
Maximizing Revenue under Market Shrinkage and Market Uncertainty

Maria-Florina Balcan
School of Computer Science
Carnegie Mellon University
ninamf@cs.cmu.edu

Siddharth Prasad
Computer Science Department
Carnegie Mellon University
sprasad2@cs.cmu.edu

Tuomas Sandholm
Computer Science Department
Carnegie Mellon University
Optimized Markets, Inc.
Strategic Machine, Inc.
Strategy Robot, Inc.
sandholm@cs.cmu.edu

Abstract

A shrinking market is a ubiquitous challenge faced by various industries. In this paper we formulate the first formal model of shrinking markets in multi-item settings, and study how mechanism design and machine learning can help preserve revenue in an uncertain, shrinking market. Via a sample-based learning mechanism, we prove the first guarantees on how much revenue can be preserved by truthful multi-item, multi-bidder auctions (for limited supply) when only a random unknown fraction of the population participates in the market. We first present a general reduction that converts any sufficiently rich auction class into a randomized auction robust to market shrinkage. Our main technique is a novel combinatorial construction called a *winner diagram* that concisely represents all possible executions of an auction on an uncertain set of bidders. Via a probabilistic analysis of winner diagrams, we derive a general possibility result: a sufficiently rich class of auctions always contains an auction that is robust to market shrinkage and market uncertainty. Our result has applications to important practically-constrained settings such as auctions with a limited number of winners. We then show how to efficiently learn an auction that is robust to market shrinkage by leveraging practically-efficient routines for solving the winner determination problem.

1 Introduction

A shrinking market with uncertain buyer participation is a natural phase of products' and services' lifecycles. Current examples of great importance include media consumers—known as cord cutters—who cancel cable-TV subscriptions in favor of streaming services [29], a thinning customer base for department stores due to online retailers like Amazon [11, 16], and reduced capacities for restaurants during the COVID-19 pandemic [43]. In this paper we study how mechanism design can help preserve revenue in this ubiquitous challenge of a shrinking market, specifically for combinatorial auctions for limited supply. The seller has m indivisible items to allocate to a set S of n bidders. The bidders can express how much they value each possible bundle $b \subseteq \{1, \dots, m\}$ of items. Combinatorial auctions have had wide reach in practice, from strategic sourcing to spectrum auctions to estate auctions. Cramton et al. [10] provide a survey of various aspects of combinatorial auctions. The design of revenue-maximizing combinatorial auctions in multi-item, multi-bidder settings is an elusive and difficult problem that has spurred a long and active line of research combining techniques from economics, artificial intelligence, and theoretical computer science. This is still largely an open question and a very active research area.

In this paper we start a new strand within that topic area, namely the study of shrinking markets with uncertain buyer participation. We introduce the first formal model of market shrinkage in multi-item settings, and prove the first revenue guarantees. Specifically, we show how much revenue can be preserved when only a random unknown fraction of the set S of bidders participates in the market.

1.1 Summary of the contributions of this paper

We present the first formal analysis of how much revenue can be preserved in a shrinking market, for multi-item settings. More precisely, there is a set S of n bidders that is known to the mechanism designer. Each bidder participates in the market independently with probability p , but the valuations of the bidders who participate in the market, denoted by $S_0 \subseteq S$, are unknown (what is known is that they belong to S). We present a learning-based method for designing a mechanism that satisfies the first known revenue-preservation guarantees in this setting.

In Section 2 we provide a formal description of the problem setting, including formal definitions of the auction settings we consider. We precisely show how to reckon with subtleties that arise when auctions are run among a shrunken market of unknown size.

In Section 3 we provide and discuss a simple example of a market where reduced competition in a shrinking market drives revenue to a lower threshold than one might expect. We furthermore show that if bidders's valuation functions can depend on what items other bidders receive, there exist scenarios in which only an exponentially small (in the number of items) fraction of the revenue obtainable by even the vanilla Vickrey-Clarke-Groves (VCG) auction on S can be guaranteed on a random subset of bidders, even if a large fraction of the market shows up. For example, if 50 items are for sale and each bidder shows up independently with 90% probability, our construction yields a maximum expected revenue of roughly 7% of the VCG revenue on S . If 100 items are for sale, at most 0.52% of the VCG revenue on S can be guaranteed.

Our main theorem is the following revenue guarantee obtained via a sample-based learning algorithm. Delineability is a structural assumption introduced by Balcan et al. [3] satisfied by nearly all commonly studied auction classes. $W_{\mathcal{M}}(S)$ denotes the maximum welfare achievable by mechanisms in \mathcal{M} , k is a term that depends on the number of winners in any mechanism in \mathcal{M} , and γ is a constant that depends on S . Rev_M denotes the revenue function induced by M . All notation and formal definitions are in the following sections. Sections 4 and 5 are dedicated to proving Theorem 1.1.

Theorem 1.1. *Let \mathcal{M} be (d, h) -delineable class of mechanisms. A mechanism $\widetilde{M} \in \mathcal{M}$ such that*

$$\mathbb{E}[\text{Rev}_{\widetilde{M}}(S_0)] \geq \Omega\left(\frac{p^2}{k^{1+\log_{1/\gamma}(4/p)}}\right) W_{\mathcal{M}}(S) - \varepsilon$$

with probability at least $1 - \delta$ can be computed in $NhT + (Nh)^{O(d)}$ time, where T is the time required to generate any given hyperplane witnessing delineability of any mechanism in \mathcal{M} and $N = O\left(\frac{d \log(dh)}{\varepsilon^2} \log\left(\frac{1}{\delta}\right)\right)$.

In Section 4 we prove that $\sup_{M \in \mathcal{M}} \mathbb{E}[\text{Rev}_M(S_0)] \geq \Omega\left(\frac{p^2}{k^{1+\log_{1/\gamma}(4/p)}}\right) W_{\mathcal{M}}(S)$, which is the major technical contribution of this paper. Our main technique is the analysis of a novel combinatorial structure we construct called a *winner diagram*, which is a graph that concisely captures all possible executions of an auction on an uncertain set of bidders. Via a probabilistic method argument that randomizes over a subgraph of the winner diagram, we arrive at a general possibility result: *if \mathcal{M} is a sufficiently rich class of mechanisms, there always exists an $M \in \mathcal{M}$ that is robust to uncertainty/shrinkage in the market.* This implies our bound on $\sup \mathbb{E}[\text{Rev}_M(S_0)]$. We primarily focus on the case where bidders participate in the market independently with probability p , but show how to generalize our results to any distribution over submarkets. Our bound is a parameterized guarantee that has interesting applications to practically motivated auction constraints: (1) limiting the number of winners and (2) bundling constraints on items.

In Section 5 we present a learning algorithm to compute a mechanism \widetilde{M} such that $\mathbb{E}[\text{Rev}_{\widetilde{M}}(S_0)] \geq \sup_{M \in \mathcal{M}} \mathbb{E}[\text{Rev}_M(S_0)] - \varepsilon$ with high probability, which proves Theorem 1.1. Our algorithm exploits geometric structure and a linear-programming approach over hyperplane arrangements. We show the run-time of our procedure is computationally tractable for a specific auction class by leveraging practically-efficient routines for solving the winner determination problem.

1.2 Related work

Shrinking markets and uncertainty Shrinking markets have been studied by various researchers in the context of oil companies [47], cable TV [1, 29], labor markets [23], telecom markets [31], housing markets [24], and in combinatorial settings including a thinning customer base for department stores due to online retailers like Amazon [11, 16] and reduced capacities for restaurants during the COVID-19 pandemic [43]. Most of this existing research is extremely domain specific, and provides advisory content based on historical observations, data, and general economic knowledge. We introduce the first formal model of market shrinkage in multi-item settings, and prove the first known guarantees for how much revenue can be preserved in a shrinking market. Our guarantee on the revenue preserved in a shrinking market provides a positive contrast to recent work of Dobzinski and Uziely [12], who study the effect of market shrinkage on revenue loss. They show that even in the case of selling a single item to n buyers with known valuation distributions, the absence of a single buyer with a fixed “low” value can surprisingly result in a (multiplicative) revenue loss of $\frac{1}{e+1}$ (in expectation). We tackle the significantly more complex multi-item setting. Furthermore, our main results are prior-free (in that they are tailored to the specific set S of bidders and do not require bidders to come from a distribution) and thus provide a strong positive contrast to this negative result.

Our results can also be viewed from the perspective of an uncertain market, since at the point of the mechanism design the subset of bidders that participates in the market is unknown. Mechanism design with uncertainty about bidder valuations has previously studied [28, 45], but to the best of our knowledge the prior-free setting for combinatorial auctions has not been considered.

Revenue in combinatorial auctions We prove guarantees when the seller limits the number of winners and when the seller places bundling constraints on the items (for example, by enforcing that certain items must be sold together). Reasons for limiting the number of winners include: (1) avoiding the logistical hassle of having a large number of winners – this constraint is commonly used in sourcing auctions [21, 36, 37, 41] and (2) increasing competition to boost revenue, as is studied by Roughgarden et al. [34] (though in a different setting than ours). Kroer and Sandholm [25] show that the VCG auction run with bundling constraints can yield significant revenue gains.

Furthermore, catalyzed by the seminal work of Bulow and Klemperer [7], a recent line of work studies the competition complexity of auctions, and provides results that compare the optimal (expected) revenue to the revenue of mechanisms like VCG when the number of bidders is augmented [6, 8, 14, 15, 34]. Competition complexity results can be seen as tackling the “opposite” situation of a growing market. However, existing work in this area uses the very different objective of expected revenue over buyer valuation distributions, and moreover has only tackled restrictive classes of valuation functions (such as unit-demand and additive valuations). Also related is the work of Rastegari et al. [32], who study settings where revenue can counterintuitively be boosted by dropping bidders.

