

Bicriteria Multidimensional Mechanism Design with Side Information

Maria-Florina Balcan*

Siddharth Prasad†

Tuomas Sandholm‡

Abstract

We develop a versatile new methodology for multidimensional mechanism design that incorporates side information about agent types with the bicriteria goal of generating high social welfare and high revenue simultaneously. Side information can come from a variety of sources—examples include advice from a domain expert, predictions from a machine-learning model trained on historical agent data, or even the mechanism designer’s own gut instinct—and in practice such sources are abundant. In this paper we adopt a prior-free perspective that makes no assumptions on the correctness, accuracy, or source of the side information. First, we design a meta-mechanism that integrates input side information with an improvement of the classical VCG mechanism. The welfare, revenue, and incentive properties of our meta-mechanism are characterized by a number of novel constructions we introduce based on the notion of a *weakest competitor*, which is an agent that has the smallest impact on welfare. We then show that our meta-mechanism—when carefully instantiated—simultaneously achieves strong welfare and revenue guarantees that are parameterized by errors in the side information. When the side information is highly informative and accurate, our mechanism achieves welfare and revenue competitive with the total social surplus, and its performance decays continuously and gradually as the quality of the side information decreases. Finally, we apply our meta-mechanism to a setting where each agent’s type is determined by a constant number of parameters. Specifically, agent types lie on constant-dimensional subspaces (of the potentially high-dimensional ambient type space) that are known to the mechanism designer. We use our meta-mechanism to obtain the first known welfare and revenue guarantees in this setting.

1 Introduction

Mechanism design is a high-impact branch of economics and computer science that studies the implementation of socially desirable outcomes among strategic self-interested agents. Major real-world use cases include combinatorial auctions (*e.g.*, strategic sourcing, radio spectrum auctions), matching markets (*e.g.*, housing allocation, ridesharing), project fundraisers, and many more. The two most commonly studied objectives in mechanism design are *welfare maximization* and *revenue maximization*. In many settings, welfare maximization, or *efficiency*, is achieved by the classic Vickrey-Clarke-Groves (VCG) mechanism [18, 27, 49]. Revenue maximization is a much more elusive problem that is only understood in very special cases. The seminal work of Myerson [41] characterized the revenue-optimal mechanism for the sale of a single item in the Bayesian setting, but it is not even known how to optimally sell two items to multiple buyers. It is known that

*School of Computer Science, Carnegie Mellon University. ninamf@cs.cmu.edu

†Computer Science Department, Carnegie Mellon University. spurasad2@cs.cmu.edu

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

welfare and revenue are generally competing objectives and optimizing one can come at the great expense of the other [1, 4, 7, 22, 32].

In this paper we study how *side information* (or *predictions*) about the agents can help with *bicriteria* optimization of both welfare and revenue. Side information can come from a variety of sources that are abundantly available in practice such as predictions from a machine-learning model trained on historical agent data, advice from domain experts, or even the mechanism designer’s own gut instinct. Mechanism design approaches that exploit the proliferation of agent data have in particular witnessed a great deal of success both in theory [8, 10, 37, 40] and in practice [23, 24, 45, 51]. In contrast to the typical Bayesian approach to mechanism design that views side information through the lens of a prior distribution over agents, we adopt a prior-free perspective that makes no assumptions on the correctness, accuracy, or source of the side information. A nascent line of work (that is part of a larger agenda on augmenting algorithms with machine-learned predictions [39]) has begun to examine the challenge of exploiting predictions (of *a priori* unknown quality) when agents are self-interested, but only for fairly specific problem settings [3, 13–15, 25, 52]. We contribute to this line of work with a general side-information-dependent meta-mechanism for a wide swath of multidimensional mechanism design problems that aim for high social welfare and high revenue.

Here we provide a few examples of the forms of side information we consider in various multidimensional mechanism design scenarios. A formal description of the model is in Section 2. (1) *The owner of a new coffee shop sets prices based on the observation that most customers are willing to pay \$9 for a coffee and a croissant, and are willing to pay \$5 for either item individually.* (2) *A real-estate agent believes that a particular buyer values a high-rise condominium with a city view three times more than one on the first floor. Alternately, the seller might know for a fact that the buyer values the first property three times more than the second based on set factors such as value per square foot.* (3) *A homeowner association is raising funds for the construction of a new swimming pool within a townhome complex. Based on the fact that a particular resident has a family with children, the association estimates that this resident is likely willing to contribute at least \$300 if the pool is opened within a block of the resident’s house but only \$100 if outside a two-block radius.* These are all examples of side information available to the mechanism designer that may or may not be useful or accurate. Our methodology allows us to derive welfare and revenue guarantees under different assumptions on the veracity of the side information. We study two slightly different settings: one in which the side information can be completely bogus (Section 4), and another in which the side information is constrained to be valid (Section 5).

1.1 Our contributions

Our main contribution is a versatile meta-mechanism that integrates side information about agent types with the bicriteria goal of simultaneously optimizing welfare and revenue.

In Section 2 we formally define the components of multidimensional mechanism design with side information. The abstraction of multidimensional mechanism design is a rich language that allows our theory to apply to many real-world settings including combinatorial auctions, matching markets, project fundraisers, and more—we expand on this list of examples further in Section 2. We also present the weakest-competitor VCG mechanism introduced by Krishna and Perry [33] and prove that it is revenue-optimal among all efficient mechanisms in the prior-free setting (extending their work which was in the Bayesian setting for a fixed known prior).

In Section 3 we present our meta-mechanism for mechanism design with side information. It generalizes the mechanism of Krishna and Perry [33]. We introduce the notion of a weakest-competitor set and a weakest-competitor hull, which are constructions that are crucial to understanding the payments and incentive properties of our meta-mechanism.

In Section 4 we prove that our meta-mechanism—when carefully instantiated—achieves strong welfare and revenue guarantees that are parameterized by errors in the side information. Our mechanism works by independently expanding the input predictions, where the expansion radius for each prediction is drawn randomly from a logarithmic discretization of the diameter of the ambient type space. Our mechanism achieves the efficient welfare OPT and revenue *at least* $\Omega(\text{OPT}/\log H)$ when the side information is highly informative and accurate, where H is an upper bound on any agent’s value for any outcome. Its revenue approaches OPT if its initialization parameters are chosen wisely. Its performance decays gradually as the quality of the side information decreases (whereas naïve approaches suffer from huge discontinuous drops in performance). *Prior-free efficient welfare OPT , or total social surplus, is the strongest possible benchmark for both welfare and revenue.* Finally, we extend our methods to more general, more expressive side information languages.

In Section 5 we use our meta-mechanism to derive new results in a setting where each agent’s type is determined by a constant number of parameters. Specifically, agent types lie on constant-dimensional subspaces (of the potentially high-dimensional ambient type space) that are known to the mechanism designer. *For example, in the condominium example from the introduction, an agent’s value per square foot might completely determine her value for each property.* When each agent’s true type is known to lie in a particular k -dimensional subspace of the ambient type space, we show how to use our meta-mechanism to guarantee revenue at least $\Omega(\text{OPT}/k(\log H)^k)$ while simultaneously guaranteeing welfare at least $\text{OPT}/\log H$.

Traditionally it is known that welfare and revenue are at odds and maximizing one objective comes at the expense of the other. Our results show that side information can help mitigate this difficulty.

1.2 Related work

Side information in mechanism design. Various mechanism design settings have been studied under the assumption that some form of public side information is available. Medina and Vassilvitskii [37] study single-item (unlimited supply) single-bidder posted-price auctions with approximate bid predictions. Devanur et al. [21] study the sample complexity of (single-parameter) auctions when the mechanism designer receives a distinguishing signal for each bidder. More generally, the field of *algorithms with predictions* aims to improve the quality of classical algorithms when (potentially machine-learned) predictions about the solution are available. This is an extremely active area of research (Mitzenmacher and Vassilvitskii [39] provide a survey of work in this area). There have been recent explicit connections of the algorithms-with-predictions paradigm to mechanism design in specific settings such as strategic scheduling, facility location, online Nash social welfare maximization, and single-leg revenue management [3, 13–15, 25]. Most related to our work, Xu and Lu [52] study the design of high-revenue auctions to sell a (single copy of a) single item to multiple bidders when the mechanism designer has access to point predictions on the bidders’ values for the items. Unlike our approach which heavily exploits randomization, they focus on *deterministic* prediction-augmented modifications of a second-price auction. An important drawback of determinism is that revenue guarantees do not decay continuously as prediction quality degrades. For agents with value capped at H there is an error threshold after which, in the worst case, only a $1/H$ -fraction of revenue can be guaranteed (this is not even competitive with a vanilla second-price auction). Xu and Lu [52] prove that such a drop in revenue is unavoidable by deterministic mechanisms. Finally, our setting is distinct from, but similar to in spirit, work that uses public attributes to perform market segmentation to improve revenue [8, 11].

Welfare-revenue tradeoffs in auctions. The relationship between welfare and revenue in Bayesian

auctions has been widely studied since the seminal work of Bulow and Klemperer [16]. Daskalakis and Pierrakos [20], Hartline and Roughgarden [30] quantify welfare-revenue tradeoffs for second-price auctions with reserve prices in the single item setting. Diakonikolas et al. [22] study the Pareto frontier between efficiency and revenue in the single-item setting. They show that understanding the Pareto frontier is in general intractable, but derive polynomial-time approximation schemes to approximate the Pareto curve when there are two bidders. Anshelevich et al. [4] study welfare-revenue tradeoffs in large markets, Aggarwal et al. [2] study the efficiency of revenue-optimal mechanisms, and Abhishek and Hajek [1] study the efficiency loss of revenue-optimal mechanisms.

Constant-parameter mechanism design. Revenue-optimal mechanism design for settings where each agent's type space is of a constant dimension has been studied previously in certain specific settings. Single-parameter mechanism design is a well-studied topic dating back to the seminal work of Myerson [41], who (1) characterized the set of all truthful allocation rules and (2) derived the Bayesian optimal auction based on virtual values (a quantity that is highly dependent on knowledge of the agents' value distributions). Archer and Tardos [5] also characterize the set of allocation rules that can be implemented truthfully in the single-parameter setting, and use this to derive polynomial-time mechanisms with strong revenue approximation guarantees in various settings. Kleinberg and Yuan [32] prove revenue guarantees for a variety of single-parameter settings that depend on distributional parameters. Constrained buyers with values determined by two parameters have also been studied [35, 42].

Revenue of combinatorial auctions for limited supply. Our mechanism when agent types lie on known linear subspaces of low degree can be seen as a generalization of the well-known logarithmic revenue approximation that is achieved by a second-price auction with a random reserve price in the single-item setting [26]. Similar revenue approximations have been derived in multi-item settings for various classes of bidder valuation functions such as unit-demand [29], additive [34, 47], and subadditive [9, 17]. To the best of our knowledge, no previous techniques handle agent types on low-dimensional subspaces. Furthermore, our results are not restricted to combinatorial auctions unlike most previous research.

2 Problem formulation, example applications, and the weakest-competitor VCG mechanism

We consider a general multidimensional mechanism design setting with a finite allocation space Γ and n agents. Θ_i is the type space of agent i . Agent i 's true private type $\theta_i \in \Theta_i$ determines her value $v(\theta_i, \alpha)$ for allocation $\alpha \in \Gamma$. We will interpret Θ_i as a subset of \mathbb{R}^Γ , so $\theta_i[\alpha] = v(\theta_i, \alpha)$. We use $\boldsymbol{\theta} \in \times_{i=1}^n \Theta_i$ to denote a profile of types and $\boldsymbol{\theta}_{-i} \in \Theta_{-i} := \times_{j \neq i} \Theta_j$ to denote a profile of types excluding agent i . We now introduce our model of side information. For each agent, the mechanism designer receives a subset of the type space predicting that the subset contains the agent's true yet-to-be-revealed type. Formally, the mechanism designer receives additional information about each agent in the form of a refinement of each agent's type space, given by $\tilde{\Theta}_1 \subseteq \Theta_1, \dots, \tilde{\Theta}_n \subseteq \Theta_n$. These refinements postulate that the true type of bidder i is actually contained in $\tilde{\Theta}_i$ (though the mechanism designer does not necessarily know whether or not these predictions are valid). We refer to the $\tilde{\Theta}_i$ as *side-information sets* or simply *predictions*. To simplify exposition, we assume that prior to receiving side information the mechanism designer has no differentiating information about the agents' types, that is, $\Theta_1 = \dots = \Theta_n$. Let Θ denote this common *ambient type space*. We assume $0 \in \Theta$. This assumption is without loss of generality—all definitions and theorems in this paper hold relative to whatever the ambient type space of agent i is.

A mechanism with side information is specified by an allocation rule $\alpha(\boldsymbol{\theta}; \tilde{\Theta}_1, \dots, \tilde{\Theta}_n) \in \Gamma$ and a payment rule $p_i(\boldsymbol{\theta}; \tilde{\Theta}_1, \dots, \tilde{\Theta}_n) \in \mathbb{R}$ for each agent i . We assume agents have quasilinear utilities. A mechanism is *incentive compatible* if $\theta_i \in \operatorname{argmax}_{\theta'_i \in \Theta_i} \theta_i[\alpha(\theta'_i, \boldsymbol{\theta}_{-i}; \tilde{\Theta}_1, \dots, \tilde{\Theta}_n)] - p_i(\theta'_i, \boldsymbol{\theta}_{-i}; \tilde{\Theta}_1, \dots, \tilde{\Theta}_n)$ holds for all $i, \theta_i \in \Theta_i, \boldsymbol{\theta}_{-i} \in \Theta_{-i}, \tilde{\Theta}_1 \subseteq \Theta_i, \dots, \tilde{\Theta}_n \subseteq \Theta_n$, that is, agents are incentivized to report their true type regardless of what other agents report and regardless of the side information used by the mechanism (this definition is equivalent to the usual notion of dominant-strategy incentive compatibility and simply stipulates that side information ought to be used in an incentive compatible manner). A mechanism is *individually rational* if $\theta_i[\alpha(\theta_i, \boldsymbol{\theta}_{-i}; \tilde{\Theta}_1, \dots, \tilde{\Theta}_n)] - p_i(\theta_i, \boldsymbol{\theta}_{-i}; \tilde{\Theta}_1, \dots, \tilde{\Theta}_n) \geq 0$ holds for all $i, \theta_i, \boldsymbol{\theta}_{-i}, \tilde{\Theta}_1, \dots, \tilde{\Theta}_n$. We will analyze a variety of randomized mechanisms that randomize over incentive compatible and individually rational mechanisms. Such randomized mechanisms are thus incentive compatible and individually rational in the strongest possible sense (as supposed to weaker in-expectation incentive compatibility/individual rationality). *An important note: no assumptions are made on the veracity of $\tilde{\Theta}_i$; agent i 's misreporting space is the full ambient type space Θ_i .*

Example applications

Our model of side information within the rich language of multidimensional mechanism design allows us to capture a variety of different problem scenarios where both welfare and revenue are desired objectives. We list a few examples of different multidimensional mechanism settings along with examples of different varieties of side information sets.

