Increasing Revenue in Efficient Combinatorial Auctions by Learning to Generate Artificial Competition

Maria-Florina Balcan[1], **Siddharth Prasad**[1], **Tuomas Sandholm**[1, 2]
[1]School of Computer Science, Carnegie Mellon University
[2]Strategy Robot, Inc.; Strategic Machine, Inc.; Optimized Markets, Inc.
ninamf@cs.cmu.edu, sprasad2@cs.cmu.edu, sandholm@cs.cmu.edu

Abstract

The design of multi-item, multi-bidder auctions involves a delicate balancing act of economic objectives, bidder incentives, and real-world complexities. *Efficient* auctions, that is, auctions that allocate items to maximize total bidder value, are practically desirable since they promote the most economically beneficial use of resources. Arguably the biggest drawback of efficient auctions, however, is their potential to generate very low revenue. In this work, we show how the auction designer can inject *competition* into the auction to boost revenue while striving to maintain efficiency. First, we invent a new auction family that enables the auction designer to specify competition in a precise, expressive, and interpretable way. We then introduce a new model of bidder behavior and individual rationality to understand how bidders act when prices are too competitive. Next, under our bidder behavior model, we use our new competitive auction class to derive the globally revenue-optimal efficient auction under two different knowledge models for the auction designer: knowledge of full bidder value distributions and knowledge of bidder value quantiles. Finally, we study a third knowledge model for the auction designer: knowledge of historical bidder valuation data. In this setting we present sample and computationally efficient learning algorithms that find high-revenue probably-efficient competitive auctions from bidder data. Our learning algorithms are *instance adaptive* and can be run in parallel across bidders, unlike most prior approaches to data-driven auction design.

1 Introduction

The design of combinatorial auctions is an intricate problem spanning theory and practice. We study the design of *competition* in *efficient* (welfare-maximizing) combinatorial auctions. Efficiency in large high-stakes auctions is important. For example, Cramton (2013) advises countries to focus on efficiency in spectrum auctions since that will yield the best and most competitive use of the allocated resources, leading to higher long-term revenues and an overall better economic state. The classic efficient auction for multi-item settings is the Vickrey (1961)-Clarke (1971)-Groves (1973) (VCG) auction, but VCG is rarely used in practice due to its unacceptably low revenue. We introduce a new class of auctions that augment VCG prices with auctioneer-specified levels of *competition*. We show that this new auction class offers the flexibility and expressive power to meaningfully boost revenue under three different auctioneer knowledge models: (i) knowledge of full bidder valuation distributions, (ii) knowledge of bidder valuation quantiles, and (iii) knowledge of historical bidder valuation data.

Our primary research question is: *how can the auction designer use additional knowledge to boost revenue via enhanced competition while striving to run an efficient auction?* If efficiency, *incentive compatibility (IC)* (each bidder is weakly best off bidding her true values no matter how other bidders act), and *individual rationality (IR)* (each bidder is motivated to participate in the auction via a guarantee of non-negative utility) are constraints of the auction design, the *weakest-competitor VCG (WCVCG)* auction of Balcan, Prasad, and Sandholm (2023) is revenue optimal. So, a more competitive auction that implements the efficient allocation and attempts to boost revenues beyond WCVCG necessarily runs the risk of determining that a bidder should pay more than her winning bid price. Such a bidder could respond in one of two ways to that situation. She could decline the offer, which would force the auctioneer to keep her winning items unsold and result in an economically inefficient allocation. But if the overcharge is not too significant, she might accept the offer—violating her individual rationality constraint—leading to the efficient allocation to be realized. In both cases an economically desirable aspect of the auction design is compromised: either the auctioneer settles for a less-than-efficient allocation, or the auctioneer accepts that a bidder was overcharged (potentially eroding bidder trust and opening the door for further unwanted negotiation). We model this behavior, design new kinds of *competitive* auctions that are sensitive to this behavior, and show how those auctions can increase revenue without violating individual rationality nor efficiency too frequently.

Summary of Contributions and Related Work

Competitive VCG auctions We introduce a new family of auctions, $\boldsymbol{f}$-VCG auctions, that gives the auction designer the expressive ability to specify precisely, for each bidder, an artificial competitor to drive competitive prices. These auctions have the feature that the auction parameters for a bidder—her *competitor*—can depend on the revealed bids of all other bidders.

Bidder behavior and individual rationality We introduce a model of bidders who are amenable to being overcharged past their winning bid price: $a(p, \kappa)$ is the probability that a bidder who bid p for a particular bundle accepts a counteroffer for the same bundle at a price of $p + \kappa > p$. For example, a television company that bid \$10 million for broadcasting rights might be willing to pay an extra \$10000 to satisfy the competitive requirements of the auction and win the rights instead of dropping out altogether. In light of this bidder model we pose a weaker—but arguably more sensible from the auctioneer's perspective—individual rationality requirement: informally, an auction is (π, κ)-IR for a bidder if (i) the set of bidder types that the auction overcharges has probability mass $\leq 1 - \pi$ and (ii) no bidder type is ever overcharged by more than κ.

An alternate widely-studied relaxation of IR is *Bayesian-IR (B-IR)*, which demands that a bidder's utility is non-negative only in expectation over the other bidders' values. The revenue-optimal auction subject to efficiency, IC, and B-IR is the Bayesian weakest-competitor VCG auction of Krishna and Perry (1998). We argue that our notion of (π, κ)-IR has several advantages over B-IR as an auction-design desideratum for at least the following reasons. First, the decision of whether or not to participate is made significantly easier for the bidders. A B-IR auction requires a bidder to understand the value distributions of other bidders, and that understanding should match the auctioneer's own understanding—a strong common knowledge assumption. In contrast, a (π, κ)-IR auction only requires bidders to reason about whether they are willing to accept an overcharge by $\$\kappa$ and thus provides bidders a greater degree of transparency. Second, a (π, κ)-IR auction is more favorable to risk-averse bidders than a B-IR auction which can lead to high overcharges (even if with low probability). B-IR auctions can indeed result in arbitrarily high overcharges (this is the case with the famous B-IR auction of Crémer and McLean (1988); see Bikhchandani (2010)) while (π, κ)-IR auctions have an explicit cap κ on overcharge. Third, B-IR auctions can overcharge bidders with much higher frequency than (π, κ)-IR auctions which have an explicit cap $1 - \pi$ on overcharge frequency (Example 3.5). Fourth, (π, κ)-IR is a flexible enough participation model to capture forms of auctioneer knowledge other than a full value distribution. In Section 3 we study a knowledge model involving quantiles. Here, the appropriate participation constraint is a "robust" (π, κ)-IR constraint. B-IR, on the other hand, is incompatible with the quantile knowledge model.

Revenue-optimal efficient auctions When counteroffers are restricted to be close enough to the bid price so that bidders accept the overcharge, we derive the revenue-optimal auction subject to efficiency, IC, and (π, κ)-IR. The revenue-optimal auction belongs to our new family of $\boldsymbol{f}$-VCG auctions. We study two auctioneer knowledge models: full bidder value distributions and bidder value quantiles. We define the appropriate notion of (π, κ)-IR and derive the revenue-optimal efficient auction for both knowledge models.

Sample and computationally efficient learning The third auctioneer knowledge model we study is sample access to historical bidder data. We derive a general learning framework to find revenue-maximizing $\boldsymbol{f}$-VCG auctions when bidder behavior is prescribed by their overcharge acceptance probability. When overcharges are sufficiently small such that efficiency can be ensured, our learning algorithms output nearly globally revenue-optimal efficient auctions subject to (π, κ)-IR. We then show how to learn high-revenue, probably-efficient $\boldsymbol{f}$-VCG auctions subject to ex-post IR when bidders *never accept overcharges* (the standard bidder assumption in auction design). In both of these important settings we show how our learning algorithms can be efficiently implemented with a winner determination oracle.

An important and unique feature of our learning framework for competition is that the algorithms are *instance adaptive* and parallelize across bidders. In all prior work, the auction parameter optimization is done based on the training data before the test instance is drawn. In our approach, the auction parameters for a particular bidder are chosen *based on the test-time revealed bids of all other bidders*, and parameter optimization across bidders can be done in parallel.

Related work The idea of a weakest competitor in VCG auctions originates from Krishna and Perry (1998), but even earlier works considered "worst-off" types first in the context of bilateral trade by Myerson and Satterthwaite (1983) and subsequently in more general trading mechanisms by Cramton, Gibbons, and Klemperer (1987). Weakest-competitor VCG had not been further studied after Krishna and Perry (1998) until Balcan, Prasad, and Sandholm (2023), who devised a prior-free version of weakest-competitor VCG and explicitly quantified welfare and revenue in terms of the quality of the auctioneer's information about the bidders. Their model of information was distribution-free (similar models have also been studied in different contexts by Hyafil and Boutilier (2004) and Chiesa, Micali, and Zhu (2015)), and their auctions were not efficient. Our goal in this paper is to use additional *distributional knowledge*—via a full prior, quantiles, or data—to boost prices while striving to maintain efficiency.

In Section 3 bidder types can be arbitrarily correlated, and the revenue-optimal choice of competitor for each bidder depends heavily on the revealed types of all other bidders. The interdependent values model (Milgrom and Weber 1982) is thematically similar in that one bidder's private value can be influenced by the others. In our setting bidders themselves have no inherent uncertainty about their private values—it is the auctioneer who can refine his knowledge about a bidder after seeing the revealed types of everybody else. In fact, efficiency might be impossible to achieve when a bidder's own understanding of her value is correlated to other bidders (Dasgupta and Maskin 2000; Jehiel et al. 2006).

Increasing competition (and thus revenue) by recruiting additional bidders has been studied starting with Bulow and Klemperer (1996). Our approach gives the auction designer the flexibility to express artificial competition. In high-stakes applications like sourcing or spectrum recruiting additional bidders might not be possible. Another class of auctions that, indirectly, boost revenue while maintaining efficiency are core-selecting auctions (Day and Raghavan 2007). However, such auctions are not IC (Goeree and Lien 2016; Othman and Sandholm 2010). Our auctions are IC and relax IR (Section 3) and sometimes efficiency (Section 4). Finally, Sandholm (2013) used *phantom bids* in sourcing auctions to optimize the decision of what items to procure through other means—a form of competition that is different from our approach since it directly affects the final allocation. Our competitive auctions solely drive prices.

Finally, our learning algorithms in Section 4 are *instance adaptive* and parallelize across bidders unlike prior approaches to data-driven auction design. We situate our work within that literature in Section 4.

2 Problem Formulation, $\boldsymbol{f}$-VCG Auctions, and Our Bidder Model

There is a set $N = \{1, \dots, n\}$ of bidders and a set $M = \{1, \dots, m\}$ of indivisible items. In a *combinatorial auction*, each bidder $i \in N$ submits bids on distinct bundles of items. Let $B_i \subseteq 2^M$ be the set of bundles bidder i bids on, so, for each $S \in B_i$, bidder i submits her value $v_i(S)$ which is the maximum amount she would be willing to pay to win bundle S ($v_i(\emptyset) = 0$). Let $v_i = (v_i(S))_{S \in B_i} \in \mathbb{R}^{|B_i|}_{\geq 0}$ denote bidder i's submitted bids, let $\boldsymbol{v} = (v_1, \dots, v_n)$ denote the full valuation profile, and let $\boldsymbol{v}_{-i}$ denote the valuation profile excluding bidder i. We use the XOR bidding language (Sandholm 2002; Nisan 2000), under which a bid vector v_i is implicitly extended to a full valuation vector in $\mathbb{R}^{2^m}$ by $v_i(S) = \max\{0, \max_{T \in B_i : T \subseteq S} v_i(T)\}$, and the seller never allocates to bidder i a bundle that i did not explicitly bid for (here we also assume free disposal, that is, $T \subseteq S \implies v_i(T) \leq v_i(S)$). Let $\Gamma \subseteq \times_{i \in N} B_i$ denote the set of feasible allocations, that is, the set of partitions $\boldsymbol{S} = (S_1, \dots, S_n)$ of M with $S_i \in B_i$ for each i and $S_i \cap S_j = \emptyset$ for each i, j.

