyers have general valuations, thus generalizing Morgenstern and Roughgarden’s analysis, which covered buyers with additive valuations.² Finally, for the class of affine maximizer auctions with single-unit supply, Balcan et al. gave a generalization bound of $O\left(\frac{n^{m+2}(H_{\bar{w}}H_v+H_\lambda)}{H_{\underline{w}}}\sqrt{\frac{m\log n}{N}}\left(\frac{n\hat{H}_v(nH_{\bar{w}}+H_\lambda)}{H_{\underline{w}}}\right) + \sqrt{n^m\log N}\right) + U\sqrt{\frac{\log(1/\delta)}{N}}$, where $H_{\bar{w}}, H_{\underline{w}}, H_v, H_\lambda$, and \hat{H}_v are constants bounding the range of the parameter search space [Balcan et al., 2016]. The bound in this paper is significantly simpler, applies to the multi-unit setting, and does not rely on the range of the parameter space. We also improve over Balcan et al.’s generalization bound of $O\left(U\sqrt{\frac{m^3\log n}{N}} + U\sqrt{\frac{\log(1/\delta)}{N}}\right)$ for mixed bundling auctions with reserve prices (MBARPs) and our proof is significantly simpler [Balcan et al., 2016].

Key challenges. A major strength of our generalization guarantees is their applicability to any algorithm that determines the optimal mechanism over the sample, a nearly optimal approximation, or any other black box procedure. For example, our results apply to any algorithm that uses samples to optimize over the classes of AMAs or virtual valuation combinatorial auctions (VVCAs), such as those developed by Sandholm and Likhodedov [Sandholm and Likhodedov, 2015]. However, generalization guarantees over the full classes of AMAs and VVCAs have proven pessimistic,

²Morgenstern and Roughgarden’s pseudo-dimension upper bounds imply Rademacher complexity upper bounds, since for any class \mathcal{F} with a finite pseudo-dimension of d and any sample \mathcal{S} of size N , the empirical Rademacher complexity is at most $O\left(\sqrt{d/N}\right)$ [Dudley, 1967].

Reference theorem and auction class \mathcal{A}	$\epsilon_{\mathcal{A}}(\delta, N)$ upper bound
Theorem 4.11: AMAs	$O\left(U\sqrt{\frac{1}{N}\prod_{i=1}^m(\kappa_i+n)}\sum_{i=1}^m\log\binom{\kappa_i+n}{n}+U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
Theorem 4.11: VVCAs	$O\left(U\sqrt{\frac{n}{N}\prod_{i=1}^m\kappa_i}\sum_{i=1}^m\log\binom{\kappa_i+n}{n}+U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
Theorem 4.11: λ -auctions	$O\left(U\sqrt{\frac{1}{N}\prod_{i=1}^m(\kappa_i+n)}\sum_{i=1}^m\log\binom{\kappa_i+n}{n}+U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
Theorem 4.14: k -sparse AMAs	$O\left(U\sqrt{\frac{n+k}{N}\left(\log(n+k)+\sum_{i=1}^m\log\binom{\kappa_i+n}{n}\right)}+U\sqrt{\frac{1}{N}\left(k\sum_{i=1}^m\log\binom{\kappa_i+n}{n}+\log\frac{1}{\delta}\right)}\right)$
Theorem 4.12: \mathcal{O} -boosted AMAs	$O\left(U\sqrt{\frac{n+ \mathcal{O} }{N}\left(\log(n+ \mathcal{O})+\sum_{i=1}^m\log\binom{\kappa_i+n}{n}\right)}+U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
Theorem 4.14: k -sparse λ -auctions	$O\left(U\sqrt{\frac{(n+k)\log(n+k)}{N}}+U\sqrt{\frac{1}{N}\left(k\sum_{i=1}^m\log\binom{\kappa_i+n}{n}+\log\frac{1}{\delta}\right)}\right)$
Theorem 4.13: \mathcal{O} -boosted λ -auctions	$O\left(U\sqrt{\frac{ \mathcal{O} \log(n \mathcal{O})}{N}}+U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
Theorem 4.15: VCG auctions with bundle reserves over \mathcal{B}	$O\left(U\sqrt{\frac{ \mathcal{B} \log(n \mathcal{B})}{N}}U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$
Theorem 4.16: MBARPs	$O\left(U\sqrt{\frac{m(\log n+m)}{N}}U\sqrt{\frac{1}{N}\log\frac{1}{\delta}}\right)$

Table 3: Generalization guarantees for auction classes. The variables κ_i for all $i \in [m]$ are upper bounds on the number of units available, as described in Section 3.

growing exponentially with the number of items [Balcan et al., 2016]. The SPM framework can ameliorate these worst-case bounds, but depends on a set of hierarchies breaking up these complex mechanism classes. The development of these hierarchies is the first challenge we address in this work. We require that these hierarchies be both economically sensible, permitting the mechanism designer to incorporate his prior knowledge about the buyers, and finely grained, allowing the designer to utilize the SPM framework and thereby determine a class in the hierarchy with the optimal tradeoff between profit and generalization.

Moreover, we favor hierarchies whose lower layers are much less complex than the full classes of AMAs and VVCAs, which raises a second key challenge: understanding the structure of these mechanism classes in order to bound their empirical Rademacher complexity. We observe that these mechanism classes are unlike many well-understood function classes in machine learning, where there is typically a simple connection between the parameter space and hypothesis space. For example, in linear or polynomial regression, there is a straight-forward mapping from the parameters of a hypothesis to its output on a given example, and a small change in parameters will lead to a predictable change in output. In our context, we discover that on a given bidding instance, the mechanism parameter space is splintered into regions such that the profit of a mechanism defined by parameters from one region is essentially unrelated to that of another. Roughly speaking, the number of significantly different profit functions over the range of parameters translates to the empirical Rademacher complexity of the mechanism class over the sample. As a result, we need to

understand the structure of the partition induced on the parameter space by the sample in order to derive our generalization guarantees and traditional methods of bounding empirical Rademacher complexity do not apply. In light of our analyses, we believe that this work is of auction-theoretic interest, since we develop novel insights into the structure of these well-studied mechanisms, as well as learning-theoretic interest, since we study function classes that diverge from those well-understood classes typically found in machine learning.

2 Related literature

The sample complexity of revenue maximization has proven to be a fruitful research area over the past several years [Elkind, 2007, Cole and Roughgarden, 2014, Huang et al., 2015, Medina and Mohri, 2014, Morgenstern and Roughgarden, 2015, Roughgarden and Schrijvers, 2016, Devanur et al., 2016, Morgenstern and Roughgarden, 2016, Balcan et al., 2016, Syrgkanis, 2017]. Unlike this work, previous work has primarily concentrated on the single-item setting, with the exception of work by Morgenstern and Roughgarden, Balcan et al., and Syrgkanis [Morgenstern and Roughgarden, 2016, Balcan et al., 2016, Syrgkanis, 2017]. Learning theory tools such as pseudo-dimension and Rademacher complexity were used by Medina and Mohri, Balcan et al., Morgenstern and Roughgarden, and Syrgkanis to prove strong guarantees [Medina and Mohri, 2014, Balcan et al., 2016, Morgenstern and Roughgarden, 2015, Morgenstern and Roughgarden, 2016, Syrgkanis, 2017]. In a similar direction, Feldman et al. and Hsu et al. have developed bounds on the sample complexity of welfare-optimal item pricing mechanisms [Feldman et al., 2015, Hsu et al., 2016]. Earlier work of Balcan et al. addressed sample complexity results for revenue maximization in unrestricted supply settings [Balcan et al., 2008].

Morgenstern and Roughgarden proved both generalization guarantees and bounded the difference between the maximum revenue achievable via a single-item t -level auction and the revenue achievable via Myerson’s optimal auction [Morgenstern and Roughgarden, 2015]. In this way, the authors study the trade-off between generalization and revenue guarantees in the single-item case, though in a worst-case sense. In contrast, we analyze this tradeoff for the multi-item, multi-unit setting and in a data-dependent, beyond-worst-case sense.

Sample-based mechanism design is closely related to *automated mechanism design*, a research area where revenue maximization in multi-item settings is a central topic. The goal is to design algorithms which take as input information about a set of buyers and return a mechanism that will extract high revenue from those buyers [Conitzer and Sandholm, 2002, Sandholm, 2003]. The input information about the buyers could be an explicit description of their priors or, as in this paper, a set of samples from their priors [Likhodedov and Sandholm, 2004, Likhodedov and Sandholm, 2005, Sandholm and Likhodedov, 2015]. Recently, automated mechanism design has been explored for applications beyond revenue maximization, such as mechanism design without money and more general assignment problems [Narasimhan and Parkes, 2016, Narasimhan et al., 2016].

Affine maximizer auctions were introduced by Roberts, who proved that they are the only *ex post* strategy-proof mechanisms over unrestricted domains of valuations [Roberts, 1979]. Lavi et al. went on to prove that under certain natural assumptions, every incentive compatible multi-item auction is an “almost” affine maximizer, i.e. an AMA for sufficiently high valuations [Lavi et al., 2003]. Lavi et al. conjecture that the “almost” qualifier is merely technical, and can be removed in future research.

There is a wealth of work on characterizing the optimal multi-item auction for restricted settings and designing mechanisms which achieve high, if not optimal revenue in specific contexts (e.g. [Daskalakis and Weinberg, 2012, Cai et al., 2012, Cai et al., 2013]). Revenue-maximizing mechanism

design complements a research area which strives to answer the question: can *simple* mechanisms achieve near-optimal revenue [Hartline and Roughgarden, 2009]? Recent work has shown that many of the mechanism classes we study achieve constant-factor approximations to the optimal revenue in certain settings, including item pricing mechanisms and grand bundle pricing mechanisms (e.g. [Chawla et al., 2007, Chawla et al., 2010, Hart and Nisan, 2012, Kleinberg and Weinberg, 2012, Babaioff et al., 2014, Cai et al., 2016]) and a special case of \mathcal{B} -pricing mechanisms known as *partition mechanisms* [Babaioff et al., 2014, Rubinstein, 2016]. We note that these approximation guarantees only hold in certain settings, such as when the buyers are unit-demand or subadditive. This highlights a strength of our results: our data-dependent bounds adapt nicely to the structure of the buyers, and we do not need to know this structure *a priori* in order to derive the improved guarantees that simple buyers afford. Further, even if the buyers do not have well-structured valuations, we can nevertheless bound generalizability.

3 Preliminaries, notation, and the mechanism hierarchies

We consider the problem of selling m heterogeneous goods to n consumers. We denote a bundle of goods as a quantity vector \mathbf{q} and we denote the i^{th} component of \mathbf{q} as $q[i]$. Accordingly, the bundle consisting only of the i^{th} item is denoted by the standard basis vector \mathbf{e}_i , where $e[i] = 1$ and $e[j] = 0$ for all $j \neq i$. In the single-unit case, $\mathbf{q} \in \{0, 1\}^m$. Each consumer $i \in [n]$ has a valuation function v_i over bundles of goods. If one bundle \mathbf{q}_0 is contained within another bundle \mathbf{q}_1 (i.e., $q_0[j] \leq q_1[j]$ for all $j \in [m]$), then $v_i(\mathbf{q}_0) \leq v_i(\mathbf{q}_1)$ and $v_i(\mathbf{0}) = 0$. We denote an allocation as $Q = (\mathbf{q}_1, \dots, \mathbf{q}_n)$ where \mathbf{q}_i is the bundle of goods that consumer i receives under allocation Q . There is an unknown distribution \mathcal{D} over buyers' values.

In the multi-unit setting, we assume that the seller has a cost function $c(Q)$ which equals the amount it costs the seller to produce the goods allocated under Q . We need to make a simple assumption about the total number of units demanded by the consumers or else the pseudo-dimension of many mechanism classes we study will be infinite. We assume that the cost function naturally caps the total number of units of each item that the producer will supply. In other words, for each item i , there is some cap κ_i such that for all valuation functions in the support of \mathcal{D} , it will cost the producer more to produce κ_i units of item i than the buyers are willing to pay. Formally, this means that there exists a vector $(\kappa_1, \dots, \kappa_m)$ such that for all valuation vectors in the support of \mathcal{D} and for all allocations $Q = (\mathbf{q}_1, \dots, \mathbf{q}_m)$, if there exists an item i such that $\sum_{j=1}^n q_j[i] > \kappa_i$, then $\sum_{j=1}^n v_j(\mathbf{q}_j) - c(Q) < 0$.

In the multi-unit setting, there are $\prod_{i=1}^m \binom{\kappa_i+n}{n}$ different allocations possible. This is because the number of ways to allocate at most κ_i unlabeled units among n consumers is $\binom{\kappa_i+n}{n}$. In the single-unit setting, the number of different allocations is $(n+1)^m$ because each of the m items can go to one of the n buyers or to no one. We denote the total number of different quantity vectors as K . In the multi-unit setting, $K = \prod_{i=1}^m (\kappa_i + 1)$ and in the single-unit setting, $K = 2^m$. We use the notation $\mathbf{v}_1 = (v_1(\mathbf{q}_1), \dots, v_1(\mathbf{q}_K))$ and $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_n)$ to denote a vector of consumer valuation functions. In the multi-item setting, we study consumers with general valuations. In the single-item setting, we study consumers with general, additive ($v_i(\mathbf{q}) = \sum_{j:q[j]=1} v_i(\mathbf{e}_j)$), and unit-demand ($v_i(\mathbf{q}) = \max_{j:q[j]=1} v_i(\mathbf{e}_j)$) valuations.