- **Combinatorial auctions:** There are m indivisible items to be allocated among n agents (or to no one). The allocation space Γ is the set of $(n+1)^m$ allocations of the items and $\theta_i[\alpha]$ is agent i 's value for the bundle of items she is allocated by α . Let X and Y denote two of the items for sale. The set $\tilde{\Theta}_i = \{\theta_i : \theta_i[\{X, Y\}] \geq 9, \theta_i[\{X\}] + \theta_i[\{Y\}] \geq 10\}$ represents the prediction that agent i 's values for X and Y individually sum up to at least \$10, and her value for the bundle is at least \$9. Here, $\tilde{\Theta}_i$ is the intersection of linear constraints.
- **Matching markets:** There are m items (*e.g.*, houses) to be matched to n buyers. The allocation space Γ is the set of matchings on the bipartite graph $K_{m,n}$ and $\theta_i[\alpha]$ is buyer i 's value for the item α assigns her. Let $\alpha_1, \alpha_2, \alpha_3$ denote three matchings that match house 1, house 2, and house 3 to agent i , respectively. The set $\tilde{\Theta}_i = \{\theta_i : \theta_i[\alpha_1] = 2 \cdot \theta_i[\alpha_2] = 0.75 \cdot \theta_i[\alpha_3]\}$ represents the information that agent i values house 1 twice as much as house 2, and $3/4$ as much as house 3. Here, $\tilde{\Theta}_i$ is the linear space given by $\operatorname{span}(\langle 1, 1/2, 4/3 \rangle)$.
- **Fundraising for a common amenity:** A multi-story office building that houses several companies is opening a new cafeteria on a to-be-determined floor and is raising construction funds. The allocation space Γ is the set of floors of the building and $\theta_i[\alpha]$ is the (inverse of the) cost incurred by building-occupant i for traveling to floor α . The set $\tilde{\Theta}_i = \{\theta_i : \|\theta_i - \theta_i^*\|_p \leq k\}$ postulates that i 's true type is no more than k away from θ_i^* in ℓ_p -distance, which might be derived from an estimate of the range of floors agent i works on based on the company agent i represents. Here, $\tilde{\Theta}_i$ is given by a (potentially nonlinear) distance constraint.
- **Bidding for a shared outcome:** A package delivery service that offers multiple delivery rates (priced proportionally) needs to decide on a delivery route to serve n customers. The allocation space Γ is the set of feasible routes and $\theta_i[\alpha]$ is agent i 's value for receiving her packages after the driving delay specified by α . For $t = 0, 1, \dots$ let α_t denote an allocation that imposes a driving delay of t time units on agent i . The set $\tilde{\Theta}_i = \{\theta_i : \theta_i[\alpha_0] \geq 500, \theta_i[\alpha_{t+1}] \geq$

$f_t(\theta_i[\alpha_t]) \forall t\}$ is the prediction that agent i is willing to pay \$500 to receive her packages as soon as possible, and is at worst a time discounter determined by (potentially nonlinear) discount functions f_t . Here, the complexity of $\tilde{\Theta}_i$ is determined by the f_t .

In our results we will assume that $\Theta = [0, H]^\Gamma$ and thereby impose an upper bound H on any agent's value for any allocation. This cap H is the only problem-specific parameter in our results. In the above four bulleted examples H represents the maximum value any agent has for the grand bundle of items, any available house, the cafeteria opening on her floor, and receiving her packages with no delay, respectively.

The weakest-competitor VCG mechanism

The VCG mechanism can generally be highly suboptimal when it comes to revenue [7, 38, 48] (and conversely mechanisms that shoot for high revenue can be highly welfare suboptimal). However, if efficiency is enforced as a constraint of the mechanism design, then the *weakest-competitor VCG mechanism* introduced by Krishna and Perry [33] is in fact revenue optimal (they call it the generalized VCG mechanism). While VCG payments are based on participation externalities, the payments of the weakest-competitor VCG mechanism are based on agents being replaced by *weakest competitors* who have the smallest impact on welfare. This approach yields a strict revenue improvement over the vanilla VCG mechanism. Krishna and Perry [33] proved that the Bayesian version of the weakest-competitor VCG mechanism is revenue optimal among all efficient, incentive compatible, and individually rational mechanisms. The (prior-free) weakest-competitor VCG mechanism works as follows. Given reported types $\theta_1, \dots, \theta_n$, it uses the efficient allocation $\alpha^* = \operatorname{argmax}_{\alpha \in \Gamma} \sum_{i=1}^n \theta_i[\alpha]$. The payments are given by

$$p_i(\theta_1, \dots, \theta_n) = \min_{\tilde{\theta}_i \in \Theta_i} \left(\max_{\alpha \in \Gamma} \sum_{j \neq i} \theta_j[\alpha] + \tilde{\theta}_i[\alpha] \right) - \sum_{j \neq i} \theta_j[\alpha^*].$$

If $0 \in \Theta_i$, p_i is simply the vanilla VCG payment. Krishna and Perry [33] prove the following result in the Bayesian setting, which we reproduce in a stronger prior-free form for completeness.

Theorem 2.1. *The weakest-competitor VCG auction is incentive compatible and individually rational. Furthermore, it is the revenue maximizing mechanism among all mechanisms that are efficient, incentive compatible, and individually rational.*

Proof. Weakest-competitor VCG is incentive compatible for the same reason that VCG is incentive compatible: the minimization in the payment formula (the *pivot* term) is independent of bidder i 's reported type. Concretely, if α' is the welfare-maximizing allocation when bidder i reports θ'_i , bidder i 's utility from reporting θ'_i is $\sum_{j=1}^n \theta_j[\alpha'] - \min_{\tilde{\theta}_i \in \Theta_i} (\max_{\alpha} \sum_{j \neq i} \theta_j[\alpha] + \tilde{\theta}_i[\alpha])$, which is maximized at $\alpha' = \alpha^*$ (which proves incentive compatibility). Furthermore, for each i , $\sum_{j=1}^n \theta_j[\alpha^*] - \min_{\tilde{\theta}_i \in \Theta_i} (\max_{\alpha} \sum_{j \neq i} \theta_j[\alpha] + \tilde{\theta}_i[\alpha]) \geq \sum_{j=1}^n \theta_j[\alpha^*] - \max_{\alpha} \sum_{j=1}^n \theta_j[\alpha] = 0$, which proves individual rationality. The proof that weakest-competitor VCG is revenue optimal follows from the revenue equivalence theorem; the necessary ingredients may be found in the monograph by Vohra [50]. Let $p_i(\boldsymbol{\theta})$ be the weakest-competitor VCG payment rule, and let $p'_i(\boldsymbol{\theta})$ be any other payment rule that also implements the efficient allocation rule. By revenue equivalence, for each i , there exists $h_i(\boldsymbol{\theta}_{-i})$ such that $p'_i(\theta_i, \boldsymbol{\theta}_{-i}) = p_i(\theta_i, \boldsymbol{\theta}_{-i}) + h_i(\boldsymbol{\theta}_{-i})$. Suppose $\boldsymbol{\theta}$ is a profile of types such that p'_i generates strictly greater revenue than p_i , that is, $\sum_{i=1}^n p'_i(\boldsymbol{\theta}) > \sum_{i=1}^n p_i(\boldsymbol{\theta})$. Equivalently $\sum_{i=1}^n p_i(\theta_i, \boldsymbol{\theta}_{-i}) + h_i(\boldsymbol{\theta}_{-i}) > \sum_{i=1}^n p_i(\theta_i, \boldsymbol{\theta}_{-i})$. Thus, there exists i^* such that $h_{i^*}(\boldsymbol{\theta}_{-i^*}) > 0$. Now, let $\theta_{i^*} = \operatorname{argmin}_{\theta'_{i^*} \in \Theta_{i^*}} \max_{\alpha \in \Gamma} \sum_{j \neq i} \theta_j[\alpha] + \theta'_{i^*}[\alpha]$ be the *weakest competitor* with respect to $\boldsymbol{\theta}_{-i^*}$.

If weakest-competitor VCG is run on the type profile $(\tilde{\theta}_{i^*}, \boldsymbol{\theta}_{-i^*})$, the agent with type $\tilde{\theta}_{i^*}$ pays their value for the efficient allocation. In other words, the individual rationality constraint is binding for $\tilde{\theta}_{i^*}$. Since $h_{i^*}(\boldsymbol{\theta}_{-i^*}) > 0$, p'_i violates individual rationality, which completes the proof. \square

Our meta-mechanism (Section 3) is a generalization of the weakest-competitor VCG mechanism that uses side information sets rather than the ambient type space to determine payments. (Misreporting is not limited to side information sets.) Our meta-mechanism relaxes efficiency in order to use the side information to boost revenue. The notion of a *weakest competitor* highlighted by the weakest-competitor VCG mechanism is a key ingredient in our study of side information.

3 Weakest-competitor sets and our meta-mechanism

In this section we present our meta-mechanism for multidimensional mechanism design with side information. Our meta-mechanism generalizes the weakest-competitor VCG mechanism. We begin by introducing some new constructions based on the concept of a weakest competitor. These constructions are the key ingredients in understanding the role of side information in our meta-mechanism. Let $\theta \preceq \theta'$ if $\theta[\alpha] \leq \theta'[\alpha]$ for all $\alpha \in \Gamma$. Let $\theta \preccurlyeq \theta'$ if $\theta[\alpha] \leq \theta'[\alpha]$ for all $\alpha \in \Gamma$ and there exists $\alpha' \in \Gamma$ with $\theta[\alpha'] < \theta'[\alpha']$. Let $\theta \prec \theta'$ if $\theta[\alpha] < \theta'[\alpha]$ for all $\alpha \in \Gamma$. We assume $\Theta_i = \Theta = [0, H]^\Gamma$ for all i , that is, all agents share a common ambient type space with no up-front differentiating information.

Definition 3.1. The *extended weakest-competitor set* of a closed set $\tilde{\Theta}_i$, denoted by $\overline{WC}(\tilde{\Theta}_i)$, is the subset of all weakest competitors in $\tilde{\Theta}_i$ over all possible type profiles of the other agents. Formally,

$$\overline{WC}(\tilde{\Theta}_i) := \left\{ \underset{\tilde{\theta}_i \in \tilde{\Theta}_i}{\operatorname{argmin}} \left(\max_{\alpha \in \Gamma} \sum_{j \neq i} \theta_j[\alpha] + \tilde{\theta}_i[\alpha] \right) : \boldsymbol{\theta}_{-i} \in \Theta_{-i} \right\} \subseteq \operatorname{bd}(\tilde{\Theta}_i).$$

The *weakest-competitor set* of $\tilde{\Theta}_i$, denoted by $WC(\tilde{\Theta}_i)$, is the subset of $\overline{WC}(\tilde{\Theta}_i)$ where ties in the argmin are broken by discarding any θ' in the argmin such that there exists θ also in the argmin with $\theta \preccurlyeq \theta'$. We call members of both $\overline{WC}(\tilde{\Theta}_i)$ and $WC(\tilde{\Theta}_i)$ *weakest competitors* and say $\hat{\theta}_i$ is a *weakest competitor relative to $\boldsymbol{\theta}_{-i}$* if $\hat{\theta}_i \in \operatorname{argmin}_{\tilde{\theta}_i \in \tilde{\Theta}_i} \max_{\alpha \in \Gamma} \sum_{j \neq i} \theta_j[\alpha] + \tilde{\theta}_i[\alpha]$.

The weakest-competitor set is a natural notion of a “lower bound” set of types corresponding to a given predicted type set. From the perspective of the weakest-competitor VCG mechanism, the payment of an agent with true type in $\tilde{\Theta}_i$ only depends on $WC(\tilde{\Theta}_i)$ and not on $\tilde{\Theta}_i$. Motivated by this observation, we define the weakest-competitor hull, which can be viewed as a “weakest-competitor relaxation” of a given set $\tilde{\Theta}_i$.

Definition 3.2. The *weakest-competitor hull* of $\tilde{\Theta}_i$, denoted by $WCH(\tilde{\Theta}_i)$, is the maximal set S such that $WC(S) = WC(\tilde{\Theta}_i)$ (no $T \supset S$ satisfies $WC(T) = WC(\tilde{\Theta}_i)$). Equivalently,

$$WCH(\tilde{\Theta}_i) = \bigcup_{\tilde{\Theta}_i : WC(\tilde{\Theta}_i) = WC(\tilde{\Theta}_i)} \tilde{\Theta}_i.$$

We now give simpler characterizations of weakest-competitor sets and weakest-competitor hulls without explicit reference to the mechanics of the weakest-competitor VCG mechanism.

Theorem 3.3. Let $\Theta = [0, H]^\Gamma$ and let $\tilde{\Theta} \subseteq \Theta$ be closed and connected. Then

$$\overline{WC}(\tilde{\Theta}) = \left\{ \theta \in \tilde{\Theta} : \left\{ \theta' \in \tilde{\Theta} : \theta' \prec \theta \right\} = \emptyset \right\}, \quad WC(\tilde{\Theta}) = \left\{ \theta \in \tilde{\Theta} : \left\{ \theta' \in \tilde{\Theta} : \theta' \preccurlyeq \theta \right\} = \emptyset \right\},$$

and

$$\text{WCH}(\tilde{\Theta}) = \left\{ \theta \in \Theta : \exists \theta' \in \tilde{\Theta} \text{ s.t. } \theta \succeq \theta' \right\}$$

is the upwards closure of $\tilde{\Theta}$.