Type spaces Before bids are submitted, bidder i's valuation v_i (her *type*) is her own private information. What is known to the auction designer is the *type space* of the bidders, denoted by $\Theta \subseteq \times_{i \in N} \mathbb{R}^{2^m}$, and that $\boldsymbol{v} \in \Theta$. Given realized bids $\boldsymbol{v}_{-i}$ of all bidders but i, let $\Theta_i(\boldsymbol{v}_{-i}) := \{\hat{v}_i : (\hat{v}_i, \boldsymbol{v}_{-i}) \in \Theta\}$ be the type space of bidder i. In other words, after seeing the submitted bids $\boldsymbol{v}_{-i}$ of all other bidders, the auction designer knows $v_i \in \Theta_i(\boldsymbol{v}_{-i})$. Importantly, what the auction designer knows about bidder i can be influenced by the *revealed bids of all other bidders*. For example, consider a seller of a painting who does not know if the painting is a forgery or not. What the seller does know is that the bidders, who are all art collectors, have similar values for the painting (either high or low). Then, if bidders $2, \dots, n$ all bid high on the painting, the seller knows that bidder 1 also has a high value before seeing her actual bid. We use the terms type and value/valuation interchangeably to refer to a bid vector v_i.

Auction design desiderata and the VCG auction An *efficient* auction selects the *efficient (welfare-maximizing) allocation* $\boldsymbol{S}^* := \mathrm{argmax}_{\boldsymbol{S} \in \Gamma} \sum_{j \in N} v_j(S_j)$. Let $w(\boldsymbol{v}) = \max_{S \in \Gamma} \sum_{j \in N} v_j(S_j)$ and let $w(\boldsymbol{v}|N \setminus i) = \sum_{j \neq i} v_j(S_j^*)$. An auction is *(dominant-strategy) incentive compatible (IC)* if each bidder's utility (her value for her winning bundle in the auction minus her payment) is weakly maximized by reporting her true type, no matter what the other bidders report. An auction is *(ex-post) individually rational (IR)* if each bidder's utility for truthful bidding is non-negative, no matter what the other bidders report. The classic efficient auction due to Vickrey (1961), Clarke (1971), and Groves (1973) (VCG) gives bidder i her winning bundle S_i^* in the efficient allocation $\boldsymbol{S}^*$, and charges her a payment of $p_i^{\mathsf{VCG}}(\boldsymbol{v}) := w(0, \boldsymbol{v}_{-i}) - w(\boldsymbol{v}|N \setminus i)$. VCG is IC and IR.

$\boldsymbol{f}$-VCG auctions We now define our new auction family: $\boldsymbol{f}$-VCG auctions. For a tuple of functions $\boldsymbol{f} = (f_1, \dots, f_n)$, $f_i : \times_{j \neq i} \mathbb{R}^{2^m} \to \mathbb{R}^{2^m}$, the $\boldsymbol{f}$-VCG auction (1) elicits bidders' types $\boldsymbol{v} = (v_1, \dots, v_n)$, (2) selects the efficient allocation $\boldsymbol{S}^*$ achieving welfare $w(\boldsymbol{v})$, and (3) offers bidder i her winning bundle S_i^* for a price of $p_i^{\boldsymbol{f}}(\boldsymbol{v}) = w(f_i(\boldsymbol{v}_{-i}), \boldsymbol{v}_{-i}) - w(\boldsymbol{v}|N \setminus i)$. In step (3) if $v_i(S_i^*) \geq p_i^{\boldsymbol{f}}(\boldsymbol{v})$, bidder i is required to pay $p_i^{\boldsymbol{f}}(\boldsymbol{v})$ (this prevents equilibria other than truthful bidding where low bidders overbid). Otherwise if $v_i(S_i^*) < p_i^{\boldsymbol{f}}(\boldsymbol{v})$, bidder i can choose to accept the higher payment (violating her IR constraint) or exit the auction altogether (leading to an inefficient allocation with S_i^* unsold). All $\boldsymbol{f}$-VCG auctions are IC since they are Groves mechanisms (that is, the *pivot* term $w(f_i(\boldsymbol{v}_{-i}), \boldsymbol{v}_{-i})$ has no dependence on bidder i's revealed type), and have the natural interpretation of $f_i(\boldsymbol{v}_{-i})$ outputting an artificial *competitor* for bidder i. The $\boldsymbol{f}$-auctions $f_i = 0$, $f_i = \mathrm{argmin}_{\widetilde{v}_i \in \Theta_i} \mathbb{E}_{\boldsymbol{v}_{-i}}[w(\widetilde{v}_i, \boldsymbol{v}_{-i})]$, and $f_i(\boldsymbol{v}_{-i}) = \mathrm{argmin}_{\widetilde{v}_i \in \Theta_i(\boldsymbol{v}_{-i})} w(\widetilde{v}_i, \boldsymbol{v}_{-i})$ are vanilla VCG, Bayesian WCVCG (Krishna and Perry 1998), and WCVCG (Balcan, Prasad, and Sandholm 2023), respectively. Let $p_i^{\widetilde{v}_i} = w(\widetilde{v}_i, \boldsymbol{v}_{-i}) - w(\boldsymbol{v}|N \setminus i)$ be the price when the competitor $\widetilde{v}_i$ is directly specified.

Overcharges, competition, bidder behavior An $\boldsymbol{f}$-VCG auction risks incurring an *overcharge* of $o_i^{\boldsymbol{f}}(\boldsymbol{v}) := p_i^{\boldsymbol{f}}(\boldsymbol{v}) - v_i(S_i^*) > 0$ for bidder i. Let $a(p, \kappa) \in [0, 1]$ be the probability that a bidder who wins bundle S_i^* with bid price $p = v_i(S_i^*)$ accepts a counteroffer for the same bundle at price $p + \kappa$. We say bidder i is *overcharged* by κ if $o_i^{\boldsymbol{f}}(\boldsymbol{v}) = \kappa > 0$, regardless of whether she accepts or not. We have $o_i^{\boldsymbol{f}}(\boldsymbol{v}) = w(f(\boldsymbol{v}_{-i}), \boldsymbol{v}_{-i}) - w(v_i, \boldsymbol{v}_{-i})$, so bidder i is overcharged if and only if $w(v_i, \boldsymbol{v}_{-i}) < w(f(\boldsymbol{v}_{-i}), \boldsymbol{v}_{-i})$, that is, she is not competitive enough. Let $\mathsf{pay}_i^{\boldsymbol{f}}(\boldsymbol{v}) = p_i^{\boldsymbol{f}}(\boldsymbol{v})(\mathbf{1}[o_i^{\boldsymbol{f}}(\boldsymbol{v}) \leq 0] + a(v_i(S_i^*), o_i^{\boldsymbol{f}}(\boldsymbol{v}))\mathbf{1}[o_i^{\boldsymbol{f}}(\boldsymbol{v}) > 0])$ be bidder i's expected payment in the $\boldsymbol{f}$-VCG auction. Let $o_i^{\widetilde{v}_i}(\boldsymbol{v}) = p_i^{\widetilde{v}_i}(\boldsymbol{v}) - v_i(S_i^*) = w(\widetilde{v}_i, \boldsymbol{v}_{-i}) - w(v_i, \boldsymbol{v}_{-i})$ and $\mathsf{pay}_i^{\widetilde{v}_i}(\boldsymbol{v}) = p_i^{\widetilde{v}_i}(\boldsymbol{v})(\mathbf{1}[o_i^{\widetilde{v}_i}(\boldsymbol{v}) \leq 0] + a(v_i(S_i^*), o_i^{\widetilde{v}_i}(\boldsymbol{v}))\mathbf{1}[o_i^{\widetilde{v}_i}(\boldsymbol{v}) > 0])$ be the overcharge and expected payment, respectively, when the competitor $\widetilde{v}_i$ is directly specified. Finally, let $C(\widetilde{v}_i; \boldsymbol{v}_{-i}) = \{v_i \in \Theta_i(\boldsymbol{v}_{-i}) : w(v_i, \boldsymbol{v}_{-i}) \geq w(\widetilde{v}_i, \boldsymbol{v}_{-i})\}$ be the set of types competitive with $\widetilde{v}_i$ given $\boldsymbol{v}_{-i}$.

3 Revenue-Optimal Efficient Auctions

We study two sources of additional bidder information available to the auction designer: knowledge of a full value distribution and knowledge of value quantiles consistent with an unknown value distribution. In this section we assume overcharges are small enough to always be accepted. This allows us to guarantee efficiency of our auctions and derive revenue-optimal efficient auctions subject to relaxed IR.

Definition 3.1. Fix $\boldsymbol{v}_{-i}$. We say κ is an *acceptable* overcharge for bidder i if $a(v_i(S_i), \kappa) = 1$ for all $v_i \in \Theta_i(\boldsymbol{v}_{-i}), S_i \in B_i, S_i \neq \emptyset$. Let v_i^κ, which we call the κ-*competitor*, denote a bidder type such that $w(v_i^\kappa, \boldsymbol{v}_{-i}) = \kappa + \min_{\widetilde{v}_i \in \Theta_i(\boldsymbol{v}_{-i})} w(\widetilde{v}_i, \boldsymbol{v}_{-i})$.

If an $\boldsymbol{f}$-VCG auction only generates acceptable overcharges, $\mathsf{pay}_i^{\boldsymbol{f}}(\boldsymbol{v}) = p_i^{\boldsymbol{f}}(\boldsymbol{v})$ and the auction is efficient.

To situate our results, we first restate the revenue optimality result of Balcan, Prasad, and Sandholm (2023) in terms of $\boldsymbol{f}$-VCG auctions. In all results, D is a Borel probability distribution on Θ.

Theorem 3.2 (Balcan, Prasad, and Sandholm (2023))**.** *Let Θ be compact and convex. Let D be any distribution on Θ. The $\boldsymbol{f}$-VCG auction $f_i(\boldsymbol{v}_{-i}) = \mathrm{argmin}_{\widetilde{v}_i \in \Theta_i(\boldsymbol{v}_{-i})} w(\widetilde{v}_i, \boldsymbol{v}_{-i})$ maximizes $\mathbb{E}_{\boldsymbol{v}\sim D}[\mathsf{pay}_i]$ for each i, and is thus revenue optimal, subject to efficiency, IC, and IR.*

Knowledge Model 1: Bidder Value Distributions

We formally define our IR relaxation, (π, κ)-IR, in the distributional knowledge model where the auction designer knows the bidder valuation distribution D over Θ.

Definition 3.3 ((π,κ)-IR, full value distribution)**.** An auction is (π, κ)-*IR* with respect to D if for each bidder i $\Pr_{\boldsymbol{v}\sim D}[o_i(\boldsymbol{v}) > 0] \leq 1 - \pi$ and $o_i(\boldsymbol{v}) \leq \kappa$ for all $\boldsymbol{v} \in \Theta$.

We now characterize the revenue-optimal auction subject to efficiency, IC, and (π, κ)-IR. It can be written as an $\boldsymbol{f}$-VCG auction. Full proofs are in App. A.

Theorem 3.4. *Let Θ be a compact and convex type space. Let D be a distribution supported on Θ and let μ be the corresponding probability measure. Let $\mu_{\boldsymbol{v}_{-i}}$ be the conditional measure over $v_i \in \Theta_i(\boldsymbol{v}_{-i})$. Let $v_i^\pi \in \Theta_i(\boldsymbol{v}_{-i})$ be such that $\mu_{\boldsymbol{v}_{-i}}(C(v_i^\pi; \boldsymbol{v}_{-i})) = \pi$ and let κ be an acceptable overcharge. The $\boldsymbol{f}$-VCG auction defined by*

$$f_i(\boldsymbol{v}_{-i}) = \begin{cases} v_i^\pi & \textit{if } w(v_i^\pi, \boldsymbol{v}_{-i}) \leq w(v_i^\kappa, \boldsymbol{v}_{-i}) \\ v_i^\kappa & \textit{otherwise} \end{cases}$$

maximizes $\mathbb{E}_{\boldsymbol{v}\sim D}[\mathsf{pay}_i]$ for each i, and is thus revenue optimal, subject to efficiency, IC, and (π, κ)-IR w.r.t. D.