Learning for auction design Our algorithm for designing a mechanism robust to uncertain market shrinkage is a sample-based learning algorithm. Such *sample-based automated mechanism design* methods were introduced by Sandholm and Likhodedov [26, 27, 42], with the first generalization guarantees for machine learning based mechanism design appearing in Balcan et al. [2]. Machine learning for auction design has since grown into a rich field blending theory and practice [3, 4, 13, 30]. Unlike most prior work on machine learning for auction design, our setting is prior-free and does not have access to samples. In our learning algorithm, the mechanism designer produces samples himself that simulate the shrunken market, which is along the lines of the learning-within-an-instance paradigm of Balcan et al. [5]. For simplicity, we use the term *sample-based* as a catch-all for any algorithm that uses samples, prior-free or not. Furthermore, while our algorithm is a sample-based, the techniques we use to prove our main guarantee are fundamentally new. The aforementioned literature seeks to find a near-optimal mechanism from samples—in contrast our goal is to prove concrete revenue guarantees on the optimal mechanism.

2 Problem formulation

Combinatorial auctions A seller has m indivisible items to allocate among a set S of n bidders. Each bidder is described by her combinatorial valuation function $v_i : 2^{\{1, \dots, m\}} \rightarrow \mathbb{R}_{\geq 0}$. We assume

bidders' valuations satisfy *free-disposal*, that is, $v_i(b) \leq v_i(b')$ for all $b \subseteq b' \subseteq \{1, \dots, m\}$, and $v_i(\emptyset) = 0$. For an allocation α , $v_i(\alpha)$ is bidder i 's value for the bundle she receives under α .

A mechanism M specifies, given a set of bidders S , an allocation α of the items and a payment p_i to be collected from each bidder $i \in S$. The welfare of M when run on S is $W_M(S) = \sum_{i \in S} v_i(\alpha)$, and its revenue on S is $\text{Rev}_M(S) = \sum_{i \in S} p_i$. M is *incentive compatible* if no bidder i may strictly increase her utility (value minus payment) by misreporting her true valuation function v_i to M . M is *individually rational* if $v_i(\alpha) \geq p_i$ for all i . All mechanisms considered in this paper are assumed to be incentive compatible and individually rational. The classical Vickrey-Clarke-Groves (VCG) auction [9, 17, 48] uses an efficient allocation α^* (one that maximizes $\sum_{i \in S} v_i(\alpha)$), and charges bidder i a payment of $\max_{\alpha} \sum_{j \neq i} v_j(\alpha) - \sum_{j \neq i} v_j(\alpha^*)$. The VCG auction is truthful and individually rational. It is well known that the VCG auction can yield low revenue in many scenarios.

Market-size uncertainty In our model, the mechanism designer has full knowledge of the entire population of bidders S (described by their valuation functions). An unknown random subset of S participates in the market. We write $S_0 \sim_p S$ to denote a subset S_0 that is sampled from S by including each bidder in S_0 independently with probability p . More generally, for a distribution D over 2^S , we write $S_0 \sim_D S$ to denote a random subset of S chosen according to D . We are interested in what happens to the maximum revenue achievable when only a random fraction of the set S of bidders participates in the auction, that is, $\sup_{M \in \mathcal{M}} \mathbb{E}_{S_0 \sim_D S} [\text{Rev}_M(S_0)]$.

Since a variable group of bidders of variable size can participate in the auction mechanisms we run, we require the important assumption that auctions in \mathcal{M} can be run on variable-size sets of bidders in a well-defined manner. Various well-studied classes of auctions satisfy this property: examples include the class of VCG auctions with reserve prices, λ -auctions [22], affine-maximizer auctions [33], and various population-size-independent auctions in Balcan et al. [5]. We assume that the mechanism designer knows the valuations of the bidders in S to begin with. So, each bidder can be thought of having an identity (for example, “the bidder who values apples at x and oranges at y ”, or “the bidder with valuation function v_4 ”), and the mechanism designer knows the identities/valuations v_1, \dots, v_n of all bidders in S . An allocation, formally, is a mapping from items to bidder identities.

The sequence of mechanism design and revelation in our setting is the same as in the standard mechanism design setting. Specifically, the mechanism design (computation of a mechanism from \mathcal{M}) takes place before the bidders in the shrunken market are asked to reveal their valuations. This is important for incentive compatibility, that is, for motivating the bidders to reveal their true valuations. If the design/choice of $M \in \mathcal{M}$ is allowed to be based on the revealed valuations, the auction might not be incentive compatible. (The class of second price auctions with reserve prices for a single item serves as an illustrative example. Choosing the reserve price to maximize revenue after the shrunken market is revealed clearly violates incentive compatibility.) Because the designer does not know exactly which bidders are in the shrunken market S_0 , the designer has uncertainty about the valuations of the bidders. He only knows that they belong to S .

Assumptions on \mathcal{M} and S For any $S' \subseteq S$, let $W_{\mathcal{M}}(S') = \max_{M \in \mathcal{M}} W_M(S')$ and $\text{Rev}_{\mathcal{M}}(S') = \max_{M \in \mathcal{M}} \text{Rev}_M(S')$. Let $\text{win}_M(S')$ denote the set of bidders in S' that win a nonempty bundle of items per M . Let $\text{win}_{\mathcal{M}}(S')$ denote the set of bidders in S' that win a nonempty bundle of items per the mechanism in \mathcal{M} achieving $W_{\mathcal{M}}(S')$. The following two assumptions are the most critical ones.

Welfare submodularity For all $S_1, S_2 \subseteq S$, $W_{\mathcal{M}}(S_1) + W_{\mathcal{M}}(S_2) \geq W_{\mathcal{M}}(S_1 \cup S_2) + W_{\mathcal{M}}(S_1 \cap S_2)$.

Winner monotonicity For all $S'' \subseteq S' \subseteq S$ and all $i \in S''$, $i \in \text{win}_{\mathcal{M}}(S') \implies i \in \text{win}_{\mathcal{M}}(S'')$.

Suppose $W_{\mathcal{M}}(S') = W_{VCG}(S')$ for any $S' \subseteq S$, that is, \mathcal{M} is sufficiently rich to be able to allocate items efficiently (as is the case with all mechanisms in the hierarchies discussed by Balcan et al. [3, 5]). Then, welfare submodularity implies winner monotonicity [19]. If the valuation functions of bidders in S satisfy the gross-substitutes property, then both welfare submodularity and winner monotonicity hold [18, 19, 49].

The final assumptions stipulate that \mathcal{M} is a sufficiently rich class of mechanisms. We assume that $\text{Rev}_{\mathcal{M}}(S') = W_{\mathcal{M}}(S') = W_{VCG}(S')$ for all $S' \subseteq S$, and further that \mathcal{M} satisfies the following “global VCG-like” property: $\text{Rev}_{\mathcal{M}}(S')$ depends only on $\text{win}_{\mathcal{M}}(S')$ and $\text{win}_{\mathcal{M}}(S' \setminus \{i\})$ for each $i \in \text{win}_{\mathcal{M}}(S')$. In words, these conditions stipulate the following: (1) \mathcal{M} is sufficiently rich such that in a non-truthful full-information setting, \mathcal{M} can always extract the full social surplus

$W_{\mathcal{M}}(S') = W_{VCG}(S')$ as revenue and (2) for any S' , the payments collected by the revenue-maximizing mechanism M that achieves $\text{Rev}_{\mathcal{M}}(S')$ depend only on $W_{\mathcal{M}}(S)$ and the maximum welfare $W_{\mathcal{M}}(S \setminus \{i\})$ achievable when each bidder drops out.

As a concrete example, if S is a set of bidders with gross-substitutes valuations, then the class of λ -auctions and the class of affine-maximizer auctions satisfy all of the above properties (we prove this fact in Appendix A, and provide more details about these auction classes).

3 Revenue loss can be drastic

At first glance it might appear that the expected revenue preserved by a mechanism M when each bidder participates independently with probability p should simply be $p \cdot \text{Rev}_{\mathcal{M}}(S)$ (or more if one thinks of revenue as having diminishing returns in the number of bidders). This intuition is indeed accurate if $\text{Rev}_{\mathcal{M}}$ is a submodular function (which captures the diminishing returns property). However, revenue can shrink by more than this when mechanisms in \mathcal{M} do not have submodular revenue. One reason for greater revenue loss is reduced competition among buyers. For example, suppose there are m items and $2m$ bidders, where bidder i for $1 \leq i \leq m$ has valuation $v_i(b) = c$ if $i \in b$ and $v_i(b) = 0$ otherwise, and bidder $m+i$ for $1 \leq i \leq m$ has valuation $v_{m+i}(b) = c - \varepsilon/m$ if $i \in b$ and $v_{m+i}(b) = 0$ otherwise (bidders have combinatorial valuations in this example, so valuation functions only depend on the bundle of items received). The VCG auction will allocate item $i \in \{1, \dots, m\}$ to bidder i . The payment collected from bidder i will be $c - \varepsilon/m$, which is the second highest value for item i . The revenue from VCG is thus $mc - \varepsilon = W(S) - \varepsilon$. Now, suppose each bidder participates in the auction independently with probability p . The expected revenue can be computed by breaking it up across items:

$$\mathbb{E}[\text{Rev}_{VCG}(S_0)] = \sum_{i=1}^m \mathbb{E}[\text{Revenue from item } i] = \sum_{i=1}^m p^2(c - \varepsilon/m) = p^2(W(S) - \varepsilon).$$

The third equality is due to the fact that VCG generates nonzero revenue from item i if and only if both bidders i and $m+i$ participate, since if at most one of them shows up there is no competition for that item. So $\mathbb{E}[\text{Rev}_{VCG}(S_0)] = p^2 \cdot \text{Rev}_{VCG}(S)$.