Proof. Let i denote the index of the agent under consideration with type space $\tilde{\Theta}$. Let $\theta \in \tilde{\Theta}$ be a point such that there exists $\theta' \in \tilde{\Theta}$ with $\theta' \prec \theta$. Then,

$$\max_{\alpha \in \Gamma} \sum_{j \neq i} \theta_j[\alpha] + \theta'[\alpha] < \max_{\alpha \in \Gamma} \sum_{j \neq i} \theta_j[\alpha] + \theta[\alpha]$$

for all $\theta_{-i} \in \Theta_{-i}$. So $\theta \notin \overline{\text{WC}}(\tilde{\Theta})$, which shows that $\overline{\text{WC}}(\tilde{\Theta}) \subseteq \{\theta \in \tilde{\Theta} : \{\theta' \in \tilde{\Theta} : \theta' \prec \theta\} = \emptyset\}$. To show the reverse containment, let $\theta \in \tilde{\Theta}$ be such that $\{\theta' \in \tilde{\Theta} : \theta' \prec \theta\} = \emptyset$. Consider any $\theta_{-i} = (\theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_n)$ such that

$$\sum_{j \neq i} \theta_j[\alpha_1] + \theta[\alpha_1] = \sum_{j \neq i} \theta_j[\alpha_2] + \theta[\alpha_2] = \dots = \sum_{j \neq i} \theta_j[\alpha_{|\Gamma|}] + \theta[\alpha_{|\Gamma|}].$$

The existence of such a θ_{-i} can be shown explicitly as follows. Let $j \neq i$ be arbitrary. For all $k \notin \{i, j\}$ set $\theta_k = (0, \dots, 0)$. Without loss of generality relabel the allocations in Γ such that $\theta[\alpha_1] \geq \theta[\alpha_2] \geq \dots \geq \theta[\alpha_{|\Gamma|}]$. Then, set

$$\theta_j = (0, \theta[\alpha_1] - \theta[\alpha_2], \dots, \theta[\alpha_1] - \theta[\alpha_{|\Gamma|}]) \in [0, H]^{\Gamma}.$$

Then, θ minimizes

$$\max \left\{ \sum_{j \neq i} \theta_j[\alpha_1] + \theta[\alpha_1], \sum_{j \neq i} \theta_j[\alpha_2] + \theta[\alpha_2], \dots, \sum_{j \neq i} \theta_j[\alpha_{|\Gamma|}] + \theta[\alpha_{|\Gamma|}] \right\}$$

since any θ' that attains a strictly smaller value must satisfy $\theta' \prec \theta$ (and no such θ' exists, by assumption). So $\theta \in \overline{\text{WC}}(\tilde{\Theta})$, which proves the reverse containment. The characterizations of WC and WCH follow immediately. \square

For example, in the single-dimensional case we have $\overline{\text{WC}}([\underline{\theta}, \bar{\theta}]) = \text{WC}([\underline{\theta}, \bar{\theta}]) = \{\underline{\theta}\}$, and $\text{WCH}([\underline{\theta}, \bar{\theta}]) = [\underline{\theta}, H]$. Figure 1 displays example weakest-competitor sets and weakest-competitor hulls in a two-dimensional type space. We list a few additional properties of WC and WCH that are all immediate consequences of Theorem 3.3.

Proposition 3.4. *Let $\tilde{\Theta} \subseteq \Theta = [0, H]^{\Gamma}$ be closed and connected. Then the following properties hold:*

1. $\text{WC}(\text{WCH}(\tilde{\Theta})) = \text{WC}(\tilde{\Theta})$.
2. $\text{WCH}(\text{WC}(\tilde{\Theta})) = \text{WCH}(\tilde{\Theta})$.
3. *Idempotence:* $\text{WC}(\text{WC}(\tilde{\Theta})) = \text{WC}(\tilde{\Theta})$ and $\text{WCH}(\text{WCH}(\tilde{\Theta})) = \text{WCH}(\tilde{\Theta})$.

We now present our meta-mechanism, which we denote by \mathcal{M} . It uses the efficient allocation, but that allocation is enjoyed only by the subset of agents able to compete with the weakest competitors in the side information set. \mathcal{M} then implements the weakest-competitor payments on

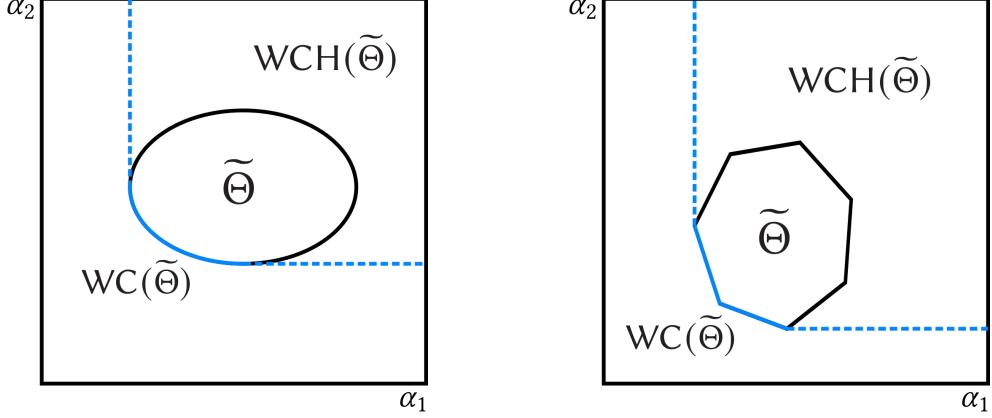


Figure 1: Example weakest-competitor hulls in a two-dimensional type space ($|\Gamma| = 2$). $\overline{WC}(\tilde{\Theta})$ is depicted in solid and dashed blue, $WC(\tilde{\Theta})$ is depicted in solid blue, and $WCH(\tilde{\Theta})$ is the region enclosed by $\overline{WC}(\tilde{\Theta})$.

those agents. The input subsets $\tilde{\Theta}_1, \dots, \tilde{\Theta}_n$ represent the side information/predictions given to the mechanism designer that postulate that $\theta_i \in \tilde{\Theta}_i$.

Meta-mechanism \mathcal{M}

Input: subsets $\tilde{\Theta}_1, \dots, \tilde{\Theta}_n \subseteq \Theta$ given to mechanism designer.

- Based on $\tilde{\Theta}_1, \dots, \tilde{\Theta}_n$, come up with $\hat{\Theta}_1, \dots, \hat{\Theta}_n$.

Agents asked to reveal types $\theta_1, \dots, \theta_n$.

- Let

$$\alpha^* = \operatorname{argmax}_{\alpha \in \Gamma} \sum_{i=1}^n \theta_i[\alpha]$$

and for each i let

$$p_i = \min_{\tilde{\theta}_i \in WC(\hat{\Theta}_i)} \left(\max_{\alpha \in \Gamma} \sum_{j \neq i} \theta_j[\alpha] + \tilde{\theta}_i[\alpha] \right) - \sum_{j \neq i} \theta_j[\alpha^*].$$

- Let

$$\mathcal{I} = \{i : \theta_i[\alpha^*] - p_i \geq 0\}.$$

If agent $i \notin \mathcal{I}$, i does not participate and receives zero utility (zero value and zero payment).^a If agent $i \in \mathcal{I}$, i enjoys allocation α^* and pays p_i .

^aOne practical consideration is that this step might require a more nuanced implementation of an “outside option” for agents to be indifferent between participating and being excluded versus not participating at all. (We do not pursue this highly application-specific issue in this work.) In auction and matching settings this step is standard and innocuous; the agent simply receives no items.

Meta-mechanism \mathcal{M} generates welfare equal to $\sum_{i \in \mathcal{I}} \theta_i[\alpha^*]$ and revenue equal to $\sum_{i \in \mathcal{I}} p_i$. \mathcal{M} does not specify how to set $\hat{\Theta}_1, \dots, \hat{\Theta}_n$ based on $\tilde{\Theta}_1, \dots, \tilde{\Theta}_n$ (hence the “meta” label). This challenge is the subject of the later sections where we will describe, based on the setting, how to set the $\hat{\Theta}_i$.

in order to generate high welfare and high revenue. We now establish the incentive properties of \mathcal{M} .

Theorem 3.5. *\mathcal{M} is incentive compatible and individually rational.*

Proof. \mathcal{M} is incentive compatible for the exact same reason weakest-competitor VCG is incentive compatible (Theorem 2.1). Individual rationality is an immediate consequence of how \mathcal{M} is defined; all agents with potential individual-rationality violations (those not in \mathcal{I}) do not participate and receive zero utility. \square

Next we show that the weakest-competitor hull precisely captures the set of agent types that never violate individual rationality. This consideration does not arise in the weakest-competitor VCG mechanism since in that setting misreporting is limited to the set used in the weakest-competitor minimization, and hence individual rationality is never violated. In our setting, we make no assumptions on the veracity of the sets $\widehat{\Theta}_i$ and must therefore reckon with the possibility that an agent is unable to compete with the weakest competitors in $\text{WC}(\widehat{\Theta}_i)$.

Theorem 3.6. *Let θ_i denote the true type of agent i and $\boldsymbol{\theta}_{-i}$ the reported types of the other agents. Let $\widehat{\Theta}_1, \dots, \widehat{\Theta}_n$ denote the side information sets used by \mathcal{M} . Then*

$$i \in \mathcal{I} \text{ for all } \boldsymbol{\theta}_{-i} \iff \theta_i \in \text{WCH}(\widehat{\Theta}_i).$$

Proof. Let θ_i denote the true type of agent i and let $\boldsymbol{\theta}_{-i} = (\theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_n)$ denote the reported types of the other agents. Suppose $\theta_i \notin \text{WCH}(\widehat{\Theta}_i)$. Then, there exists $\widetilde{\theta}_i \in \text{WC}(\widehat{\Theta}_i)$ such that $\widetilde{\theta}_i \succ \theta_i$, and there exists $\boldsymbol{\theta}_{-i}$ such that $\widetilde{\theta}_i \in \operatorname{argmin}_{\widetilde{\theta}_i \in \widehat{\Theta}_i} \max_{\alpha \in \Gamma} \sum_{j \neq i} \theta_j[\alpha] + \widetilde{\theta}_i[\alpha]$ is a weakest competitor relative to $\boldsymbol{\theta}_{-i}$ (the existence of $\boldsymbol{\theta}_{-i}$ follows from the same reasoning as in the proof of Theorem 3.3). As $\widetilde{\theta}_i \succ \theta_i$, agent i 's overall utility will be negative. The utility is unchanged and remains negative if $\widetilde{\theta}_i$ is replaced by $\theta_i^* \in \text{WC}(\widehat{\Theta}_i)$ that is also a weakest competitor relative to $\boldsymbol{\theta}_{-i}$. So we have shown there exists $\boldsymbol{\theta}_{-i}$ such that $i \notin \mathcal{I}$.

Conversely suppose $\theta_i \in \text{WCH}(\widehat{\Theta}_i)$. Then, there exists $\theta'_i \in \text{WC}(\widehat{\Theta}_i)$ such that $\theta_i \succeq \theta'_i$. Let $\boldsymbol{\theta}_{-i}$ be arbitrary. Agent i 's utility is $\sum_{j=1}^n \theta_j[\alpha^*] - \min_{\widetilde{\theta}_i \in \text{WC}(\widehat{\Theta}_i)} (\max_{\alpha} \sum_{j \neq i} \theta_j[\alpha] + \widetilde{\theta}_i[\alpha]) \geq \sum_{j=1}^n \theta_j[\alpha^*] - (\max_{\alpha} \sum_{j \neq i} \theta_j[\alpha] + \theta'_i[\alpha]) \geq \sum_{j=1}^n \theta_j[\alpha^*] - (\max_{\alpha} \sum_{j \neq i} \theta_j[\alpha] + \theta_i[\alpha]) = 0$, so $i \in \mathcal{I}$, as desired. \square

The key takeaway from Theorem 3.6 is that i is guaranteed to participate in \mathcal{M} regardless of the other agents' types if and only if $\theta_i \in \text{WCH}(\widehat{\Theta}_i)$. We will capitalize on this observation when we derive revenue guarantees for \mathcal{M} , since if $\widehat{\Theta}_i$ is derived from a high quality prediction/side information set $\widetilde{\Theta}_i$, one would expect that $\theta_i \in \text{WCH}(\widehat{\Theta}_i)$. In particular, the welfare of \mathcal{M} is at least $\sum_{i: \theta_i \in \text{WCH}(\widehat{\Theta}_i)} \theta_i[\alpha^*]$ and its revenue is at least $\sum_{i: \theta_i \in \text{WCH}(\widehat{\Theta}_i)} p_i$.

3.1 Computational considerations

Before we proceed to our main analyses of the key properties and guarantees of \mathcal{M} , we briefly discuss its computational complexity. We consider the special case where the side-information sets are polytopes. Let $\text{size}(\widehat{\Theta}_i)$ denote the encoding size of the constraints defining $\widehat{\Theta}_i$.

Theorem 3.7. *Let $\widehat{\Theta}_i$ be a polytope. Agent i 's payment p_i in the execution of \mathcal{M} can be computed in $\text{poly}(|\Gamma|, \text{size}(\widehat{\Theta}_i), n)$ time. Furthermore, determining membership in $\text{WCH}(\widehat{\Theta}_i)$ can be done in $\text{poly}(|\Gamma|, \text{size}(\widehat{\Theta}_i))$ time.*

Proof. The weakest competitor in $\widehat{\Theta}_i$ relative to θ_{-i} is the solution $\widetilde{\theta}_i \in \mathbb{R}^\Gamma$ to the linear program

$$\min \left\{ \gamma : \begin{array}{l} \widetilde{\theta}_i[\alpha] + \sum_{j \neq i} \theta_j[\alpha] \leq \gamma \quad \forall \alpha \in \Gamma, \\ \widetilde{\theta}_i \in \widehat{\Theta}_i, \gamma \geq 0 \end{array} \right\}$$

with $|\Gamma| + 1$ variables and $|\Gamma| + \text{size}(\widehat{\Theta}_i)$ constraints. Generating the first set of constraints requires computing $\sum_{j \neq i} \theta_j[\alpha]$ for each $\alpha \in \Gamma$, which takes time $\leq n|\Gamma|$.

Checking membership of θ_i in $\text{WCH}(\widehat{\Theta}_i)$ is equivalent to checking feasibility of a polytope

$$\theta_i \in \text{WCH}(\widehat{\Theta}_i) \iff \left\{ \widetilde{\theta}_i : \widetilde{\theta}_i \in \widehat{\Theta}_i, \theta_i[\alpha] \geq \widetilde{\theta}_i[\alpha] \quad \forall \alpha \in \Gamma \right\} \neq \emptyset$$

defined by $\text{size}(\widehat{\Theta}_i) + |\Gamma|$ constraints. \square

More generally, the complexity of the above two mathematical programs is determined by the complexity of constraints needed to define $\widehat{\Theta}_i$: for example, if $\widehat{\Theta}_i$ is a convex set then they are convex programs. Naturally, a major caveat of this brief discussion on computational complexity is that $|\Gamma|$ can be very large (for example, $|\Gamma|$ is exponential in combinatorial auctions).