Proof sketch. The competitor v_i^κ provides a cap on the welfare induced by the competitor chosen by f_i since no type in $\Theta_i(\boldsymbol{v}_{-i})$ can be overcharged by more than κ. If the competitive set of types $C(v_i^\kappa; \boldsymbol{v}_{-i})$ has measure at least π, v_i^κ is the "most competitive" competitor since a competitor with higher welfare (thus increasing the $\boldsymbol{f}$-VCG auction payment) would cause a set of types with measure $> 1 - \pi$ to be overcharged. Otherwise if $C(v_i^\kappa; \boldsymbol{v}_{-i})$ has measure less than π, we must pick a weaker competitor to reduce the probability of overcharge. That weaker competitor is precisely the v_i^π such that $C(v_i^\pi; \boldsymbol{v}_{-i})$ has measure π (which we show always exists). An application of revenue equivalence (Vohra 2011, Theorem 4.3.1) proves global revenue optimality among all efficient, IC, and (π, κ)-IR auctions. $\square$

We show via an example that B-IR auctions, specifically the Bayesian weakest-competitor auction of Krishna and Perry (1998), can overcharge with high frequency, giving further credence to our approach of optimal efficient auction design subject to an overcharge frequency cap.

Example 3.5. Consider an auction with two items A and B and two bidders $i \in \{1, 2\}$. The type space is $\Theta = \Theta_1 \times \Theta_2$ with $\Theta_i = \{(v_i(A), v_i(B)) \in \mathbb{R}^2_{\geq 0} : v_i(A) + v_i(B) = 1\}$ for both bidders (so, implicitly, $v_i(AB) = 0$ for both bidders). Suppose both bidders' valuations are distributed uniformly and independently over the type space. The Bayesian weakest competitors prescribed by Krishna and Perry (1998) are chosen before true values are revealed. The Bayesian weakest competitor for bidder 1 (and identically for bidder 2) is the valuation $\widetilde{v}_1 = (\widetilde{v}_1(A), \widetilde{v}_1(B))$ that minimizes $\mathbb{E}_{v_2}[w(\widetilde{v}_1, v_2)]$, which is $\widetilde{v}_1 = (1/2, 1/2)$ (the calculation showing this is in App. A). Suppose now that the realized type of bidder 2 is $(v_2(A), v_2(B)) = (1, 0)$, so bidder 2 wins item A and bidder 1 wins item B. According to the Bayesian weakest competitor, bidder 1 is charged $(1 + 1/2) - 1 = 1/2$, so whenever $v_1(B) < 1/2$, bidder 1 is overcharged. So, there is a 50% probability that bidder 1 is overcharged.

Looking to Theorem 3.4, the auction designer chooses the competitor v_i^π for bidder 1 after having seen bidder 2's revealed type of $(1, 0)$. For an overcharge probability of $1 - \pi$, that weakest competitor is $v_i^\pi = (1 - \pi, \pi)$. The Bayesian weakest competitor is $v_i^{1/2}$ which induces an impractically high overcharge rate of 50%.

Knowledge Model 2: Bidder Value Quantiles

We now study a knowledge model where the auctioneer has less knowledge than a full bidder value distribution. In the quantile knowledge model, we assume some underlying *unknown* value distribution, but the auction designer knows quantiles corresponding to the distribution. Formally, for each bidder i, the auctioneer possesses a sequence of sets (that can depend on the revealed types of the other bidders) $\{\Theta_i^\pi(\boldsymbol{v}_{-i})\}_{0<\pi\leq 1}$ with $\Theta_i^\pi(\boldsymbol{v}_{-i}) \supseteq \Theta_i^{\pi'}(\boldsymbol{v}_{-i})$ for any $\pi \geq \pi'$ and $\Theta_i^1(\boldsymbol{v}_{-i}) = \Theta_i(\boldsymbol{v}_{-i})$. This sequence of *quantiles* represents the knowledge that $v_i \in \Theta_i^\pi(\boldsymbol{v}_{-i})$ with probability π given the bid profile $\boldsymbol{v}_{-i}$ of all other bidders. A distribution D over Θ is *consistent* with the quantiles if $\Pr_{\hat{\boldsymbol{v}}\sim D}[\hat{v}_i \in \Theta_i^\pi(\boldsymbol{v}_{-i}) | \hat{\boldsymbol{v}}_{-i} = \boldsymbol{v}_{-i}] = \pi$. The notion of (π, κ)-IR in the quantile knowledge model is a robust version of the distributional version.

Definition 3.6 ((π,κ)-IR, quantiles)**.** An auction is (π, κ)-*IR* with respect to $\{\Theta_i^\pi\}$ if for each bidder i $\sup_{\hat{D} \text{ consistent with } \{\Theta_i^\pi\}} \Pr_{\boldsymbol{v}\sim\hat{D}}[o_i(\boldsymbol{v}) > 0] \leq 1 - \pi$ and $o_i(\boldsymbol{v}) \leq \kappa$ for all $\boldsymbol{v} \in \Theta$.

Theorem 3.7. *Let Θ be a compact and convex type space. Let $\{\Theta_i^\pi(\boldsymbol{v}_{-i})\}$ be a sequence of quantiles such that (i) the set-valued function $\pi \mapsto \Theta_i^\pi(\boldsymbol{v}_{-i})$ is continuous and (ii) the map $\pi \mapsto \min_{\widetilde{v}_i \in \Theta_i^\pi(\boldsymbol{v}_{-i})} w(\widetilde{v}_i, \boldsymbol{v}_{-i})$ is decreasing in π. Let D be any distribution supported on Θ consistent with $\{\Theta_i^\pi(\boldsymbol{v}_{-i})\}$. Let $v_i^\pi = \mathrm{argmin}_{\widetilde{v}_i \in \Theta_i^\pi} w(\widetilde{v}_i, \boldsymbol{v}_{-i})$ be the weakest competitor in quantile $\Theta_i^\pi(\boldsymbol{v}_{-i})$ and let κ be an ac-*

ceptable overcharge. The $\boldsymbol{f}$-VCG auction defined by

$$f_i(\boldsymbol{v}_{-i}) = \begin{cases} v_i^\pi & \text{if } w(v_i^\pi, \boldsymbol{v}_{-i}) \leq w(v_i^\kappa, \boldsymbol{v}_{-i}) \\ v_i^\kappa & \text{otherwise} \end{cases}$$

maximizes $\mathbb{E}_{\boldsymbol{v}\sim D}[\mathsf{pay}_i]$ *for each* i*, and is thus revenue optimal, subject to efficiency, IC, and* (π, κ)*-IR w.r.t.* $\{\Theta^\pi\}$*.*

Proof sketch. We construct a worst-case distribution $\hat{D}$ (let $\hat{\mu}$ be the corresponding probability measure) that achieves the supremum in the definition of (π, κ)-IR for any $\boldsymbol{f}$-VCG auction. Then, as in the proof of Theorem 3.4, if $\hat{\mu}(C(v_i^\kappa; \boldsymbol{v}_{-i})) \geq \pi$, v_i^κ is the optimal competitor since it satisfies (π, κ)-IR and we cannot overcharge by more than κ. Otherwise if $\hat{\mu}(C(v_i^\kappa; \boldsymbol{v}_{-i})) < \pi$, that is, the overcharge probability is $> 1 - \pi$, we must pick a weaker competitor to reduce the probability of overcharge. That competitor is the v_i^π such that $\hat{\mu}(C(v_i^\pi; \boldsymbol{v}_{-i})) = \pi$, which, based on our construction, is precisely the weakest competitor v_i^π that minimizes $w(v_i^\pi, \boldsymbol{v}_{-i})$ over $v_i^\pi \in \Theta_i^\pi(\boldsymbol{v}_{-i})$. To prove global optimality we apply revenue equivalence as in Theorem 3.4.

The construction of the worst-case measure $\hat{\mu}$ is simple: it is supported on a set of weakest competitors of the form $\{v_i^\pi : v_i^\pi = \operatorname{argmin}_{\widetilde{v}_i \in \Theta_i^\pi(\boldsymbol{v}_{-i})} w(\widetilde{v}_i, \boldsymbol{v}_{-i}), \pi \in (0, 1]\}$ and is defined to be consistent with the quantiles as $\hat{\mu}(\{v_i^\pi : \pi \in (\pi_1, \pi_2)\}) = \pi_2 - \pi_1$ for all $0 \leq \pi_1 < \pi_2 \leq 1$. □

Let us emphasize that (for both knowledge models) in a $(99\%, \kappa)$-IR auction, only 1% of bidder types ever have to deal with issues of overcharge and participation. 99% of the time the auction is perfectly efficient, IC, IR, and enjoys improved revenues. The auctioneer sets π and κ to strike a balance between risk of overcharging weak bidders and enjoying increased revenue from the large majority of bidders.

We derived the globally revenue optimal efficient auction for *acceptable* overcharges. Acceptability ensured that a (π, κ)-IR auction remained efficient. Otherwise it is unlikely that a concise global revenue optimality guarantee exists since, without an efficiency constraint, that would solve revenue-optimal multi-item auction design—a major open question—as a special case. In Section 4 we use a data-driven approach to design $\boldsymbol{f}$-VCG auctions that are nearly revenue-optimal for the class of $\boldsymbol{f}$-VCG auctions (but not globally revenue optimal) for general overcharges. Before that, we discuss how our results generalize beyond auctions.

Beyond Auctions: General Mechanism Design

Our results so far are not specific to combinatorial auctions and hold in a more general multidimensional mechanism design setting as in Balcan, Prasad, and Sandholm (2023). In that setting, Γ is a finite set of outcomes and an agent's type is a vector $v_i \in \mathbb{R}^\Gamma$ indexing her value for each outcome. The chief issue that must be addressed when applying our methodology to other settings is non-participation due to overcharge. What does non-participation mean, and what are its consequences, in the mechanism design setting of interest? In auctions, a non-participating agent receives no items. In other settings, for example public projects where the final outcome involves a resource shared by agents, non-participation might not be as naturally implementable.

4 Learning to Generate Competition

In the previous section we studied two different knowledge models for the auction designer: knowledge of the bidders' value distributions and knowledge of quantiles consistent with the bidders' value distributions. In practice, access to an exact prior is unrealistic, and fine-grained knowledge of quantiles as in the continuity requirement in Theorem 3.7 might be impractical. In this section we study a third, realistic, knowledge model: access to historical bidder data.

First we establish the formal setting. Our setup mirrors how combinatorial auctions are run in practice. We then prove our main learning guarantees for independently distributed bidder values (this is the standard assumption in mechanism design; we discuss challenges to extending our approach to correlated bidders) and provide learning algorithms. We then study the computational complexity of the learning algorithms. Throughout, we situate our results within the broader context of data-driven auction design.

Bidder valuations In practice a full valuation vector cannot be communicated due to its exponential length. Instead, the auction designer alleviates this issue by placing one of two restrictions on bidder valuations: (i) bidder i is restricted to submit bids on a set $B_i \subseteq 2^M$ of predetermined bundles or (ii) bidder i is allowed to submit bids on at most b bundles of her choice. Let $\mathsf{supp}(v_i) = \{S \subseteq M : v_i(S) > 0\}$ denote the *supported bids* of a valuation vector. We refer to valuation functions supported on B_i as B_i*-valuations* and valuation functions with support size $\leq b$ as b*-valuations*. In this section, for simplicity, we assume that bidder i submits a B_i-valuation function where B_i is set by the auction designer (we handle b-valuations in App. B). This is a practical requirement in combinatorial auctions to alleviate communication costs and the computational cost of winner determination (*e.g.,* spectrum auctions in the UK and Canada employed the XOR language with explicit bid limits of 4000 and 500, respectively (Ausubel and Baranov 2017)).