Furthermore, if bidders' valuations are *allocational*, that is, $v_i(\alpha)$ can depend on what items other bidders receive, revenue loss can be even more dramatic.

Theorem 3.1. *For any $\varepsilon > 0$ there exists a set S of bidders with allocational valuations such that for $S_0 \sim_p S$, $\mathbb{E}[\text{Rev}_{\mathcal{M}}(S_0)] \leq p^{m/2} \cdot (\text{Rev}_{VCG}(S) + 2\varepsilon) + \varepsilon$ for any individually rational auction M .*

The exponential revenue decay in the number of items means that even if the shrunken market is large in expectation, the revenue loss can be dramatic. For example, if 50 items are for sale and each bidder shows up independently with 90% probability, our construction shows that any auction can guarantee only at most roughly 7% of the VCG revenue on S . If 100 items are for sale, at most 0.52% of the VCG revenue on S can be guaranteed.

4 Main guarantee on preserved revenue

We now present our main revenue guarantee when each bidder participates in the auction independently with probability p . For a set of bidders $S' \subseteq S$, let $\omega(S') = \text{win}_{\mathcal{M}}(S') \cup (\cup_{i \in S'} \text{win}_{\mathcal{M}}(S' \setminus \{i\}))$ be the set of bidders in S' whose valuations determine $\text{Rev}_{\mathcal{M}}(S')$. Define an equivalence relation \equiv on subsets of S by $S_1 \equiv S_2$ if and only if $\omega(S_1) = \omega(S_2)$. Let $\varphi(S') = \frac{1}{n} \sum_{i=1}^n W_{\mathcal{M}}(S' \setminus \{i\})$. φ serves as a potential function in the proof of the following theorem and represents the average max-welfare of S' when a uniformly random bidder in S drops out.

Mechanism \mathcal{A}

- (1) Let S_1, \dots, S_ℓ be an enumeration of the equivalence classes with $\varphi(S_i) > \frac{p}{4} W_{\mathcal{M}}(S)$.
- (2) Let M_1, \dots, M_ℓ denote the mechanisms that achieve $\text{Rev}_{\mathcal{M}}(S_1), \dots, \text{Rev}_{\mathcal{M}}(S_\ell)$.
- (3) Choose M uniformly at random from $\{M_1, \dots, M_\ell\}$, and run M .

The main challenge in analyzing this mechanism is bounding ℓ . Before we do that, we analyze the revenue guarantee it satisfies in terms of ℓ .

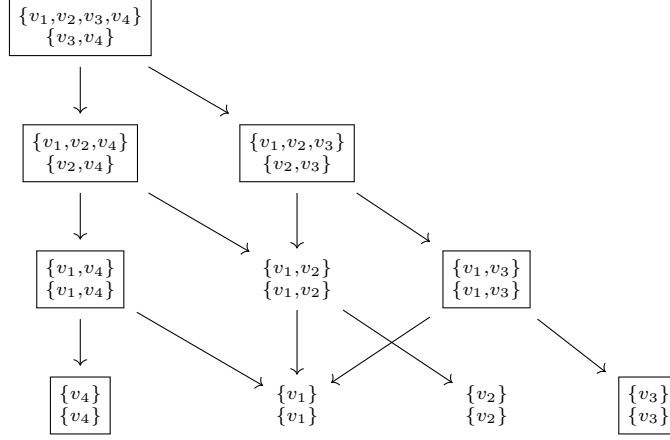


Figure 1: A winner diagram representing a second-price auction with a single item and four bidders with valuations $S = \{v_1 = 1, v_2 = 2, v_3 = 4, v_4 = 8\}$. At each node, the top set S' is the set of remaining bidders, and the bottom set is the set of bidders $\omega(S')$ that actually determine revenue. Boxed nodes represent heavy equivalence classes for $p = 8/9$, which is the subgraph of the winner diagram \mathcal{A} randomizes over.

Lemma 4.1. *Let $|S| \geq 2$. For $S_0 \sim_p S$, $\mathbb{E}[\text{Rev}_{\mathcal{A}}(S_0)] \geq \Omega(p^2/\ell)W_{\mathcal{M}}(S)$.*

Proof. By definition of ω , if $S_1 \equiv S_2$, then $W_{\mathcal{M}}(S_1) = W_{\mathcal{M}}(S_2)$, $\varphi(S_1) = \varphi(S_2)$, and $\text{Rev}_{\mathcal{M}}(S_1) = \text{Rev}_{\mathcal{M}}(S_2)$ (and the maximum revenue is achieved by the same $M \in \mathcal{M}$ for both sets). Call a set of bidders $S' \subseteq S$ *heavy* if $\varphi(S') > \frac{p}{4}W_{\mathcal{M}}(S)$. If S_0 is heavy, there is $M \in \{M_1, \dots, M_\ell\}$ such that $\text{Rev}_M(S_0) = W_{\mathcal{M}}(S_0)$, so $\mathbb{E}_{\mathcal{A}}[\text{Rev}_{\mathcal{A}}(S_0)] \geq \frac{1}{\ell}W(S_0) \geq \frac{1}{\ell}\varphi(S_0) > \frac{p/4}{\ell}W_{\mathcal{M}}(S)$. Let H denote the event that S_0 is heavy. Then,

$$\mathbb{E}_{\mathcal{A}}[\mathbb{E}_{S_0}[\text{Rev}_{\mathcal{A}}(S_0)]] = \mathbb{E}_{S_0}[\mathbb{E}_{\mathcal{A}}[\text{Rev}_{\mathcal{A}}(S_0)]] \geq \mathbb{E}_{S_0}[\mathbb{E}_{\mathcal{A}}[\text{Rev}_{\mathcal{A}}(S_0)|H]] \cdot \Pr(H) \geq \frac{p/4}{\ell}W(S) \cdot \Pr(H).$$

We now derive a lower bound on $\Pr(H)$. We have

$$\mathbb{E}_{S_0}[\varphi(S_0)] = \mathbb{E}_{S_0}[\frac{1}{n} \sum_{i=1}^n W_{\mathcal{M}}(S_0 \setminus \{i\})] = \mathbb{E}_{i \sim S}[\mathbb{E}_{S_0}[W_{\mathcal{M}}(S_0 \setminus \{i\})]] \geq \frac{p}{2}W_{\mathcal{M}}(S)$$

where in the final inequality we use the fact that $|S| \geq 2$ and that $W_{\mathcal{M}}$ is submodular, and so by Hartline et al. [20], $\mathbb{E}_{S_0}[W_{\mathcal{M}}(S_0)] \geq pW_{\mathcal{M}}(S)$. By Markov's inequality on the (nonnegative) random variable $W_{\mathcal{M}}(S) - \varphi(S_0)$,

$$\Pr(S_0 \text{ is heavy}) \geq \frac{(p/2)W_{\mathcal{M}}(S) - (p/4)W_{\mathcal{M}}(S)}{W_{\mathcal{M}}(S) - (p/4)W_{\mathcal{M}}(S)} = \frac{p/4}{1-p/4} \geq \frac{p}{4}.$$

Substituting this into our previous bound yields $\mathbb{E}_{\mathcal{A}}[\mathbb{E}_{S_0}[\text{Rev}_{\mathcal{A}}(S_0)]] \geq \frac{p^2}{16\ell}W_{\mathcal{M}}(S)$, as desired. \square

We now bound the number of heavy equivalence classes ℓ . In order to do this, we introduce the notion of a *winner diagram*, which is a subgraph of the Hasse diagram of S . The winner diagram for S is the following directed graph \mathcal{G} : each node is labeled $(S', \omega(S'))$ for some subset $S' \subset S$. The root node is labeled $(S, \omega(S))$. The children of node $(S', \omega(S'))$ are given by $(S' \setminus \{i\}, \omega(S' \setminus \{i\}))$ for each $i \in \omega(S')$. Figure 1 illustrates the winner diagram corresponding to a second-price auction for a single item with four bidders. Winner monotonicity will allow us to show that \mathcal{G} contains a node that represents every equivalence class of \equiv .

Lemma 4.2. *\mathcal{G} contains all equivalence classes of \equiv .*

Proof. Let $S^* \subseteq S$ be a set of bidders that arises as a winner set, that is, $S^* = \omega(S'')$ for some $S'' \supseteq S^*$. The set $S' \supset S'' \supset S^*$ is *maximal* for S^* if $\omega(S') = S^*$ and $\omega(S' \cup \{i\}) \neq S^*$ for every $i \notin S'$. We show that for a given winner set of bidders S^* , there is a unique maximal set of bidders $S' \supseteq S^*$ such that $\omega(S') = S^*$. Initialize $S' = S^*$, and greedily add bidders from S to S' while $\omega(S') = S^*$ does not change. Due to winner monotonicity, if $i \notin \omega(S')$, then $i \notin \omega(S' \cup \{j\})$ for

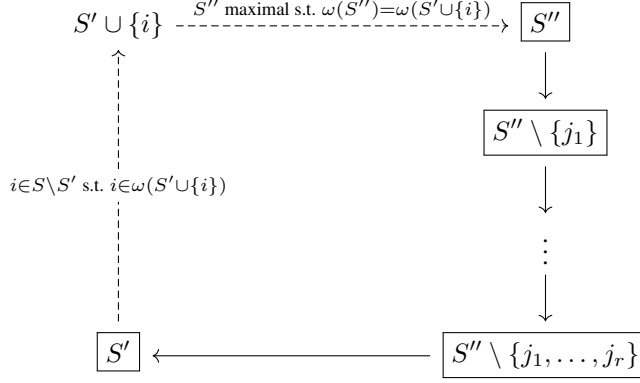


Figure 2: Illustration of the inductive step in Lemma 4.2. Boxed sets correspond to representative elements of equivalence classes in \mathcal{G} . Solid arrows represent directed edges in \mathcal{G} from parent to child.

any bidder j . Hence, the order in which bidders are added by the greedy procedure does not matter, and therefore the final set S' is the unique maximal set for S^* . Let the representative element of each equivalence class $[(S', \omega(S'))]$ be the one such that S' is maximal for $\omega(S')$.