4 Main guarantees of the mechanism in terms of prediction quality

In this section we prove our main guarantees on our meta-mechanism \mathcal{M} in terms of the quality of the side information $\widehat{\Theta}_1, \dots, \widehat{\Theta}_n$. We will largely refer to the side information as *predictions* in this section to emphasize that $\widehat{\Theta}_i$ could be wildly incorrect/inaccurate. To state our results we need the following notation which will be used throughout the remainder of the paper. Given agent types $\theta_1, \dots, \theta_n$, let α_{opt} denote the efficient allocation among the n agents and let $\text{OPT} = \max_{\alpha \in \Gamma} \sum_{i=1}^n \theta_i[\alpha] = \sum_{i=1}^n \theta_i[\alpha_{\text{opt}}]$ denote the welfare of the efficient allocation (also called the total social surplus). For a subset $S \subseteq \{1, \dots, n\}$ of agents, let $\text{OPT}_S = \sum_{i \in S} \theta_i[\alpha_{\text{opt}}]$ be the welfare generated by the efficient allocation restricted to agents in S . Let $\text{VCG}(S)$ denote the revenue of the vanilla VCG mechanism when run among the agents in S . Let $\text{VCG}(\beta) = \mathbb{E}[\text{VCG}(S)]$ where $S \subseteq \{1, \dots, n\}$ is sampled by including each agent in S independently with probability β . In general, VCG is not revenue monotonic [43], that is $S \subseteq T \not\Rightarrow \text{VCG}(S) \leq \text{VCG}(T)$, so $\text{VCG}(\beta)$ need not be increasing in β , but there are various sufficient conditions for revenue monotonicity.¹

The following calculation shows the revenue of \mathcal{M} can be related to the efficient welfare restricted to the subset of agents with valid predictions. We incur a loss term equal to the sum of ℓ_∞ -Hausdorff distances from each type θ_i to $\text{WC}(\widehat{\Theta}_i)$. The Hausdorff distance between θ_i and $\text{WC}(\widehat{\Theta}_i)$ is defined as $d_H(\theta_i, \text{WC}(\widehat{\Theta}_i)) := \max_{\widehat{\theta}_i \in \text{WC}(\widehat{\Theta}_i)} \|\theta_i - \widehat{\theta}_i\|_\infty$.

Lemma 4.1. *Run \mathcal{M} with $\widehat{\Theta}_1, \dots, \widehat{\Theta}_n$. Let $S = \{i : \theta_i \in \text{WCH}(\widehat{\Theta}_i)\}$ where θ_i is the true and revealed type of agent i (so $S \subseteq \mathcal{I}$). Then, for $i \in S$, \mathcal{M} satisfies*

$$p_i \geq \theta_i[\alpha_{\text{opt}}] - d_H(\theta_i, \text{WC}(\widehat{\Theta}_i)),$$

and

$$\text{revenue} \geq \text{OPT}_S - \sum_{i \in S} d_H(\theta_i, \text{WC}(\widehat{\Theta}_i)).$$

¹For example, if efficient welfare is a submodular set function over the agents, then VCG is revenue monotonic due to a result of Ausubel and Milgrom [6], and in this case $\text{VCG}(\beta)$ is increasing in β . In the combinatorial auction setting, a sufficient condition for efficient welfare to be submodular is that bidders have gross-substitutes valuations over items [28, 53].

Proof. Let p_i denote the payment collected from agent $i \in S$. Let θ_i^* be the weakest competitor in $\widehat{\Theta}_i$ with respect to θ_{-i} . The utility for agent i under \mathcal{M} is

$$\begin{aligned}
\theta_i[\alpha_{\text{opt}}] - p_i &= \sum_{j=1}^n \theta_j[\alpha_{\text{opt}}] - \min_{\widehat{\theta}_i \in \widehat{\Theta}_i} \left(\max_{\alpha \in \Gamma} \sum_{j \neq i} \theta_j[\alpha] + \widehat{\theta}_i[\alpha] \right) \\
&= \sum_{j=1}^n \theta_j[\alpha_{\text{opt}}] - \left(\max_{\alpha \in \Gamma} \sum_{j \neq i} \theta_j[\alpha] + \theta_i^*[\alpha] \right) \\
&\leq \sum_{j=1}^n \theta_j[\alpha_{\text{opt}}] - \left(\sum_{j \neq i} \theta_j[\alpha_{\text{opt}}] + \theta_i^*[\alpha_{\text{opt}}] \right) \\
&= \theta_i[\alpha_{\text{opt}}] - \theta_i^*[\alpha_{\text{opt}}] \\
&\leq \max_{\widehat{\theta}_i \in \text{WC}(\widehat{\Theta}_i)} \|\theta_i - \widehat{\theta}_i\|_{\infty} = d_H(\theta_i, \text{WC}(\widehat{\Theta}_i)),
\end{aligned}$$

as required. The revenue guarantee follows by summing up the bound for p_i over all $i \in S$. \square

Truncating the proof of Lemma 4.1 yields the following important bound, which is a more direct lower bound on p_i in terms of the weakest competitor's value for the efficient allocation (versus the bound in terms of Hausdorff distances of Lemma 4.1).

Lemma 4.2. *Let $i \in S$ (adopting the same notation and setup as Lemma 4.1). Then*

$$p_i \geq \widetilde{\theta}_i[\alpha_{\text{opt}}]$$

where $\widetilde{\theta}_i$ is the weakest competitor in $\widehat{\Theta}_i$ with respect to θ_{-i} .

Measuring the error of a prediction

Before instantiating \mathcal{M} with specific rules to determine the $\widehat{\Theta}_i$ from the $\widetilde{\Theta}_i$, we define our main notions of error in the side information/predictions, which are motivated by Lemma 4.1.

Definition 4.3. The *invalidity* of a prediction $\widetilde{\Theta}_i$, denoted by γ_i^V , is the distance from the true type θ_i of agent i to $\text{WCH}(\widetilde{\Theta}_i)$:

$$\gamma_i^V := d(\theta_i, \text{WCH}(\widetilde{\Theta}_i)) = \min_{\widetilde{\theta}_i \in \text{WCH}(\widetilde{\Theta}_i)} \|\theta_i - \widetilde{\theta}_i\|_{\infty}.$$

Definition 4.4. The *inaccuracy* of a prediction $\widetilde{\Theta}_i$ is the quantity

$$\gamma_i^A := d_H(\theta_i, \text{WC}(\widetilde{\Theta}_i)) = \max_{\widetilde{\theta}_i \in \text{WC}(\widetilde{\Theta}_i)} \|\theta_i - \widetilde{\theta}_i\|_{\infty}.$$

We say that a prediction $\widetilde{\Theta}_i$ is *valid* if $\gamma_i^V = 0$, that is, $\theta_i \in \text{WCH}(\widetilde{\Theta}_i)$. We say that a prediction is *perfect* if $\gamma_i^A = 0$ or, equivalently, $\text{WC}(\widetilde{\Theta}_i) = \{\theta_i\}$. If a prediction is perfect, then it is also valid. Our main results will depend on these error measures. Figure 2 illustrates these notions of prediction error for two different prediction sets.

Consistency and robustness. We say a mechanism is (a, b) -consistent and (c, d) -robust if when predictions are perfect it satisfies $\mathbb{E}[\text{welfare}] \geq a \cdot \text{OPT}$, $\mathbb{E}[\text{revenue}] \geq b \cdot \text{OPT}$, and satisfies $\mathbb{E}[\text{welfare}] \geq c \cdot \text{OPT}$, $\mathbb{E}[\text{revenue}] \geq d \cdot \text{VCG}$ independent of the prediction quality. High consistency and robustness ratios are in fact trivial to achieve, and we will thus largely not be too concerned with

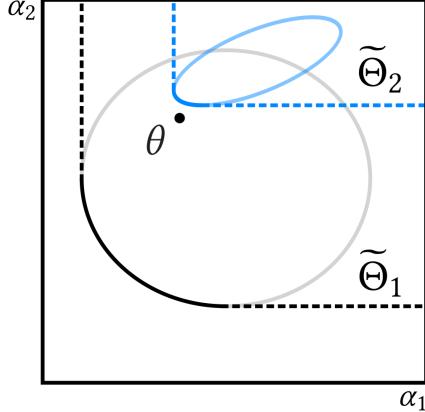


Figure 2: The prediction $\tilde{\Theta}_1$ (depicted in gray) is valid ($\gamma^V = 0$), but is highly inaccurate. The prediction $\tilde{\Theta}_2$ (depicted in light blue) is invalid ($\gamma^V > 0$), but is more accurate than $\tilde{\Theta}_1$. A small expansion of $\tilde{\Theta}_2$ would yield a valid and highly accurate prediction.

these measures—our main goal is to design high-performance mechanisms that degrade gracefully as the prediction errors increase. Indeed, the trivial mechanism that discards all side information with probability β and trusts the side information completely with probability $1-\beta$ is $(1, 1-\beta)$ -consistent and (β, β) -robust, but suffers from huge discontinuous drops in performance even when predictions are nearly perfect. That is, with probability β set $(\hat{\Theta}_1, \dots, \hat{\Theta}_n) = (\Theta, \dots, \Theta)$ and with probability $1-\beta$ set $(\hat{\Theta}_1, \dots, \hat{\Theta}_n) = (\text{WCH}(\tilde{\Theta}_1), \dots, \text{WCH}(\tilde{\Theta}_n))$, and then run \mathcal{M} . Let $V = \{i : \theta_i \in \text{WCH}(\tilde{\Theta}_i)\}$ be the set of valid predictions. From Lemma 4.1, we have $\mathbb{E}[\text{welfare}] = \beta \cdot \text{OPT} + (1-\beta) \cdot \text{OPT}_V$ and $\mathbb{E}[\text{revenue}] \geq \beta \cdot \text{VCG}(1) + (1-\beta) \cdot (\text{OPT}_V - \sum_{i \in V} \gamma_i^A)$, and thus obtain $(1, 1-\beta)$ -consistency and (β, β) -robustness. This approach suffers from a major issue: its revenue drops drastically the moment predictions are invalid ($\gamma_i^V > 0$). In particular, if predictions are highly accurate but very slightly invalid (such as the blue prediction in Figure 2), this approach completely misses out on any payments from such agents and drops to the revenue of VCG (which can be drastically smaller than OPT). But, a tiny expansion of these predictions would have sufficed to increase revenue significantly and perform competitively with OPT. One simple approach is to set $\hat{\Theta}_i$ to be an expansion of $\tilde{\Theta}_i$ by a parameter η_i with some probability, and discard the prediction with complementary probability. If $\gamma_i^V \leq \eta_i$ for all i , then such a mechanism would perform well. The main issue with such an approach is that the moment $\gamma_i^V > \eta_i$, our expansion by η_i fails to capture the true type θ_i and the performance drastically drops. Our approach in the next subsection essentially selects the η_i randomly from a suitable discretization of the ambient type space to be able to capture θ_i with reasonable probability.

4.1 Random expansion mechanism

Our main guarantee in this section will depend on H , which is an upper bound on any agent's value for any allocation. This is the only problem-domain-specific parameter in our results, and is naturally interpreted based on the setting (as in the examples we gave in Section 2). We have $H = \text{diam}(\Theta) := \max_{\theta_1, \theta_2 \in \Theta} \|\theta_1 - \theta_2\|_\infty$ is the ℓ_∞ -diameter of Θ . For a point θ , let $\mathcal{B}(\theta, r) = \{\theta' : \|\theta - \theta'\|_\infty \leq r\}$ be the closed ℓ_∞ -ball centered at θ with radius r . For a set $\tilde{\Theta}$, let $\mathcal{B}(\tilde{\Theta}, r) = \cup_{\tilde{\theta} \in \tilde{\Theta}} \mathcal{B}(\tilde{\theta}, r) = \{\theta' : \exists \tilde{\theta} \in \tilde{\Theta} \text{ s.t. } \|\theta' - \tilde{\theta}\|_\infty \leq r\}$ denote the ℓ_∞ -expansion of $\tilde{\Theta}$ by r . For $\zeta_i \geq 0$, $\lambda_i > 0$, and $K_i := \lceil \log_2((H - \zeta_i)/\lambda_i) \rceil$, let $\mathcal{M}_{\zeta, \lambda}$ denote the mechanism that for each i independently

sets

$$\widehat{\Theta}_i = \mathcal{B} \left(\widetilde{\Theta}_i, \zeta_i + 2^{k_i} \cdot \lambda_i \right), \text{ where } k_i \sim \text{unif. } \{0, 1, \dots, K_i\}.$$

$\mathcal{M}_{\zeta, \lambda}$ in essence performs a doubling search with initial hop ζ_i and λ_i controlling how fine-grained the search proceeds. We now prove and discuss our main welfare and revenue guarantees on $\mathcal{M}_{\zeta, \lambda}$. Define $\log^+ : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ by $\log^+(x) = 0$ if $x \leq 0$ and $\log^+(x) = \max\{0, \log(x)\}$ if $x > 0$.