Data-driven auction design The auction designer in our setting has access to K independently and identically distributed (IID) samples $V = \{\boldsymbol{v}^{(1)}, \ldots, \boldsymbol{v}^{(K)}\}$ drawn from an unknown distribution D supported on Θ. We assume bidders' type spaces and type distributions are independent, that is, $\Theta = \Theta_1 \times \cdots \times \Theta_n$ and $D = D_1 \times \cdots \times D_n$ have product structures. So, $\Theta_i = \Theta_i(\boldsymbol{v}_{-i})$ is independent of the revealed types of the other agents and the conditional distribution over bidder i's type given $\boldsymbol{v}_{-i}$ is just D_i. As discussed above, D_i is a distribution over B_i-valuations, that is, the type space of bidder i is of the form $\Theta_i \subseteq \{v_i \in [0, H]^{2^m} : \mathsf{supp}(v_i) = B_i\}$ where H is an upper bound on any bid.

Overcharge acceptance probability We assume that the probability of accepting an overcharge only depends on the overcharge: $a(\kappa) = a(p, \kappa)$. For example, if there are known appraisal values on the items being auctioned, it might be reasonable to assume some bid-independent probability of overcharge acceptance. This (stylized) assumption is solely for technical ease of exposition; without it our bounds would only change slightly to depend on the structure of $a(p, \kappa)$.

Learning Guarantees and Algorithms for Independent Bidder Types

Even with independent bidder types, our learning algorithms choose a competitor for bidder i that is highly dependent on $\boldsymbol{v}_{-i}$. For a dataset $V = \{\boldsymbol{v}^{(1)}, \ldots, \boldsymbol{v}^{(K)}\}$ of type profiles, define $V_i = \{v_i^{(1)}, \ldots, v_i^{(K)}\}$ to be the dataset of bidder-i types. Since bidders are independently distributed, each V_i is an IID dataset from D_i. Let $\mathsf{OPT}_i^{\boldsymbol{f}}(\pi, \kappa)$ denote the optimal payment $\mathbb{E}_{\boldsymbol{v}\sim D}[\mathsf{pay}_i]$ of any (π, κ)-IR $\boldsymbol{f}$-VCG auction and let $\mathsf{OPT}_i(\pi, \kappa)$ denote the globally optimal payment of any efficient, IC, and (π, κ)-IR mechanism (achieved by the $\boldsymbol{f}$-VCG auction of Theorem 3.4 for acceptable κ).

We now present our main learning guarantees. Let $\mathcal{F}^{\mathsf{price}}(B_i) = \{p_i^{\widetilde{v}_i} : \Theta \to [0, H] : \mathsf{supp}(\widetilde{v}_i) = B_i\}$ and $\mathcal{F}^{\mathsf{pay}}(B_i) = \{\mathsf{pay}_i^{\widetilde{v}_i} : \Theta \to [0, H] : \mathsf{supp}(\widetilde{v}_i) = B_i\}$ be the collection of price and payment functions, respectively, parameterized by B_i-competitor $\widetilde{v}_i$. We bound the intrinsic complexity as measured by *pseudodimension* of these function families in order to prove our learning guarantees. The pseudodimension of a class of functions $\mathcal{F} = \{f : \Theta \to [0, H]\}$ (defined in App. B), denoted by $\mathsf{Pdim}(\mathcal{F})$, is a standard learning-theoretic measure of complexity for real-valued functions. Full proofs from this section and learning theory background are in App. B.

Lemma 4.1. $\mathsf{Pdim}(\mathcal{F}^{\mathsf{price}}(B_i))$ *and* $\mathsf{Pdim}(\mathcal{F}^{\mathsf{pay}}(B_i))$ *are at most* $O(|B_i| \log |B_i|)$.

Proof sketch. Fix $\boldsymbol{v}$. We exhibit a decomposition of competitor space $\mathbb{R}^{B_i}$ (treating $\widetilde{v}_i$ as a tunable parameter) by $O(B_i^2)$ hyperplanes into regions such that within each region (i) whether or not v_i is competitive with (and thus not overcharged by) $\widetilde{v}_i$ is invariant and (ii) p_i and pay_i are linear as functions of $\widetilde{v}_i$. The result of Balcan et al. (2023) allows us to turn this structure into a pseudodimension bound. □

Let $\varepsilon(K, \delta) = O(H\sqrt{(|B_i| \log |B_i| + \ln(1/\delta))/K})$. The following corollary, which is a consequence of standard results from learning theory (see App. B), shows that ε controls the error between empirical payment and expected payment uniformly over all possible competitors.

Corollary 4.2. *Fix* $\boldsymbol{v}$. *With probability* $\geq 1-\delta$ *over the draw of dataset* $V = \{\boldsymbol{v}^{(1)}, \ldots, \boldsymbol{v}^{(K)}\}$, *the following quantities are at most* $\varepsilon(K, \delta/n)$ *for all* i *and all* B_i*-valuations* $\widetilde{v}_i$.

- $|\frac{1}{K}\sum_{\ell=1}^{K} p_i^{\widetilde{v}_i}(v_i^{(\ell)}, \boldsymbol{v}_{-i}) - \mathbb{E}_{v_i\sim D_i}[p_i^{\widetilde{v}_i}(v_i, \boldsymbol{v}_{-i})]|$
- $|\frac{1}{K}\sum_{\ell=1}^{K} \mathsf{pay}_i^{\widetilde{v}_i}(v_i^{(\ell)}, \boldsymbol{v}_{-i}) - \mathbb{E}_{v_i\sim D_i}[\mathsf{pay}_i^{\widetilde{v}_i}(v_i, \boldsymbol{v}_{-i})]|$
- $|\frac{|\{\ell : o_i^{\widetilde{v}_i}(v_i^{(\ell)}, \boldsymbol{v}_{-i}) > 0\}|}{K} - \Pr_{v_i\sim D_i}[o_i^{\widetilde{v}_i}(v_i, \boldsymbol{v}_{-i}) > 0]|$

The above uniform convergence bounds are, for each bidder i, over a transformed training set of the form $(v_i^1, \boldsymbol{v}_{-i}), \ldots, (v_i^K, \boldsymbol{v}_{-i})$ for each bidder i. This is a form of *instance-adaptive* learning since we use the test-time revealed bids $\boldsymbol{v}_{-i}$ to (i) define the training set for bidder i and (ii) optimize the auction parameters, namely the competitor $\widetilde{v}_i$, for bidder i (as we show in Theorems 4.3 and 4.4). This is markedly different from prior approaches to data-driven auction design, for example by Balcan, Sandholm, and Vitercik (2023) and references within, where in order for the learned auction to be IC, the auction parameters are set before the test instance is seen. Some prior work tackles the *unlimited supply* setting by learning prices "within an instance" from other bidders' revealed bids (Baliga and Vohra 2003; Balcan et al. 2005), but limited supply (our setting) is more challenging (Balcan, Prasad, and Sandholm 2021).

We now translate these generalization guarantees into concrete learning algorithms. The most general result for any overcharge acceptance function $a(\kappa)$ is Theorem B.2 in App. B. It outputs an empirical-payment-maximizing competitor subject to empirical overcharge constraints. Here, we present algorithms for two pertinent cases. The first case is for acceptable κ—here the revenue-optimal efficient (π, κ)-IR auction is given by Theorem 3.4. The second case is for bidders who *do not accept overcharges*, that is, $a(\kappa) = 0$ for all κ. Here we obtain high-revenue learned auctions that are exactly ex-post IR and probably efficient.

Acceptable overcharges: nearly revenue-optimal efficient auctions The learning algorithm defining $f_i(\boldsymbol{v}_{-i})$ outputs a competitor either from the dataset V_i or defaults to a κ-competitor v_i^{κ} based on empirical overcharge frequency.

Theorem 4.3. *Let* κ *be acceptable and let the underlying type space* Θ *be compact and convex. Given an IID dataset* $V = \{\boldsymbol{v}^{(1)}, \ldots, \boldsymbol{v}^{(K)}\}$ *define the following* $\boldsymbol{f}$*-VCG auction:*

$$f_i(\boldsymbol{v}_{-i}) = \operatorname*{argmax}_{\widetilde{v}_i \in V_i \cup \{v_i^{\kappa}\}} w(\widetilde{v}_i, \boldsymbol{v}_{-i})$$

$$\text{s.t.} \quad \frac{|\{\ell : o_i^{\widetilde{v}_i}(v_i^{(\ell)}, \boldsymbol{v}_{-i}) > 0\}|}{K} \leq 1 - \pi + \varepsilon(K, \delta/n)$$

$$w(\widetilde{v}_i, \boldsymbol{v}_{-i}) \leq w(v_i^{\kappa}, \boldsymbol{v}_{-i}).$$

The resulting auction is efficient and, with probability $\geq 1 - \delta$ *over the draw of* V, $\mathbb{E}_{\boldsymbol{v}\sim D}[\mathsf{pay}_i^{\boldsymbol{f}}(\boldsymbol{v})] \geq \mathsf{OPT}_i(\pi, \kappa) - 2\varepsilon(K, \delta/n)$ *for all* i *and is thus nearly revenue-optimal, and is* $(\pi - 2\varepsilon(K, \delta/n), \kappa)$*-IR.*

High-revenue probably-efficient $\boldsymbol{f}$-VCG auctions We apply our techniques to the setting where bidders do not accept overcharges ($a(\kappa) = 0$ for all $\kappa > 0$). In other words, bidders' IR constraints must be satisfied. In this case, the only way to increase revenue is to sacrifice efficiency. As discussed previously, this is the standard model of bidders in auction design. We learn revenue-maximizing auctions within the class of $\boldsymbol{f}$-VCG auctions subject to a efficiency constraint: let $\mathsf{OPT}_i^{\boldsymbol{f}}(\pi)$ denote the optimal payment $\mathbb{E}_{\boldsymbol{v}\sim D}[\mathsf{pay}_i^{\boldsymbol{f}}] = \mathbb{E}_{\boldsymbol{v}\sim D}[p_i^{\boldsymbol{f}} \cdot \mathbf{1}[o_i^{\boldsymbol{f}}(\boldsymbol{v}) \leq 0]]$ of any $\boldsymbol{f}$-VCG auction such that $\Pr_{\boldsymbol{v}\sim D}[o_i^{\boldsymbol{f}}(\boldsymbol{v}) < 0] < 1 - \pi$ (so π is the probability bidder i is sold her winning bundle in the efficient allocation). We no longer need to consider a κ-competitor since there are no IR violations.

Theorem 4.4. *Assume bidders do not accept overcharges. Given an IID dataset* $V = \{\boldsymbol{v}^{(1)}, \ldots, \boldsymbol{v}^{(K)}\}$ *define the following* $\boldsymbol{f}$*-VCG auction:* $f_i(\boldsymbol{v}_{-i})$ *outputs*

$$\operatorname*{argmax}_{\widetilde{v}_i \in V_i} \frac{1}{K}\sum_{\ell=1}^{K} p_i^{\widetilde{v}_i}(v_i^{(\ell)}, \boldsymbol{v}_{-i}) \mathbf{1}[o_i^{\widetilde{v}_i}(v_i^{(\ell)}, \boldsymbol{v}_{-i}) \leq 0]$$

$$\text{s.t.} \quad \frac{|\{\ell : o_i^{\widetilde{v}_i}(v_i^{(\ell)}, \boldsymbol{v}_{-i}) > 0\}|}{K} \leq 1 - \pi + \varepsilon(K, \delta/n).$$

The resulting auction is IR and, with probability $\geq 1 - \delta$ *over the draw of* V*, for all bidders* i*:* $\mathbb{E}_{\boldsymbol{v}\sim D}[\mathsf{pay}_i^{\boldsymbol{f}}(\boldsymbol{v})] \geq \mathsf{OPT}_i^{\boldsymbol{f}}(\pi) - 2\varepsilon(K, \delta/n)$ *and* i *is sold her winning bundle in the efficient allocation with probability* $\geq \pi - 2\varepsilon(K, \delta)$.