We say that $\text{profit}_A(\mathbf{v})$ is the profit of a mechanism A on the valuation vector \mathbf{v} . Denoting the payment of any one consumer i under mechanism A given valuation vector \mathbf{v} as $p_{i,A}(\mathbf{v})$ and the outcome as $Q_A(\mathbf{v})$, we have that $\text{profit}_A(\mathbf{v}) = \sum_{i=1}^n p_{i,A}(\mathbf{v}) - c(Q_A(\mathbf{v}))$. Throughout this work, we study mechanisms parameterized by a vector $\mathbf{p} \in \mathbb{R}^d$ for some d , and we refer to the profit of the mechanism defined by \mathbf{p} on the valuation vector \mathbf{v} as $\text{profit}_{\mathbf{p}}(\mathbf{v})$. If we fix \mathbf{v} and consider the

profit as a function of \mathbf{p} , then we use the notation $\text{profit}_v(\mathbf{p})$. Finally, we define $U_{\mathcal{D},\mathcal{A}}$ to be an upper bound on the profit achievable by any mechanism in the class \mathcal{A} over the support of \mathcal{D} . For the sake of readability, we drop the subscript when the mechanism class and distribution are clear from context.

Many of the auction classes that we consider have a design based on the classic *Vickrey-Clarke-Groves mechanism (VCG)* [Vickrey, 1961, Clarke, 1971, Groves, 1973]. The VCG allocates the items such that the social welfare of the bidders, that is, the sum of their valuations for the allocated items, is maximized. Each winning bidder then pays her bid minus a “rebate” equal to the increase in welfare attributable to her presence in the auction. We note that every auction in the classes we study is incentive compatible, so we may assume that the bids equal the bidders’ valuations.

3.1 Mechanism classes

We now define the mechanism families in the hierarchies we study. See Figures 1 and 2 for the hierarchical organization of the mechanism classes, together with the papers that introduced each family.

3.1.1 Multi-unit pricing mechanisms

First, we describe the pricing mechanisms we study, beginning with multi-unit mechanisms. All of these mechanisms can utilize either *anonymous* or *non-anonymous* prices. Under an anonymous pricing mechanism each consumer is subject to the same prices as every other consumer. Under a non-anonymous pricing mechanism different consumers are subject to different prices. This typically results in higher profit.

Non-linear pricing mechanisms. This paper studies non-linear pricing under the *bundling interpretation* described in Chapter 4.3 of Wilson’s book on non-linear pricing [Wilson, 1993]. The mechanism designer sets a price per quantity vector \mathbf{q} denoted $p(\mathbf{q})$. Consumer j will purchase the bundle that maximizes $v_j(\mathbf{q}) - p(\mathbf{q})$.

Additively decomposable non-linear pricing mechanisms. There exists m functions $p_i : [\kappa_i] \rightarrow \mathbb{R}$ for all $i \in [m]$ such that for every quantity vector \mathbf{q} , $p(\mathbf{q}) = \sum_{i=1}^m p_i(q[i])$.

Two-part tariffs. A two-part tariff is a non-linear pricing scheme for a single good made up of two parts: 1) A fixed, up-front fee charged to each consumer who buys at least one unit of the good and 2) A price per unit bought. We denote a menu of M two-part tariffs as $\{(p_0^1, p_1^1), \dots, (p_0^M, p_1^M)\}$, where (p_0^i, p_1^i) is the i^{th} entry on the menu with a fixed fee p_0^i and a price per unit p_1^i . Each consumer can choose his payment plan among any of the M menu entries. Since there is only one item, we denote the cap on the total number of units the producer will supply as κ .

3.1.2 Single-unit pricing mechanisms

For single-unit pricing mechanisms, we assume that there is some fixed but arbitrary ordering on the consumers such that the first consumer in the ordering arrives first at the marketplace and buys the bundle of goods that maximizes his utility, then the next consumer in the ordering arrives at the marketplace and buys the bundle of remaining goods that maximizes his utility, and so on. We assume that this ordering over bidders is known to the mechanism designer. It is important for the mechanism designer to know this ordering in the single-unit setting because the order in which

heterogeneous consumers arrive at the marketplace could affect profit drastically. In this setting, we assume the seller does not have a cost function.

Bundle pricing mechanisms. First, we define bundle pricing mechanisms with non-anonymous reserve prices, which we refer to as $(\mathcal{B}_1, \dots, \mathcal{B}_n)$ -pricing mechanisms, where $\mathcal{B}_1, \dots, \mathcal{B}_n$ are sets of bundles. Each bundle (or equivalently, quantity vector) $\mathbf{q} \in \mathcal{B}_i$ has a buyer-specific reserve price $p_i(\mathbf{q})$. We require that the singleton set \mathbf{e}_j is in \mathcal{B}_i for each item $j \in [m]$ and each buyer $i \in [n]$ so that if $\mathbf{q} \notin \mathcal{B}_i$, then we define $p_i(\mathbf{q}) = \sum_{j:q[j]=1} p_i(\mathbf{e}_j)$. In the case of anonymous reserve prices, we refer to the mechanism as a \mathcal{B} -pricing mechanism, where $\mathcal{B} = \mathcal{B}_1 = \dots = \mathcal{B}_n$ and $p_1 = \dots = p_n$.

Item pricing mechanisms. Item pricing mechanisms are a special case of bundle pricing mechanisms where only the bundles consisting of a single item receive prices.

3.1.3 Multi-unit auction classes

Next we define the multi-unit auction classes we study in this work.

Affine maximizer auctions (AMAs). An AMA A is defined by a set of weights per bidder $w_j \in \mathbb{R}_{>0}$ and boosts per allocation $\lambda(Q) \in \mathbb{R}$. These parameters allow the mechanism designer to multiplicatively boost any bidder's bids by their corresponding weight and to increase the likelihood that any allocation Q is returned as the output of an auction by increasing $\lambda(Q)$. More concretely, the allocation Q^* of an AMA A is the one which maximizes the weighted social welfare, i.e. $Q^* = \operatorname{argmax} \left\{ \sum_{j=1}^n w_j v_j(\mathbf{q}_j) + \lambda(Q) - c(Q) \right\}$. The payment function of A has the same form as the VCG payment rule, with the parameters factored in to ensure incentive compatibility. For all $j \in [n]$, the payments are

$$p_{j,A}(\mathbf{v}) = \frac{1}{w_j} \left[\sum_{\ell \neq j} w_\ell v_\ell(\mathbf{q}_\ell^{-j}) + \lambda(Q^{-j}) - c(Q^{-j}) - \left(\sum_{\ell \neq j} w_\ell v_\ell(\mathbf{q}_\ell^*) + \lambda(Q^*) - c(Q^*) \right) \right],$$

where $Q^{-j} = \operatorname{argmax} \left\{ \sum_{\ell \neq j} w_\ell v_\ell(\mathbf{q}_\ell) + \lambda(Q) - c(Q) \right\}$.

λ -auctions. A λ -auction is an AMA where $w_1 = \dots = w_n = 1$.

\mathcal{O} -boosted AMAs and λ -auctions. Let \mathcal{O} be a set of allocations. The set of \mathcal{O} -boosted AMAs (respectively, λ -auctions) consists of all AMAs (respectively, λ -auctions) where only outcomes in \mathcal{O} are boosted. In other words, if $\lambda(Q) > 0$, then it must be that $Q \in \mathcal{O}$.

k -sparse AMAs and λ -auctions. The set of k -sparse AMAs (respectively, λ -auctions) consists of all AMAs (respectively, λ -auctions) where at most k outcomes are boosted. Notice that if $\mathcal{A}_{\mathcal{O}}$ is the set of \mathcal{O} -boosted AMAs for some set \mathcal{O} , then the set of k -sparse AMAs is equal to the union of all $\mathcal{A}_{\mathcal{O}}$ where $|\mathcal{O}| \leq k$.

Virtual valuation combinatorial auctions (VVCAs). VVCAs are a subset of AMAs. The defining characteristic of a VVCA is that each $\lambda(Q)$ is split into n terms such that $\lambda(Q) = \sum_{i=1}^n \lambda_i(Q)$ where $\lambda_i(Q) = c_{i,\mathbf{q}}$ for all allocations Q that give Bidder i exactly bundle \mathbf{q} .

3.1.4 Single-unit auction classes

As in the case of single-unit pricing mechanisms, we assume the seller does not have a cost function.

VCG auctions with anonymous bundle reserves. Each auction is defined by a set of bundles \mathcal{B} and each bundle $\mathbf{q} \in \mathcal{B}$ has a reserve price $p(\mathbf{q})$. We require that \mathcal{B} contains the singleton bundle \mathbf{e}_j for all items $j \in [m]$. This way, if a bundle \mathbf{q} is not in \mathcal{B} , then we define the reserve price of \mathbf{q} to be $p(\mathbf{q}) = \sum_{j:q[j]=1} p(\mathbf{e}_j)$. For an outcome Q , let \mathbf{q}_Q be the bundle of items not allocated. The allocation Q^* of this auction is the one that maximizes $\sum_{i=1}^n v_i(\mathbf{q}_i) + p(\mathbf{q}_Q)$. Bidder j pays $\sum_{i \neq j} v_i(\mathbf{q}_i^{-j}) + p(\mathbf{q}_{Q^{-j}}) - \left(\sum_{i \neq j} v_i(\mathbf{q}_i^*) + p(\mathbf{q}_{Q^*}) \right)$, where Q^{-j} is the allocation that maximizes $\sum_{i \neq j} v_i(\mathbf{q}_i) + p(\mathbf{q}_Q)$.

Mixed bundling auctions with reserve prices (MBARPs). The class of mixed bundling auctions (MBAs) is parameterized by a constant $c \geq 0$ which can be seen as a discount for any bidder who receives the grand bundle. Formally, the c -MBA is the λ -auction with $\lambda(Q) = c$ if some bidder receives the grand bundle in allocation Q and 0 otherwise.³ MBARPs are identical to MBAs though with *reserve prices*. In a single-item VCG auction (i.e. second price auction) with a reserve price, the item is only sold if the highest bidder's bid exceeds the reserve price, and the winner must pay the maximum of the second highest bid and the reserve price. To generalize this intuition to the multi-item case, we enlarge the set of agents to include the seller, who is now Bidder 0 and whose valuation for a set of items is the set's reserve price. Working with this expanded set of agents, the bidder weights are all 1 and the λ terms are the same as in the standard MBA setup. Importantly, the seller makes no payments, no matter her allocation. More formally, given a vector of valuation functions \mathbf{v} , the MBARP allocation is $Q^* = \operatorname{argmax} \{ \sum_{i=0}^n v_i(\mathbf{q}_i) + \lambda(Q) \}$. For each $i \in \{1, \dots, n\}$, Bidder i 's payment is

$$p_{A,i}(\mathbf{v}) = \sum_{j \in \{0, \dots, n\} \setminus \{i\}} v_j(\mathbf{q}_j^{-i}) + \lambda(Q^{-i}) - \sum_{j \in \{0, \dots, n\} \setminus \{i\}} v_j(\mathbf{q}_j^*) - \lambda(Q^*),$$

where

$$Q^{-i} = \operatorname{argmax} \left\{ \sum_{j \in \{0, \dots, n\} \setminus \{i\}} v_j(\mathbf{q}_j) + \lambda(Q) \right\}.$$

3.2 Learning theory background

In this work, we are concerned with *generalization guarantees* for *uniformly learnable* classes of mechanisms. We begin with a formal definition of these concepts.

Definition 3.1 (Uniformly learnable). *A class \mathcal{A} is uniformly learnable if there exists a function $\epsilon_{\mathcal{A}}(N, \delta)$ such that for any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the draw of a sample of size N from the distribution over buyers' values \mathcal{D} , for any mechanism A in \mathcal{A} , the difference between the average profit of A over the sample and its expected profit over \mathcal{D} is at most $\epsilon_{\mathcal{A}}(N, \delta)$.*

The function $\epsilon_{\mathcal{A}}(N, \delta)$ is known as a *generalization guarantee* for learning over the class \mathcal{A} . Learning theorists have developed many tools to help derive generalization guarantees such as *empirical Rademacher complexity* and *pseudo-dimension* which quantify the "complexity" of a class of

³Balcan et al. proved that if \mathcal{A} is the class of MBAs, then $\epsilon_{\mathcal{A}}(N, \delta) \leq O\left(U \sqrt{\frac{1}{N} \log \frac{1}{\delta}}\right)$ [Balcan et al., 2016].

functions. First, we formally define empirical Rademacher complexity in terms of mechanism profit functions and then provide a more intuitive notion of the quantity that it measures. Throughout this paper, we slightly abuse notation and use \mathcal{A} both to refer to a class of mechanisms and the class of those mechanisms' profit functions.