Theorem 4.5 (Welfare bound). *The welfare of $\mathcal{M}_{\zeta, \lambda}$ satisfies*

$$\mathbb{E}[\text{welfare}] \geq \max \left\{ 1 - \max_i \frac{\lceil \log_2((\gamma_i^V - \zeta_i)/\lambda_i) \rceil}{1 + \lceil \log_2((H - \zeta_i)/\lambda_i) \rceil}, \frac{1}{1 + \lceil \max_i \log_2((H - \zeta_i)/\lambda_i) \rceil} \right\} \text{OPT}.$$

Proof. For each agent i , let k_i^* be the smallest $k \in \{0, \dots, K_i\}$ such that $\theta_i \in \text{WCH}(\mathcal{B}(\widetilde{\Theta}_i, \zeta_i + 2^k \cdot \lambda_i))$. Equivalently, k_i^* is the minimal k such that $\gamma_i^V \leq \zeta_i + 2^k \lambda_i$. So $k_i^* = \lceil \log_2((\gamma_i^V - \zeta_i)/\lambda_i) \rceil$. We have, using the fact that $\theta_i \in \widehat{\Theta}_i \implies i \in \mathcal{I}$ (Theorem 3.6),

$$\mathbb{E}[\text{welfare}] = \mathbb{E} \left[\sum_{i=1}^n \theta_i[\alpha_{\text{opt}}] \cdot \mathbf{1}(i \in \mathcal{I}) \right] \geq \mathbb{E} \left[\sum_{i=1}^n \theta_i[\alpha_{\text{opt}}] \cdot \mathbf{1}(\theta_i \in \widehat{\Theta}_i) \right] = \sum_{i=1}^n \theta_i[\alpha_{\text{opt}}] \cdot \Pr(\theta_i \in \widehat{\Theta}_i)$$

and

$$\Pr(\theta_i \in \widehat{\Theta}_i) = \Pr(k_i \geq k_i^*) = 1 - \Pr(k_i < k_i^*) = 1 - \frac{k_i^*}{1 + K_i} = 1 - \frac{\lceil \log_2((\gamma_i^V - \zeta_i)/\lambda_i) \rceil}{1 + \lceil \log_2((H - \zeta_i)/\lambda_i) \rceil}.$$

Therefore

$$\mathbb{E}[\text{welfare}] \geq \left(1 - \max_i \frac{\lceil \log_2((\gamma_i^V - \zeta_i)/\lambda_i) \rceil}{1 + \lceil \log_2((H - \zeta_i)/\lambda_i) \rceil} \right) \cdot \text{OPT}.$$

If all predictions are valid, we get $\mathbb{E}[\text{welfare}] = \text{OPT}$. The other term of the bound follows from $\Pr(\theta_i \in \widehat{\Theta}_i) \geq \Pr(k_i = k_i^*) = \frac{1}{1+K_i}$. \square

Theorem 4.6 (Revenue bound 1). *Let $\rho_i = 2(\gamma_i^V - \zeta_i)\mathbf{1}(\zeta_i + \lambda_i < \gamma_i^V) + \lambda_i\mathbf{1}(\zeta_i + \lambda_i \geq \gamma_i^V)$. The revenue of $\mathcal{M}_{\zeta, \lambda}$ satisfies*

$$\mathbb{E}[\text{revenue}] \geq \frac{1}{1 + \lceil \max_i \log_2((H - \zeta_i)/\lambda_i) \rceil} \left(\text{OPT} - \sum_{i=1}^n (\gamma_i^A + \zeta_i + \rho_i) \right).$$

Proof. Let k_i^* be defined as in the proof of Theorem 4.5. We compute expected revenue by computing $\mathbb{E}[p_i]$ for each agent i . Let $S = \{i : \theta_i \in \widehat{\Theta}_i\}$ be the (random) set of agents with valid predictions post expansion. We have

$$\mathbb{E}[p_i] \geq \mathbb{E}[p_i \mid k_i = k_i^*] \cdot \Pr(k_i = k_i^*) = \frac{1}{1 + K_i} \cdot \mathbb{E}[p_i \mid k_i = k_i^*].$$

Now $k_i = k_i^* \implies i \in S$, so we may apply the payment bound of Lemma 4.1:

$$\begin{aligned} \mathbb{E}[p_i \mid k_i = k_i^*] &\geq \mathbb{E} \left[\theta_i[\alpha_{\text{opt}}] - d_H(\theta_i, \text{WC}(\mathcal{B}(\widetilde{\Theta}_i, \zeta_i + 2^{k_i^*} \lambda_i))) \mid k_i = k_i^* \right] \\ &= \theta_i[\alpha_{\text{opt}}] - d_H(\theta_i, \text{WC}(\mathcal{B}(\widetilde{\Theta}_i, \zeta_i + 2^{k_i^*} \lambda_i))). \end{aligned}$$

Next, we bound $d_H(\theta_i, \text{WC}(\mathcal{B}(\widetilde{\Theta}_i, \zeta_i + 2^{k_i^*} \lambda_i)))$. Let $\widetilde{\theta}_i \in \text{WC}(\mathcal{B}(\widetilde{\Theta}_i, \zeta_i + 2^{k_i^*} \lambda_i))$ be arbitrary. By Lemma A.1 (the statement and proof are in Appendix A), $\widetilde{\theta}_i \in \text{WC}(\mathcal{B}(\text{WC}(\widetilde{\Theta}_i), \zeta_i + 2^{k_i^*} \lambda_i))$, so there

exists $\theta'_i \in \text{WC}(\tilde{\Theta}_i)$ such that $\|\tilde{\theta}_i - \theta'_i\|_\infty \leq \zeta_i + 2^{k_i^*} \lambda_i$. Moreover, $\|\theta_i - \theta'_i\|_\infty \leq \gamma_i^A$ by definition of γ_i^A . The triangle inequality therefore yields

$$\|\theta_i - \tilde{\theta}_i\|_\infty \leq \|\tilde{\theta}_i - \theta'_i\|_\infty + \|\theta_i - \theta'_i\|_\infty \leq \gamma_i^A + \zeta_i + 2^{k_i^*} \lambda_i,$$

so, as $\tilde{\theta}_i \in \text{WC}(\mathcal{B}(\tilde{\Theta}_i, \zeta_i + 2^{k_i^*} \lambda_i))$ was arbitrary, $d_H(\theta_i, \text{WC}(\mathcal{B}(\tilde{\Theta}_i, r_i^*))) \leq \gamma_i^A + \zeta_i + 2^{k_i^*} \lambda_i$. Now, we claim $2^{k_i^*} \lambda_i \leq \rho_i$. To show this, consider two cases. If $\zeta_i + \lambda_i \geq \gamma_i^V$, that is, $k_i^* = 0$, then $2^{k_i^*} \lambda_i = \lambda_i = \rho_i$. If $\zeta_i + \lambda_i < \gamma_i^V$, then $k_i^* > 0$ and we have $\zeta_i + 2^{k_i^*-1} \lambda_i < \gamma_i^V \leq \zeta_i + 2^{k_i^*} \lambda_i$, so $2^{k_i^*} \lambda < 2(\gamma_i^V - \zeta_i) = \rho_i$. So

$$d_H(\theta_i, \text{WC}(\mathcal{B}(\tilde{\Theta}_i, \zeta_i + 2^{k_i^*} \lambda_i))) \leq \gamma_i^A + \zeta_i + \rho_i.$$

Finally, we have

$$\mathbb{E}[p_i] \geq \frac{1}{1+K_i} (\theta_i[\alpha_{\text{opt}}] - (\gamma_i^A + \zeta_i + \rho_i)),$$

so

$$\begin{aligned} \mathbb{E}[\text{revenue}] &= \mathbb{E}\left[\sum_{i=1}^n p_i\right] = \sum_{i=1}^n \mathbb{E}[p_i] \geq \sum_{i=1}^n \frac{1}{1+K_i} (\theta_i[\alpha_{\text{opt}}] - (\gamma_i^A + \zeta_i + \rho_i)) \\ &\geq \frac{1}{1 + \lceil \max_i \log_2(\frac{H-\zeta_i}{\lambda_i}) \rceil} \left(\text{OPT} - \sum_{i=1}^n (\gamma_i^A + \zeta_i + \rho_i) \right), \end{aligned}$$

as desired. \square

Theorem 4.7 (Revenue bound 2). *The revenue of $\mathcal{M}_{\zeta, \lambda}$ satisfies*

$$\mathbb{E}[\text{revenue}] \geq \left(1 - \max_i \frac{\lceil \log_2((\gamma_i^V - \zeta_i)/\lambda_i) \rceil}{1 + \lceil \log_2((H - \zeta_i)/\lambda_i) \rceil}\right) \left(\text{OPT} - \sum_{i=1}^n (\gamma_i^A + \zeta_i)\right) - \sum_{i=1}^n \frac{4H}{1 + \lceil \log_2((H - \zeta_i)/\lambda_i) \rceil}.$$

Proof. Let k_i^* be defined as in Theorems 4.5 and 4.6. We bound $\mathbb{E}[p_i]$ similarly to the approach in the proof of Theorem 4.6, but account for all possible values of k_i (rather than only conditioning on $k_i = k_i^*$). If $k_i < k_i^*$, then agent i does not participate and pays nothing. We have

$$\begin{aligned} \mathbb{E}[p_i] &= \sum_{k=k_i^*}^{K_i} \mathbb{E}[p_i | k_i = k] \cdot \Pr(k_i = k) \geq \frac{1}{1+K_i} \sum_{k=k_i^*}^{K_i} (\theta_i[\alpha_{\text{opt}}] - d_H(\theta_i, \text{WC}(\mathcal{B}(\tilde{\Theta}_i, \zeta_i + 2^k \cdot \lambda_i)))) \\ &\geq \frac{1}{1+K_i} \sum_{k=k_i^*}^{K_i} (\theta_i[\alpha_{\text{opt}}] - (\gamma_i^A + \zeta_i + 2^k \cdot \lambda_i)) \\ &= \left(1 - \frac{k_i^*}{1+K_i}\right) (\theta_i[\alpha_{\text{opt}}] - \gamma_i^A - \zeta_i) - \frac{\lambda_i}{1+K_i} \sum_{k=k_i^*}^{K_i} 2^k \\ &\geq \left(1 - \frac{k_i^*}{1+K_i}\right) (\theta_i[\alpha_{\text{opt}}] - \gamma_i^A - \zeta_i) - \frac{\lambda_i 2^{K_i+1}}{1+K_i} \end{aligned}$$

where in the second inequality we have used the bound $d_H(\theta_i, \text{WC}(\mathcal{B}(\tilde{\Theta}_i, \zeta_i + 2^k \lambda_i))) \leq \gamma_i^A + \zeta_i + 2^k \lambda_i$, which was derived in the proof of Theorem 4.6. We have $2^{K_i+1} \leq \lambda_i (4(H - \zeta_i)/\lambda_i) \leq 4H$. Substituting and summing over agents yields the desired revenue bound. \square

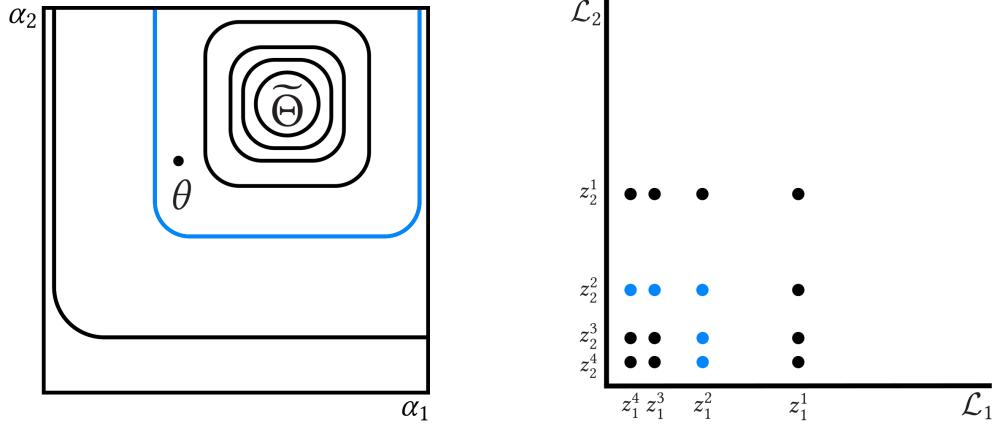


Figure 3: **Left:** An illustration of the sets $\mathcal{B}(\tilde{\Theta}, \zeta_i + 2^{k_i} \lambda_i)$ that the random expansion mechanism $\mathcal{M}_{\zeta, \lambda}$ chooses from for a particular agent with true type θ and side-information set/prediction $\tilde{\Theta}$. The blue expansion is the set $\mathcal{B}(\tilde{\Theta}, \zeta_i + 2^{k^*} \lambda_i)$. **Right:** An illustration of the weakest competitors $\tilde{\theta}(\ell_1, \dots, \ell_k)$ the mechanism \mathcal{M}_k randomizes over for a particular agent when $k = 2$ and $H = 2^4$. The weakest competitors with $(\ell_1, \dots, \ell_k) \in W_2$ are displayed in blue. (In both figures we drop the subscript of i indicating the agent index.)

First, consider constant ζ, λ . Our welfare guarantee degrades from OPT to $\Omega(\text{OPT}/\log H)$ as the invalidity of the predictions increase. Revenue bound 1 degrades from $\Omega(\text{OPT}/\log H)$ as both the invalidity and inaccuracy of the predictions increase. Revenue bound 2 illustrates that we can obtain significantly better performance if the parameters ζ, λ are chosen appropriately. In particular, for any ε , $\zeta_i = \gamma_i^V$ and $\lambda_i \leq O((H - \zeta_i)/2^{H/\varepsilon})$ yields $\mathbb{E}[\text{revenue}] \geq \text{OPT} - \sum_{i=1}^n (\gamma_i^A + \gamma_i^V + \varepsilon)$, a bound that is only additively worse than the total social surplus (and recovers the total social surplus as $\lambda_i \downarrow 0$ if the predictions are perfect). This bound degrades gradually as ζ_i, λ_i deviate.

To summarize, if the side information is of very high quality, the best parameters ζ, λ nearly recover the total social surplus OPT as welfare *and* revenue, and revenue degrades gradually as the chosen parameters ζ, λ worsen. If the side information is of questionable quality, the best parameters ζ, λ still obtain OPT as welfare, with revenue suffering additively by the prediction errors. As the parameter selection worsens, welfare and revenue degrade to $\Omega(\text{OPT}/\log H)$ with revenue suffering the same additive loss. Effective parameters can be, for example, learned from data [31].

Consistency and robustness. In this discussion we assume constant $\zeta, \lambda > 0$. Assuming revenue monotonicity, since $\theta_i \in \hat{\Theta}_i$ with probability at least $\Omega(1/\log H)$, the revenue of our mechanism is never worse than $\text{VCG}(\Omega(1/\log H))$. Thus, $\mathcal{M}_{\zeta, \lambda}$ is $(1, \Omega(1/\log H))$ -consistent (but if ζ, λ are chosen per the previous discussion, approaches the optimal $(1, 1)$ -consistency) and, assuming revenue monotonicity, $(\Omega(1/\log H), \text{VCG}(\Omega(1/\log H)))$ -robust. If VCG revenue is submodular, the revenue robustness ratio is $\geq \Omega(1/\log H)$ (but in general VCG revenue can shrink by more than this ratio [12]). In contrast to the trivial approach that either trusted the side information completely or discarded predictions completely, our random expansion approach does not suffer from large discontinuous drops in welfare nor revenue.

4.2 More expressive forms of side information

In this subsection we establish two avenues for richer and more expressive side information languages. The first deals with uncertainty and the second with joint multi-agent predictions.