Setting $\pi = 0$ corresponds to "pure" revenue maximization subject to IC and IR within the class of $\boldsymbol{f}$-VCG auctions with no other constraints. On the other hand, $\pi = 0.99$ corresponds to the revenue-maximizing $\boldsymbol{f}$-VCG auction that retains each bidder with probability at least 99%.

Challenges posed by correlation in bidder types The assumption of independent bidder types is critical to the above empirical payment maximization algorithms; each V_i is independent and consists of IID draws of bidder i's type and therefore we can optimize over the dataset $(v_i^{(1)}, \boldsymbol{v}_{-i}), \ldots, (v_i^{(K)}, \boldsymbol{v}_{-i})$ for each bidder without introducing any correlation. Without independence, the dataset can be completely uninformative about bidder i's test-time type. To illustrate, consider the extreme scenario where the type space for bidder i implied by the test-time revealed types of all other bidders $\boldsymbol{v}_{-i}$ is completely disjoint from the samples, that is, $\Theta_i(\boldsymbol{v}_{-i}) \cap V_i = \emptyset$. Then, V_i gives the auction designer absolutely no information about the conditional distribution over v_i given $\boldsymbol{v}_{-i}$. Tackling this challenge, possibly via out-of-distribution learning (Ben-David et al. 2010), is a compelling direction for future work since most real-world settings involve correlation.

Computational Considerations

A feature of the $\boldsymbol{f}$-VCG auctions above is that the competitor $f_i(\boldsymbol{v}_{-i})$ is determined via a search over the set V_i of historical bids for i and the κ-competitor. Furthermore they are sample efficient: the number of samples required to meet a prescribed error bound ε is $O(\frac{H^2}{\varepsilon^2}(|B_i| \log |B_i| + \ln(n/\delta)))$. This is in stark contrast with other combinatorial auction formats (*e.g.,* affine maximizer auctions (Roberts 1979)) for which empirical revenue maximization requires exponentially many samples and is computationally intractable (Balcan, Sandholm, and Vitercik 2023) (an approach via hyperplane arrangements has been explored in some restricted settings (Balcan, Prasad, and Sandholm 2021, 2022)).

We determine the computational complexity of our learning algorithms given a *winner-determination oracle* that on input $\boldsymbol{v}$ outputs $w(\boldsymbol{v})$ and the efficient allocation $\boldsymbol{S}^*$. Winner determination is NP-complete but can be efficiently implemented in practice via custom search algorithms (Sandholm et al. 2005; Sandholm 2006) or by integer programming. First, assuming type spaces described by linear constraints, we show how to compute a κ-competitor.

Theorem 4.5. *Given as input* $\boldsymbol{v}_{-i}$ *and a polynomial number of linear constraints defining* $\Theta_i(\boldsymbol{v}_{-i})$*, a* κ*-competitor* v_i^κ *with* $w(v_i^\kappa, \boldsymbol{v}_{-i}) = \kappa + \min_{\widetilde{v}_i \in \Theta_i(\boldsymbol{v}_{-i})} w(\widetilde{v}_i, \boldsymbol{v}_{-i})$ *can be computed with a polynomial number of calls to a winner-determination oracle and additional polynomial run time.*

Proof sketch. Balcan, Prasad, and Sandholm (2023) give a linear program (LP) to compute *weakest competitors* (zero-competitors in our terminology). Their LP enumerates all feasible allocations Γ in its constraint set. We show that a separation oracle for that LP can be implemented with a single call to a winner determination oracle. So the weakest competitor $\widetilde{v}_i$ that minimizes $w(\widetilde{v}_i, \boldsymbol{v}_{-i})$ can be computed via the ellipsoid algorithm (Grotschel, Lovász, and Schrijver 1993). Extending to a κ-competitor is straightforward. □

To find the empirical-payment-maximizing competitor $f_i(\boldsymbol{v}_{-i})$ one needs to call the winner determination oracle to compute $w(v_i^{(\ell)}, \boldsymbol{v}_{-i})$ for each $v_i^{(\ell)} \in V_i$.

Theorem 4.6. *The competitor* $f_i(\boldsymbol{v}_{-i})$ *in Theorems 4.3 and 4.4 can be computed with polynomial calls to a winner determination oracle and additional polynomial run time.*

Finally, observe that the $\boldsymbol{f}$-VCG auction computation can be parallelized across (independent) bidders. The empirical-payment-maximization algorithm to compute $f_i(\boldsymbol{v}_{-i})$ for different bidders uses completely disjoint portions of the dataset, and is an independent computation for each bidder. This has not been the case even in modern approaches to data-driven auction design via, for example, deep learning (Dütting et al. 2019; Curry, Sandholm, and Dickerson 2023; Duan et al. 2023).

5 Conclusions and Future Research

We showed how to inject artificial competition into combinatorial auctions to accomplish greater revenue when efficiency is a constraint of the auction design. While the weakest-competitor VCG mechanism of Balcan, Prasad, and Sandholm (2023) (Krishna and Perry (1998)) poses a revenue barrier for efficient, IC, and IR (B-IR) auctions, we showed that under a relaxed participation model for bidders we can nonetheless make fruitful progress. We derived the revenue optimal auction subject to efficiency, IC, and a relaxed notion of IR that involved auctioneer-set caps on overcharge frequency and magnitude, for different auctioneer knowledge models. Our new auction class, $\boldsymbol{f}$-VCG auctions, provided a unified language of artificial competition and contained the revenue optimal auctions in all the above settings. Finally, we gave sample and computationally efficient *instance-adaptive* learning algorithms that parallelize across bidders in a data-driven auction design setting.

There are a number of important theoretical and practical extensions needed to develop a more complete landscape of competitive efficient auctions. First, extensions of and more realistic versions of our bidder participation model are needed. While our (stylized) model takes a first step towards understanding how a bidder would respond to competitive prices, a more nuanced model that ties together bidder uncertainty, rationality, and attitudes towards risk is needed. Another important direction is to understand how our methodology helps move prices closer to *the core*, an important group-fairness criterion in real-world combinatorial auctions. Finally, the broader idea of auction parameter optimization that *separates* across bidders and uses the revealed types of other bidders merits deeper investigation. Current combinatorial auctions do not have this property and a more thorough understanding of when it can be exploited might lead to new and better designs.

6 Acknowledgments

This material is based on work supported by the NSF under grants IIS-1901403, CCF-1733556, RI-2312342, and RI-1901403, the ARO under award W911NF2210266, the Vannevar Bush Faculty Fellowship ONR N00014-23-1-2876, NIH award A240108S001, and a Simons Investigator award.

References

Anthony, M.; and Bartlett, P. 1999. *Neural network learning: Theoretical foundations*, volume 9. Cambridge university press.

Ausubel, L. M.; and Baranov, O. 2017. A Practical Guide to the Combinatorial Clock Auction. *Economic Journal*, 127(605).

Balcan, M.-F.; Blum, A.; Hartline, J. D.; and Mansour, Y. 2005. Mechanism design via machine learning. In *IEEE Symposium on Foundations of Computer Science (FOCS)*. IEEE.

Balcan, M.-F.; Prasad, S.; and Sandholm, T. 2021. Learning Within an Instance for Designing High-Revenue Combinatorial Auctions. In *International Joint Conference on Artificial Intelligence (IJCAI)*.

Balcan, M.-F.; Prasad, S.; and Sandholm, T. 2022. Maximizing Revenue Under Market Shrinkage and Market Uncertainty. In *Conference on Neural Information Processing Systems (NeurIPS)*.

Balcan, M.-F.; Prasad, S.; and Sandholm, T. 2023. Bicriteria Multidimensional Mechanism Design with Side Information. In *Conference on Neural Information Processing Systems (NeurIPS)*.

Balcan, M.-F.; Sandholm, T.; and Vitercik, E. 2023. Generalization guarantees for multi-item profit maximization: Pricing, auctions, and randomized mechanisms. *Operations Research*.

Baliga, S.; and Vohra, R. 2003. Market research and market design. *Advances in theoretical Economics*, 3(1).

Ben-David, S.; Blitzer, J.; Crammer, K.; Kulesza, A.; Pereira, F.; and Vaughan, J. W. 2010. A theory of learning from different domains. *Machine learning*, 79: 151–175.

Bikhchandani, S. 2010. Information acquisition and full surplus extraction. *Journal of Economic Theory*, 145(6): 2282–2308.

Bulow, J.; and Klemperer, P. 1996. Auctions versus negotiations. *The American Economic Review*.

Chiesa, A.; Micali, S.; and Zhu, Z. A. 2015. Knightian analysis of the Vickrey mechanism. *Econometrica*, 83(5): 1727–1754.

Clarke, E. H. 1971. Multipart pricing of public goods. *Public Choice*.

Cramton, P. 2013. Spectrum auction design. *Review of industrial organization*, 42: 161–190.

Cramton, P.; Gibbons, R.; and Klemperer, P. 1987. Dissolving a partnership efficiently. *Econometrica: Journal of the Econometric Society*, 615–632.

Crémer, J.; and McLean, R. P. 1988. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica: Journal of the Econometric Society*, 1247–1257.

Curry, M.; Sandholm, T.; and Dickerson, J. 2023. Differentiable economics for randomized affine maximizer auctions. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (IJCAI)*.

Dasgupta, P.; and Maskin, E. 2000. Efficient Auctions. *Quarterly Journal of Economics*, 115: 341–388.

Day, R.; and Raghavan, S. 2007. Fair Payments for Efficient Allocations in Public Sector Combinatorial Auctions. *Management Science*, 53(9): 1389–1406.

Duan, Z.; Sun, H.; Chen, Y.; and Deng, X. 2023. A scalable neural network for DSIC affine maximizer auction design. *Advances in Neural Information Processing Systems*, 36.

Dütting, P.; Feng, Z.; Narasimhan, H.; Parkes, D.; and Ravindranath, S. S. 2019. Optimal auctions through deep learning. In *International Conference on Machine Learning (ICML)*, 1706–1715. PMLR.

Goeree, J. K.; and Lien, Y. 2016. On the impossibility of core-selecting auctions. *Theoretical economics*, 11(1): 41–52.

Grotschel, M.; Lovász, L.; and Schrijver, A. 1993. *Geometric Algorithms and Combinatorial Optimizations*. Springer-Verlag.

Groves, T. 1973. Incentives in Teams. *Econometrica*.

Hyafil, N.; and Boutilier, C. 2004. Regret minimizing equilibria and mechanisms for games with strict type uncertainty. In *Proceedings of the 20th conference on Uncertainty in Artificial Intelligence (UAI)*.

Jehiel, P.; Meyer-ter Vehn, M.; Moldovanu, B.; and Zame, W. R. 2006. The limits of ex post implementation. *Econometrica*, 74(3): 585–610.

Krishna, V.; and Perry, M. 1998. Efficient mechanism design. *Available at SSRN 64934*.

Milgrom, P.; and Weber, R. 1982. A theory of auctions and competitive bidding. *Econometrica*, 50: 1089–1122.

Myerson, R.; and Satterthwaite, M. 1983. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28: 265–281.

Nisan, N. 2000. Bidding and Allocation in Combinatorial Auctions. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, 1–12. Minneapolis, MN.

Othman, A.; and Sandholm, T. 2010. Envy Quotes and the Iterated Core-Selecting Combinatorial Auction. In *AAAI Conference on Artificial Intelligence (AAAI)*.