Definition 3.2 (Empirical Rademacher complexity). *The empirical Rademacher complexity of \mathcal{A} with respect to the sample $\mathcal{S} = \{\mathbf{v}^1, \dots, \mathbf{v}^N\}$ is defined as*

$$\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) = \mathbb{E}_{\sigma} \left[\sup_{A \in \mathcal{A}} \frac{1}{N} \sum_{i=1}^N \sigma_i \cdot \mathbf{profit}_A(\mathbf{v}^i) \right],$$

where $\sigma = (\sigma_1, \dots, \sigma_N)^{\top}$, with $\sigma_i \sim U(\{-1, 1\})$.

Intuitively, the supremum measures, for a given sample \mathcal{S} and Rademacher vector σ , the maximum correlation between $\mathbf{profit}_A(\mathbf{v}^i)$ and σ_i over all $A \in \mathcal{A}$. Taking the expectation over σ , we can then say that the empirical Rademacher complexity of \mathcal{A} measures the ability of profit functions from \mathcal{A} (when applied to a fixed sample \mathcal{S}) to fit random noise. We are able to derive strong sample complexity bounds with empirical Rademacher complexity. For example, the following bound is well-known.

Theorem 3.3 ([Shalev-Shwartz and Ben-David, 2014]). *Suppose that for any sample \mathcal{S} of size N , $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq M$ for some $M \in \mathbb{R}$. Then*

$$\epsilon_{\mathcal{A}}(N, \delta) \leq 2M + 4U \sqrt{\frac{2}{N} \ln \frac{4}{\delta}}.$$

Moreover, for a sample \mathcal{S} , suppose $\mathbf{profit}_{\hat{A}} \in \mathcal{A}$ is a profit function with the maximum average value over \mathcal{S} and \mathbf{profit}_{A^*} is a function with maximum expected value with respect to the distribution \mathcal{D} . Then with probability at least $1 - \delta$,

$$\mathbb{E}_{\mathbf{v} \sim \mathcal{D}} [\mathbf{profit}_{A^*}(\mathbf{v})] - \mathbb{E}_{\mathbf{v} \sim \mathcal{D}} [\mathbf{profit}_{\hat{A}}(\mathbf{v})] \leq 2\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) + 5U \sqrt{\frac{2}{N} \ln \frac{8}{\delta}}.$$

The *pseudo-dimension* of a mechanism class \mathcal{A} is another means of analyzing the complexity of \mathcal{A} , and thereby deriving useful generalization guarantees. To define pseudo-dimension, let $\mathcal{S} = \{\mathbf{v}^1, \dots, \mathbf{v}^N\}$ be a sample drawn from \mathcal{D} and let $(r_1, \dots, r_N) \in \mathbb{R}^N$ be a set of *targets*. We say that (r_1, \dots, r_N) *witnesses* the shattering of \mathcal{S} by \mathcal{A} if for all $T \subseteq \mathcal{S}$, there exists some mechanism $A_T \in \mathcal{A}$ such that for all $\mathbf{v}^i \in T$, $\mathbf{profit}_{A_T}(\mathbf{v}^i) \leq r_i$ and for all $\mathbf{v}^i \notin T$, $\mathbf{profit}_{A_T}(\mathbf{v}^i) > r_i$. If there exists some \mathbf{r} that witnesses the shattering of \mathcal{S} by \mathcal{A} , then we say that \mathcal{S} is *shatterable* by \mathcal{A} . Finally, the pseudo-dimension of \mathcal{A} is the size of the largest set that is shatterable by \mathcal{A} . The following theorem is well-known.

Theorem 3.4 ([Dudley, 1967]). *If the pseudo-dimension of a class \mathcal{A} is d , then for any sample \mathcal{S} of size N , $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) = O\left(U \sqrt{\frac{d}{N}}\right)$.*

4 Generalization guarantees

In this section, we uncover structural connections between the mechanism classes that we study which allow us to abstractly reason about their Rademacher complexity. The profit function depends first and foremost on the allocation made to each of the buyers. In the case of pricing

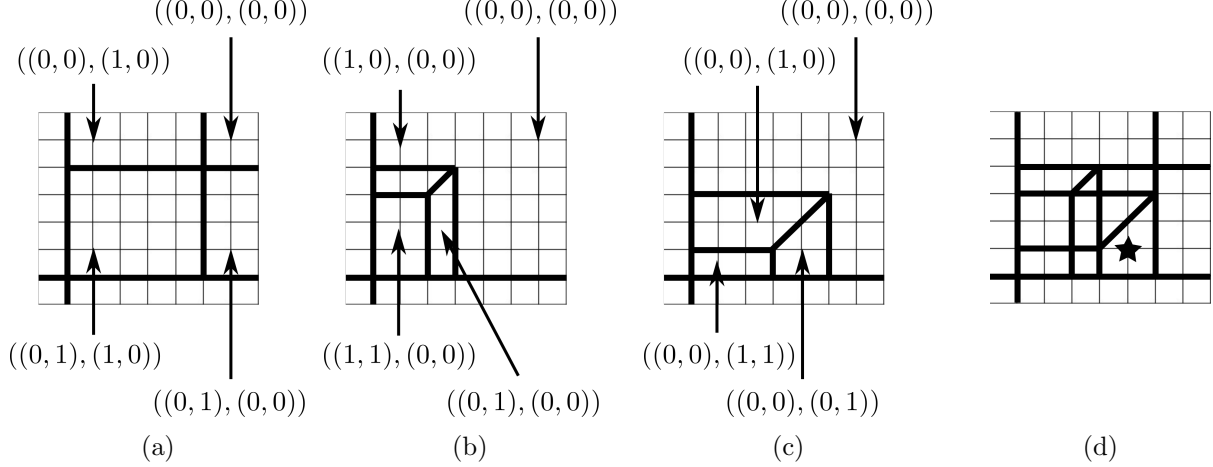


Figure 3: An illustration of the partition induced on $\mathbb{R}_{\geq 0}^2$ by the scenario described in Example 4.1.

mechanisms, profit is completely determined once these allocations are set. Meanwhile, the auction classes we study are variants on the classic VCG mechanism, so the profit also depends on the n allocations that would be made without each buyer's participation in turn. Once all of these $n + 1$ allocations are fixed, the payment function is simple — it is a linear function of the auction's parameters and the buyers' valuations for each of the $n + 1$ allocations. The challenge of bounding the complexity of these classes of mechanisms therefore comes down to understanding the connection between a mechanism's parameters and the resulting allocation, or in the case of auctions, the resulting $n + 1$ relevant allocations on a given valuation vector. We can often partition the mechanism parameter space into a finite number of regions over which these allocations are fixed for a given valuation vector and over which the profit is a fixed linear function. The following example illustrates the partition induced on $\mathbb{R}_{\geq 0}^2$ by a specific valuation profile and the VCG auction with anonymous item reserve prices.

Example 4.1. *Suppose that there are two bidders and two items for sale in the single-unit setting. Bidder 1 has the valuation function $v_1((1,0)) = 3$, $v_1((0,1)) = 4$, and $v_1((1,1)) = 6$. Bidder 2 has the valuation function $v_2((1,0)) = 5$, $v_2((0,1)) = 3$, and $v_2((1,1)) = 6$. We analyze the class of VCG auctions with anonymous item reserves, so the auction's parameters are in $\mathbb{R}_{\geq 0}^2$. First, we show that $\mathbb{R}_{\geq 0}^2$ can be partitioned into 4 regions so that if $(p((1,0)), p((0,1)))$ and $(p'((1,0)), p'((0,1)))$ are from the same region, the outcome of the resulting VCG auction will be the same. To this end, notice that the allocation $((0,1), (1,0))$, where Bidder 1 receives item 2 and Bidder 2 receives item 1, will be the allocation of any auction so long as $v_1((0,1)) + v_2((1,0)) \geq v_1(\mathbf{q}_1) + v_2(\mathbf{q}_2) + p(\mathbf{q}_Q)$ for any allocation Q . Simple calculations show that this will be the case so long as $p((1,0)) \leq 5$ and $p((0,1)) \leq 4$. Along the same lines, the allocation will be $((0,1), (0,0))$ so long as $p((1,0)) \geq 5$ and $p((0,1)) \leq 4$ and the allocation will be $((0,0), (1,0))$ so long as $p((1,0)) \leq 5$ and $p((0,1)) \geq 4$. Otherwise, the allocation is $((0,0), (0,0))$. This is illustrated by Figure 3(a), where $p((1,0))$ scales along the x -axis and $p((0,1))$ scales along the y -axis.*

In Figure 3(b) we perform the same analysis in the case where Bidder 2 is not present in the auction and Figure 3(c) shows the same analysis when Bidder 1 is not present in the auction. Finally, Figure 3(d) displays the overlay of these three partitions. By construction, the profit in each region is a linear function of the prices. For example, consider the region marked by a star. From Figure 3(a), we know that the allocation of auction defined by any price pair from this region is $((0,1), (1,0))$. From Figure 3(b), the allocation without Bidder 2's participation is $((0,0), (0,0))$,

and from Figure 3(c), the allocation without Bidder 1's participation is $((0, 0), (0, 1))$. Therefore, Bidder 1 pays $v_2((0, 1)) + p((1, 0)) - v_2((1, 0)) = p((1, 0)) - 2$ and Bidder 2 pays $p((1, 0)) + p((0, 1)) - v_1((0, 1)) = p((1, 0)) + p((0, 1)) - 4$, so the profit in this region is $2p((1, 0)) + p((0, 1)) - 6$, a linear function of $(p((1, 0)), p((0, 1)))$.

In this section, we show that the empirical Rademacher complexity of a class of mechanisms can be bounded by a function of both the number of regions an arbitrary sample induces on the parameter space and the dimension of the parameter space. We find that mechanisms that have been deemed “simple” in the literature often admit partitions that are small and easy to characterize over a low-dimensional parameter space. We also demonstrate that the complexity of the partition induced by a sample, and thus the empirical Rademacher complexity of a class of mechanisms over that sample, is intricately bound to the complexity of the buyers. Since empirical Rademacher complexity is a measurement the mechanism designer can make on the sample, he can count or bound the number of regions induced by the sample on the parameter space and calculate his generalization guarantee accordingly. In this way, he does not need to make any assumption about whether the distribution is over simple buyers or not.

We now present our main empirical Rademacher complexity guarantee, which we then instantiate for a wide variety of mechanisms. These bounds immediately imply the generalization guarantee upper bounds listed in Tables 1, 2, and 3 by Theorem 3.3.

Theorem 4.2. *Let \mathcal{F} be a class of mechanism profit functions such that each mechanism is parameterized by a vector $\mathbf{p} \in \mathcal{X} \subseteq \mathbb{R}^d$ and let $\mathcal{S} = \{\mathbf{v}^1, \dots, \mathbf{v}^N\}$ be a sample of valuation vectors. Suppose that there are at most r regions partitioning \mathcal{X} such that on each sample \mathbf{v}^i , $\text{profit}_{\mathbf{v}^i}(\mathbf{p})$ is linear as \mathbf{p} ranges over a single region R , i.e. $\text{profit}_{\mathbf{v}^i}(\mathbf{p}) = \mathbf{u}_R^i \cdot \mathbf{p} + a_R$ for some $\mathbf{u}_R^i \in \mathbb{R}^d$ and $a_R \in \mathbb{R}$. Then $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{F}) \leq O\left(U\sqrt{(\log r + d \log d)/N}\right)$.*

Proof. Suppose that N is the size of the largest shatterable set, so $Pdim(\mathcal{F}) = N$. As we see in Lemma 4.3, it must be that $N \leq d \log N + \log(rd)$. We can deduce that $N \leq 4d \log(2d) + 2 \log(rd) = O(d \log d + \log r)$ by using the fact that for $c \geq 1$ and $b > 0$, if $x \leq c \log x + b$, then $x \leq 4c \log(2c) + 2b$ [Shalev-Shwartz and Ben-David, 2014], and setting $x = N$, $c = d$, and $b = \log(rd)$. Therefore, $Pdim(\mathcal{F}) \leq O(d \log d + \log r)$. The Rademacher complexity bound follows from Theorem 3.4. \square

Lemma 4.3. *Let \mathcal{F} be a class of mechanism profit functions such that each mechanism is parameterized by a vector $\mathbf{p} \in \mathcal{X} \subseteq \mathbb{R}^d$ and let $\mathcal{S} = \{\mathbf{v}^1, \dots, \mathbf{v}^N\}$ be a shatterable sample of valuation vectors. Suppose that there are at most r regions partitioning \mathcal{X} such that on each sample \mathbf{v}^i , $\text{profit}_{\mathbf{v}^i}(\mathbf{p})$ is linear as \mathbf{p} ranges over a single region R , i.e. $\text{profit}_{\mathbf{v}^i}(\mathbf{p}) = \mathbf{u}_R^i \cdot \mathbf{p} + a_R$ for some $\mathbf{u}_R^i \in \mathbb{R}^d$ and $a_R \in \mathbb{R}$. Then $N \leq d \log N + \log(rd)$*

Proof. Since \mathcal{S} is shatterable, there must be a set of N targets r_1, \dots, r_N that witness the shattering of \mathcal{S} by \mathcal{F} . In other words, for all $T \subseteq [N]$, there exists a parameter vector \mathbf{p}_T such that $\text{profit}_{\mathbf{v}^i}(\mathbf{p}_T) \leq r_i$ if $i \in T$ and $\text{profit}_{\mathbf{v}^i}(\mathbf{p}_T) > r_i$ if $i \notin T$. We refer to the set of these 2^N special vectors \mathbf{p}_T as P . Since the regions R partition the parameter space \mathcal{X} , we know that each parameter vector in P comes from some region R . We will now count the maximum number of vectors in P from a single arbitrary region R .