Uncertainty. We now show that the techniques we have developed so far readily extend to an even larger more expressive form of side information that allows one to express varying degrees of uncertainty. A *side information structure* corresponding to agent i is given by a partition (A_1^i, \dots, A_m^i) of the ambient type space Θ_i into disjoint sets, probabilities $\mu_1^i, \dots, \mu_m^i \geq 0$; $\sum_j \mu_j^i = 1$ corresponding to each partition element, and for each partition element an optional probability density function f_j^i ; $\int_{A_j^i} f_j^i = 1$. The side information structure represents (1) a belief over what partition element A_j^i the true type θ_i lies in and (2) if a density is specified, the precise nature of uncertainty over the true type within A_j^i . Our model of side information sets $\tilde{\Theta}_i$ considered earlier in the paper corresponds to the partition $(\tilde{\Theta}_i, \Theta_i \setminus \tilde{\Theta}_i)$ with $\mu(\tilde{\Theta}_i) = 1$ and no specified densities. The richer model allows side information to convey finer-grained beliefs; for example one can express quantiles of certainty, precise distributional beliefs, and arbitrary mixtures of these.

Our notions of prediction error (invalidity and inaccuracy) can be naturally generalized. We define $\bar{\gamma}_i^V = \sum_j \mu_j^i \gamma_i^V(A_j^i; f_j^i)$ and $\bar{\gamma}_i^A = \sum_j \mu_j^i \gamma_i^A(A_j^i; f_j^i)$, where $\gamma_i^V(A_j^i; f_j^i) = d(\theta_i, \text{WCH}(A_j^i))$ if $f_j^i = \text{None}$ and $\gamma_i^V(A_j^i; f_j^i) = \mathbb{E}_{\tilde{\theta}_i \sim f_j^i} [d(\theta_i, \text{WCH}(\{\tilde{\theta}_i\}))]$ if f_j^i is a well-defined density. Similarly $\gamma_i^A(A_j^i; f_j^i) = d_H(\theta_i, \text{WC}(A_j^i))$ if $f_j^i = \text{None}$ and $\gamma_i^A(A_j^i; f_j^i) = \mathbb{E}_{\tilde{\theta}_i \sim f_j^i} [d(\theta_i, \tilde{\theta}_i)]$ otherwise.

Our generalized version of $M_{\zeta, \lambda}$ first samples a partition element A_j^i according to $(\mu_1^i, \dots, \mu_m^i)$, and draws $k_i \sim \text{unif. } \{0, \dots, K_i\}$ where K_i is defined as before. If $f_j^i = \text{None}$, it sets $\tilde{\Theta}_i = \mathcal{B}(A_j^i, \zeta_i + 2^{k_i} \lambda_i)$. Otherwise, it samples $\tilde{\theta}_i \sim f_j^i$ and sets $\tilde{\Theta}_i = \mathcal{B}(\{\tilde{\theta}_i\}, \zeta_i + 2^{k_i} \lambda_i)$. Versions of Theorems 4.5, 4.6, and 4.7 carry forward with $\bar{\gamma}_i^V$ and $\bar{\gamma}_i^A$ as the error measures. In Appendix B we provide the derivations, and also take the expressive power one step further by specifying a probability space over Θ_i ; its σ -algebra captures the granularity of knowledge being conveyed and its probability measure captures the uncertainty.

Joint side information. So far, side information has been independent across agents. Specifically, the mechanism designer receives sets $\tilde{\Theta}_i \subseteq \Theta_i$ for each agent i postulating that $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n) \in \tilde{\Theta}_1 \times \dots \times \tilde{\Theta}_n$. We show that our techniques extend to a more expressive form of side information that allows one to express predictions involving multiple agents. Let $\tilde{\Theta} = \tilde{\Theta}_1 \times \dots \times \tilde{\Theta}_n$. The mechanism designer receives as side information a set $\tilde{\Theta} \subseteq \Theta$ postulating that $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n) \in \tilde{\Theta}$. Given an agent i , a side information set $\tilde{\Theta} \subseteq \Theta$, and $\overline{\Theta}_{-i} \subseteq \times_{j \neq i} \Theta_j$, let

$$\text{proj}_i(\tilde{\Theta}; \overline{\Theta}_{-i}) = \left\{ \tilde{\theta}_i \in \Theta_i : \exists \boldsymbol{\theta}_{-i} \in \overline{\Theta}_{-i} \text{ s.t. } (\tilde{\theta}_i, \boldsymbol{\theta}_{-i}) \in \tilde{\Theta} \right\}$$

be the i th *projection* of $\tilde{\Theta}$ with respect to $\overline{\Theta}_{-i}$. The projection is the set of types for agent i consistent with $\tilde{\Theta}$ given that the realizations of the other agents' true types are contained in $\overline{\Theta}_{-i}$.

First, if $\tilde{\Theta}$ is known to be a valid prediction, that is, the true type profile $\boldsymbol{\theta}$ is guaranteed a priori to lie in $\tilde{\Theta}$ (equivalently, the joint misreporting space is limited to $\tilde{\Theta}$), we generalize the weakest-competitor VCG mechanism as follows. First, agents are asked to reveal their true types $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n)$. The allocation used is

$$\alpha^* = \underset{\alpha \in \Gamma}{\text{argmax}} \sum_{i=1}^n \theta_i[\alpha]$$

and agent i pays

$$p_i = \min_{\tilde{\theta}_i \in \text{proj}_i(\tilde{\Theta}; \boldsymbol{\theta}_{-i})} \left(\max_{\alpha \in \Gamma} \sum_{j \neq i} \theta_j[\alpha] + \tilde{\theta}_i[\alpha] \right) - \sum_{j \neq i} \theta_j[\alpha^*].$$

This generalized weakest-competitor VCG mechanism is IC, IR, and revenue optimal subject to efficiency in the joint information setting for the same reason that weakest-competitor VCG is in the independent information setting.

More generally, given side information set $\tilde{\Theta}$, our random expansion mechanism can be generalized as follows. Agents first reveal their true types $\theta_1, \dots, \theta_n$. For each agent i , independently set $\hat{\Theta}_i = \mathcal{B}(\text{proj}_i(\tilde{\Theta}, \theta_{-i}), \zeta_i + 2^{k_i} \lambda_i)$. The same guarantees we derived previously hold, with appropriately modified quality measures: invalidity is $\gamma_i^V = d(\theta_i, \text{WCH}(\text{proj}_i(\tilde{\Theta}, \theta_{-i})))$ and inaccuracy is $\gamma_i^A = d_H(\theta_i, \text{WC}(\text{proj}_i(\tilde{\Theta}, \theta_{-i})))$. An important idea highlighted by these mechanisms for joint side information is that the true types of all agents *other than* i can be heavily utilized in determining p_i . This model of side information affords us significantly more expressive power than the agent-independent model that has a product structure considered previously in this paper. For example, the mechanism designer might know that sum of the valuations of two customers for a cup of coffee exceeds a particular threshold, but does not know who has the higher value. Joint side information allows us to express such a belief precisely, and allows the mechanism designer to refine his beliefs on one agent based on the realized true type of the other agent. This was not possible in our previous framework.

We conclude this section with the observation that these mechanisms can loosely be interpreted as prior-free quantitative analogues of the seminal total-surplus-extraction Bayesian mechanism of Crémer and McLean [19] for correlated agents (generalized to infinite type spaces by McAfee and Reny [36]). This is an interesting connection to explore further in future research.

5 Constant-parameter agents: types on low-dimensional subspaces

In this section we show how the theory we have developed so far can be used to derive new revenue approximation results when the mechanism designer knows that each agent's type belongs to some low-dimensional subspace of \mathbb{R}^Γ (these subspaces can be different for each agent).

This is a slightly different setup from the previous sections. So far, we have assumed that $\Theta_i = \Theta$ for all i , that is, there is an ambient type space that is common to all the agents. Side information sets $\tilde{\Theta}_i$ are given as input to the mechanism designer, with no assumptions on quality/correctness (and our guarantees in Section 4 were parameterized by quality). Here, we assume the side information that each agent's type lies in a particular subspace is guaranteed to be valid. Two equivalent ways of stating this setup are (1) that Θ_i is the corresponding subspace for agent i and the mechanism designer receives no additional prediction set $\tilde{\Theta}_i$ or (2) $\Theta_i = \Theta$ for all i , $\tilde{\Theta}_i = \Theta \cap U_i$ where U_i is a subspace of \mathbb{R}^Γ , and the mechanism designer has the additional guarantee that $\theta_i \in U_i$ (so $\tilde{\Theta}_i$ is a *valid* side-information set). We shall use the language of the second interpretation.

In this setting, while the side information is valid, the inaccuracy errors γ_i^A of the sets $\tilde{\Theta}_i = \Theta \cap U_i$ can be too large to meaningfully use our previous guarantees. Nevertheless, we show how to fruitfully use the information provided by the subspaces U_1, \dots, U_n within the framework of our meta-mechanism. We begin by analyzing the most informative case: when agent types lie on lines. We then generalize our reasoning to any k -dimensional subspace that admits an orthonormal basis in the non-negative orthant of \mathbb{R}^Γ . In this section we assume $\Theta = [1, H]^\Gamma$, thereby imposing a lower bound of 1 on agent values (this choice of lower bound is not important, but the knowledge of some lower bound is needed in our analysis).

5.1 Agent types that lie on lines

Suppose for each i , the mechanism designer knows that θ_i lies in a line given by $U_i = \{\lambda u_i : \lambda \in \mathbb{R}\}$, for some $u_i \in \mathbb{R}_{\geq 0}^\Gamma$. Let $\mathcal{L}_i = U_i \cap [0, H]^\Gamma$ be the line segment that is the portion of this line that lies in $[0, H]^\Gamma$. We assume that $H = 2^a$ for some positive integer a .

Let $\tilde{\theta}_i^0$ be the endpoint of \mathcal{L}_i with $\|\tilde{\theta}_i^0\|_\infty = H$ (the other endpoint of \mathcal{L}_i is the origin). Let $\tilde{\theta}_i^1 = \tilde{\theta}_i^0/2$ be the midpoint of \mathcal{L}_i , and for $\ell = 2, \dots, \log_2 H$ let $\tilde{\theta}_i^\ell = \tilde{\theta}_i^{\ell-1}/2$ be the midpoint of $0\tilde{\theta}_i^{\ell-1}$. So $\|\tilde{\theta}_i^{\log_2 H}\|_\infty = 1$. We terminate the halving of the line segment \mathcal{L}_i after $\log_2 H$ steps due to the assumption that $\theta_i \in [1, H]^\Gamma$.

We can now specify our instantiation of \mathcal{M} , which we denote by \mathcal{M}_1 . For each i we independently set

$$\hat{\Theta}_i = \overline{\tilde{\theta}_i^{\ell_i} \tilde{\theta}_i^0} \text{ where } \ell_i \sim_{\text{unif.}} \{1, \dots, \log_2 H\}.$$

The weakest-competitor sets of the segments used in \mathcal{M}_1 are precisely the points in our logarithmic discretizations of the segments \mathcal{L}_i , that is, $\text{WC}(\tilde{\theta}_i^{\ell_i} \tilde{\theta}_i^0) = \{\tilde{\theta}_i^\ell\}$. We could equivalently instantiate \mathcal{M}_1 by simply setting $\hat{\Theta}_i = \{\tilde{\theta}_i^\ell\}$. We show that in this special case of types belonging to known rays, \mathcal{M}_1 satisfies effectively the same guarantees as the general mechanism in the previous section, but without the loss associated with prediction error. The key observation is that by Lemma 4.2, we can lower bound the payment p_i of agent i by the value $\tilde{\theta}_i[\alpha_{\text{opt}}]$ of the corresponding weakest competitor. This is a more direct lower bound than the one we used previously in terms of inaccuracy γ_i^A given by Lemma 4.1.

Theorem 5.1. \mathcal{M}_1 satisfies

$$\mathbb{E}[\text{welfare}] \geq \frac{1}{\log_2 H} \cdot \text{OPT}$$

and

$$\mathbb{E}[\text{revenue}] \geq \frac{1}{2 \log_2 H} \cdot \text{OPT}.$$

Proof. We have $\mathbb{E}[\text{welfare}] \geq \mathbb{E}[\sum_{i=1}^n \theta_i[\alpha_{\text{opt}}] \cdot \mathbf{1}(\theta_i \in \hat{\Theta}_i)] = \sum_{i=1}^n \theta_i[\alpha_{\text{opt}}] \cdot \Pr(\theta_i \in \hat{\Theta}_i) \geq \frac{1}{\log_2 H} \cdot \text{OPT}$. We now prove the revenue guarantee. Let $\ell_i^* = \min\{\ell : \theta_i \succeq \tilde{\theta}_i^\ell\}$. By how the $\tilde{\theta}_i^\ell$ are defined, we have $\tilde{\theta}_i^{\ell_i^*} \succeq \frac{1}{2} \theta_i$. We have

$$\mathbb{E}[p_i] \geq \mathbb{E}[p_i \mid \ell_i = \ell_i^*] \cdot \Pr(\ell_i = \ell_i^*) = \frac{1}{\log_2 H} \cdot \mathbb{E}[p_i \mid \ell_i = \ell_i^*].$$

Let $S = \{i : \theta_i \in \hat{\Theta}_i\}$. We have $\ell_i = \ell_i^* \implies i \in S$. Therefore, conditioned on $\ell_i = \ell_i^*$, Lemma 4.2 yields $p_i \geq \tilde{\theta}_i^{\ell_i^*}[\alpha_{\text{opt}}]$. As $\tilde{\theta}_i^{\ell_i^*} \succeq \frac{1}{2} \theta_i$, we have

$$\mathbb{E}[p_i \mid \ell_i^* = \ell_i] \geq \tilde{\theta}_i^{\ell_i^*}[\alpha_{\text{opt}}] \geq \frac{1}{2} \theta_i[\alpha_{\text{opt}}].$$

Finally,

$$\mathbb{E}[\text{revenue}] = \sum_{i=1}^n \mathbb{E}[p_i] \geq \frac{1}{2 \log_2 H} \cdot \sum_{i=1}^n \theta_i[\alpha_{\text{opt}}] = \frac{1}{2 \log_2 H} \cdot \text{OPT},$$

as desired. \square

5.2 Generalization to agent types that lie in subspaces

Suppose for each i , the mechanism designer knows that θ_i lies in a k -dimensional subspace $U_i = \text{span}(u_{i,1}, \dots, u_{i,k})$ of \mathbb{R}^Γ where each $u_{i,j} \in \mathbb{R}_{\geq 0}^\Gamma$ lies in the non-negative orthant and $\{u_{i,1}, \dots, u_{i,k}\}$ is an orthonormal basis for U_i . As in the previous section, we assume $H = 2^a$ for some positive integer a .