We prove the lemma by backwards induction on the size of the representative set S' of any equivalence class. The base case of $|S'| = n$ is immediate since the root $(S, \omega(S))$ is the only node for which the representative set has size n . For the inductive step suppose that \mathcal{G} contains a node for every equivalence class for which the representative set is of size at least n' . Let $(S', \omega(S'))$ be the representative of an equivalence class with $|S'| = n' - 1$. Let $i \notin S'$ be a bidder such that $i \in \omega(S' \cup \{i\})$. Such an i exists due to winner monotonicity: if $i \in \omega(S)$, then $i \in \omega(S' \cup \{i\})$, since $S' \cup \{i\} \subset S$. Let S'' be the maximal set such that $\omega(S'') = \omega(S' \cup \{i\})$. We have $|S''| \geq |S' \cup \{i\}| > |S'|$, so by the induction hypothesis \mathcal{G} contains a node labelled $(S'', \omega(S''))$. Now, there must exist a bidder $j_1 \in S'' \setminus S'$ such that $j_1 \in \omega(S'')$. If not, adding all the bidders in $S'' \setminus S'$ to S' would not introduce any new winners, that is, $\omega(S' \cup (S'' \setminus S')) = \omega(S'') = \omega(S')$, which contradicts the maximality of S' for $\omega(S')$. Therefore, the node $(S'', \omega(S''))$ has a child

$$(S'' \setminus \{j_1\}, \omega(S'' \setminus \{j_1\}))$$

($S'' \setminus \{j_1\}$ is maximal due to winner monotonicity). We may now find a bidder $j_2 \in S'' \setminus \{j_1\} \setminus S'$ such that $j_2 \in \omega(S'' \setminus \{j_1\} \setminus S')$ for the same reason as before. Continuing in this fashion yields a path from $(S'', \omega(S''))$ to $(S', \omega(S'))$, so $(S', \omega(S')) \in \mathcal{G}$, as desired. \square

Combining Lemmas 4.1 and 4.2 yields our main guarantee. Let $\gamma = \max_{S', i \in \omega(S')} \frac{\varphi(S' \setminus \{i\})}{\varphi(S')}$ and let $k = \max_{S'} |\omega(S')|$. We have $k \leq 2m$. The parameter γ measures the smallest decrease in φ between any two levels of \mathcal{G} , which we use to control the depth of nodes considered by our main mechanism \mathcal{A} . We stipulate that $\gamma < 1$. In Appendix C, we discuss how to remove this assumption and replace γ with an appropriate parameter that is unconditionally strictly less than 1.

Theorem 4.3. *Let $|S| \geq 2$ and $S_0 \sim_p S$. We have*

$$\mathbb{E}_{\mathcal{A}}[\mathbb{E}_{S_0}[\text{Rev}_{\mathcal{A}}(S_0)]] \geq \Omega\left(\frac{p^2}{k^{1+\log_{1/\gamma}(4/p)}}\right) W_{\mathcal{M}}(S).$$

If \mathcal{M} consists of revenue-monotonic mechanisms the slightly improved bound $\Omega\left(\frac{p^2}{k^{\log_{1/\gamma}(4/p)}}\right) W_{\mathcal{M}}(S)$ holds. In particular, there exist mechanisms in \mathcal{M} achieving the above guarantees in expectation.

Proof. Let \mathcal{G}' denote the restriction of the winner diagram \mathcal{G} to nodes representing heavy equivalence classes. Each node of \mathcal{G}' has out-degree at most $k = \max_{S'} |\omega(S')|$, and the depth of \mathcal{G}' is at most $\log_{1/\gamma}\left(\frac{W_{\mathcal{M}}(S)}{(p/4)W_{\mathcal{M}}(S)}\right) = \log_{1/\gamma}(4/p)$ since φ decreases by a factor of at least γ when passing from a parent node to a child node (and \mathcal{G} is truncated at nodes that are not heavy). Hence the number of nodes in \mathcal{G}' is at most $k^{1+\log_{1/\gamma}(4/p)}$. If mechanisms in \mathcal{M} are revenue monotonic, then we may modify \mathcal{A} to randomize only over mechanisms corresponding to nodes of \mathcal{G}' with out-degree 0. The

number of such nodes is at most $k^{\log_{1/\gamma}(4/p)}$. By Lemma 4.2, we may substitute this quantity for ℓ in Lemma 4.1, which completes the proof. \square

We have thus shown, via an application of the probabilistic method, that if \mathcal{M} is a sufficiently rich mechanism class, there always exists $M \in \mathcal{M}$ that is robust to uncertainty in the market. Thus, $\sup_{M \in \mathcal{M}} \mathbb{E}[\text{Rev}_M(S_0)] \geq \Omega\left(\frac{p^2}{k^{1+\log_{1/\gamma}(4/p)}}\right) W_{\mathcal{M}}(S)$, with a slight improvement under revenue monotonicity.

In Section 3 we showed that if bidders can have allocational valuations there exist scenarios where only an exponentially small fraction of revenue can be guaranteed. Theorem 4.3 also holds in this setting under a stronger version of winner monotonicity, which we show in Appendix C.

4.1 Applications

The dependence of Theorem 4.3 on $\max_{S'} |\omega(S')|$ allows us to derive interesting families of guarantees when the seller places practical constraints on the auction setting. The first is a constraint on the number of winners, and the second is a bundling constraint that favors allocations that sell certain items together. Reasons for limiting the number of winners include: (1) avoiding the logistical hassle of having a large number of winners – this constraint is commonly used in sourcing auctions [21, 36, 37, 41] and (2) increasing competition to boost revenue, as is studied in Roughgarden et al. [34] (though in a different setting than ours). Kroer and Sandholm [25] show that even the vanilla VCG auction run with bundling constraints can yield significant revenue gains compared to VCG with no bundling constraints.

4.1.1 Limiting the number of winners

Suppose \mathcal{M} is a class of mechanisms such that $|\text{win}_M(S)| \leq n_0$ for all $M \in \mathcal{M}$ such that $W_{\mathcal{M}}$ is submodular, winner-monotonic, and satisfies the global-VCG-like property discussed previously. The proofs of all previous theorems go through with this constraint taken into account, with parameters modified correspondingly. Let φ, γ be defined as previously.

Theorem 4.4. *Let $|S| \geq 2$ and $S_0 \sim p$, and let \mathcal{M} be a class of mechanisms that sell to at most n_0 bidders. Then, there exists $M \in \mathcal{M}$ such that*

$$\mathbb{E}[\text{Rev}_M(S_0)] \geq \Omega\left(\frac{p^2}{(2n_0)^{1+\log_{1/\gamma}(4/p)}}\right) W_{\mathcal{M}}(S).$$

In practical settings the auction designer might limit the number of bidders that can win a nonempty bundle of items – and in such cases n_0 can potentially be treated as a constant relative to m and n .

In Appendix C, we consider VCG-like auctions that favor allocations that bundle certain items together. For such auctions, we obtain a guarantee in terms of the relevant bundling constraints.

4.2 General distribution over submarkets

So far we have stated our guarantees under the assumption that each bidder participates in the auction independently with probability p . When bidders participated independently with probability p , welfare submodularity was required to ensure that $\mathbb{E}[W_{\mathcal{M}}(S_0)] \geq pW_{\mathcal{M}}(S)$. In Appendix C we give a more general guarantee in terms of $\mathbb{E}[W_{\mathcal{M}}(S_0)]$ which only requires winner monotonicity.

5 How to choose an auction

Computing the mechanism $M \in \mathcal{M}$ that achieves the revenue guarantee of Theorem 4.3 can be accomplished by searching over the set $\{M_1, \dots, M_\ell\}$ that \mathcal{A} randomizes over, but this would potentially be a highly-inefficient procedure. Moreover, \mathcal{A} itself is not computationally-efficient: determining the heavy sets of bidders, and determining the mechanisms M_1, \dots, M_ℓ that are revenue maximizing for the heavy sets is an exhaustive procedure that would require enumerating over a potentially exponential number of subsets of S .

A more natural way for the mechanism designer to arrive at a mechanism is to learn from samples, which ensures that the mechanism designer uses the auction that (nearly) optimizes the expected preserved revenue, which could be significantly higher than what Theorem 4.3 guarantees.

We give a learning algorithm that the mechanism designer can use to learn a mechanism $\widetilde{M} \in \mathcal{M}$ that achieves an expected revenue of nearly $\sup_{M \in \mathcal{M}} \mathbb{E}[\text{Rev}_M(S_0)]$. Our algorithm is similar in spirit to the learning-within-an-instance paradigm of Balcan et al. [5]. To describe the algorithm, we require the following structural notion of mechanism *delineability* introduced by Balcan et al. [3].

Definition 5.1 (Balcan et al. [3]). Let $d, h \in \mathbb{N}$. A mechanism M is (d, h) -*delineable* if (1) each mechanism $M \in \mathcal{M}$ can be parameterized by a d -dimensional vector $\theta \in \mathbb{R}^d$ and (2) for any set $S' \subseteq S$ of bidder valuations, there is a set \mathcal{H} of at most h hyperplanes in \mathbb{R}^d such that $\text{Rev}_\theta(S)$ is linear in θ over any given connected component of $\mathbb{R}^d \setminus \mathcal{H}$.