Fix a region R and a single sample \mathbf{v}^i with its corresponding witness. For a single sample \mathbf{v}^i , we know that $\text{profit}_{\mathbf{v}^i}(\mathbf{p}) = \mathbf{u}_R^i \cdot \mathbf{p} + a_R$. The set of all \mathbf{p} such that $\mathbf{u}_R^i \cdot \mathbf{p} + a_R$ is less than its witness r_i form a halfspace defined by the hyperplane $\mathbf{u}_R^i \cdot \mathbf{p} + a_R = r_i$. There exists one such hyperplane per sample, leading to a total of N hyperplanes. These hyperplanes induce a partition of R consisting of at most dN^d cells [Buck, 1943]. By construction, as we range \mathbf{p} over any one

cell, for all $i \in [N]$, $\text{profit}_{\mathbf{v}^i}(\mathbf{p})$ can only be either less than its witness or greater than its witness. Therefore, at most one parameter vector in P can come from each cell, so the maximum number of vectors in P from R is dN^d . We chose R to be arbitrary, so this bound holds for all r regions. Therefore, $2^N = |P| \leq rdN^d$, which means that $N \leq d \log N + \log(rd)$. \square

We observe that the partitions induced by the mechanism classes we study can be described as *hyperplane arrangements*. In particular, let \mathbf{v} be a fixed valuation vector and consider any mechanism class featured in this paper that is defined by a parameter vector $\mathbf{p} \in \mathcal{X}$. We show that we can identify a set \mathcal{H} of hyperplanes such that the regions induced on \mathcal{X} over which $\text{profit}_{\mathbf{v}}(\mathbf{p})$ is linear are the connected components of $\mathcal{X} \setminus \mathcal{H}$. Buck proved that if \mathcal{X} is a subset of \mathbb{R}^d , then the number of connected components is bounded by $\sum_{i=1}^d \binom{|\mathcal{H}|}{i} \leq d|\mathcal{H}|^d$ [Buck, 1943]. We use this fact to prove the following corollary of Theorem 4.2 in this context.

Corollary 4.4. *Suppose that for every $\mathbf{v} \in \mathcal{S}$, there are at most t hyperplanes that partition \mathcal{X} into regions such that $\text{profit}_{\mathbf{v}}(\mathbf{p})$ is linear over any given region. Then $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{F}) \leq O\left(U\sqrt{d \log(dt)/N}\right)$.*

Proof. First, each sample in \mathcal{S} comes with at most t hyperplanes which partition \mathcal{X} into regions where the corresponding profit function is linear. We denote the partition corresponding to \mathbf{v}^i by \mathcal{P}_i . Next, the union of all Nt hyperplanes also partition \mathcal{X} , and we refer to this partition as \mathcal{P} . Notice that for each region $R \in \mathcal{P}$ and each index $i \in [N]$, R is fully contained within a single region in \mathcal{P}_i . Therefore, for each region in $R \in \mathcal{P}$ and for all $i \in [N]$, $\text{profit}_{\mathbf{v}^i}(\mathbf{p})$ is a fixed linear function. This means that $r = |\mathcal{P}| \leq d(Nt)^d$ [Buck, 1943]. By Lemma 4.3 we see that if \mathcal{S} is shatterable, then $N \leq d \log N + \log(rd) \leq d \log N + 2 \log d + d \log(Nt) = 2d \log N + 2 \log d + d \log t$. Recall that for $c \geq 1$ and $b > 0$, if $x \leq c \log x + b$, then $x \leq 4c \log(2c) + 2b$ [Shalev-Shwartz and Ben-David, 2014]. We set $x = N$, $c = 2d$, and $b = 2 \log d + d \log t$ and derive that $N \leq 8d \log(4d) + 2(2 \log d + d \log t) = O(d \log(dt))$. \square

We now use these results to prove bounds on the empirical Rademacher complexity of an array of mechanism classes.

4.1 Multi-unit pricing mechanisms

We begin by applying Theorem 4.2 to multi-unit pricing mechanisms. Our most general result is for non-linear pricing mechanisms.

Theorem 4.5. *Let \mathcal{A} and \mathcal{A}' be the classes of non-linear pricing mechanisms with anonymous and non-anonymous prices, respectively. Then for any sample \mathcal{S} of size N ,*

$$\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(U\sqrt{\frac{1}{N}\prod_{i=1}^m \kappa_i \left(\log n + \sum_{i=1}^m \log \kappa_i\right)}\right)$$

and

$$\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}') \leq O\left(U\sqrt{\frac{n}{N}\prod_{i=1}^m \kappa_i \left(\log n + \sum_{i=1}^m \log \kappa_i\right)}\right).$$

Proof. In the case of anonymous prices, any non-linear pricing mechanism is defined by $\prod_{i=1}^m (\kappa_i + 1)$ parameters because that is the number of different bundles. Consumer j will prefer the bundle corresponding to the quantity vector \mathbf{q} over the bundle corresponding to the quantity vector \mathbf{q}' if

$v_j(\mathbf{q}) - p(\mathbf{q}) \geq v_j(\mathbf{q}') - p(\mathbf{q}')$. Therefore, there are at most $\prod_{i=1}^m (\kappa_i + 1)^2$ hyperplanes determining each buyer's preferred bundle — one hyperplane per pair of bundles. This means that there are a total of $n \prod_{i=1}^m (\kappa_i + 1)^2$ hyperplanes in $(\prod_{i=1}^m (\kappa_i + 1))$ -dimensional space such that in any one region induced by these hyperplanes, the demand bundles of all n consumers are fixed and profit is linear in the prices of these n bundles. By Corollary 4.4, with $d = \prod_{i=1}^m (\kappa_i + 1)$ and $t = n \prod_{i=1}^m (\kappa_i + 1)^2$, we have that $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(U \sqrt{\frac{1}{N} \prod_{i=1}^m \kappa_i (\log n + \sum_{i=1}^m \log \kappa_i)}\right)$.

In the case of non-anonymous prices, the same argument holds, except that every non-linear pricing mechanism is defined by $n \prod_{i=1}^m (\kappa_i + 1)$ parameters — one parameter per bundle per consumer. Therefore, $d = n \prod_{i=1}^m (\kappa_i + 1)$, and the result follows from Corollary 4.4. \square

We achieve improved bounds when the non-linear prices decompose additively over the items (see the definition of additively decomposable non-linear pricing mechanisms in Section 3.1.1).

Theorem 4.6. *Let \mathcal{A} and \mathcal{A}' be the classes of additively decomposable non-linear pricing mechanisms with anonymous and non-anonymous prices, respectively. Then for any sample \mathcal{S} of size N ,*

$$\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(U \sqrt{\frac{1}{N} \sum_{i=1}^m \kappa_i \left(\log n + \sum_{i=1}^m \log \kappa_i\right)}\right)$$

and

$$\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}') \leq O\left(U \sqrt{\frac{n}{N} \sum_{i=1}^m \kappa_i \left(\log n + \sum_{i=1}^m \log \kappa_i\right)}\right).$$

Proof. In the case of anonymous prices, any additively decomposable non-linear pricing mechanism is defined by $\sum_{i=1}^m (\kappa_i + 1)$ parameters. As in the proof of Theorem 4.5, there are a total of $n \prod_{i=1}^m (\kappa_i + 1)^2$ hyperplanes such that in any one region induced by these hyperplanes, the demand bundles of all n consumers are fixed and profit is linear in the prices of these n bundles. By Corollary 4.4, with $d = \sum_{i=1}^m (\kappa_i + 1)$ and $t = n \prod_{i=1}^m (\kappa_i + 1)^2$, we have that $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(U \sqrt{\frac{1}{N} \sum_{i=1}^m \kappa_i (\log n + \sum_{i=1}^m \log \kappa_i)}\right)$.

In the case of non-anonymous prices, the same argument holds, except that every non-linear pricing mechanism is defined by $n \sum_{i=1}^m (\kappa_i + 1)$ parameters — one parameter per item, quantity, and consumer tuple. Therefore, $d = n \sum_{i=1}^m (\kappa_i + 1)$, and the result follows from Corollary 4.4. \square

Finally, we apply Theorem 4.2 to menus of two-part tariffs.

Theorem 4.7. *Let \mathcal{A} and \mathcal{A}' be the classes of length- M menus of two-part tariffs with anonymous and non-anonymous prices, respectively. Then for any sample \mathcal{S} of size N , $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(U \sqrt{\frac{M}{N} \log(n\kappa M)}\right)$ and $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}') \leq O\left(U \sqrt{\frac{nM}{N} \log(n\kappa M)}\right)$.*

Proof. In the case of anonymous prices, every length- M menu of two-part tariffs is defined by $2M$ parameters: the fixed fee and unit price for each of the M menu entries. Consumer j will choose the quantity q and menu entry (p_0^i, p_1^i) that maximizes $v_j(q) - (p_0^i \cdot \mathbf{1}_{q>0} + p_1^i q)$. Therefore, the quantity q and menu entry that she chooses is determined by $(\kappa M)^2$ hyperplanes of the form $v_j(q) - (p_0^i \cdot \mathbf{1}_{q>0} + p_1^i q) \geq v_j(q') - (p_0^k \cdot \mathbf{1}_{q'>0} + p_1^k q')$. In total, there are $n(M\kappa)^2$ hyperplanes that determine the menu entry and quantity demanded by all n buyers, over which profit is linear in the fixed fees and unit prices. Therefore, with $d = 2M$ and $t = n(M\kappa)^2$, by Corollary 4.4, $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(U \sqrt{\frac{M}{N} \log(n\kappa M)}\right)$.

In the case of non-anonymous reserve prices, the same argument holds, except that every length- M menu of two-part tariffs is defined by $2nM$ parameters: for each buyer, we must set the fixed fee and unit price for each of the M menu entries. Therefore, $d = 2nM$, and the result follows from Corollary 4.4. \square

4.2 Single-unit pricing mechanisms

We begin by deriving bounds for the simplest type of single-unit pricing mechanisms: item pricing mechanisms when the buyers have additive valuations. When we say that the buyers are additive (respectively, unit-demand), we mean that the support of the distribution of buyers' values is over additive (respectively, unit-demand) valuations.

Theorem 4.8. *Let \mathcal{A} and \mathcal{A}' be the classes of item pricing mechanisms with anonymous prices and non-anonymous prices, respectively. If the buyers are additive, then for any sample \mathcal{S} of size N , $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(U\sqrt{\frac{m \log m}{N}}\right)$ and $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}') \leq O\left(U\sqrt{\frac{nm \log(nm)}{N}}\right)$.*

Proof. For anonymous prices, we appeal to Corollary 4.4. We know that an item i will be bought under the valuation vector \mathbf{v}^k if there exists a buyer j such that $v_j^k(\mathbf{e}_i) \geq p(\mathbf{e}_i)$. For a sample \mathbf{v}^k , let $v^k(\mathbf{e}_i)$ be the maximum amount that any buyer values item i , i.e. $v^k(\mathbf{e}_i) = \max_j v_j^k(\mathbf{e}_i)$. Item i is bought under the valuation vector \mathbf{v}^k if and only if $v^k(\mathbf{e}_i) \geq p(\mathbf{e}_i)$. Therefore, for each sample, there are at most m hyperplanes determining the set of items bought. On a single region induced by these hyperplanes, profit is a linear function of the prices. Therefore, by Corollary 4.4 with $d = m$ and $t = m$, we have that $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(\sqrt{\frac{m \log m}{N}}\right)$.

The proof for anonymous prices easily generalizes to the case with non-anonymous prices and we include it below for completeness. We know that an item i will be bought under valuation vector \mathbf{v}^k if there exists a buyer j such that $v_j^k(\mathbf{e}_i) \geq p_j(\mathbf{e}_i)$. Therefore, for each sample, there are at most nm hyperplanes determining the set of items bought. On a single region induced by these hyperplanes, profit is a linear function of the prices. Therefore, by Corollary 4.4 with $d = nm$ and $t = nm$, we have that $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}') \leq O\left(U\sqrt{\frac{nm \log(nm)}{N}}\right)$. \square

A similar analysis allows us to derive the Rademacher complexity of item pricing mechanisms when the buyers have unit-demand valuations.

Theorem 4.9. *Let \mathcal{A} and \mathcal{A}' be the classes of item pricing mechanisms with anonymous prices and non-anonymous prices, respectively. If the buyers are unit-demand, then for any sample \mathcal{S} of size N , $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(U\sqrt{\frac{m \log(nm)}{N}}\right)$ and $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}') \leq O\left(U\sqrt{\frac{nm \log(nm)}{N}}\right)$.*

Proof. For unit-demand buyers with anonymous item pricing, each buyer i will buy the item j that maximizes $v_i(\mathbf{e}_j) - p(\mathbf{e}_j)$. Therefore, the outcome of the mechanism is determined by the $n\binom{m}{2}$ hyperplanes $v_i(\mathbf{e}_j) - p(\mathbf{e}_j) = v_i(\mathbf{e}_k) - p(\mathbf{e}_k)$ for all $i \in [n]$ and $j, k \in [m]$. With $d = m$ and $t = O(nm^2)$, Corollary 4.4 guarantees the theorem statement. For unit-demand buyers with non-anonymous prices, the only difference is that $d = nm$, and the theorem holds by Corollary 4.4. \square

For the generalized class of bundle pricing mechanisms, we prove the following bounds.