We generalize the single-dimensional analysis from the previous subsection in the design of our mechanism here. Let $\mathcal{L}_{i,j} = \{\lambda u_{i,j} : \lambda \geq 0\} \cap [0, H]^\Gamma$ be the line segment that is the portion of the ray generated by $u_{i,j}$ that lies in $[0, H]^\Gamma$. Let $y_{i,j}$ be the endpoint of $\mathcal{L}_{i,j}$ with $\|y_{i,j}\|_\infty = H$ (the other endpoint of $\mathcal{L}_{i,j}$ is the origin). Let $z_{i,j}^1 = y_{i,j}/2$ be the midpoint of $\mathcal{L}_{i,j}$, and for $\ell = 2, \dots, \log_2 H$ let $z_{i,j}^\ell = z_{i,j}^{\ell-1}/2$ be the midpoint of $\overline{0z_{i,j}^{\ell-1}}$. So $\|z_{i,j}^{\log_2 H}\|_\infty = 1$. We terminate the halving of $\mathcal{L}_{i,j}$ after $\log_2 H$ steps due to the assumption that $\theta_i \in [1, H]^\Gamma$. For every k -tuple $(\ell_1, \dots, \ell_k) \in \{1, \dots, \log_2 H\}^k$, let

$$\tilde{\theta}_i(\ell_1, \dots, \ell_k) = \sum_{j=1}^k z_{i,j}^{\ell_j}.$$

Furthermore, let $W_\ell = \{(\ell_1, \dots, \ell_k) \in \{1, \dots, \log_2 H\}^k : \min_j \ell_j = \ell\}$. The sets $W_1, \dots, W_{\log_2 H}$ partition $\{1, \dots, \log_2 H\}^k$ into levels, where W_ℓ is the set of points with ℓ_∞ -distance $H/2^\ell$ from the origin. We can now specify our instantiation of \mathcal{M} , which we denote by \mathcal{M}_k . For each i , we independently set

$\widehat{\Theta}_i = \{\tilde{\theta}_i(\ell_{i,1}, \dots, \ell_{i,k})\}$ where $\ell_i \sim_{\text{unif.}} \{1, \dots, \log_2 H\}$ and $(\ell_{i,1}, \dots, \ell_{i,k}) \sim_{\text{unif.}} W_{\ell_i}$.

Figure 3 is an illustration of the mechanism in the case of $k = 2$.

Theorem 5.2. \mathcal{M}_k satisfies

$$\mathbb{E}[\text{welfare}] \geq \frac{1}{\log_2 H} \cdot \text{OPT}$$

and

$$\mathbb{E}[\text{revenue}] \geq \frac{1}{2k(\log_2 H)^k} \cdot \text{OPT}.$$

Proof. We have $\mathbb{E}[\text{welfare}] \geq \sum_{i=1}^n \theta_i[\alpha_{\text{opt}}] \cdot \Pr(\theta_i \in \widehat{\Theta}_i) \geq \sum_{i=1}^n \theta_i[\alpha_{\text{opt}}] \cdot \Pr(\ell_i = \log_2 H) = \frac{1}{\log_2 H} \cdot \text{OPT}$ (since $\theta_i \succeq \tilde{\theta}_i(\log_2 H, \dots, \log_2 H)$). The proof of the revenue guarantee relies on the following key claim: for each agent i , there exists $\ell_{i,1}^*, \dots, \ell_{i,k}^* \in \{1, \dots, \log_2 H\}$ such that $\tilde{\theta}(\ell_{i,1}^*, \dots, \ell_{i,k}^*) \succeq \frac{1}{2}\theta_i$. To show this, let θ_i^j denote the projection of θ_i onto u_j , so $\theta_i = \sum_{j=1}^k \theta_i^j$ since $\{u_{i,1}, \dots, u_{i,k}\}$ is an orthonormal basis. Let $\ell_{i,j}^* = \min\{\ell : \theta_i^j \succeq z_{i,j}^\ell\}$. Then, $z_{i,j}^{\ell_{i,j}^*} \succeq \frac{1}{2}\theta_i^j$, so

$$\tilde{\theta}(\ell_{i,1}^*, \dots, \ell_{i,k}^*) = \sum_{j=1}^k z_{i,j}^{\ell_{i,j}^*} \succeq \sum_{j=1}^k \frac{1}{2}\theta_i^j = \frac{1}{2}\theta_i.$$

We now bound the expected payment of agent i as in the previous results. Let $\ell_i^* = \min_j \ell_{i,j}^*$. We have

$$\begin{aligned} \mathbb{E}[p_i] &\geq \mathbb{E}[p_i \mid (\ell_{i,1}, \dots, \ell_{i,k}) = (\ell_{i,1}^*, \dots, \ell_{i,k}^*)] \cdot \Pr((\ell_{i,1}, \dots, \ell_{i,k}) = (\ell_{i,1}^*, \dots, \ell_{i,k}^*)) \\ &= \frac{1}{|W_{\ell_i^*}| \log_2 H} \cdot \mathbb{E}[p_i \mid (\ell_{i,1}, \dots, \ell_{i,k}) = (\ell_{i,1}^*, \dots, \ell_{i,k}^*)] \end{aligned}$$

$$\begin{aligned}
&\geq \frac{1}{\log_2 H((\log_2 H)^k - (\log_2 H - 1)^k)} \cdot \mathbb{E}[p_i \mid (\ell_{i,1}, \dots, \ell_{i,k}) = (\ell_{i,1}^*, \dots, \ell_{i,k}^*)] \\
&\geq \frac{1}{k(\log_2 H)^k} \cdot \mathbb{E}[p_i \mid (\ell_{i,1}, \dots, \ell_{i,k}) = (\ell_{i,1}^*, \dots, \ell_{i,k}^*)]
\end{aligned}$$

since the probability of obtaining the correct weakest competitor $\tilde{\theta}(\ell_{i,1}^*, \dots, \ell_{i,k}^*)$ can be written as the probability of drawing the correct “level” $\ell_i^* \in \{1, \dots, \log_2 H\}$ times the probability of drawing the correct weakest competitor within the correct level $W_{\ell_i^*}$. We bound the conditional expectation as in the proof of Theorem 5.1. By Lemma 4.2,

$$\mathbb{E}[p_i \mid (\ell_{i,1}, \dots, \ell_{i,k}) = (\ell_{i,1}^*, \dots, \ell_{i,k}^*)] \geq \tilde{\theta}_i(\ell_{i,1}^*, \dots, \ell_{i,k}^*)[\alpha_{\text{opt}}] \geq \frac{1}{2} \cdot \theta_i[\alpha_{\text{opt}}].$$

Finally,

$$\mathbb{E}[\text{revenue}] = \sum_{i=1}^n \mathbb{E}[p_i] \geq \frac{1}{2k(\log_2 H)^k} \cdot \sum_{i=1}^n \theta_i[\alpha_{\text{opt}}] = \frac{1}{2k(\log_2 H)^k} \cdot \text{OPT},$$

as desired. \square

As mentioned in Section 1, the mechanisms in this section can be viewed as a generalization of the well-known $\log H$ revenue approximation in the single-item limited-supply setting that is achieved by a second-price auction with a random reserve price chosen uniformly from $\{H/2, H/4, H/8, \dots, 1\}$ [26]. Our results apply not only to combinatorial auctions but to general multidimensional mechanism design problems such as the examples presented in Section 2.

6 Conclusions and future research

We developed a versatile new methodology for multidimensional mechanism design that incorporates side information about agent types with the bicriteria goal of generating high social welfare and high revenue simultaneously. We designed a side-information-dependent meta-mechanism. This mechanism generalizes the weakest-competitor VCG mechanism of Krishna and Perry [33]. Careful instantiations of our meta-mechanism simultaneously achieved strong welfare and revenue guarantees that were parameterized by errors in the side information, and additionally proved to be fruitful in a setting where each agent’s type lies on a constant-dimensional subspace (of the potentially high-dimensional ambient type space) that is known to the mechanism designer.

There are many new interesting research directions that stem from our work. First, how far off are our mechanisms from the welfare-versus-revenue Pareto frontier? The weakest-competitor VCG mechanism is one extreme point, but what does the rest of the frontier look like? One possible approach here would be to extend our theory beyond VCG to the larger class of affine maximizers (which are known to contain high-revenue mechanisms)—we provide some initial ideas in Appendix C. Another important facet that we have largely ignored is computational complexity. The computations in our mechanism involving weakest competitors scale with the description complexity of $\tilde{\Theta}_i$ (e.g., the number of constraints, the complexity of constraints, and so on). An important question here is to understand the computational complexity of our mechanisms as a function of the differing (potentially problem-specific) language structures used to describe the side-information sets $\tilde{\Theta}_i$. In particular, the classes of side-information sets that are accurate, natural/interpretable, and easy to describe might depend on the specific mechanism design domain. Expressive bidding languages for combinatorial auctions have been extensively studied with massive impact in practice [45, 46]. Can a similar methodology for side information be developed in

conjunction with the results of this paper? Another natural direction is to extend our techniques to derive even stronger welfare and revenue guarantees when there is a known prior distribution on the agent types. All of our results apply to the Bayesian setting, but knowledge of the prior (or even sample access to the prior) could yield stronger guarantees. Another interesting direction along this vein is to generalize our mechanisms to depend on a known prior over prediction errors. Finally, the weakest-competitor variant of VCG of Krishna and Perry [33] is a strict improvement over the vanilla VCG mechanism, yet it appears to not have been further studied nor applied since its discovery. The weakest-competitor paradigm highlighted by that work and our results could potentially have important applications in the field of economics and computation more broadly.

Acknowledgements

This material is based on work supported by the NSF under grants IIS-1901403, CCF-1733556, CCF-1910321, and SES-1919453, and by the ARO under award W911NF2210266. S. Prasad thanks Morgan McCarthy for helpful discussions about real-world use cases of multidimensional mechanism design and Misha Khodak for detailed feedback on an earlier draft.

References

- [1] Vineet Abhishek and Bruce Hajek. Efficiency loss in revenue optimal auctions. In *49th IEEE Conference on Decision and Control (CDC)*. IEEE, 2010.
- [2] Gagan Aggarwal, Gagan Goel, and Aranyak Mehta. Efficiency of (revenue-)optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*, 2009.
- [3] Priyank Agrawal, Eric Balkanski, Vasilis Gkatzelis, Tingting Ou, and Xizhi Tan. Learning-augmented mechanism design: Leveraging predictions for facility location. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*, 2022.
- [4] Elliot Anshelevich, Koushik Kar, and Shreyas Sekar. Pricing to maximize revenue and welfare simultaneously in large markets. In *International Conference on Web and Internet Economics (WINE)*. Springer, 2016.
- [5] Aaron Archer and Éva Tardos. Truthful mechanisms for one-parameter agents. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, 2001.
- [6] Lawrence Ausubel and Paul Milgrom. Ascending auctions with package bidding. *The BE Journal of Theoretical Economics*, 2002.
- [7] Lawrence Ausubel and Paul Milgrom. The lovely but lonely Vickrey auction. *Combinatorial auctions*, 2006.
- [8] 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.
- [9] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Item pricing for revenue maximization. In *ACM Conference on Electronic Commerce (EC)*, 2008.
- [10] 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.

- [11] Maria-Florina Balcan, Siddharth Prasad, and Tuomas Sandholm. Efficient algorithms for learning revenue-maximizing two-part tariffs. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI)*, 2020.
- [12] Maria-Florina Balcan, Siddharth Prasad, and Tuomas Sandholm. Maximizing revenue under market shrinkage and market uncertainty. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [13] Eric Balkanski, Vasilis Gkatzelis, and Xizhi Tan. Strategyproof scheduling with predictions. In *Innovations in Theoretical Computer Science (ITCS)*, 2023.
- [14] Santiago Balseiro, Christian Kroer, and Rachitesh Kumar. Single-leg revenue management with advice. *arXiv preprint arXiv:2202.10939*, 2022.
- [15] Siddhartha Banerjee, Vasilis Gkatzelis, Artur Gorokh, and Billy Jin. Online nash social welfare maximization with predictions. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2022.
- [16] Jeremy Bulow and Paul Klemperer. Auctions versus negotiations. *The American Economic Review*, 1996.
- [17] Tanmoy Chakraborty, Zhiyi Huang, and Sanjeev Khanna. Dynamic and nonuniform pricing strategies for revenue maximization. *SIAM Journal on Computing*, 2013.
- [18] Ed H. Clarke. Multipart pricing of public goods. *Public Choice*, 1971.
- [19] Jacques Crémer and Richard P McLean. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica: Journal of the Econometric Society*, pages 1247–1257, 1988.
- [20] Constantinos Daskalakis and George Pierrakos. Simple, optimal and efficient auctions. In *International Workshop on Internet and Network Economics (WINE)*. Springer, 2011.
- [21] Nikhil R Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. The sample complexity of auctions with side information. In *ACM Symposium on Theory of Computing (STOC)*, 2016.
- [22] Ilias Diakonikolas, Christos Papadimitriou, George Pierrakos, and Yaron Singer. Efficiency-revenue trade-offs in auctions. In *International Colloquium on Automata, Languages, and Programming (ICALP)*, 2012.
- [23] 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.
- [24] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review*, 97(1):242–259, 2007.
- [25] Vasilis Gkatzelis, Kostas Kollias, Alkmini Sgouritsa, and Xizhi Tan. Improved price of anarchy via predictions. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*, 2022.
- [26] Andrew V. Goldberg, Jason D. Hartline, and Andrew Wright. Competitive auctions and digital goods. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2001.