Theorem 1.1. *Let \mathcal{M} be (d, h) -delineable class of mechanisms. A mechanism $\widetilde{M} \in \mathcal{M}$ such that*

$$\mathbb{E}[\text{Rev}_{\widetilde{M}}(S_0)] \geq \Omega\left(\frac{p^2}{k^{1+\log_{1/\gamma}(4/p)}}\right) W_{\mathcal{M}}(S) - \varepsilon$$

with probability at least $1 - \delta$ can be computed in $NhT + (Nh)^{O(d)}$ time, where T is the time required to generate any given hyperplane witnessing delineability of any mechanism in \mathcal{M} and $N = O\left(\frac{d \log(dh)}{\varepsilon^2} \log\left(\frac{1}{\delta}\right)\right)$.

Proof. We design a mechanism \widetilde{M} such that $\mathbb{E}[\text{Rev}_{\widetilde{M}}(S_0)] \geq \sup_{M \in \mathcal{M}} \mathbb{E}[\text{Rev}_M(S_0)] - \varepsilon$ with high probability. The theorem statement then follows from Theorem 4.3. Our algorithm is based on the framework of empirical risk minimization from machine learning. The mechanism designer samples $S_1, \dots, S_N \subseteq S$ independently and identically according to distribution D on 2^S . (We assume for simplicity that sampling according to D can be done in a computationally efficient manner. If bidders participate independently with probability p , then the mechanism designer simply needs to flip N coins of bias p for each of the n bidders in S .) The auction used will be the one that maximizes empirical revenue $\widetilde{M} = \arg\max_{M \in \mathcal{M}} \frac{1}{N} \sum_{t=1}^N \text{Rev}_M(S_t)$. Balcan et al. [3] show that $N = O\left(\frac{d \log(dh)}{\varepsilon^2} \log\left(\frac{1}{\delta}\right)\right)$ samples suffice to guarantee that the expected revenue of \widetilde{M} is ε -close to optimal with probability at least $1 - \delta$ over the draw of S_1, \dots, S_N .

We now determine the computational complexity of maximizing empirical revenue. Our algorithm exploits similar geometric intuition that was used by Balcan et al. [3] to derive the above sample complexity guarantee. A similar approach has been used in other settings as well [4, 5].

For each $S_t \in \{S_1, \dots, S_n\}$, let \mathcal{H}_t denote the set of at most h hyperplanes witnessing (d, h) -delineability of \mathcal{M} , and let $\mathcal{H} = \cup_t \mathcal{H}_t$, so $|\mathcal{H}| \leq Nh$. The number of connected components of $\mathbb{R}^d \setminus \mathcal{H}$ is at most $|\mathcal{H}|^d \leq (Nh)^d$. Each connected component is a convex polyhedron that is the intersection of at most $|\mathcal{H}|$ halfspaces. Representations of these regions as 0/1 constraint-vectors of length \mathcal{H} (a 0 in entry $h \in \mathcal{H}$ corresponds to one side of h , a 1 corresponds to the other side) can be computed in $\text{poly}(|\mathcal{H}|^d)$ time using standard techniques [46]. Empirical revenue is linear as a function of θ in each connected component due to delineability, so the parameter θ that maximizes empirical revenue within a given component can be found by solving a linear program that involves d variables and at most $|\mathcal{H}|$ constraints, which can be done in $\text{poly}(|\mathcal{H}|, d)$ time. \square

Our algorithm has a run-time that is exponential in the number of parameters d required to describe mechanisms in \mathcal{M} . In Appendix D, we study a class of sparse λ -auctions that can be described by a constant number of parameters. By leveraging practically-efficient routines for winner determination [35, 39, 40] (a generalization of the problem of computing welfare-maximizing allocations), we show how our empirical revenue maximization is computationally tractable for this setting (in particular, the run-time T of computing the hyperplanes witnessing delineability is in terms of the run-time of winner determination).

6 Conclusions and future research

Our work is the first to formally study the problem of preserving revenue in a shrinking market. We gave a sample-based learning algorithm to design a mechanism that is robust to shrinkage

and uncertainty in the market. The crux of our analysis was a new combinatorial construction we introduced called a winner diagram. Our model of a shrinking market is simple and natural, and can serve as a testbed for many exciting mechanism design questions.

There are several open questions and new interesting research directions that stem from this study. The most immediate question is to derive tight bounds on revenue loss. There is a gap between the bound of Theorem 4.3 and the $(1 - p^2)$ -fraction revenue loss of the simple example of a market with competition. Where does the true answer lie? Another interesting, and seemingly more difficult, setting is the one where the mechanism designer does not know the distribution D over 2^S beforehand. Can he still arrive at an auction that is robust to the shrinking market? If the mechanism designer knows that each bidder participates independently with probability p , but does not know p , is it still possible to design a robust auction? Finally, we believe that the combinatorial bidder structure uncovered by our notion of a winner diagram could have interesting applications to other areas in mechanism design. While our analysis required a number of assumptions on the set of bidders, it would be interesting to extend the concept of a winner diagram to prove more general results with weaker assumptions. It would be interesting to extend our techniques to understand market shrinkage in other settings including objectives beyond revenue, other auction classes, and unlimited supply.

Acknowledgments and Disclosure of Funding

This material is based on work supported by the National Science Foundation under grants CCF1733556, CCF-1910321, IIS-1901403, and SES-1919453, the Defense Advanced Research Projects Agency under cooperative agreement HR00112020003, an AWS Machine Learning Research Award, an Amazon Research Award, a Bloomberg Research Grant, and a Microsoft Research Faculty Fellowship.

References

- [1] Nicole P Aliloupour. The impact of technology on the entertainment distribution market: The effects of Netflix and Hulu on cable revenue. 2016.
- [2] Maria-Florina Balcan, Avrim Blum, Jason D Hartline, and Yishay Mansour. Mechanism design via machine learning. In *IEEE Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2005.
- [3] Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik. A general theory of sample complexity for multi-item profit maximization. In *ACM Conference on Economics and Computation (EC)*, 2018.
- [4] Maria-Florina Balcan, Siddharth Prasad, and Tuomas Sandholm. Efficient algorithms for learning revenue-maximizing two-part tariffs. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2020.
- [5] Maria-Florina Balcan, Siddharth Prasad, and Tuomas Sandholm. Learning within an instance for designing high-revenue combinatorial auctions. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2021.
- [6] Hedyeh Beyhaghi and S Matthew Weinberg. Optimal (and benchmark-optimal) competition complexity for additive buyers over independent items. In *ACM Symposium on Theory of Computing (STOC)*, 2019.
- [7] Jeremy Bulow and Paul Klemperer. Auctions versus negotiations. *The American Economic Review*, 86(1):180, 1996.
- [8] Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *ACM Symposium on Theory of Computing (STOC)*, pages 311–320, 2010.
- [9] Ed H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [10] Peter Cramton, Yoav Shoham, and Richard Steinberg. *Combinatorial Auctions*. MIT Press, 2006.
- [11] Michael A Cusumano. Amazon and Whole Foods: follow the strategy (and the money). *Communications of the ACM*, 60(10):24–26, 2017.

- [12] Shahar Dobzinski and Nitzan Uziely. Revenue loss in shrinking markets. In *ACM Conference on Economics and Computation (EC)*, 2018.
- [13] Paul Dütting, Zhe Feng, Harikrishna Narasimhan, David Parkes, and Sai Srivatsa Ravindranath. Optimal auctions through deep learning. In *International Conference on Machine Learning (ICML)*, pages 1706–1715. PMLR, 2019.
- [14] Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S Matthew Weinberg. The competition complexity of auctions: A Bulow-Klemperer result for multi-dimensional bidders. In *ACM Conference on Economics and Computation (EC)*, 2017.
- [15] Michal Feldman, Ophir Friedler, and Aviad Rubinstein. 99% revenue via enhanced competition. In *ACM Conference on Economics and Computation (EC)*, 2018.
- [16] Maris Goldmanis, Ali Hortaçsu, Chad Syverson, and Önsel Emre. E-commerce and the market structure of retail industries. *The Economic Journal*, 120(545):651–682, 2010.
- [17] Theodore Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [18] Faruk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87:95–124, 1999.
- [19] Mingyu Guo. VCG redistribution with gross substitutes. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2011.
- [20] Jason Hartline, Vahab Mirrokni, and Mukund Sundararajan. Optimal marketing strategies over social networks. In *Proceedings of the 17th International Conference on World Wide Web (WWW)*, pages 189–198, 2008.
- [21] Gail Hohner, John Rich, Ed Ng, Grant Reid, Andrew J. Davenport, Jayant R. Kalagnanam, Ho Soo Lee, and Chae An. Combinatorial and quantity-discount procurement auctions benefit Mars, Incorporated and its suppliers. *Interfaces*, 33(1):23–35, 2003.
- [22] Philippe Jehiel, Moritz Meyer-Ter-Vehn, and Benny Moldovanu. Mixed bundling auctions. *Journal of Economic Theory*, 127(1):494–512, 2007.
- [23] Randall S Jones and Haruki Seitani. Labour market reform in Japan to cope with a shrinking and ageing population. *OECD Economic Department Working Papers*, (1568), 2019.
- [24] Kohei Kawai, Masatomo Suzuki, and Chihiro Shimizu. Shrinkage in Tokyo’s central business district: Large-scale redevelopment in the spatially shrinking office market. *Sustainability*, 11(10):2742, 2019.
- [25] Christian Kroer and Tuomas Sandholm. Computational bundling for auctions. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2015. Extended version: Carnegie Mellon University, Computer Science Department technical report CMU-CS-13-111, 2013.
- [26] Anton Likhodedov and Tuomas Sandholm. Methods for boosting revenue in combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 232–237, San Jose, CA, 2004.
- [27] Anton Likhodedov and Tuomas Sandholm. Approximating revenue-maximizing combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Pittsburgh, PA, 2005.
- [28] Giuseppe Lopomo, Luca Rigotti, and Chris Shannon. Uncertainty in mechanism design. *arXiv preprint arXiv:2108.12633*, 2021.
- [29] Victor J Massad. Understanding the cord-cutters: An adoption/self-efficacy approach. *International Journal on Media Management*, 20(3):216–237, 2018.
- [30] Jamie Morgenstern and Tim Roughgarden. Learning simple auctions. In *Conference on Learning Theory (COLT)*, pages 1298–1318. PMLR, 2016.
- [31] Ioannis Neokosmidis, Theodoros Rokkas, Dimitris Xydias, Antonino Albanese, Muhammad Shuaib Siddiqui, Carlos Colman-Meixner, and Dimitra Simeonidou. Are 5g networks and the neutral host model the solution to the shrinking telecom market. In *IFIP International Conference on Artificial Intelligence Applications and Innovations*, pages 70–77. Springer, 2018.
- [32] Baharak Rastegari, Anne Condon, and Kevin Leyton-Brown. Revenue monotonicity in deterministic, dominant-strategy combinatorial auctions. *Artif. Intell.*, 175(2):441–456, 2011.