Theorem 4.10. *Let \mathcal{A} be the class of \mathcal{B} -pricing mechanisms and let \mathcal{A}' be the class of $(\mathcal{B}_1, \dots, \mathcal{B}_n)$ -pricing mechanisms. Then for any sample \mathcal{S} of size N , $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(U\sqrt{\frac{|\mathcal{B}|(m+\log(n|\mathcal{B}|))}{N}}\right)$ and $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}') \leq O\left(U\sqrt{\frac{\sum_{i=1}^n |\mathcal{B}_i|(m+\log(n\sum_{i=1}^n |\mathcal{B}_i|))}{N}}\right)$.*

Proof. Every \mathcal{B} -pricing mechanism is defined by $|\mathcal{B}|$ parameters: the $|\mathcal{B}|$ reserve prices. Let $\mathbf{p} \in \mathbb{R}^{|\mathcal{B}|}$ be a vector of the reserve prices. Without loss of generality, suppose that Buyer 1 is the first consumer to choose which bundle to buy, then Buyer 2, and so on, until Buyer n has bought his desired bundle. Given a vector of valuations \mathbf{v} , Buyer 1's desired bundle will be \mathbf{q} so long as $v_1(\mathbf{q}) - p(\mathbf{q}) \geq v_1(\mathbf{q}') - p(\mathbf{q}')$ for all other bundles \mathbf{q}' . Therefore, this decision is based on $\binom{2^m}{2}$ hyperplanes: one for every pair of bundles. Similarly, each Buyer i 's desired bundle will be based on $\binom{2^m}{2}$ hyperplanes: $v_i(\mathbf{q}) - p(\mathbf{q}) = v_i(\mathbf{q}') - p(\mathbf{q}')$. (Technically, this decision may be based on fewer hyperplanes for all buyers after Buyer 1, since not all items may be available to them when they choose which bundle to buy.) By Corollary 4.4, with $t = O(n2^{2m})$ and $d = |\mathcal{B}|$ we have that $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(U\sqrt{\frac{|\mathcal{B}|\log(n|\mathcal{B}|2^{2m})}{N}}\right) = O\left(U\sqrt{\frac{|\mathcal{B}|(m+\log(n|\mathcal{B}|))}{N}}\right)$.

If there are non-anonymous reserve prices, then the same reasoning holds, except now there are $\sum_{i=1}^n |\mathcal{B}_i|$ parameters defining each mechanism. Therefore,

$$\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}') \leq O\left(U\sqrt{\frac{\sum_{i=1}^n |\mathcal{B}_i|(m+\log(n\sum_{i=1}^n |\mathcal{B}_i|))}{N}}\right),$$

as claimed. \square

4.3 Multi-unit auctions

We now cover the classes of AMAs and VVCAs, which are the only classes with an exponential Rademacher complexity upper bound. This is the best we can hope for, since Balcan et al. also prove an exponential lower bound [Balcan et al., 2016].⁴

Theorem 4.11. *Let \mathcal{A} be the class of AMAs, let \mathcal{A}' be the class of VVCAs, and let \mathcal{A}'' be the class of λ -auctions. For any sample \mathcal{S} of size N ,*

1. $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(U\sqrt{\frac{1}{N}\prod_{i=1}^m \binom{\kappa_i+n}{n} \sum_{i=1}^m \log \binom{\kappa_i+n}{n}}\right)$,
2. $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}') \leq O\left(U\sqrt{\frac{1}{N} \cdot n \prod_{i=1}^m \kappa_i \sum_{i=1}^m \log \binom{\kappa_i+n}{n}}\right)$, and
3. $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}'') \leq O\left(U\sqrt{\frac{1}{N}\prod_{i=1}^m \binom{\kappa_i+n}{n} \sum_{i=1}^m \log \binom{\kappa_i+n}{n}}\right)$.

Proof. Every AMA is defined by $n + \prod_{i=1}^m \binom{\kappa_i+n}{n} \leq 2\prod_{i=1}^m \binom{\kappa_i+n}{n}$ parameters since there are n bidder weights and at most $\prod_{i=1}^m \binom{\kappa_i+n}{n}$ allocation boosts. Given a vector \mathbf{v} of valuations, an allocation $Q = (\mathbf{q}_1, \dots, \mathbf{q}_n)$ will be the allocation of the AMA so long as $\sum_{i=1}^n w_i v_i(\mathbf{q}_i) + \lambda(Q) - c(Q) \geq \sum_{i=1}^n w_i v_i(\mathbf{q}'_i) + \lambda(Q') - c(Q')$ for all allocations $Q' = (\mathbf{q}'_1, \dots, \mathbf{q}'_n) \neq Q$. Since the number of different allocations is at most $\prod_{i=1}^m \binom{\kappa_i+n}{n}$, the allocation of the auction on \mathbf{v} is defined by at

⁴Specifically, Balcan et al. prove that exponentially-many samples are required for uniform convergence over the classes of AMAs and VVCAs. This implies that the Rademacher complexity of these classes is also exponential.

most $\prod_{i=1}^m \binom{\kappa_i+n}{n}^2$ hyperplanes. Similarly, the allocations Q^{-1}, \dots, Q^{-n} are also determined by at most $\prod_{i=1}^m \binom{\kappa_i+n}{n}^2$ hyperplanes in $(2 \prod_{i=1}^m \binom{\kappa_i+n}{n})$ -dimensional space. Therefore, by Corollary 4.4, setting $d = 2 \prod_{i=1}^m \binom{\kappa_i+n}{n}$ and $t = (n+1) \binom{\kappa_i+n}{n}^2$ we have that

$$\begin{aligned} \widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) &\leq O \left(U \sqrt{\frac{1}{N} \cdot \prod_{i=1}^m \binom{\kappa_i+n}{n} \log \left((n+1) \prod_{i=1}^m \binom{\kappa_i+n}{n}^3 \right)} \right) \\ &\leq O \left(U \sqrt{\frac{1}{N} \cdot \prod_{i=1}^m \binom{\kappa_i+n}{n} \log \prod_{i=1}^m \binom{\kappa_i+n}{n}^4} \right) \\ &\leq O \left(U \sqrt{\frac{1}{N} \cdot \prod_{i=1}^m \binom{\kappa_i+n}{n} \sum_{i=1}^m \log \binom{\kappa_i+n}{n}} \right). \end{aligned}$$

For the class of VVCAs, the same argument holds except that every VVCA is defined by $n+n \prod_{i=1}^m (\kappa_i+1)$ parameters, since there are n bidders weights and $n \prod_{i=1}^m (\kappa_i+1)$ bidder-specific bundle boosts $c_{i,\mathbf{q}}$. Similarly, for the class of λ -auctions, the same argument holds except there are zero bidder weights and $\prod_{i=1}^m \binom{\kappa_i+n}{n}$ allocation boosts. Therefore, $d = O(n \prod_{i=1}^m \kappa_i)$ for VVCAs and $d = \prod_{i=1}^m \binom{\kappa_i+n}{n}$ for λ -auctions, and we may again apply Corollary 4.4 to arrive at the theorem statement. \square

Next, we analyze the class of \mathcal{O} -boosted AMAs.

Theorem 4.12. *Let \mathcal{A} be the class of \mathcal{O} -boosted AMAs. Then for any sample \mathcal{S} of size N ,*

$$\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O \left(U \sqrt{\frac{(n+|\mathcal{O}|)(\log(n+|\mathcal{O}|)+m \log n)}{N}} \right).$$

Proof. Every \mathcal{O} -boosted AMA is defined by $n+|\mathcal{O}|$ parameters, since there are n bidders weights and $|\mathcal{O}|$ allocation boosts. Fix some valuation vector \mathbf{v} . We claim that the allocation of any \mathcal{O} -boosted AMA is determined by at most $(n+1) \prod_{i=1}^m \binom{\kappa_i+n}{n}^2$ hyperplanes, where $\prod_{i=1}^m \binom{\kappa_i+n}{n}$ is the number of different allocations. To see why this is, notice that the VCG allocation will be the AMA allocation by default unless there exists some $Q^j = (\mathbf{q}_1^j, \dots, \mathbf{q}_n^j)$ such that $\sum w_i v_i (\mathbf{q}_i^j) + \lambda(Q^j) - c(Q^j) \geq \sum w_i v_i (\mathbf{q}_i^k) + \lambda(Q^k) - c(Q^k)$ for all allocations $Q^k = (\mathbf{q}_1^k, \dots, \mathbf{q}_n^k)$. This decision governing which of the $\prod_{i=1}^m \binom{\kappa_i+n}{n}$ possible allocations will be the AMA allocation is defined by the $\prod_{i=1}^m \binom{\kappa_i+n}{n}^2$ hyperplanes, one per pair of distinct allocations Q^j and Q^k . We now give a precise characterization of the hyperplane corresponding to an arbitrary pair Q^j and Q^k , which depends on whether both allocations are in \mathcal{O} , neither are in \mathcal{O} , or just one is in \mathcal{O} . Therefore, there are three cases:

1. If Q^j and Q^k are both in \mathcal{O} , then Q^k will not be allocated by any AMA where the parameters $w_1, \dots, w_n, \lambda(Q^j)$, and $\lambda(Q^k)$ are such that $\sum w_i v_i (\mathbf{q}_i^j) + \lambda(Q^j) - c(Q^j) > \sum w_i v_i (\mathbf{q}_i^k) + \lambda(Q^k) - c(Q^k)$. This corresponds to a hyperplane in the following way. We will use the notation \mathbf{v}^j to denote the $(n+1)$ -dimensional vector consisting of each bidder's value for the allocation Q^j , i.e., $\mathbf{v}^j = (v_1(\mathbf{q}_1^j), \dots, v_n(\mathbf{q}_n^j))$. Also, we use the notation $\mathbf{v}^j \circ \mathbf{e}_j$ to denote the vector of valuations \mathbf{v}^j concatenated with the standard basis vector $\mathbf{e}_j \in \mathbb{R}^{|\mathcal{O}|}$, which implies that $(\mathbf{w}, \lambda) \cdot \mathbf{v}^j \circ \mathbf{e}_j = \sum w_i v_i (\mathbf{q}_i^j) + \lambda(Q^j)$. This notation allows us to conclude that the set of

parameter vectors (\mathbf{w}, λ) (where $\lambda \in \mathbb{R}^{|\mathcal{O}|}$) under which Q^k is never allocated under the corresponding AMA is the halfspace $\{(\mathbf{w}, \lambda) \mid (\mathbf{w}, \lambda) \cdot \mathbf{v}^j \circ \mathbf{e}_j - c(Q^j) > (\mathbf{w}, \lambda) \cdot \mathbf{v}^k \circ \mathbf{e}_k - c(Q^k)\}$. This halfspace's separating hyperplane is defined by the equation $(\mathbf{w}, \lambda) \cdot (\mathbf{v}^j \circ \mathbf{e}_j - \mathbf{v}^k \circ \mathbf{e}_k) = c(Q^j) - c(Q^k)$.

2. Without loss of generality, suppose that Q^k is in \mathcal{O} and Q^j is not. Then Q^k will not be allocated so long as the parameters w_1, \dots, w_n , and $\lambda(Q^k)$ are such that $\sum w_i v_i(\mathbf{q}_i^j) - c(Q^j) > \sum w_i v_i(\mathbf{q}_i^k) + \lambda(Q^k) - c(Q^k)$ (since Q^j is not in \mathcal{O} , $\lambda(Q^j) = 0$). In other words, the set of parameter vectors (\mathbf{w}, λ) where Q^k is never allocated are in the halfspace $\{(\mathbf{w}, \lambda) \mid (\mathbf{w}, \lambda) \cdot \mathbf{v}^j \circ \mathbf{0} - c(Q^j) > (\mathbf{w}, \lambda) \cdot \mathbf{v}^k \circ \mathbf{e}_k - c(Q^k)\}$. This halfspace's separating hyperplane is defined by the equation $(\mathbf{w}, \lambda) \cdot (\mathbf{v}^j \circ \mathbf{0} - \mathbf{v}^k \circ \mathbf{e}_k) = c(Q^j) - c(Q^k)$.
3. If neither Q^j nor Q^k are in \mathcal{O} , then by the same reasoning as in the above two cases, one of the $\prod_{i=1}^m \binom{\kappa_i+n}{n}^2$ hyperplanes defining which parameters lead to which AMA allocations on the valuation vector \mathbf{v} is $(\mathbf{w}, \lambda) \cdot (\mathbf{v}^j \circ \mathbf{0} - \mathbf{v}^k \circ \mathbf{0}) = c(Q^j) - c(Q^k)$.