Roberts, K. 1979. The characterization of implementable social choice rules. In Laffont, J.-J., ed., *Aggregation and Revelation of Preferences*.

Sandholm, T. 2002. Algorithm for Optimal Winner Determination in Combinatorial Auctions. *Artificial Intelligence*, 135: 1–54.

Sandholm, T. 2006. Optimal Winner Determination Algorithms. In Cramton, P.; Shoham, Y.; and Steinberg, R., eds., *Combinatorial Auctions*, 337–368. MIT Press. Chapter 14.

Sandholm, T. 2013. Very-Large-Scale Generalized Combinatorial Multi-Attribute Auctions: Lessons from Conducting $60 Billion of Sourcing. In Neeman, Z.; Roth, A.; and Vulkan, N., eds., *Handbook of Market Design*. Oxford University Press.

Sandholm, T.; Suri, S.; Gilpin, A.; and Levine, D. 2005. CABOB: A Fast Optimal Algorithm for Winner Determination in Combinatorial Auctions. *Management Science*.

Vickrey, W. 1961. Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*.

Vohra, R. V. 2011. *Mechanism design: a linear programming approach*, volume 47. Cambridge University Press.

A Omitted Proofs From Section 3

Proof of Theorem 3.4. For $\pi \in (0,1]$ let

$$L_\pi(\boldsymbol{v}_{-i}) = \underset{L\subseteq\Theta_i(\boldsymbol{v}_{-i})}{\text{argmax}} \left\{ w(\widetilde{v}_i, \boldsymbol{v}_{-i}) : \begin{array}{l} \mu_{\boldsymbol{v}_{-i}}(L) = \pi, \\ \widetilde{v}_i = \text{argmin}_{\hat{v}_i\in L}\, w(\hat{v}_i, \boldsymbol{v}_{-i}) \end{array} \right\},$$

that is, $L_\pi(\boldsymbol{v}_{-i}) \subseteq \Theta_i(\boldsymbol{v}_{-i})$ is the set of probability mass π with the strongest weakest competitor. For a candidate weakest competitor $\hat{v}_i$, consider the set $C(\hat{v}_i; \boldsymbol{v}_{-i}) = \{v_i \in \Theta_i(\boldsymbol{v}_{-i}) : w(v_i, \boldsymbol{v}_{-i}) \geq w(\hat{v}_i, \boldsymbol{v}_{-i})\}$. $C(\hat{v}_i, \boldsymbol{v}_{-i})$ is precisely the set of types v_i in $\Theta_i(\boldsymbol{v}_{-i})$ that are *competitive with* $\hat{v}_i$, that is, types v_i that are not overcharged by $p_i^{\hat{v}_i}(\cdot, \boldsymbol{v}_{-i})$. There are three steps to the proof. First, we show there exists a type $\widetilde{v}_i$ such that $C(\widetilde{v}_i; \boldsymbol{v}_{-i})$ has measure π. Next, we show that $L_\pi(\boldsymbol{v}_{-i}) = C(\widetilde{v}_i; \boldsymbol{v}_{-i})$, which explicitly characterizes $L_\pi(\boldsymbol{v}_{-i})$ in terms of competitive types (this is an alternate characterization to the one in the theorem statement in the main body of the paper that was solely based on competitive sets). Finally, an application of revenue equivalence in the style of Krishna and Perry (1998); Balcan, Prasad, and Sandholm (2023) allows us to establish payment optimality.

Let $\underline{v}_i = \text{argmin}_{\hat{v}_i\in\Theta_i(\boldsymbol{v}_{-i})}\, w(\hat{v}_i, \boldsymbol{v}_{-i})$ and $\overline{v}_i = \text{argmax}_{\hat{v}_i\in\Theta_i(\boldsymbol{v}_{-i})}\, w(\hat{v}_i, \boldsymbol{v}_{-i})$ be the weakest and strongest competitors in $\Theta_i(\boldsymbol{v}_{-i})$, respectively (both exist due to compactness of $\Theta_i(\boldsymbol{v}_{-i})$). We have $C(\underline{v}_i; \boldsymbol{v}_{-i}) = \Theta_i(\boldsymbol{v}_{-i})$ so $\mu_{\boldsymbol{v}_{-i}}(C(\underline{v}_i; \boldsymbol{v}_{-i})) = 1$. We now argue that $\mu_{\boldsymbol{v}_{-i}}(C(\overline{v}_i; \boldsymbol{v}_{-i})) = 0$. For $S_i \in B_i$ let S_{-i} be the allocation restricted to $N \setminus i$ that maximizes welfare subject to the constraint that bidder i wins S_i. We have

$$\begin{aligned} C(\overline{v}_i; \boldsymbol{v}_{-i}) &= \{v_i : w(v_i, \boldsymbol{v}_{-i}) = w(\overline{v}_i, \boldsymbol{v}_{-i})\} \\ &= \bigcup_{S_i^*\in B_i} \left\{ v_i : \begin{array}{l} v_i(S_i^*) + \sum_{j\neq i} v_j(S_j^*) \geq v_i(S_i') + \sum_{j\neq i} v_j(S_j')\ \forall\, S_i' \in B_i \setminus S_i^*, \\ v_i(S_i^*) + \sum_{j\neq i} v_j(S_j^*) = w(\overline{v}_i, \boldsymbol{v}_{-i}) \end{array} \right\} \end{aligned}$$

where each set in the (finite) union is of measure zero since the second constraint demands the zero probability event that $v_i(S_i^*)$ take on the particular fixed value of $w(\overline{v}_i, \boldsymbol{v}_{-i}) - \sum_{j\neq i} v_j(S_j^*)$. So $C(\overline{v}_i; \boldsymbol{v}_{-i})$ is itself of measure zero. As $\Theta_i(\boldsymbol{v}_{-i})$ is convex (and thus connected), continuity of $\mu_{\boldsymbol{v}_{-i}}(C(\cdot; \boldsymbol{v}_{-i}))$ and the intermediate value theorem imply the existence of $\widetilde{v}_i$ with $\mu_{\boldsymbol{v}_{-i}}(C(\widetilde{v}_i; \boldsymbol{v}_{-i})) = \pi$. Fix this $\widetilde{v}_i$. We claim that $L_\pi(\boldsymbol{v}_{-i}) = C(\widetilde{v}_i; \boldsymbol{v}_{-i})$. For the sake of contradiction, suppose that $L_\pi(\boldsymbol{v}_{-i}) = L \neq C(\widetilde{v}_i; \boldsymbol{v}_{-i})$, and let v_i' be the weakest competitor of L. So $\mu_{\boldsymbol{v}_{-i}}(L) = \pi$ and $w(v_i', \boldsymbol{v}_{-i}) > w(\widetilde{v}_i, \boldsymbol{v}_{-i})$. Since the weakest competitor v_i' of L generates strictly more welfare than $\widetilde{v}_i$, the set of types competitive with v_i' is a strict subset of the set of types competitive with $\widetilde{v}_i$, that is, $C(v_i'; \boldsymbol{v}_{-i}) \subset C(\widetilde{v}_{-i}; \boldsymbol{v}_{-i})$, which means $\mu_{\boldsymbol{v}_{-i}}(C(v_i'; \boldsymbol{v}_{-i})) < \pi$. But as $L \subseteq C(v_i'; \boldsymbol{v}_{-i})$, this is a contradiction.

We now use the revenue equivalence theorem and the above characterization to prove payment optimality. The key intuition is that $f_i(\boldsymbol{v}_{-i})$ outputs a competitor that makes either the κ-constraint or the π-constraint of (π, κ)-IR tight. Therefore, greater payment cannot be obtained without violating relaxed-IR. Formally, suppose $p_i'(\boldsymbol{v})$ is an alternate payment rule that implements the efficient allocation, is IC, and $\mathbb{E}_{\boldsymbol{v}\sim D}[p_i'(\boldsymbol{v})] > \mathbb{E}_{\boldsymbol{v}\sim D}[p_i^{\boldsymbol{f}}(\boldsymbol{v})]$. By revenue equivalence (Vohra 2011, Theorem 4.3.1), there exists a function $h_i(\boldsymbol{v}_{-i})$ such that $p_i'(\boldsymbol{v}) = p_i^{\boldsymbol{f}}(\boldsymbol{v}) + h_i(\boldsymbol{v}_{-i})$. So

$$\mathop{\mathbb{E}}_{\boldsymbol{v}\sim D}[p_i^{\boldsymbol{f}}(\boldsymbol{v}) + h_i(\boldsymbol{v}_{-i})] > \mathop{\mathbb{E}}_{\boldsymbol{v}\sim D}[p_i^{\boldsymbol{f}}(\boldsymbol{v})],$$

which means there exists a particular $\boldsymbol{v}_{-i}$ such that $h_i(\boldsymbol{v}_{-i}) > 0$. Fix this $\boldsymbol{v}_{-i}$, and let $\widetilde{v}_i = f_i(\boldsymbol{v}_{-i})$ (so $\widetilde{v}_i \in \{v_i^\kappa, v_i^\pi\}$). If $\widetilde{v}_i = v_i^\kappa$, the weakest competitor $\underline{v}_i$ of $\Theta_i(\boldsymbol{v}_{-i})$ is overcharged by exactly κ by $p_i^{\boldsymbol{f}}(\underline{v}_i, \boldsymbol{v}_{-i})$. That weakest competitor is therefore overcharged by more than κ by p_i', violating the κ-constraint in (π, κ)-IR. Else if $\widetilde{v}_i = v_i^\pi$, v_i^π's utility is zero (that is, her IR constraint is tight) under $p_i^{\boldsymbol{f}}(v_i^\pi, \boldsymbol{v}_{-i})$. Therefore, v_i^π is overcharged by $p_i'(v_i^\pi, \boldsymbol{v}_{-i})$, and more importantly by continuity of $p_i^{\boldsymbol{f}}(\cdot, \boldsymbol{v}_{-i})$ there is a sufficiently-small open ball centered at v_i^π such that all types in that ball are overcharged by p_i'. Since the measure of types less competitive than v_i^π is exactly $1-\pi$, a set of types of measure $> 1-\pi$ is overcharged by p_i'. So p_i' violates the π-constraint of (π, κ)-IR in this case. □

Bayesian weakest-competitor in Example 3.5. Let E denote the event that the weakest competitor $\widetilde{v}_1$ wins item A, so $E = \{\widetilde{v}_1(A) \geq v_2(A)\}$ and $\Pr(E) = \widetilde{v}_1(A)$. By definition of the type space, the weakest competitor wins B if and only if event E does not occur. For a given $\widetilde{v}_1$,

$$\begin{aligned} \mathop{\mathbb{E}}_{v_2}[w(\widetilde{v}_1, v_2)] &= \mathop{\mathbb{E}}_{v_2}[\widetilde{v}_1(A)\cdot \mathbf{1}[E] + \widetilde{v}_1(B)\cdot(1-\mathbf{1}[E]) + v_2(A)\cdot(1-\mathbf{1}[E]) + v_2(B)\cdot\mathbf{1}[E]] \\ &= \widetilde{v}_1(A)^2 + (1-\widetilde{v}_1(A))^2 + (1-\widetilde{v}_1(A))\cdot\frac{\widetilde{v}_1(A)+1}{2} + \frac{\widetilde{v}_1(A)^2}{2} \end{aligned}$$

which is minimized at $\widetilde{v}_1(A) = 1/2$, as claimed. □

Proof of Theorem 3.7. Suppose there exists a prior distribution D consistent with the quantiles and an alternate payment rule p_i' that generates strictly more payment than the $\boldsymbol{f}$-VCG auction defined in the theorem statement, that is, $\mathbb{E}_{\boldsymbol{v}\sim D}[p_i'(\boldsymbol{v})] > \mathbb{E}_{\boldsymbol{v}\sim D}[p_i^{\boldsymbol{f}}(\boldsymbol{v})]$, and p_i' implements the efficient allocation and is IC. We will show that there exists a distribution $\hat{D}$ consistent with the quantiles such that under p_i', $\Pr_{\boldsymbol{v}\sim\hat{D}}[o_i'(\boldsymbol{v})] > 1-\pi$. First, by revenue equivalence (Vohra 2011, Theorem 4.3.1), there exists a function $h_i(\boldsymbol{v}_{-i})$ such that $p_i'(\boldsymbol{v}) = p_i^{\boldsymbol{f}}(\boldsymbol{v}) + h_i(\boldsymbol{v}_{-i})$. So we have $\mathbb{E}_{\boldsymbol{v}\sim D}[p_i^{\boldsymbol{f}}(\boldsymbol{v}) + h_i(\boldsymbol{v}_{-i})] > \mathbb{E}_{\boldsymbol{v}\sim D}[p_i^{\boldsymbol{f}}(\boldsymbol{v})]$, which means there must exist a particular $\boldsymbol{v}_{-i}$ such that $h_i(\boldsymbol{v}_{-i}) > 0$. Fix this $\boldsymbol{v}_{-i}$. We next construct the promised worst-case measure $\hat{\mu}$.