- [33] Kevin Roberts. The characterization of implementable social choice rules. In J-J Laffont, editor, *Aggregation and Revelation of Preferences*. North-Holland Publishing Company, 1979.
- [34] Tim Roughgarden, Inbal Talgam-Cohen, and Qiqi Yan. Robust auctions for revenue via enhanced competition. *Operations Research*, 2020.
- [35] Tuomas Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135:1–54, January 2002.
- [36] Tuomas Sandholm. Expressive commerce and its application to sourcing: How we conducted \$35 billion of generalized combinatorial auctions. *AI Magazine*, 28(3):45–58, 2007.
- [37] Tuomas Sandholm. Very-large-scale generalized combinatorial multi-attribute auctions: Lessons from conducting \$60 billion of sourcing. In Zvika Neeman, Alvin Roth, and Nir Vulkan, editors, *Handbook of Market Design*. Oxford University Press, 2013.
- [38] Tuomas Sandholm and Anton Likhodedov. Automated design of revenue-maximizing combinatorial auctions. *Operations Research*, 63(5):1000–1025, 2015.
- [39] Tuomas Sandholm and Subhash Suri. BOB: Improved winner determination in combinatorial auctions and generalizations. *Artificial Intelligence*, 145:33–58, 2003.
- [40] Tuomas Sandholm, Subhash Suri, Andrew Gilpin, and David Levine. CABOB: A fast optimal algorithm for winner determination in combinatorial auctions. *Management Science*, 51(3): 374–390, 2005. Early version in IJCAI-01.
- [41] Tuomas Sandholm, David Levine, Michael Concordia, Paul Martyn, Rick Hughes, Jim Jacobs, and Dennis Begg. Changing the game in strategic sourcing at Procter & Gamble: Expressive competition enabled by optimization. *Interfaces*, 36(1):55–68, 2006.
- [42] Tuomas Sandholm, Anton Likhodedov, and Andrew Gilpin. Automated design of revenue-maximizing combinatorial auctions. *Operations Research*, 63(5):1000–1025, 2015. Special issue on Computational Economics. Subsumes and extends over a AAAI-05 paper and a AAAI-04 paper.
- [43] Hyoung Ju Song, Jihwan Yeon, and Seoki Lee. Impact of the COVID-19 pandemic: Evidence from the US restaurant industry. *International Journal of Hospitality Management*, 92:102702, 2021.
- [44] Pingzhong Tang and Tuomas Sandholm. Mixed-bundling auctions with reserve prices. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2012.
- [45] Taiki Todo. Mechanism design with uncertainty. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI)*, 2020.
- [46] Csaba D Tóth, Joseph O’Rourke, and Jacob E Goodman. *Handbook of Discrete and Computational Geometry*. Chapman and Hall/CRC, 2017.
- [47] Thijs Van de Graaf. Battling for a shrinking market: oil producers, the renewables revolution, and the risk of stranded assets. In *The Geopolitics of Renewables*, pages 97–121. Springer, 2018.
- [48] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.
- [49] Makoto Yokoo, Yuko Sakurai, and Shigeo Matsubara. The effect of false-name bids in combinatorial auctions: New fraud in Internet auctions. *Games and Economic Behavior*, 46(1):174–188, 2004.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#)
 - (b) Did you describe the limitations of your work? [\[Yes\]](#) We described the limitations of our work in the Conclusions and future research section
 - (c) Did you discuss any potential negative societal impacts of your work? [\[N/A\]](#)

- (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
- 2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes]
- 3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
- 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
- 5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Omitted results from Section 2

A λ -auction is an incentive-compatible parameterized generalization of the VCG auction where the mechanism designer may specify additive boosts to specific allocations. λ -auctions were introduced by Jehiel et al. [22]. The class of λ -auctions is a rich class of auctions and has been studied towards designing high-revenue combinatorial auctions [3, 38]. We use the notation $W(\alpha) = \sum_{i=1}^n v_i(\alpha)$, and $W(S) = \max_{\alpha} W(\alpha)$. Formally, a λ -auction run among n buyers is specified by a vector $\lambda \in \mathbb{R}^{(n+1)^m}$ indexed by the $(n+1)^m$ possible allocations. The overall allocation chosen is

$$\alpha^* = \operatorname{argmax}_{\alpha} W(\alpha) + \lambda(\alpha)$$

and bidder i is charged a payment of

$$\max_{\alpha} (W_{-i}(\alpha) + \lambda(\alpha)) - (W_{-i}(\alpha^*) + \lambda(\alpha^*)).$$

An affine-maximizer auction [33] is a generalization of a λ -auction with multiplicative bidder-specific weights. We do not define them here, but all subsequent discussions on λ -auctions apply to affine maximizers as well.

The following proposition shows that λ -auctions are rich enough (and hence so are affine maximizers) to extract the entire social surplus as revenue if the bidders' valuations are known beforehand.

Proposition A.1. *For any set of bidders S , there exists a λ -auction, in the full-information setting with no incentive-compatibility constraints, with revenue equal to $W(S)$.*

Proof. Let α denote the efficient allocation among bidders in S and let α_{-i} denote the efficient allocation among bidders in $S \setminus \{i\}$. We show that the λ -auction with $\lambda(\alpha) = 0$, $\lambda(\alpha_{-i}) = W(\alpha) - W_{-i}(\alpha_{-i})$, and $\lambda(\beta) = -\infty$ for all other allocations β collects a payment of $v_i(\alpha)$ from each bidder, and thus extracts a revenue of $W(S)$. First, note that α is the overall allocation used since $W(\alpha) + \lambda(\alpha) \geq W(\alpha_{-i}) + \lambda(\alpha_{-i})$ for each i . To show that α_{-i} is the allocation used when bidder i is absent, observe that

$$\begin{aligned} W_{-i}(\alpha_{-i}) + \lambda(\alpha_{-i}) &= W(\alpha) \\ &\geq W(\alpha) - v_i(\alpha) \\ &= W_{-i}(\alpha) + \lambda(\alpha) \end{aligned}$$

and

$$\begin{aligned} W_{-i}(\alpha_{-i}) + \lambda(\alpha_{-i}) &= W(\alpha) \\ &\geq W(\alpha) - v_i(\alpha_{-j}) \\ &= W(\alpha) + v_j(\alpha_{-j}) - v_i(\alpha_{-j}) \\ &= W_{-i}(\alpha_{-j}) + W(\alpha) - W_{-j}(\alpha_{-j}) \\ &= W_{-i}(\alpha_{-j}) + \lambda(\alpha_{-j}) \end{aligned}$$

for any $j \neq i$. (We use the fact that if α allocates nothing to i , then $v_i(\alpha) = 0$.)

Thus the allocations used by this λ -auction are precisely the VCG allocations. The payment of bidder i is therefore $(W_{-i}(\alpha_{-i}) + \lambda(\alpha_{-i})) - (W_{-i}(\alpha) + \lambda(\alpha)) = W(\alpha) - (W(\alpha) - v_i(\alpha)) = v_i(\alpha)$, and so the total revenue is $W(\alpha)$. \square

As a consequence of the above proof, λ -auctions and affine-maximizer auctions both satisfy the global-VCG-like property discussed in Section 2.

Since a variable group of bidders of variable size can participate in the λ -auctions we run, it is important to formalize how to distinguish between allocations since the λ -auction adds “boosts” to allocations specifically. We assume that the mechanism designer knows the valuations of the bidders in S to begin with. So, each bidder can be thought of having an identity (for example, “the bidder who values apples at x and oranges at y ”, or “the bidder with valuation function v_4 ”), and the mechanism designer knows the identities/valuations v_1, \dots, v_n of all bidders in S . An allocation, formally, is a mapping from items to bidder identities. Traditionally λ -auctions assume that the number of bidders is fixed, so allocations are usually interpreted as mappings from items to the position/index of a given

bidder. So, we, too, will consider λ -auctions run among all bidders, so an auction is parameterized by a $(n + 1)^m$ -dimensional vector that specifies boosts for allocations among all n identities in S . However, our auctions also have to be well defined in the case of a shrinking market that has only a subset of the bidders. We address that as follows. If a λ -auction is run among a subset of bidder identities S_0 and chooses an allocation that allocates an item to a bidder identity not in S_0 , we assume that the seller keeps that item.