These $\prod_{i=1}^m \binom{\kappa_i+n}{n}^2$ hyperplanes split the parameter space into regions so that for any two AMAs defined by parameters in the same region, the allocation will be the same. Specifically, given a parameter vector (\mathbf{w}, λ) , there must be some allocation Q^j such that $\sum w_i v_i(\mathbf{q}_i^j) + \lambda(Q^j) - c(Q^j) \geq \sum w_i v_i(\mathbf{q}_i^k) + \lambda(Q^k) - c(Q^k)$ for all allocations Q^k , and therefore Q^j is the allocation of the corresponding AMA (or an allocation with the same weighted social welfare, which implies the same profit, and thus we are unconcerned with specific tie-breaking rules). By construction, for any set of parameters in the same region as (\mathbf{w}, λ) , the allocation of the corresponding AMA will also be Q^j .

By a similar argument, it is straightforward to see that $\prod_{i=1}^m \binom{\kappa_i+n}{n}^2$ hyperplanes determine the allocation of any AMA in this restricted space without any one bidder's participation. This leads us to a total of $(n+1) \prod_{i=1}^m \binom{\kappa_i+n}{n}^2$ hyperplanes which partition the space of \mathcal{O} -boosted AMA parameters in a way such that for any two parameter vectors in the same cell, the auction allocations are the same, as are the allocations without any one bidder's participation.

Therefore, by Corollary 4.4, setting $d = n + |\mathcal{O}|$ and $t = (n+1) \prod_{i=1}^m \binom{\kappa_i+n}{n}^2 = O\left(\prod_{i=1}^m \binom{\kappa_i+n}{n}^3\right)$, we have that $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(U \sqrt{\frac{1}{N}(n+|\mathcal{O}|) (\log(n+|\mathcal{O}|) + \sum_{i=1}^m \log \binom{\kappa_i+n}{n})}\right)$. \square

A similar proof technique admits the following theorem.

Theorem 4.13. *Let \mathcal{A} be the class of \mathcal{O} -boosted λ -auctions. Then for any sample \mathcal{S} of size N ,*
 $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(U \sqrt{\frac{1}{N} \cdot |\mathcal{O}| \log(n|\mathcal{O}|)}\right)$.

Proof. This proof is very similar to that of Theorem 4.12. However, we claim that the allocation of any \mathcal{O} -boosted λ -auction is determined by at most $(n+1)(|\mathcal{O}|+1)^2$ hyperplanes. This is because without the bidder weights, the VCG allocation is the only unboosted allocation that has any chance of being the allocation of the \mathcal{O} -boosted λ -auction. Therefore, there are only $(|\mathcal{O}|+1)^2$ hyperplanes determining the allocation of the λ -auction, and the same number of hyperplanes determine the allocation of the λ -auction in this restricted space without any one bidder's participation. Therefore, by Corollary 4.4, setting $d = |\mathcal{O}|$ and $t = (n+1)(|\mathcal{O}|+1)^2$, we have that $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(U \sqrt{\frac{1}{N} \cdot |\mathcal{O}| \log((n+1)|\mathcal{O}|(|\mathcal{O}|+1)^2)}\right) = O\left(U \sqrt{\frac{1}{N} \cdot |\mathcal{O}| \log(n|\mathcal{O}|)}\right)$. \square

We use Theorems 4.12 and 4.13 to prove the following generalization guarantees regarding k -sparse AMAs and λ -auctions.

Theorem 4.14. *Let \mathcal{A} be the class of k -sparse AMAs and let \mathcal{A}' be the class of k -sparse λ -auctions. Then $\epsilon_{\mathcal{A}}(N, \delta)$ is at most*

$$O \left(U \sqrt{\frac{n+k}{N} \left(\log(n+k) + \sum_{i=1}^m \log \binom{\kappa_i+n}{n} \right)} + U \sqrt{\frac{1}{N} \left(k \sum_{i=1}^m \log \binom{\kappa_i+n}{n} + \log \frac{1}{\delta} \right)} \right)$$

and $\epsilon_{\mathcal{A}'}(N, \delta)$ is at most

$$O \left(U \sqrt{\frac{(n+k) \log(n+k)}{N}} + U \sqrt{\frac{1}{N} \left(k \sum_{i=1}^m \log \binom{\kappa_i+n}{n} + \log \frac{1}{\delta} \right)} \right).$$

Proof. We prove this theorem by effectively spreading the confidence parameter δ over all \mathcal{O} -boosted AMAs where $|\mathcal{O}| = k$ of which there are at most $\alpha := \prod_{i=1}^m \binom{\kappa_i+n}{n}^k$. Let $\mathcal{A}_{\mathcal{O}}$ be the set of \mathcal{O} -boosted AMAs for an arbitrary set of allocations \mathcal{O} and let \mathcal{A}_k be the set of all k -sparse AMAs. By Theorem 4.12 and the generalization guarantee stated in Section 3, we know that for all sets of allocations \mathcal{O} such that $|\mathcal{O}| \leq k$, the probability that there exists an auction $A \in \mathcal{A}_{\mathcal{O}}$ such that $|\text{profit}_{\mathcal{D}}(A) - \text{profit}_{\mathcal{S}}(A)| \geq \epsilon_{\mathcal{A}_{\mathcal{O}}}(N, \delta/\alpha)$ is at most δ/α , where

$$\epsilon_{\mathcal{A}_{\mathcal{O}}}(N, \delta/\alpha) \leq O \left(U \sqrt{\frac{n+k}{N} \left(\log(n+k) + \sum_{i=1}^m \log \binom{\kappa_i+n}{n} \right)} + U \sqrt{\frac{1}{N} \log \frac{\alpha}{\delta}} \right).$$

Notice that $\epsilon_{\mathcal{A}_{\mathcal{O}}}(N, \delta/\alpha) = \epsilon_{\mathcal{A}_k}(N, \delta)$, where $\epsilon_{\mathcal{A}_k}(N, \delta/\alpha)$ equals

$$O \left(U \sqrt{\frac{n+k}{N} \left(\log(n+k) + \sum_{i=1}^m \log \binom{\kappa_i+n}{n} \right)} + U \sqrt{\frac{1}{N} \left(k \sum_{i=1}^m \log \binom{\kappa_i+n}{n} + \log \frac{1}{\delta} \right)} \right)$$

since $O(\log \frac{\alpha}{\delta}) = O(k \sum_{i=1}^m \log \binom{\kappa_i+n}{n} + \log \frac{1}{\delta})$.

Next, let $B_{\mathcal{S}}$ be the bad event where there exists $A \in \mathcal{A}_k$ such that $|\text{profit}_{\mathcal{D}}(A) - \text{profit}_{\mathcal{S}}(A)| \geq \epsilon_{\mathcal{A}_k}(N, \delta)$. We want to show that $\Pr_{\mathcal{S} \sim \mathcal{D}^N}[B_{\mathcal{S}}] < \delta$. To this end, let $B_{\mathcal{S}, \mathcal{O}}$ be the event where there exists an auction $A \in \mathcal{A}_{\mathcal{O}}$ such that $|\text{profit}_{\mathcal{D}}(A) - \text{profit}_{\mathcal{S}}(A)| \geq \epsilon_{\mathcal{A}_{\mathcal{O}}}(N, \delta/\alpha)$. Again, by Theorem 4.12, we have that $\Pr_{\mathcal{S} \sim \mathcal{D}^N}[B_{\mathcal{S}, \mathcal{O}}] < \delta/\alpha$. Since $\mathcal{A}_k = \cup_{\mathcal{O}: |\mathcal{O}|=k} \mathcal{A}_{\mathcal{O}}$, a union bound ensures that $\Pr[B_{\mathcal{S}}] \leq \sum_{\mathcal{O}: |\mathcal{O}|=k} \Pr_{\mathcal{S} \sim \mathcal{D}^N}[B_{\mathcal{S}, \mathcal{O}}] < \alpha \delta/\alpha = \delta$, as desired.

The bound on $\epsilon_{\mathcal{A}'}(N, \delta)$ follows by the exact same reasoning from Theorem 4.13. \square

4.4 Single-unit auctions

For the class of VCG auctions with bundle reserve prices, we prove the following theorem.

Theorem 4.15. *Let \mathcal{A} be the class of VCG auctions with anonymous reserve prices over \mathcal{B} . Then for any sample \mathcal{S} of size N , $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O \left(U \sqrt{\frac{|\mathcal{B}|(m+\log(n|\mathcal{B}|))}{N}} \right)$.*

Proof. The VCG auction with anonymous item reserve prices over \mathcal{B} is defined by $|\mathcal{B}|$ parameters: the $|\mathcal{B}|$ reserve prices. Let $\mathbf{p} \in \mathbb{R}^{|\mathcal{B}|}$ be a vector of the reserve prices. Given a vector of valuations

\mathbf{v} , any allocation allocating in total the items in a bundle \mathbf{q} will partition them among the bidders in such a way that the social welfare is maximized. Let $Q^{\mathbf{q}} = (\mathbf{q}_1^{\mathbf{q}}, \dots, \mathbf{q}_n^{\mathbf{q}})$ be this allocation. This will be the allocation of the auction if it maximizes $\sum_{i=1}^n v_i(\mathbf{q}_i^{\mathbf{q}}) + p(\mathbf{1} - \mathbf{q})$ since $\mathbf{1} - \mathbf{q}$ is the bundle of items not allocated under $Q^{\mathbf{q}}$. Therefore, the allocation of the auction is determined by at most $\binom{2^m}{2}$ hyperplanes: $\sum_{i=1}^n v_i(\mathbf{q}_i^{\mathbf{q}}) + p(\mathbf{1} - \mathbf{q}) = \sum_{i=1}^n v_i(\mathbf{q}_i^{\mathbf{q}'}) + p(\mathbf{1} - \mathbf{q}')$ for all bundles $\mathbf{q}, \mathbf{q}' \subseteq \{0, 1\}^m$. Similarly, at most $\binom{2^m}{2}$ hyperplanes determine the allocation without any one bidder's participation, so there are at most $(n+1)2^{2m}$ relevant hyperplanes per sample. Given a fixed set of the $n+1$ allocations $Q^*, Q^{-1}, \dots, Q^{-n}$, the profit of the VCG auction is a fixed linear function. Therefore, by Corollary 4.4, with $t = O(n2^{2m})$, we have that $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(U\sqrt{\frac{|\mathcal{B}|\log(n|\mathcal{B}|2^{2m})}{N}}\right) = O\left(U\sqrt{\frac{|\mathcal{B}|(m+\log(n|\mathcal{B}|))}{N}}\right)$. \square

We conclude by bounding the Rademacher complexity of the class of MBARPs.

Theorem 4.16. *Let \mathcal{A} be the set of MBARPs and let \mathcal{S} be a sample of size N . Then $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(U\sqrt{\frac{m(\log n+m)}{N}}\right)$.*

Proof. Fix $\mathbf{v}^t \in \mathcal{S}$. First, for each bundle $\mathbf{q} \in \{0, 1\}^m$, let $\mathcal{O}_{\mathbf{q}}$ be the set of allocations where exactly the elements of \mathbf{q} are allocated, and let

$$Q^{\mathbf{q}} = \operatorname{argmax}_{Q \in \mathcal{O}_{\mathbf{q}}} \left\{ \sum_{i=1}^n v_i^t(\mathbf{q}_i) \right\}.$$

Notice that regardless of the reserve prices, if \mathbf{q} is comprised of the items allocated in the allocation of an MBARP, then $Q^{\mathbf{q}}$ will be the allocation. After all, if (r_1, \dots, r_m) are the reserve prices of an arbitrary MBARP, then it will always be the case that

$$\sum_{i=1}^n v_i^t(\mathbf{q}_i^{\mathbf{q}}) + \sum_{j:q^{\mathbf{q}}[j]=0} r_j \geq \sum_{i=1}^n v_i^t(\mathbf{q}_i^{\mathbf{q}'}) + \sum_{j:q'^{\mathbf{q}}[j]=0} r_j$$

for any allocation $Q' = (\mathbf{q}_1^{\mathbf{q}'}, \dots, \mathbf{q}_n^{\mathbf{q}'}) \in \mathcal{O}_{\mathbf{q}}$ by definition of $Q^{\mathbf{q}}$.