The construction of the worst-case measure $\hat{\mu}$ is simple: it is supported on a set of weakest competitors of the form[1]

$$\left\{ v_i^{\pi} : v_i^{\pi} = \operatorname*{argmin}_{\widetilde{v}_i \in \Theta_i^{\pi}(\boldsymbol{v}_{-i})} w(\widetilde{v}_i, \boldsymbol{v}_{-i}), \pi \in (0,1] \right\}$$

and is defined to be consistent with the quantiles as $\hat{\mu}(\{v_i^{\pi} : \pi \in [\pi_1, \pi_2]\}) = \pi_2 - \pi_1$ for all $0 < \pi_1 < \pi_2 \le 1$. The key property of this distribution is that if $\pi_1 < \pi_2$, $w(v_i^{\pi_1}, \boldsymbol{v}_{-i}) > w(v_i^{\pi_2}, \boldsymbol{v}_{-i})$, that is, the weakest competitor in quantile π_2 cannot compete with the weakest competitor in quantile π_1, so the $\hat{\mu}$-measure of types that cannot compete with v_i^{π} is precisely $1-\pi$ (this shows that $\hat{\mu}$ achieves the supremum in the definition of (π,κ)-IR for any distribution D consistent with the quantiles). Another key fact is that the map $\omega : (0,1] \to \mathbb{R}_{\ge 0}$ defined by $\omega(\pi) = w(v_i^{\pi}, \boldsymbol{v}_{-i})$ is continuous (this is a consequence of Berge's Maximum Theorem and the fact that the set-valued function $\pi \mapsto \Theta_i^{\pi}$ is continuous).

Now, as in the proof of Theorem 3.4, if $\hat{\mu}(C(v_i^{\kappa}; \boldsymbol{v}_{-i})) \ge \pi$, v_i^{κ} is the optimal competitor since it satisfies (π,κ)-IR and we cannot overcharge by more than κ. Otherwise if $\hat{\mu}(C(v_i^{\kappa}; \boldsymbol{v}_{-i})) < \pi$, that is, the overcharge probability is $> 1-\pi$, we must pick a weaker competitor to reduce the probability of overcharge. That competitor is the v_i^{π} such that $\hat{\mu}(C(v_i^{\pi}; \boldsymbol{v}_{-i})) = \pi$, which is precisely the weakest competitor v_i^{π} that minimizes $w(v_i^{\pi}, \boldsymbol{v}_{-i})$ over $v_i^{\pi} \in \Theta_i^{\pi}(\boldsymbol{v}_{-i})$.

Finally, consider the alternate payment rule p_i', and let $\widetilde{v}_i = f_i(\boldsymbol{v}_{-i})$ (so $\widetilde{v}_i \in \{v_i^{\pi}, v_i^{\kappa}\}$). If $\widetilde{v}_i = v_i^{\kappa}$, the weakest competitor $\underline{v}_i = v_i^1$ of $\Theta_i(\boldsymbol{v}_{-i})$ is overcharged by exactly κ by $p_i^{\boldsymbol{f}}(\underline{v}_i, \boldsymbol{v}_{-i})$. That weakest competitor is therefore overcharged by more than κ by p_i', violating the κ-constraint in (π,κ)-IR. Else if $\widetilde{v}_i = v_i^{\pi}$, v_i^{π}'s IR constraint is tight when using payment rule $p_i^{\boldsymbol{f}}$. So p_i' overcharges v_i^{π}, and more importantly, by continuity of ω, there exists ε sufficiently small such that for all $\pi' \in (\pi-\varepsilon, \pi+\varepsilon)$, $v_i^{\pi'}$ is overcharged by p_i'. So the $\hat{\mu}$-probability mass of types being overcharged is more than $1-\pi$, so p_i' violates the π-constraint of (π,κ)-IR. □

B Learning Theory Background and Omitted Proofs From Section 4

Definition B.1 (Pseudodimension)**.** The *pseudodimension* of a class of real valued functions $\mathcal{F} = \{f : \Theta \to \mathbb{R}\}$, denoted by $\mathsf{Pdim}(\mathcal{F})$, is the largest positive integer d such that there exist d inputs $\boldsymbol{v}^1, \ldots, \boldsymbol{v}^d \in \Theta$ and d thresholds $r_1, \ldots, r_d \in \mathbb{R}$ such that

$$|\{(\text{sign}(f(\boldsymbol{v}^1) - r_1), \ldots, \text{sign}(f(\boldsymbol{v}^d) - r_d)) : f \in \mathcal{F}\}| = 2^d.$$

Uniform convergence The pseudodimension of a class of real valued functions $\mathcal{F} = \{f : \Theta \to [0,H]\}$ with bounded range controls the rate of convergence of the difference between the empirical value over an IID dataset from Θ and the expected value over the underlying distribution, uniformly over all functions in $\mathcal{F}$, for any distribution. Formally, (see, *e.g.*, Anthony and Bartlett (1999)) for any $K \in \mathbb{N}$, $\delta \in (0,1)$, and any D supported on Θ,

$$\Pr_{\boldsymbol{v}^1,\ldots,\boldsymbol{v}^K \sim D}\left(\sup_{f\in\mathcal{F}} \left| \frac{1}{K}\sum_{\ell=1}^{K} f(\boldsymbol{v}^{\ell}) - \mathop{\mathbb{E}}_{\boldsymbol{v}\sim D}[f(\boldsymbol{v})] \right| \le \varepsilon_{\mathcal{F}}(K,\delta) \right) \ge 1-\delta$$

where

$$\varepsilon_{\mathcal{F}}(K,\delta) = O\left(H\sqrt{\frac{\mathsf{Pdim}(\mathcal{F}) + \ln\frac{1}{\delta}}{K}} \right).$$

Proof of Lemma 4.1. We prove the bounds for $\mathcal{F}^{\mathsf{price}}(B_i)$ and $\mathcal{F}^{\mathsf{pay}}(B_i)$ first. Fix $\boldsymbol{v}$. For each $S_i \in B_i$, let $S_{-i} = (S_1, \ldots, S_{i-1}, S_{i+1}, \ldots, S_n)$ denote the allocation that maximizes welfare subject to the constraint that bidder i wins S_i. Over all $S_i \in B_i$, consider the set of halfspaces in $\widetilde{v}_i \in \mathbb{R}^{B_i}$:

$$\widetilde{v}_i(S_i) + \sum_{j\ne i} v_j(S_j) \ge \widetilde{v}_i(S_i') + \sum_{j\ne i} v_j(S_j') \;\forall\, S_i' \in B_i \setminus S_i$$

[1]How ties are broken in the argmin is irrelevant. What is important is continuity of the induced welfare function which is a consequence of Berge's theorem of the maximum.

where S'_{-i} denotes the welfare maximizing allocation subject to the constraint that i wins S'_i. This set of $\leq |B_i|^2$ hyperplanes corresponding to those halfspaces partitions $\mathbb{R}^{B_i}$ into regions such that within each region, the overall efficient allocation $\boldsymbol{S}$ is fixed. Thus, within each region,

$$p_i^{\widetilde{v}_i}(\boldsymbol{v}) = \widetilde{v}_i(S_i) + \sum_{j \neq i} v_j(S_j) - \sum_{j \neq i} v_j(S_j^*)$$

is linear in $\widetilde{v}_i$. An application of the main result of Balcan, Sandholm, and Vitercik (2023) proves the pseudodimension bound for $\mathcal{F}^{\mathsf{price}}(B_i)$. To understand the structure of $\mathsf{pay}_i^{\widetilde{v}_i}$, consider the same set of halfspaces as above along with the following set of B_i additional halfspaces:

$$\widetilde{v}_i(S_i) + \sum_{j \neq i} v_j(S_j) \geq \sum_{j=1}^{n} v_j(S_j^*) \; \forall \, S_i \in B_i.$$

In each region in the previous decomposition where some fixed allocation $\boldsymbol{S}$ was efficient over all $\widetilde{v}_i$ in that region, the new halfspace creates two additional regions: in one $\widetilde{v}_i$ is less competitive than v_i and so $o_i^{\widetilde{v}_i}(\boldsymbol{v}) = 0 \implies \mathsf{pay}_i^{\widetilde{v}_i}(\boldsymbol{v}) = p_i^{\widetilde{v}_i}(\boldsymbol{v})$ and in the other $\widetilde{v}_i$ is more competitive than v_i so $\mathsf{pay}_i^{\widetilde{v}_i}(\boldsymbol{v}) = p_i^{\widetilde{v}_i}(\boldsymbol{v}) \cdot a(\kappa)$. In both cases pay is linear within each region. So in total, $O(|B_i|^2)$ hyperplanes partition $\mathbb{R}^{B_i}$ into regions such that within each region, $\mathsf{pay}_i^{\widetilde{v}_i}(\boldsymbol{v})$ is linear as a function of $\widetilde{v}_i$. The pseudodimension bound follows from Balcan, Sandholm, and Vitercik (2023). □

We now present the general form of our learning algorithm for any overcharge cap κ. We will then show how the algorithm can be implemented more simply and computationally efficiently for acceptable κ.

Theorem B.2. *Given an IID dataset $V = \{\boldsymbol{v}^{(1)}, \ldots, \boldsymbol{v}^{(K)}\}$, define the following $\boldsymbol{f}$-VCG auction:*

$$f_i(\boldsymbol{v}_{-i}) = \underset{\widetilde{v}_i \in \Theta_i}{\operatorname{argmax}} \frac{1}{K} \sum_{\ell=1}^{K} \mathsf{pay}_i^{\widetilde{v}_i}(v_i^{(\ell)}, \boldsymbol{v}_{-i})$$

$$\text{s.t.} \quad \frac{|\{\ell : o_i^{\widetilde{v}_i}(v_i^{(\ell)}, \boldsymbol{v}_{-i}) > 0\}|}{K} \leq 1 - \pi + \varepsilon(K, \delta)$$

$$w(\widetilde{v}_i, \boldsymbol{v}_{-i}) \leq w(v_i^{\kappa}, \boldsymbol{v}_{-i})$$

where v_i^{κ} is a κ-competitor. Then, with probability $\geq 1 - \delta$ over the draw of V, $\mathbb{E}_{\boldsymbol{v} \sim D}[p_i^{\boldsymbol{f}}(\boldsymbol{v})] \geq \mathsf{OPT}_i^{\boldsymbol{f}}(\pi, \kappa) - 2\varepsilon(K, \delta)$ and the resulting auction is $(\pi - 2\varepsilon(K, \delta), \kappa)$-IR.