- [27] Theodore Groves. Incentives in teams. *Econometrica*, 1973.
- [28] Faruk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic theory*, 87(1):95–124, 1999.
- [29] Venkatesan Guruswami, Jason D Hartline, Anna R Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On profit-maximizing envy-free pricing. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2005.
- [30] Jason D Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 225–234, 2009.
- [31] Misha Khodak, Maria-Florina Balcan, Ameet Talwalkar, and Sergei Vassilvitskii. Learning predictions for algorithms with predictions. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [32] Robert Kleinberg and Yang Yuan. On the ratio of revenue to welfare in single-parameter mechanism design. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce (EC)*, 2013.
- [33] Vijay Krishna and Motty Perry. Efficient mechanism design. *Available at SSRN 64934*, 1998.
- [34] Anton Likhodedov, Tuomas Sandholm, et al. Approximating revenue-maximizing combinatorial auctions. In *AAAI*, volume 5, pages 267–274, 2005.
- [35] Alexey Malakhov and Rakesh V Vohra. An optimal auction for capacity constrained bidders: a network perspective. *Economic Theory*, 2009.
- [36] R Preston McAfee and Philip J Reny. Correlated information and mechanism design. *Econometrica: Journal of the Econometric Society*, pages 395–421, 1992.
- [37] Andres Muñoz Medina and Sergei Vassilvitskii. Revenue optimization with approximate bid predictions. *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [38] Cade Metz. Facebook doesn’t make as much money as it could—on purpose. *Wired*, 2015.
- [39] Michael Mitzenmacher and Sergei Vassilvitskii. Algorithms with predictions. *Communications of the ACM*, 2022.
- [40] Jamie Morgenstern and Tim Roughgarden. Learning simple auctions. In *Conference on Learning Theory (COLT)*, 2016.
- [41] Roger Myerson. Optimal auction design. *Mathematics of operations research*, 1981.
- [42] Mallesh M Pai and Rakesh Vohra. Optimal auctions with financially constrained buyers. *Journal of Economic Theory*, 150:383–425, 2014.
- [43] Baharak Rastegari, Anne Condon, and Kevin Leyton-Brown. Revenue monotonicity in deterministic, dominant-strategy combinatorial auctions. *Artificial Intelligence*, 2011.
- [44] Kevin Roberts. The characterization of implementable social choice rules. In J-J Laffont, editor, *Aggregation and Revelation of Preferences*. North-Holland Publishing Company, 1979.
- [45] Tuomas Sandholm. Expressive commerce and its application to sourcing: How we conducted \$35 billion of generalized combinatorial auctions. *AI Magazine*, 28(3):45–45, 2007.

- [46] 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.
- [47] Tuomas Sandholm and Anton Likhodedov. Automated design of revenue-maximizing combinatorial auctions. *Operations Research*, 2015.
- [48] Hal R Varian and Christopher Harris. The VCG auction in theory and practice. *American Economic Review*, 2014.
- [49] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 1961.
- [50] Rakesh V Vohra. *Mechanism design: a linear programming approach*, volume 47. Cambridge University Press, 2011.
- [51] William E Walsh, David C Parkes, Tuomas Sandholm, and Craig Boutilier. Computing reserve prices and identifying the value distribution in real-world auctions with market disruptions. In *AAAI*, pages 1499–1502, 2008.
- [52] Chenyang Xu and Pinyan Lu. Mechanism design with predictions. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 2022.
- [53] 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.

A Omitted proof from Section 4

We prove the set-theoretic result concerning $\text{WC}(\cdot)$ and $\mathcal{B}(\cdot, r)$ used in Theorem 4.6.

Lemma A.1. $\text{WC}(\mathcal{B}(\tilde{\Theta}, r)) = \text{WC}(\mathcal{B}(\text{WC}(\tilde{\Theta}), r))$.

Proof. We first prove the forwards containment. For the sake of contradiction suppose there exists $\tilde{\theta} \in \text{WC}(\mathcal{B}(\tilde{\Theta}, r))$ such that $\tilde{\theta} \notin \text{WC}(\mathcal{B}(\text{WC}(\tilde{\Theta}), r))$. Then, there exists $\theta' \in \mathcal{B}(\text{WC}(\tilde{\Theta}), r)$ such that $\theta' \preccurlyeq \tilde{\theta}$. But

$$\text{WC}(\tilde{\Theta}) \subseteq \tilde{\Theta} \implies \mathcal{B}(\text{WC}(\tilde{\Theta}), r) \subseteq \mathcal{B}(\tilde{\Theta}, r),$$

so $\theta' \in \mathcal{B}(\tilde{\Theta}, r)$ and $\theta' \preccurlyeq \tilde{\theta}$ which contradicts the assumption that $\tilde{\theta} \in \text{WC}(\mathcal{B}(\tilde{\Theta}, r))$.

We now prove the reverse containment. For the sake of contradiction suppose there exists $\tilde{\theta} \in \text{WC}(\mathcal{B}(\text{WC}(\tilde{\Theta}), r))$ such that $\tilde{\theta} \notin \text{WC}(\mathcal{B}(\tilde{\Theta}, r))$. Then, there exists $\theta' \in \mathcal{B}(\tilde{\Theta}, r)$ such that $\theta' \preccurlyeq \tilde{\theta}$. Furthermore, if $\theta' \notin \text{WC}(\mathcal{B}(\tilde{\Theta}, r))$, there exists $\theta'' \in \text{WC}(\mathcal{B}(\tilde{\Theta}, r))$ such that $\theta'' \preccurlyeq \theta' \preccurlyeq \tilde{\theta}$ (if $\theta' \in \text{WC}(\mathcal{B}(\tilde{\Theta}, r))$, set $\theta'' = \theta'$). From the forward inclusion, we have

$$\text{WC}(\mathcal{B}(\tilde{\Theta}, r)) \subseteq \text{WC}(\mathcal{B}(\text{WC}(\tilde{\Theta}), r)) \subseteq \mathcal{B}(\text{WC}(\tilde{\Theta}), r),$$

so $\theta'' \in \mathcal{B}(\text{WC}(\tilde{\Theta}), r)$ and $\theta'' \preccurlyeq \tilde{\theta}$ which contradicts the assumption that $\tilde{\theta} \in \text{WC}(\mathcal{B}(\text{WC}(\tilde{\Theta}), r))$. \square

B An expressive language for side information

A *side information structure* is a probability space $(\Theta, \mathcal{F}, \mu)$ where the ambient type space Θ is the sample space, \mathcal{F} is a σ -algebra on Θ , and μ is a probability measure. (We suppress the agent index for brevity.)

\mathcal{F} induces a partition of Θ into equivalence classes where $\theta \equiv \theta'$ if $\mathbf{1}(\theta \in A) = \mathbf{1}(\theta' \in A)$ for all $A \in \mathcal{F}$ (so the side-information structure does not distinguish between θ and θ'). Let $A_\theta = \{\theta' : \theta \equiv \theta'\} \in \mathcal{F}$ be the equivalence class of θ . In this way the σ -algebra determines the granularity of knowledge being conveyed by the side information structure, and the probability measure $\mu : \mathcal{F} \rightarrow [0, 1]$ establishes uncertainty over this knowledge.

We define invalidity and inaccuracy of a side information structure in the natural way. Define random variables $X_i^V, X_i^A : \Theta \rightarrow \mathbb{R}_{\geq 0}$ by

$$X_i^V(\theta) = \gamma_i^V(A_\theta) = d(\theta_i, \text{WCH}(A_\theta))$$

and

$$X_i^A(\theta) = \gamma_i^A(A_\theta) = d_H(\theta_i, \text{WC}(A_\theta)).$$

X_i^V and X_i^A are \mathcal{F} -measurable since they are (by definition) constant on all atoms of \mathcal{F} (sets $A \in \mathcal{F}$ such that no nonempty $B \subsetneq A$ is in \mathcal{F}). The invalidity/inaccuracy distributions on $\mathbb{R}_{\geq 0}$ are given by

$$\Pr(a \leq X_i^V \leq b) = \mu(\{\theta \in \Theta : a \leq \gamma_i^V(A_\theta) \leq b\}) = \mu\left(\bigcup\{A_\theta : a \leq \gamma_i^V(A_\theta) \leq b\}\right).$$

The generalized version of $M_{\zeta, \lambda}$ that receives as input a side information structure for each agent i given by $(\Theta_i, \mathcal{F}_i, \mu_i)$ works as follows. It samples $\tilde{\theta}_i \sim \Theta_i$ according to (\mathcal{F}_i, μ_i) and draws $k_i \sim_{\text{unif.}} \{0, \dots, K_i\}$ where K_i is defined as before. It then sets

$$\widehat{\Theta}_i = \mathcal{B}\left(A_{\tilde{\theta}_i}, \zeta_i + 2^{k_i} \lambda_i\right).$$

Executing the same analysis in the proofs of Theorems 4.5, 4.6, and 4.7 for a fixed $\tilde{\theta}_i$, then taking expectation over the draw of $\tilde{\theta}_i$ yields similar guarantees with γ_i^V and γ_i^A replaced by $\mathbb{E}[X_i^V]$ and $\mathbb{E}[X_i^A]$, respectively. To show this, we loosen the bounds in Theorems 4.5, 4.6, and 4.7 slightly to make the multiplicative terms convex. As $X_i^V \geq 0$, we have

$$\left\lceil \log_2 \left(\frac{X_i^V - \zeta_i}{\lambda_i} \right) \right\rceil = \left\lceil \max \left(0, \log_2 \left(\frac{X_i^V - \zeta_i}{\lambda_i} \right) \right) \right\rceil \leq 1 + \log_2 \left(1 + \frac{X_i^V - \zeta_i}{\lambda_i} \right)$$

and therefore

$$\begin{aligned} \mathbb{E} \left[1 - \frac{\lceil \log_2((X_i^V - \zeta_i)/\lambda_i) \rceil}{1 + \lceil \log_2((H - \zeta_i)/\lambda_i) \rceil} \right] &\geq \mathbb{E} \left[1 - \frac{1 + \log_2(1 + (X_i^V - \zeta_i)/\lambda_i)}{1 + \lceil \log_2((H - \zeta_i)/\lambda_i) \rceil} \right] \\ &\geq 1 - \frac{1 + \log_2(1 + (\mathbb{E}[X_i^V] - \zeta_i)/\lambda_i)}{1 + \lceil \log_2((H - \zeta_i)/\lambda_i) \rceil} \end{aligned}$$

by Jensen's inequality. The corresponding versions of Theorems 4.5, 4.6, and 4.7 follow.

Theorem B.1. $\mathbb{E}[\text{welfare}] \geq \max\{(1 - \max_i \frac{1 + \log_2(1 + (\mathbb{E}[X_i^V] - \zeta_i)/\lambda_i)}{1 + \lceil \log_2((H - \zeta_i)/\lambda_i) \rceil}), \frac{1}{1 + \lceil \max_i \log_2((H - \zeta_i)/\lambda_i) \rceil}\} \text{OPT}.$

Theorem B.2 (Revenue bound 1). *Let $\rho_i = \mathbb{E}[2(X_i^V - \zeta_i)\mathbf{1}(\zeta_i + \lambda_i < X_i^V) + \lambda_i\mathbf{1}(\zeta_i + \lambda_i \geq X_i^V)]$. Then $\mathbb{E}[\text{revenue}] \geq \frac{1}{1 + \lceil \max_i \log_2((H - \zeta_i)/\lambda_i) \rceil} (\text{OPT} - \sum_{i=1}^n (\mathbb{E}[X_i^A] + \zeta_i + \mathbb{E}[\rho_i]))$.*

Theorem B.3 (Revenue bound 2). $\mathbb{E}[\text{revenue}] \geq (1 - \max_i \frac{1 + \log_2(1 + (\mathbb{E}[X_i^V] - \zeta_i)/\lambda_i)}{1 + \lceil \log_2((H - \zeta_i)/\lambda_i) \rceil})(\text{OPT} - \sum_{i=1}^n (\gamma_i^A + \zeta_i)) - \sum_{i=1}^n \frac{4H}{1 + \lceil \log_2((H - \zeta_i)/\lambda_i) \rceil}.$

C Beyond the VCG mechanism: affine maximizers

Given agent-specific multipliers $\omega = (\omega_1, \dots, \omega_n) \in \mathbb{R}_{\geq 0}$ and an allocation-based boost function $\lambda : \Gamma \rightarrow \mathbb{R}_{\geq 0}$, we define the following meta-mechanism $\mathcal{M}(\omega, \lambda)$ which is a generalization of our meta-mechanism \mathcal{M} . The mechanism designer receives as input $\tilde{\Theta}_1, \dots, \tilde{\Theta}_n$, and based on these decides on prediction sets $\hat{\Theta}_1, \dots, \hat{\Theta}_n$. The agents are then asked to reveal their true types $\theta_1, \dots, \theta_n$. The allocation used is

$$\alpha_{\omega, \lambda} = \operatorname{argmax}_{\alpha \in \Gamma} \sum_{i=1}^n \omega_i \theta_i[\alpha] + \lambda(\alpha).$$

Let

$$p_i = \frac{1}{\omega_i} \left[\min_{\tilde{\theta}_i \in \hat{\Theta}_i} \left(\max_{\alpha \in \Gamma} \sum_{j \neq i} \omega_j \theta_j[\alpha] + \omega_i \tilde{\theta}_i[\alpha] + \lambda(\alpha) \right) - \left(\sum_{j \neq i} \omega_j \theta_j[\alpha_{\omega, \lambda}] + \lambda(\alpha_{\omega, \lambda}) \right) \right].$$

Let

$$\mathcal{I} = \{i : \theta_i[\alpha_{\omega, \lambda}] - p_i \geq 0\}.$$

Agents in i enjoy allocation $\alpha_{\omega, \lambda}$ and pay p_i . Agents not in i do not participate and receive zero utility.

This mechanism is the natural generalization of the affine-maximizer mechanism [44] parameterized by ω, λ to our setting. The special case where agent misreporting is limited to $\hat{\Theta}_i$ is the natural generalization of the weakest-competitor VCG mechanism of Krishna and Perry [33] to affine-maximizer mechanisms. The following is a simple consequence of the proofs that \mathcal{M} and the affine-maximizer mechanism parameterized by ω, λ are incentive compatible and individually rational.

Theorem C.1. *For any $\omega \in \mathbb{R}_{\geq 0}^n$ and $\lambda : \Gamma \rightarrow \mathbb{R}_{\geq 0}$, $\mathcal{M}(\omega, \lambda)$ is incentive compatible and individually rational.*

Let $\text{OPT}(\omega, \lambda) = \sum_{i=1}^n \theta_i[\alpha_{\omega, \lambda}]$ be the welfare of the (ω, λ) -efficient allocation. All of the guarantees satisfied by \mathcal{M} carry over to $\mathcal{M}(\omega, \lambda)$, the only difference being the modified benchmark of $\text{OPT}(\omega, \lambda)$. Of course, $\text{OPT}(\omega, \lambda) \leq \text{OPT}$ is a weaker benchmark than the welfare of the efficient allocation. However, the class of affine maximizer mechanisms is known to achieve much higher revenue than the vanilla VCG mechanism. We leave it as a compelling open question to derive even stronger guarantees on mechanisms of the form $\mathcal{M}(\omega, \lambda)$ when the underlying affine maximizer is known to achieve greater revenue than vanilla VCG.