Next, let $R_{\mathbf{q}}^{\mathbf{v}^t}$ be the subset of \mathbb{R}^{m+1} such that if an MBARP is parameterized by $(c, r_1, \dots, r_m) \in R_{\mathbf{q}}^{\mathbf{v}^t}$, then the allocation of the MBARP on \mathbf{v}^t is $Q^{\mathbf{q}}$. This means that if $\mathbf{q} \neq \mathbf{1}$,

$$\begin{aligned} \sum_{i=1}^n v_i^t(\mathbf{q}_i^{\mathbf{q}}) + \sum_{j:q^{\mathbf{q}}[j]=0} r_j &\geq \sum_{i=1}^n v_i^t(\mathbf{q}_i^{\mathbf{q}'}) + \sum_{j:q'^{\mathbf{q}}[j]=0} r_j && \forall \mathbf{q}' \notin \{\mathbf{q}, \mathbf{1}\} \text{ and} \\ \sum_{i=1}^n v_i^t(\mathbf{q}_i^{\mathbf{q}}) + \sum_{j:q^{\mathbf{q}}[j]=0} r_j &\geq \sum_{i=1}^n v_i^t(\mathbf{q}_i^{\mathbf{1}}) + c. \end{aligned}$$

In other words, $(c, r_1, \dots, r_m) \in R_{\mathbf{q}}^{\mathbf{v}^t}$ if and only if it falls in the intersection of $2^m - 1$ halfspaces:

$$\begin{aligned} \sum_{j:q^{\mathbf{q}}[j]=0} r_j - \sum_{j:q'^{\mathbf{q}}[j]=0} r_j &\geq \sum_{i=1}^n v_i^t(\mathbf{q}_i^{\mathbf{q}'}) - v_i^t(\mathbf{q}_i^{\mathbf{q}}) && \forall \mathbf{q}' \notin \{\mathbf{q}, \mathbf{1}\} \\ \sum_{j:q^{\mathbf{q}}[j]=0} r_j - c &\geq \sum_{i=1}^n v_i^t(\mathbf{q}_i^{\mathbf{1}}) - v_i^t(\mathbf{q}_i^{\mathbf{q}}). \end{aligned}$$

Similarly, if $\mathbf{q} = \mathbf{1}$, it is not hard to see that we can write $R_1^{v^t}$ as the intersection of $2^m - 1$ halfspaces. The same holds for the MBARP without any one bidder's participation, leading to a total of $(n + 1)2^m (2^m - 1) = O(n2^{2m})$ relevant hyperplanes. Whenever these $n + 1$ allocations are fixed, the profit is a fixed linear function of the m reserve prices. By Corollary 4.4, with $t = O(n2^{2m})$ and $\mathbf{p} \in \mathbb{R}^{m+1}$, we have that $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A}) \leq O\left(U\sqrt{\frac{m \log(nm2^{2m})}{N}}\right) = O\left(U\sqrt{\frac{m(\log n + m)}{N}}\right)$. \square

5 Structural profit maximization

In this section, we provide a data-dependent methodology by which the mechanism designer can determine a mechanism class which provides him the optimal profit-generalization tradeoff. This tradeoff can be described in terms of two types of profit loss: *approximation loss* and *estimation loss*. To understand this trade-off, let $opt(\mathcal{D})$ be the mechanism which maximizes expected profit over the distribution of buyers' values \mathcal{D} . Crucially, we do not assume that $opt(\mathcal{D})$ is in the mechanism class \mathcal{A} that the mechanism designer is optimizing over. Further, for a set \mathcal{S} of N samples from \mathcal{D} , with a slight abuse of notation, let $\mathcal{A}(\mathcal{S})$ be the mechanism in \mathcal{A} which maximizes empirical profit, and let $\text{profit}_{\mathcal{D}}(\mathcal{A}(\mathcal{S}))$ be the expected profit of the mechanism $\mathcal{A}(\mathcal{S})$ over \mathcal{D} . Similarly, let $\mathcal{A}(\mathcal{D})$ be the mechanism in \mathcal{A} that maximizes expected profit over \mathcal{D} . We can write the true profit loss as the difference between the expected profit of $opt(\mathcal{D})$ and the expected profit of $\mathcal{A}(\mathcal{S})$, which decomposes as

$$\begin{aligned} & \underbrace{\text{profit}_{\mathcal{D}}(opt(\mathcal{D})) - \text{profit}_{\mathcal{D}}(\mathcal{A}(\mathcal{S}))}_{\text{true loss}} \\ = & \underbrace{\text{profit}_{\mathcal{D}}(opt(\mathcal{D})) - \text{profit}_{\mathcal{D}}(\mathcal{A}(\mathcal{D}))}_{\text{approximation loss}} + \underbrace{\text{profit}_{\mathcal{D}}(\mathcal{A}(\mathcal{D})) - \text{profit}_{\mathcal{D}}(\mathcal{A}(\mathcal{S}))}_{\text{estimation loss}}. \end{aligned}$$

In words, approximation loss is the amount of profit lost given that the optimal mechanism, $opt(\mathcal{D})$, is not in the design space \mathcal{A} , and estimation loss measures the amount of profit lost given that we do not know the distribution \mathcal{D} , but only have samples from \mathcal{D} . Structural risk minimization is a general technique used in machine learning to pin down the optimal tradeoff between estimation and approximation loss. In the case of sample-based mechanism design, we are not minimizing risk, but maximizing profit, so we refer to this process as *structural profit maximization* (SPM).

We will demonstrate this tradeoff using the abstract notion of a generalization guarantee $\epsilon_{\mathcal{A}}(N, \delta)$, as defined in Definition 3.1. For a given class \mathcal{A} , the form of $\epsilon_{\mathcal{A}}(N, \delta)$ invariably depends on a measure of the class's intrinsic complexity, which we will refer to as $\text{comp}(\mathcal{A})$. For example, c could be a bound on $\widehat{\mathcal{R}}_{\mathcal{S}}(\mathcal{A})$, in which case $\epsilon_{\mathcal{A}}(N, \delta) \leq 2\text{comp}(\mathcal{A}) + 4U\sqrt{2 \ln(4/\delta)/N}$.

Suppose that \mathcal{A} is a rich class of mechanisms which can be decomposed into a nested sequence of subclasses $\mathcal{A}_1 \subseteq \mathcal{A}_2 \subseteq \dots \subseteq \mathcal{A}_p = \mathcal{A}$. For example, if \mathcal{A} is the class of AMAs, then \mathcal{A}_k could be the class of k -sparse AMAs. For any standard complexity measure, we have that $\text{comp}(\mathcal{A}_1) \leq \text{comp}(\mathcal{A}_2) \leq \dots \leq \text{comp}(\mathcal{A}_p)$. Prior work has primarily given uniform generalization bounds without taking advantage of hierarchical structure within a mechanism class, as we illustrate in the left panel of Figure 4 with $p = 4$. On the x -axis, we chart the growth in mechanism complexity. On the y -axis, for $i = 1, 2, 3, 4$, we plot the hypothetical average revenue over an arbitrary sample of the mechanism $A_i^* \in \mathcal{A}_i$ that maximizes empirical revenue, i.e. $\text{profit}_{\mathcal{S}}(A_i^*)$. We also plot the lower bound on the expected profit of A_i^* which is equal to the empirical profit of A_i^* minus the fixed constant $\epsilon_{\mathcal{A}}(N, \delta)$ (the dash-dot line). Finally, we know that the expected profit of A_i^* falls somewhere above this lower bound and potentially below its empirical profit, or in other words, somewhere in the constant band between the solid and dash-dot line. Our goal is

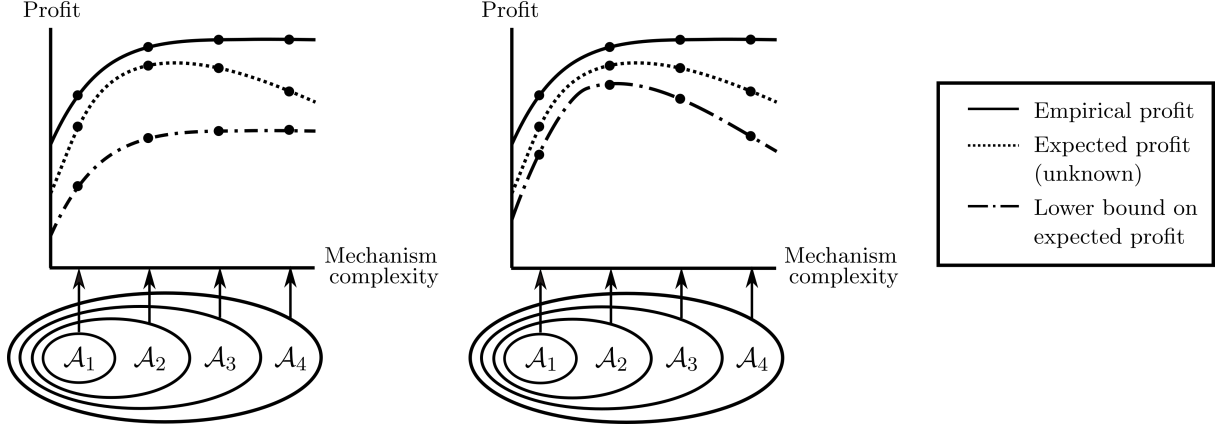


Figure 4: These charts are a theoretical illustration of the stronger bounds guaranteed via SPM (the right panel) as compared to uniform generalization guarantees (the left panel).

to maximize expected profit, but without knowing the underlying distribution, we can only choose the mechanism that maximizes the lower bound on expected profit. Crucially, the lower bound on expected profit is always increasing in this scenario and as a result, the mechanism designer may erroneously think that A_4^* is the best mechanism to field. The mechanism designer will not know when the mechanism class has grown so complex that performance on the sample no longer generalizes to the distribution, a phenomenon known as *overfitting*.

Meanwhile, our general theorem allows us to easily provide a generalization bound $\epsilon_{\mathcal{A}_i}(N, \delta \cdot w(i))$ for each individual class \mathcal{A}_i , where $w : \mathbb{N} \rightarrow [0, 1]$ is a weight function which we explain later in this section. This is illustrated in the right panel of Figure 4, where for $i = 1, 2, 3, 4$, the lower bound on the expected profit of A_i^* is equal to its empirical profit minus $\epsilon_{\mathcal{A}_i}(N, \delta \cdot w(i))$, which is proportional to the complexity of \mathcal{A}_i . Since simpler classes have lower intrinsic complexity, we will have tighter bounds for classes that are lower in the hierarchy. This lower bound is much more informative and indicates roughly when the mechanism complexity has grown so large that overfitting has occurred. By maximizing this complexity-dependent lower bound on expected profit, the mechanism designer can correctly determine that A_2^* is a better mechanism to field than A_4^* . In this way, the mechanism designer can optimize profit versus generalization. This leads us to the notion of a *non-uniformly learnable* class of mechanisms, upon which SPM can be performed.

Definition 5.1 (Non-uniformly learnable). *A class \mathcal{A} of mechanisms is non-uniformly learnable if \mathcal{A} is the countable union of a set of uniformly learnable classes $\mathcal{C} = \{\mathcal{A}_1, \mathcal{A}_2, \dots\}$ with functions $\epsilon_{\mathcal{A}_i}$ bounding the estimation loss of learning over \mathcal{A}_i . This means that given a weight function $w : \mathbb{N} \rightarrow [0, 1]$ such that $\sum_{i=1}^{\infty} w(i) = 1$, for all $\delta \in (0, 1)$ and for all distributions \mathcal{D} , with probability at least $1 - \delta$, for all mechanism classes \mathcal{A}_i such that $w(i) \neq 0$ and all mechanisms $A \in \mathcal{A}_i$, the empirical profit of A is $\epsilon_{\mathcal{A}_i}(N, \delta \cdot w(i))$ -close to its expected profit.*

In effect, the weight function allows us to spread the confidence δ across all mechanism sub-families in \mathcal{A} . Given a non-uniformly learnable class of mechanisms, structural profit maximization is the process of determining the mechanism class \mathcal{A}_i and the mechanism $A \in \mathcal{A}_i$ such that $\text{profit}_{\mathcal{S}}(A) - \epsilon_{\mathcal{A}_i}(N, \delta \cdot w(i))$ is maximized, since we know that the expected profit of A is at least this value. Further, the following theorem shows that the profit of A is close to optimal. Though this result is well-known for general learning problems, we include the proof for completeness in Appendix A.

Theorem 5.2. *Let $\mathcal{A}(\mathcal{S})$ be the mechanism that maximizes $\mathit{profit}_{\mathcal{S}}(A) - \min_{i:A \in \mathcal{A}_i} \epsilon_{\mathcal{A}_i}(N, \delta \cdot w(i))$ and let $A^* \in \mathcal{A}$ be the mechanism that maximizes $\mathit{profit}_{\mathcal{D}}(A)$. Then with probability at least $1 - 2\delta$,*

$$\mathit{profit}_{\mathcal{D}}(\mathcal{A}(\mathcal{S})) \geq \mathit{profit}_{\mathcal{D}}(A^*) - \min_{i:A^* \in \mathcal{A}_i} \epsilon_{\mathcal{A}_i}(N, \delta \cdot w(i)) - U \sqrt{\frac{1}{2N} \ln \frac{2}{\delta}}.$$

Both the structural decomposition of \mathcal{A} into subsets and the assignment w of weights to these subsets allow the mechanism designer to encode any prior knowledge he might have about the buyers. After all, the larger the weight $w(i)$ assigned to a mechanism class \mathcal{A}_i is, the larger $\delta \cdot w(i)$ is, and a larger $\delta \cdot w(i)$ implies a smaller $\epsilon_{\mathcal{A}_i}(N, \delta \cdot w(i))$, thereby implying stronger guarantees. Where might this prior knowledge come from? Oftentimes, it might come from domain expertise; perhaps the mechanism designer knows that some mechanisms are likely to be more profitable than others, so assigning a higher weight to classes containing those mechanisms will lead to better profit. To present this section's results, we will need the following notation. For an AMA A , let \mathcal{O}_A be the set of all allocations Q such that $\lambda(Q) > 0$ in A . We first derive the following generalization guarantee for the SPM hierarchy consisting of the classes of \mathcal{O} -boosted AMAs.