Proof of Theorem B.2. For a competitor $\widetilde{v}_i$, let

$$\widehat{o}(\widetilde{v}_i) = \frac{|\{\ell : o_i^{\widetilde{v}_i}(v_i^{(\ell)}, \boldsymbol{v}_{-i}) > 0\}|}{K} \text{ and } \overline{o}(\widetilde{v}_i) = \Pr_{v_i \sim D_i}[o_i^{\widetilde{v}_i}(v_i \boldsymbol{v}_{-i}) > 0]$$

be its empirical overcharge frequency and true overcharge frequency, respectively. Let

$$\widehat{\mathsf{pay}}(\widetilde{v}_i) = \frac{1}{K} \sum_{\ell=1}^{K} \mathsf{pay}_i^{\widetilde{v}_i}(v_i^{(\ell)}, \boldsymbol{v}_{-i}) \text{ and } \overline{\mathsf{pay}}(\widetilde{v}_i) = \mathbb{E}_{v_i \sim D_i}[\mathsf{pay}_i^{\widetilde{v}_i}(v_i, \boldsymbol{v}_{-i})]$$

be the empirical payment $\widetilde{v}_i$ generates and the true expected payment it generates, respectively. Let v_i^* be the optimal B_i-valuation output by the optimal (π, κ)-IR $\boldsymbol{f}$-VCG auction.

First, by Corollary 4.2 (all inequalities hold with high probability), we have

$$\widehat{o}(v_i^*) \leq \overline{o}(v_i^*) + \varepsilon(K, \delta) \leq 1 - \pi + \varepsilon(K, \delta).$$

Now, in the definition of our $\boldsymbol{f}$-VCG auction, $\widetilde{v}_i$ is chosen to maximize empirical payment subject to the constraint that empirical overcharge frequency is at most $1 - \pi + \varepsilon(K, \delta)$. Therefore,

$$\begin{aligned} \overline{\mathsf{pay}}(\widetilde{v}_i) &\geq \widehat{\mathsf{pay}}(\widetilde{v}_i) - \varepsilon(K, \delta) \\ &\geq \widehat{\mathsf{pay}}(v_i^*) - \varepsilon(K, \delta) \\ &\geq \overline{\mathsf{pay}}(v_i^*) - 2\varepsilon(K, \delta). \end{aligned}$$

Finally, the actual overcharge frequency of our auction is

$$\overline{o}(\widetilde{v}_i) \leq \widehat{o}(\widetilde{v}_i) + \varepsilon(K, \delta) \leq 1 - \pi + 2\varepsilon(K, \delta).$$

So our $\boldsymbol{f}$-VCG auction yields expected payment at worst $2\varepsilon(K, \delta)$ less than the optimal payment, and its overcharge frequency is at worst $2\varepsilon(K, \delta)$ greater than $1 - \pi$. All of the competitors considered above never overcharge by over κ. □

Proofs of Theorems 4.3 and 4.4. The proofs of Theorems 4.3 and 4.4 boil down to showing that the argmax in Theorem B.2 can be computed by searching over the dataset V_i (and the κ-competitor if needed for Theorem 4.3). We outline the argument in the context of Theorem 4.4 first. Here,

$$\mathsf{pay}_i^{\widetilde{v}_i}(v_i^{(\ell)}, \boldsymbol{v}_{-i}) = p_i^{\widetilde{v}_i}(v_i^{(\ell)}, \boldsymbol{v}_{-i}) \cdot \mathbf{1}[o_i^{\widetilde{v}_i}(v_i^{(\ell)}, \boldsymbol{v}_{-i}) \le 0] = \Big(w(\widetilde{v}_i, \boldsymbol{v}_{-i}) - (w(v_i^{(\ell)}, \boldsymbol{v}_{-i}) - v_i^{(\ell)}(S_i^{\ell}))\Big) \cdot \mathbf{1}[o_i^{\widetilde{v}_i}(v_i^{(\ell)}, \boldsymbol{v}_{-i}) \le 0]$$

where S_i^{ℓ} is $v_i^{(\ell)}$'s winning bundle in $w(v_i^{(\ell)}, \boldsymbol{v}_{-i})$. Let us relabel the dataset V_i in order of increasing welfare:

$$w(v_i^{(1)}, \boldsymbol{v}_{-i}) \le \cdots \le w(v_i^{(K)}, \boldsymbol{v}_{-i}).$$

Suppose

$$w(v_i^{(k)}, \boldsymbol{v}_{-i}) < w(\widetilde{v}_i, \boldsymbol{v}_{-i}) < w(v_i^{(k+1)}, \boldsymbol{v}_{-i}).$$

We claim that

$$\mathsf{pay}_i^{v_i^{(k+1)}}(v_i^{(\ell)}, \boldsymbol{v}_{-i}) \ge \mathsf{pay}_i^{\widetilde{v}_i}(v_i^{(\ell)}, \boldsymbol{v}_{-i}).$$

Indeed, if $\widetilde{v}_i$ induces an overcharge, so must $v_i^{(k+1)}$, so both payments are zero. If $\widetilde{v}_i$ does not induce an overcharge, that is $w(\widetilde{v}_i, \boldsymbol{v}_{-i}) \le w(v_i^{(\ell)}, \boldsymbol{v}_{-i})$, $v_i^{(k+1)}$ also does not induce an overcharge. Therefore, replacing $\widetilde{v}_i$ with $v_i^{(k+1)}$ increases payment, and the overcharge frequency constraint is not affected. So, the argmax in Theorem B.2 is attained by $\widetilde{v}_i \in V_i$, as desired.

The empirical payment maximization defining f_i in Theorem 4.3 has an even simpler form due to the assumption that κ is acceptable. The argmax has an even simpler form since for acceptable κ,

$$\mathsf{pay}_i^{\widetilde{v}_i}(v_i^{(\ell)}, \boldsymbol{v}_{-i}) = w(\widetilde{v}_i, \boldsymbol{v}_{-i}) - (w(v_i^{(\ell)}, \boldsymbol{v}_{-i}) - v_i^{(\ell)}(S_i^{\ell}))$$

(as $\widetilde{v}_i$ is restricted to have welfare no greater than $w(v_i^{\kappa}, \boldsymbol{v}_{-i})$). Therefore, as above (including v_i^{κ} in the welfare-sorted dataset), $v_i^{(k+1)}$ generates strictly higher empirical payment with the same overcharge frequency as $\widetilde{v}_i$. Therefore it suffices to restrict to $\widetilde{v}_i \in V_i \cup \{v_i^{\kappa}\}$. □

Proof of Theorem 4.5. The linear program for computing a weakest competitor in $\Theta_i(\boldsymbol{v}_{-i})$ based on the formulation in Balcan, Prasad, and Sandholm (2023) is:

$$\min_{\gamma, \widetilde{v}_i} \left\{ \gamma : \begin{array}{l} \widetilde{v}_i(S_i) + \sum_{j \neq i} v_j(S_j) \le \gamma \ \ \forall \boldsymbol{S} \in \Gamma, \\ \widetilde{v}_i \in \Theta_i(\boldsymbol{v}_{-i}) \end{array} \right\}.$$

It has $|B_i| + 1$ variables and enumerates the set of feasible allocations Γ in the constraints. Given a candidate solution (γ', v_i'), to find the most violated constraint we solve

$$w(v_i', \boldsymbol{v}_{-i}) = \max_{\boldsymbol{S} \in \Gamma} v_i'(S_i) + \sum_{j \neq i} v_j(S_j)$$

via a call to our winner determination oracle. Let $\boldsymbol{S}'$ be the welfare maximizing allocation on valuation profile $v_i', \boldsymbol{v}_{-i}$. If $\gamma' - w(v_i', \boldsymbol{v}_{-i}) < 0$ the constraint corresponding to $\boldsymbol{S}'$ in the original LP is violated. Else there are no violated constraints and γ', v_i' is optimal. The described routine, along with an additional scan over the constraints defining $\Theta_i(\boldsymbol{v}_{-i})$ serve as a separation oracle to be used within the ellipsoid algorithm. The ellipsoid algorithm makes a polynomial number of calls to the separation oracle and requires additional polynomial run time (Grotschel, Lovász, and Schrijver 1993).

Having found a weakest competitor $\widetilde{v}_i$, the following type v_i^{κ} defines a valid κ-competitor: $v_i^{\kappa}(\widetilde{S}_i) = \kappa + \widetilde{v}_i(\widetilde{S}_i)$ where $\widetilde{S}_i \in B_i$ is the weakest competitor's winning bundle in $w(\widetilde{v}_i, \boldsymbol{v}_{-i})$ and $v_i^{\kappa}(S_i') = \widetilde{v}_i(S_i')$ for all other bundles $S_i' \in B_i$. □

Bidders with b-valuations

A b-valuation v_i is one with $|\mathsf{supp}(v_i)| \le b$. As discussed in Section 4, many real world combinatorial auctions (*e.g.,* spectrum auctions) are run with an auctioneer-set limit on the number of bids any bidder can submit. We show in this section that similar learning guarantees can be derived for the class of $\boldsymbol{f}$-VCG auctions such that each f_i outputs a b-valuation.

Let

$$\mathcal{F}^{\mathsf{price}}(b) = \{p_i^{\widetilde{v}_i} : \Theta \to [0, H] : |\mathsf{supp}(\widetilde{v}_i)| \le b\}$$

and

$$\mathcal{F}^{\mathsf{pay}}(b) = \{\mathsf{pay}_i^{\widetilde{v}_i} : \Theta \to [0, H] : |\mathsf{supp}(\widetilde{v}_i)| \le b\}$$

be the collection of price and payment functions, respectively, parameterized by b-competitor $\widetilde{v}_i$.

Lemma B.3. $\mathsf{Pdim}(\mathcal{F}^{\mathsf{price}}(b))$ *and* $\mathsf{Pdim}(\mathcal{F}^{\mathsf{pay}}(b))$ *are at most* $O(bm \log b)$.

Proof. To prove the bounds for $\mathcal{F}^{\mathsf{price}}(b)$ and $\mathcal{F}^{\mathsf{pay}}(b)$ we use the facts that

$$\mathcal{F}^{\mathsf{price}}(b) = \bigcup_{B_i : |B_i| \le b} \mathcal{F}^{\mathsf{price}}(B_i) \text{ and } \mathcal{F}^{\mathsf{pay}}(b) = \bigcup_{B_i : |B_i| \le b} \mathcal{F}^{\mathsf{pay}}(B_i).$$

So, $\mathcal{F}^{\mathsf{price}}(b)$ (resp. $\mathcal{F}^{\mathsf{pay}}(b)$) is the union of $\binom{2^m}{b} < 2^{mb}$ function classes each with pseudodimension $O(b \log b)$ (what we proved in Lemma 4.1). Standard results[2] imply that $\mathsf{Pdim}(\mathcal{F}^{\mathsf{price}}(b))$ and $\mathsf{Pdim}(\mathcal{F}^{\mathsf{pay}}(b))$ are upper bounded by

$$O\left(\max\left(b\log b, \log\binom{2^m}{b} + b\log b \log\left(\log\binom{2^m}{b}/(b\log b)\right)\right)\right) = O(\max(b\log b, bm + b\log b\log m)) \le O(bm\log b).$$

□

Analogues of Theorems B.2 and 4.3 can then be derived for b-valuations. The main caveat here is that since the set of b-valuations is non-convex, any nontrivial type space of b-valuations will be non-convex. Hence the global revenue optimality theorem (Theorem 3.4) for acceptable κ does not hold since revenue equivalence need not hold for non-convex type spaces. In the analogue of 4.3 we will therefore only be able to compare to the revenue-optimal efficient $\boldsymbol{f}$-VCG auction such that f_i outputs a b-valuation, rather than the globally optimal efficient auction.

[2] See https://home.ttic.edu/~nati/Teaching/TTIC31120/2015/hw1.pdf for the needed result on VC dimension for unions of concept classes. That result carries over to pseudodimension due to the fact that pseudodimension of $\mathcal{F}$ is equivalent to the VC dimension of the collection of epigraphs of functions in $\mathcal{F}$.