B Omitted proof from Section 3

When bidder valuations can depend on what items the other bidders receive, we can construct a simple example where a random fraction of bidders participating incurs a much more significant revenue loss than the $(1 - p^2)$ fraction in our example for the combinatorial setting. Suppose there is a distinguished set of k bidders (who can be viewed as unwealthy bidders) with negligibly low valuations for each bundle. All other bidders' (who can be viewed as wealthy bidders) valuation functions are defined to be above some threshold only on allocations that give a nonempty bundle to each of the k distinguished bidders (and zero otherwise). In other words, while the wealthy bidders would like to receive items, they are not willing to participate unless unwealthy bidders are also guaranteed items. Then, for any nontrivial fraction of the revenue to be preserved, all distinguished bidders must participate in the auction, which occurs with probability p^k . Thus, any auction can preserve revenue at most $p^k \cdot W(S)$. The number of distinguished bidders k can be taken to be as large as, for example, $m/2$. We now give a formal construction of the described example. Our construction satisfies the property that even the vanilla VCG auction extracts revenue nearly equal to the entire social surplus on the full set S of bidders.

Theorem B.1. *For any $\varepsilon > 0$ there exists a set S of bidders with allocational valuations such that*

$$\sup_{M \in \mathcal{M}} \mathbb{E}_{S_0 \sim_p S} [\text{Rev}_M(S_0)] \leq p^{m/2} \cdot (\text{Rev}_{VCG}(S) + 2\varepsilon) + \varepsilon$$

for any auction class \mathcal{M} .

Proof. For each item $1 \leq i \leq m/2$ we introduce two buyers with valuations $v_{i,1}, v_{i,2}$. For each item $m/2+1 \leq j \leq m$ we introduce a single buyer with valuation v_j . For $1 \leq i \leq m/2$ valuations $v_{i,1}$ are defined by $v_{i,1}(\alpha) = c$ if bidder $(i, 1)$ is allocated item i and bidders $j = m/2+1, \dots, m$ each receive at least one item, and $v_{i,1}(\alpha) = 0$ otherwise. Valuations $v_{i,2}$ are defined by $v_{i,2}(\alpha) = c - 2\varepsilon/m$ if bidder $(i, 2)$ is allocated item i and bidders $j = m/2+1, \dots, m$ each receive at least one item, and $v_{i,2}(\alpha) = 0$ otherwise. The only requirement on the valuations of bidders $j = m/2+1, \dots, m$ is that $v_j(\alpha) \leq 2\varepsilon/m$ for all α . The VCG auction would allocate item i to bidder $(i, 1)$ for each $i = 1, \dots, m/2$, and allocate the remaining $m/2$ items to bidders $j = m/2+1, \dots, m$ such that each bidder j receives exactly one item. The welfare of this (efficient) allocation is at most $cm/2 + \varepsilon$. The revenue obtained by VCG is at least $cm/2 - \varepsilon = W(S) - 2\varepsilon$. Let S^* denote the set of small-valuation bidders $j = m/2+1, \dots, m$. If each bidder shows up independently with probability p , the expected revenue of any auction M is

$$\begin{aligned} \mathbb{E}[\text{Rev}_M(S_0)] &= \mathbb{E}[\text{Rev}_M(S_0) \mid S^* \subseteq S_0] \cdot \Pr(S^* \subseteq S_0) + \mathbb{E}[\text{Rev}_M(S_0) \mid S^* \not\subseteq S_0] \cdot \Pr(S^* \not\subseteq S_0) \\ &\leq p^{m/2} \cdot \mathbb{E}[\text{Rev}_M(S_0) \mid S^* \subseteq S_0] + \mathbb{E}[\text{Rev}_M(S_0) \mid S^* \not\subseteq S_0] \\ &\leq p^{m/2} \cdot \mathbb{E}[W(S_0) \mid S^* \subseteq S_0] + W(S^*) \\ &\leq p^{m/2} \cdot W(S) + \varepsilon, \end{aligned}$$

as desired. □

C Omitted results and proofs from Section 4

C.1 More details on γ

If $\gamma = 1$, we simply replace it with

$$\gamma := \max \left\{ \frac{\varphi(S' \setminus \{i\})}{\varphi(S')} : i \in \omega(S'), \frac{\varphi(S' \setminus \{i\})}{\varphi(S')} < 1 \right\} < 1$$

which is the worst-case non-trivial decrease in φ across any two levels of the winner diagram. We are able to enforce that $\gamma < 1$ in this way because if S' is such that $\varphi(S') = \varphi(S' \setminus \{i\})$, then the same mechanism achieves optimal revenue on both S' and $S' \setminus \{i\}$. So these nodes can be considered jointly as a single node in the winner diagram, without incurring any penalty in the number of heavy equivalence classes \mathcal{A} randomizes over.

C.2 Allocational valuations

All of our results hold when bidders have allocational valuations. The only modification is that we require a stronger version of winner monotonicity, which states that if i wins bundle b under the mechanism that achieves $W_{\mathcal{M}}(S')$, then i must win a bundle $b' \supseteq b$ under the mechanism that achieves $W_{\mathcal{M}}(S'')$ for any $i \in S'' \subseteq S'$. The weaker notion of winner monotonicity implies the stronger version when bidders have combinatorial valuations, but this is not necessarily the case when bidders have allocational valuations.

C.3 Limiting the number of winners

Proof of Theorem 4.4. Given a set $S' \subseteq S$ of bidders, if a bidder in $\text{win}(S')$ is removed, at most one new bidder can win a nonempty bundle per the efficient allocation, due to winner monotonicity and the limit on the number of bidders. Thus, $\max_{S'} |\omega(S')| \leq 2n_0$. The remainder of the proof follows from the same arguments used to prove Theorem 4.3. \square

C.4 Bundling constraints

A *bundling* is a partition of the set of items $\{1, \dots, m\}$. An allocation α *respects* a bundling ϕ if no two items in the same bundle according to ϕ are allocated to different buyers. For a set of bundlings Φ , the class of Φ -boosted λ -auctions consists of all λ -auctions satisfying $\lambda(\alpha) \geq 0$ for all α that respects a bundling in Φ and $\lambda(\alpha) = 0$ otherwise. Let $W^{\Phi}(S)$ denote the maximum welfare of any allocation that respects a bundling in Φ . Proposition A.1 holds with $W^{\Phi}(S)$ for the class of Φ -boosted λ -auctions (the λ -auction constructed in Proposition A.1 can be shifted by a constant additive factor to make all boosts nonnegative). In the following theorem statement, $\varphi(S') = \frac{1}{n} \sum_{i=1}^n W^{\Phi}(S' \setminus \{i\})$.

Theorem C.1. *Let Φ be a set of bundlings. Let S be a set of $n \geq 2$ bidders with valuations such that $W^{\Phi} : 2^S \rightarrow \mathbb{R}_{\geq 0}$ is submodular. Let $\gamma = \max_{S', i \in \omega(S')} \frac{\varphi(S' \setminus \{i\})}{\varphi(S')}$. Let m_0 be the greatest number of bundles in any bundling in Φ . Let \mathcal{M} be the class of Φ -boosted λ -auctions. Then,*

$$\sup_{\lambda \in \mathcal{M}} \mathbb{E}_{S_0 \sim_p S} [\text{Rev}_{\lambda}(S_0)] \geq \Omega \left(\frac{p^2}{(2m_0)^{1+\log_{1/\gamma}(4/p)}} \right) W^{\Phi}(S).$$

Proof. At most m_0 bidders can win a nonempty bundle of items, so $\max_{S'} |\omega(S')| \leq 2m_0$ by the same reasoning used to prove Theorem 4.4. The arguments used to prove Theorem 4.3 yield the desired bound. \square

C.5 General distribution over submarkets

Our proof techniques easily generalize to handle any distribution D over subsets of bidders since the only statistic of the distribution required is the expected welfare of a random subset of bidders $\mathbb{E}_{S_0 \sim_D S} [W_{\mathcal{M}}(S_0)]$. When bidders participated independently with probability p , submodularity of the welfare function was required to ensure that $\mathbb{E}[W_{\mathcal{M}}(S_0)] \geq pW_{\mathcal{M}}(S)$. In the following more general guarantee, which is in terms of $\mathbb{E}[W_{\mathcal{M}}(S_0)]$, we only need the more general condition of winner monotonicity.

Theorem C.2. *Let S be a set of $n \geq 2$ bidders with valuations that satisfy winner monotonicity. Let D be a distribution supported on 2^S with $\mathbb{E}_{S_0 \sim_D S} [W_{\mathcal{M}}(S_0)] = \mu \cdot W_{\mathcal{M}}(S)$. Let $\gamma = \max_{S', i \in \omega(S')} \frac{\varphi(S' \setminus \{i\})}{\varphi(S')}$ and let $k = \max_{S'} |\omega(S')|$. We have*

$$\sup_{M \in \mathcal{M}} \mathbb{E}_{S_0 \sim_D S} [\text{Rev}_M(S_0)] \geq \frac{\eta\mu}{k^{1+\log_{1/\gamma}(1/\eta\mu)}} \left(\frac{\mu - 2\eta\mu}{2(1 - \eta\mu)} \right) \cdot W_{\mathcal{M}}(S)$$

for all $0 \leq \eta \leq 1/2$.