Theorem 5.3. *Let \mathcal{A} be the class of AMAs and let w be a weight function which maps sets of allocations \mathcal{O} to \mathbb{R} such that $\sum w(\mathcal{O}) \leq 1$. Then for any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the draw of a sample of size N from \mathcal{D} , for any auction $A \in \mathcal{A}$, the difference between the average profit of A over the sample and the expected profit of A over \mathcal{D} is at most*

$$O \left(U \sqrt{\frac{1}{N} (n + |\mathcal{O}_A|) \left(\log(n + |\mathcal{O}_A|) + \sum_{i=1}^m \log \binom{\kappa_i + n}{n} \right)} + U \sqrt{\frac{1}{N} \log \frac{1}{\delta \cdot w(\mathcal{O}_A)}} \right).$$

We prove a similar theorem for the class of anonymous pricing mechanisms, by which we mean the union of all \mathcal{B} -pricing mechanisms over all sets of bundles \mathcal{B} . For a particular anonymous pricing mechanism A , let \mathcal{B}_A be the set of bundles with set prices in A .

Theorem 5.4. *Let \mathcal{A} be the class of anonymous pricing mechanisms and let w be a weight function which maps sets of bundles \mathcal{B} to $[0, 1]$ such that $\sum w(\mathcal{B}) \leq 1$. Then for any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the draw of a sample of size N from \mathcal{D} , for any mechanism $A \in \mathcal{A}$, the difference between the average profit of A over the sample and the expected profit of A over \mathcal{D} is at most*

$$O \left(U \sqrt{\frac{1}{N} |\mathcal{B}_A| (m + \log(n |\mathcal{B}_A|))} + U \sqrt{\frac{1}{N} \log \frac{1}{\delta \cdot w(\mathcal{B}_A)}} \right).$$

Theorem 5.3 follows directly from Theorem 4.12 and Theorem 5.4 follows from Theorem 4.10, since we only need to multiply the weight term with δ as it appears in both Theorem 4.12 and 4.10. Theorems 4.7, 4.13, 4.14 and 4.10 similarly imply SPM results for length- M menus of two-part tariffs, \mathcal{O} -boosted λ -auctions, k -sparse AMAs, k -sparse λ -auctions, and bundle pricing mechanisms with non-anonymous prices.

6 Conclusion

In this work, we provide a unifying framework for bounding the complexity of a wide variety of mechanism classes. We characterize structural similarities of mechanism classes ranging from simple take-it-or-leave-it pricing mechanisms to combinatorially challenging VCG-based mechanisms such as affine maximizer auctions. These similarities lead us to an overarching theorem that bounds

the empirical Rademacher complexity of these mechanism classes. Despite this theorem’s wide applicability, we match and improve over many of the generalization guarantees already provided in the sample-based mechanism design literature. This all-encompassing theorem also allows us to easily bound the complexity of finely grained mechanism hierarchies in one swoop. We then call upon the learning framework known as *structural profit maximization* in order to show how the mechanism designer can find the precise level of each hierarchy that will provide him with the optimal tradeoff between revenue and generalization.

Our work opens many directions for future exploration. A particularly interesting one is an investigation into data-driven mechanism design from a computational complexity perspective. Assuming full expressiveness, there is an inevitable tradeoff between deriving data-dependent generalization guarantees and computational complexity because scanning the input alone will take an exponential number of steps. Further, empirical Rademacher complexity can be computationally challenging to compute exactly, but it is sometimes possible to formulate a data-dependent upper bound as a convex optimization problem (e.g. [Riondato and Upfal, 2015]). A similar approach might work here. In another direction, as of yet, there has been no research into generalization guarantees for multi-dimensional sample-based mechanism design when the bidders have valuation profiles other than general, additive, or unit-demand. It would be interesting to see whether the wealth of knowledge regarding other valuation profiles, such as submodular valuations, can lead to improved generalization bounds.

References

- [Babaioff et al., 2014] Babaioff, M., Immorlica, N., Lucier, B., and Weinberg, S. M. (2014). A simple and approximately optimal mechanism for an additive buyer. In *Proceedings of the IEEE Symposium on Foundations of Computer Science (FOCS)*.
- [Balcan et al., 2008] Balcan, M.-F., Blum, A., Hartline, J., and Mansour, Y. (2008). Reducing mechanism design to algorithm design via machine learning. *Journal of Computer and System Sciences*, 74:78–89.
- [Balcan et al., 2016] Balcan, M.-F., Sandholm, T., and Vitercik, E. (2016). Sample complexity of automated mechanism design. In *Advances In Neural Information Processing Systems (NIPS)*.
- [Balcan et al., 2017] Balcan, M.-F., Sandholm, T., and Vitercik, E. (2017). Sample complexity of multi-item profit maximization. *arXiv preprint arXiv:1705.00243*.
- [Bartlett and Mendelson, 2002] Bartlett, P. L. and Mendelson, S. (2002). Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482.
- [Buck, 1943] Buck, R. C. (1943). Partition of space. *Amer. Math. Monthly*, 50:541–544.
- [Cai et al., 2012] Cai, Y., Daskalakis, C., and Weinberg, S. M. (2012). An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the Annual Symposium on Theory of Computing (STOC)*.
- [Cai et al., 2013] Cai, Y., Daskalakis, C., and Weinberg, S. M. (2013). Reducing revenue to welfare maximization: Approximation algorithms and other generalizations. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

- [Cai et al., 2016] Cai, Y., Devanur, N. R., and Weinberg, S. M. (2016). A duality based unified approach to Bayesian mechanism design. In *Proceedings of the Annual Symposium on Theory of Computing (STOC)*.
- [Chawla et al., 2007] Chawla, S., Hartline, J. D., and Kleinberg, R. (2007). Algorithmic pricing via virtual valuations. In *Proceedings of the ACM Conference on Economics and Computation (EC)*.
- [Chawla et al., 2010] Chawla, S., Hartline, J. D., Malec, D. L., and Sivan, B. (2010). Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the Annual Symposium on Theory of Computing (STOC)*.
- [Clarke, 1971] Clarke, E. H. (1971). Multipart pricing of public goods. *Public Choice*, 11:17–33.
- [Cole and Roughgarden, 2014] Cole, R. and Roughgarden, T. (2014). The sample complexity of revenue maximization. In *Proceedings of the Annual Symposium on Theory of Computing (STOC)*.
- [Conitzer and Sandholm, 2002] Conitzer, V. and Sandholm, T. (2002). Complexity of mechanism design. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*.
- [Daskalakis and Weinberg, 2012] Daskalakis, C. and Weinberg, S. M. (2012). Symmetries and optimal multi-dimensional mechanism design. In *Proceedings of the ACM Conference on Economics and Computation (EC)*.
- [Devanur et al., 2016] Devanur, N. R., Huang, Z., and Psomas, C.-A. (2016). The sample complexity of auctions with side information. In *Proceedings of the Annual Symposium on Theory of Computing (STOC)*.
- [Dudley, 1967] Dudley, R. M. (1967). The sizes of compact subsets of Hilbert space and continuity of Gaussian processes. *Journal of Functional Analysis*, 1(3):290 – 330.
- [Elkind, 2007] Elkind, E. (2007). Designing and learning optimal finite support auctions. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- [Feldman et al., 2015] Feldman, M., Gravin, N., and Lucier, B. (2015). Combinatorial auctions via posted prices. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- [Groves, 1973] Groves, T. (1973). Incentives in teams. *Econometrica*, 41:617–631.
- [Hart and Nisan, 2012] Hart, S. and Nisan, N. (2012). Approximate revenue maximization with multiple items. In *Proceedings of the ACM Conference on Economics and Computation (EC)*.
- [Hartline and Roughgarden, 2009] Hartline, J. D. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *Proceedings of the ACM Conference on Economics and Computation (EC)*.
- [Hsu et al., 2016] Hsu, J., Morgenstern, J., Rogers, R., Roth, A., and Vohra, R. (2016). Do prices coordinate markets? In *Proceedings of the Annual Symposium on Theory of Computing (STOC)*.
- [Huang et al., 2015] Huang, Z., Mansour, Y., and Roughgarden, T. (2015). Making the most of your samples. In *Proceedings of the ACM Conference on Economics and Computation (EC)*.

- [Kleinberg and Weinberg, 2012] Kleinberg, R. and Weinberg, S. M. (2012). Matroid prophet inequalities. In *Proceedings of the Annual Symposium on Theory of Computing (STOC)*.
- [Koltchinskii, 2001] Koltchinskii, V. (2001). Rademacher penalties and structural risk minimization. *IEEE Transactions on Information Theory*, 47(5):1902–1914.
- [Lavi et al., 2003] Lavi, R., Mu’Alem, A., and Nisan, N. (2003). Towards a characterization of truthful combinatorial auctions. In *Proceedings of the IEEE Symposium on Foundations of Computer Science (FOCS)*.
- [Likhodedov and Sandholm, 2004] Likhodedov, A. and Sandholm, T. (2004). Methods for boosting revenue in combinatorial auctions. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [Likhodedov and Sandholm, 2005] Likhodedov, A. and Sandholm, T. (2005). Approximating revenue-maximizing combinatorial auctions. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [Medina and Mohri, 2014] Medina, A. M. and Mohri, M. (2014). Learning theory and algorithms for revenue optimization in second price auctions with reserve. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- [Mohri et al., 2012] Mohri, M., Rostamizadeh, A., and Talwalkar, A. (2012). *Foundations of Machine Learning*. MIT Press.
- [Morgenstern and Roughgarden, 2015] Morgenstern, J. and Roughgarden, T. (2015). On the pseudo-dimension of nearly optimal auctions. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*.
- [Morgenstern and Roughgarden, 2016] Morgenstern, J. and Roughgarden, T. (2016). Learning simple auctions. In *Proceedings of the Conference on Learning Theory (COLT)*.
- [Narasimhan et al., 2016] Narasimhan, H., Agarwal, S. B., and Parkes, D. C. (2016). Automated mechanism design without money via machine learning. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- [Narasimhan and Parkes, 2016] Narasimhan, H. and Parkes, D. C. (2016). A general statistical framework for designing strategy-proof assignment mechanisms. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*.
- [Riondato and Upfal, 2015] Riondato, M. and Upfal, E. (2015). Mining frequent itemsets through progressive sampling with rademacher averages. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- [Roberts, 1979] Roberts, K. (1979). The characterization of implementable social choice rules. In Laffont, J.-J., editor, *Aggregation and Revelation of Preferences*. North-Holland Publishing Company.
- [Roughgarden and Schrijvers, 2016] Roughgarden, T. and Schrijvers, O. (2016). Ironing in the dark. In *Proceedings of the ACM Conference on Economics and Computation (EC)*.
- [Rubinstein, 2016] Rubinstein, A. (2016). On the computational complexity of optimal simple mechanisms. In *Proceedings of the ACM Conference on Innovations in Theoretical Computer Science (ITCS)*.

- [Sandholm, 2003] Sandholm, T. (2003). Automated mechanism design: A new application area for search algorithms. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*.
- [Sandholm and Likhodedov, 2015] Sandholm, T. and Likhodedov, A. (2015). Automated design of revenue-maximizing combinatorial auctions. *Operations Research*, 63(5):1000–1025.
- [Shalev-Shwartz and Ben-David, 2014] Shalev-Shwartz, S. and Ben-David, S. (2014). *Understanding machine learning: From theory to algorithms*. Cambridge University Press.
- [Syrkkanis, 2017] Syrkkanis, V. (2017). A sample complexity measure with applications to learning optimal auctions. *arXiv preprint arXiv:1704.02598*.
- [Vickrey, 1961] Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37.
- [Wilson, 1993] Wilson, R. B. (1993). *Nonlinear pricing*. Oxford University Press on Demand.

A Proof of Theorem 5.2

Proof of Theorem 5.2. By assumption, we have that

$$\begin{aligned}
\text{profit}_{\mathcal{D}}(A^*) - \text{profit}_{\mathcal{D}}(\mathcal{A}(\mathcal{S})) &= \text{profit}_{\mathcal{D}}(A^*) - \text{profit}_{\mathcal{S}}(\mathcal{A}(\mathcal{S})) \\
&\quad + \text{profit}_{\mathcal{S}}(\mathcal{A}(\mathcal{S})) - \text{profit}_{\mathcal{D}}(\mathcal{A}(\mathcal{S})) \\
&\leq \text{profit}_{\mathcal{D}}(A^*) - \text{profit}_{\mathcal{S}}(\mathcal{A}(\mathcal{S})) + \min_{i:\mathcal{A}(\mathcal{S}) \in \mathcal{A}_i} \epsilon_{\mathcal{A}_i}(N, \delta \cdot w(i)) \\
&\leq \text{profit}_{\mathcal{D}}(A^*) - \text{profit}_{\mathcal{S}}(A^*) + \min_{i:\mathcal{A}(\mathcal{S}) \in \mathcal{A}_i} \epsilon_{\mathcal{A}_i}(N, \delta \cdot w(i)) \\
&\leq U \sqrt{\frac{\ln(2/\delta)}{2N}} + \min_{i:\mathcal{A}(\mathcal{S}) \in \mathcal{A}_i} \epsilon_{\mathcal{A}_i}(N, \delta \cdot w(i)).
\end{aligned}$$

The final inequality holds because by Hoeffding’s inequality, with probability at least $1 - \delta$, $\text{profit}_{\mathcal{D}}(A^*) - \text{profit}_{\mathcal{S}}(A^*) \leq U \sqrt{\frac{\ln(2/\delta)}{2N}}$. \square