Proof. The proof is nearly identical to that of Theorem 4.3. The main modification is that $S' \subset S$ is heavy if $\varphi(S') \geq \eta\mu \cdot W_{\mathcal{M}}(S)$, and \mathcal{A} randomizes over mechanisms corresponding to sets S' with this property. Then, $\mathbb{E}_{S_0 \sim_{D,S}}[\varphi(S_0)] \geq \frac{\eta}{2} W_{\mathcal{M}}(S)$ and so Markov's inequality yields $\Pr(S_0 \text{ is heavy}) \geq \frac{\mu/2-\eta}{1-\eta}$. The remainder of the proof is identical. \square

Versions of Theorems 4.4 and Theorems C.1 for general distributions can be similarly obtained.

D Omitted results from Section 5

Learning a high-revenue λ -auction would require a number of samples on the order of $(n+1)^m$ [3]. However, for sparse λ -auctions that are restricted to boost only a constant number of allocations, we can perform sample and computationally efficient learning while satisfying a similar guarantee to the ones derived for the entire class.

D.1 λ auctions with limited boosting

Let $\Gamma \subseteq \{1, \dots, (n+1)^m\}$ be a set of allocations. The class of Γ -boosted λ -auctions consists of all λ -auctions satisfying $\lambda(\alpha) \geq 0$ for all $\alpha \in \Gamma$ and $\lambda(\alpha) = 0$ for all $\alpha \notin \Gamma$ (and can be specified by vectors in $\mathbb{R}^{|\Gamma|}$). Γ -boosted λ -auctions were introduced by Balcan, Sandholm, and Vitercik [3]. Let $W^\Gamma(S) = \max_{\alpha \in \Gamma} W(\alpha)$. We may derive a guarantee for this class of auctions analogous to Theorem 4.3. In the following theorem statement, $\varphi(S') = \frac{1}{n} \sum_{i=1}^n W^\Gamma(S' \setminus \{i\})$.

Theorem D.1. *Let S be a set of $n \geq 2$ bidders with valuations such that $W^\Gamma : 2^S \rightarrow \mathbb{R}_{\geq 0}$ is submodular. Let $\gamma = \max_{S', i \in \omega(S')} \frac{\varphi(S' \setminus \{i\})}{\varphi(S')}$ and let k be the maximum number of winners in any allocation in Γ . The class \mathcal{M} of Γ -boosted λ -auctions satisfies*

$$\sup_{\lambda \in \mathcal{M}} \mathbb{E}_{S_0 \sim_p S} [\text{Rev}_\lambda(S_0)] \geq \Omega \left(\frac{p^2}{(2k)^{1+\log_{1/\gamma}(4/p)}} \right) W^\Gamma(S).$$

D.2 Algorithm for learning an auction from samples

We now give an algorithm that the mechanism designer can use to compute a Γ -boosted λ -auction that nearly achieves an expected revenue of $\sup_{\lambda \in \mathcal{M}} \mathbb{E}_{S_0 \sim_p S} [\text{Rev}_\lambda(S_0)]$. Our algorithm leverages practically-efficient routines for solving winner determination, which is a generalization of the problem of computing welfare-maximizing allocations. The mechanism designer samples several subsets of bidders according to D , and computes the Γ -boosted λ -auction that maximizes empirical revenue over the samples.

While computing the empirical-revenue-maximizing auction is NP-hard in general, since winner determination is NP-hard, winner determination can be efficiently solved in practice [35, 39, 40]. Furthermore, when bidders have gross-substitutes valuations, winner determination can be solved in polynomial time [10]. The run-time of our algorithm is exponential only in $|\Gamma|$ but polynomial in all other problem parameters (including the run-time required to solve winner determination with m items and n bidders).

Theorem D.2. *Let \mathcal{M} be the class of Γ -boosted λ -auctions. A $\hat{\lambda} \in \mathcal{M}$ such that*

$$\mathbb{E}_{S_0 \sim S} [\text{Rev}_{\hat{\lambda}}(S_0)] \geq \sup_{\lambda \in \mathcal{M}} \mathbb{E}_{S_0 \sim S} [\text{Rev}_\lambda(S_0)] - \varepsilon$$

with probability at least $1 - \delta$ can be computed in $N(\min\{m, n\} + 1)w(m, n) + (Nn|\Gamma|)^{O(|\Gamma|)}$ time, where $w(m, n)$ is the time required to solve winner determination for n buyers with valuations over m items and $N = O\left(\frac{|\Gamma| \ln(n|\Gamma|)}{\varepsilon^2} \ln\left(\frac{1}{\delta}\right)\right)$.

Proof. The algorithm and proof are similar to Theorem 1.1. \mathcal{M} is $(|\Gamma|, O(n|\Gamma|^2))$ -delineable [3]. Now, we explicitly describe how to generate the hyperplanes witnessing delineability. For each $1 \leq t \leq N$ let α^t denote the efficient allocation among bidders in S_t . For each $1 \leq t \leq N$ and each $i \in \text{win}(S_t)$ let α^t_{-i} denote the efficient allocation among bidders in $S_t \setminus \{i\}$. For $i \notin \text{win}(S_t)$,

$\alpha_{-i}^t = \alpha^t$. Determining these allocations requires at most $N + N \cdot \min\{m, n\}$ calls to the winner determination routine, since $|\text{win}(S_t)| \leq m$. The allocation used by any Γ -boosted λ -auction on S_t is in $\Gamma \cup \{\alpha^t\}$, and the allocation used to determine the payment by bidder $i \in S_t$ by any Γ -boosted λ -auction on S_t is in $\Gamma \cup \{\alpha_{-i}^t\}$.

For each t and each pair of allocations $\alpha, \alpha' \in \Gamma \cup \{\alpha^t\}$, let $H(t, \alpha, \alpha')$ denote the hyperplane

$$\sum_{i \in S_t} v_i(\alpha) + \lambda(\alpha) = \sum_{i \in S_t} v_i(\alpha') + \lambda(\alpha').$$

For each t , each $i \in S_t$, and each pair of allocations $\alpha, \alpha' \in \Gamma \cup \{\alpha_{-i}^t\}$, let $H_{-i}(t, \alpha, \alpha')$ denote the hyperplane

$$\sum_{j \in S_t \setminus \{i\}} v_j(\alpha) + \lambda(\alpha) = \sum_{j \in S_t \setminus \{i\}} v_j(\alpha') + \lambda(\alpha').$$

Let \mathcal{H} denote the collection of these hyperplanes. The total number of such hyperplanes is at most $N(|\Gamma| + 1)^2 + Nn(|\Gamma| + 1)^2$. It is a basic combinatorial fact that \mathcal{H} partitions $\mathbb{R}^{|\Gamma|}$ into at most $|\mathcal{H}|^{|\Gamma|} \leq (N(n+1)(|\Gamma|+1)^2)^{|\Gamma|}$ regions. Each region is a convex polytope that is the intersection of at most $|\mathcal{H}|$ halfspaces. Representations of these regions as 0/1 constraint-vectors of length $|\mathcal{H}|$ can be computed in $\text{poly}(|\mathcal{H}|^{|\Gamma|})$ time using standard techniques [46]. Empirical revenue is linear as a function of λ in each region, since the allocations used by λ are constant within as λ varies in a given region. Thus, the auction maximizes empirical revenue within a given region can be found by solving a linear program that involves $|\Gamma|$ variables and at most $|\mathcal{H}|$ constraints, which can be done in $\text{poly}(|\mathcal{H}|, |\Gamma|)$ time. \square

One special case is when all allocations in Γ are given the *same* boost. Then, the parameter space is \mathbb{R} , the number of relevant regions (subintervals of \mathbb{R}) is $O(Nn|\Gamma|^2)$, and the algorithm in Theorem 1.1 has a run-time of $O(N \min\{m, n\}w(m, n) + Nn|\Gamma|^2)$. Mixed bundling auctions [22, 44] are an instance of this with $|\Gamma| = n$.

D.3 Structural revenue maximization

Let $\Gamma_1 \subset \Gamma_2$ be collections of allocations, and let \mathcal{M}_1 and \mathcal{M}_2 denote the classes of Γ_1 -boosted λ -auctions and Γ_2 -boosted λ -auctions, respectively. Suppose the mechanism designer has drawn some number of samples N , and observes that the empirical-revenue-maximizing auction λ_2 over \mathcal{M}_2 yields slightly higher revenue than the empirical-revenue-maximizing auction λ_1 over \mathcal{M}_1 , but λ_2 assigns nonzero boosts to significantly more allocations than λ_1 and is much more complex to describe. \mathcal{M}_2 is a richer auction class than \mathcal{M}_1 , so it always yields higher empirical revenue, but there is the risk that it overfits to the samples. *Structural revenue maximization* allows the mechanism designer to precisely choose between such auctions by quantifying the tradeoff between empirical revenue maximization and overfitting [2, 3]. Instead of choosing λ_2 by default, the mechanism designer should choose $\lambda_k, k \in \{1, 2\}$ that maximizes empirical revenue minus a regularization term $\varepsilon_{\mathcal{M}_k}(N, \delta)$. The correct regularization term is precisely the error term

$$\varepsilon_{\mathcal{M}_k}(N, \delta) = O\left(\frac{|\Gamma_k| \ln(n|\Gamma_k|)}{\varepsilon^2} \ln\left(\frac{1}{\delta}\right)\right)$$

in the generalization guarantee for Γ -boosted λ -auctions obtained by Balcan, Sandholm, and Vitercik [3], which is fine-tuned to the intrinsic complexity of the auction class. Structural revenue maximization can be especially useful to the mechanism designer when there is a limit on the number of samples he can draw (due to a run-time constraint, for example). In this case, he may run the exact same geometric algorithm given in Theorem 1.1 with the modified objective of empirical revenue minus the regularizer described above. In particular, the algorithm may be run over the entire class of λ -auctions, and the mechanism designer effectively learns the best set Γ of allocations to boost in order to guarantee high expected revenue while also generalizing well with high confidence.