

Regret Minimization and Separation in Multi-Bidder Multi-Item Auctions

Çağıl Koçyiğit

Luxembourg Centre for Logistics and Supply Chain Management, University of Luxembourg, Luxembourg,
cagil.kocyiigit@uni.lu

Daniel Kuhn

Risk Analytics and Optimization Chair, École Polytechnique Fédérale de Lausanne, Switzerland,
daniel.kuhn@epfl.ch

Napat Rujeerapaiboon

Department of Industrial Systems Engineering and Management, National University of Singapore, Singapore,
napat.rujeerapaiboon@nus.edu.sg

We study a robust auction design problem with a minimax regret objective, where a seller seeks a mechanism for selling multiple items to multiple bidders with additive values. The seller knows that the bidders' values range over a box uncertainty set but has no information on their probability distribution. The robust auction design model we study requires no distributional information except for upper bounds on the bidders' values for each item. This model is relevant if there is no trust-worthy distributional information or if any distributional information is costly or time-consuming to acquire. We propose a mechanism that sells each item separately via a second price auction with a random reserve price and prove that this mechanism is optimal using duality techniques from robust optimization. We then interpret the auction design problem as a zero-sum game between the seller, who chooses a mechanism, and a fictitious adversary or 'nature,' who chooses the bidders' values from within the uncertainty set with the aim to maximize the seller's regret. We characterize the Nash equilibrium of this game analytically when the bidders are symmetric. The Nash strategy of the seller coincides with the optimal separable second price auction, whereas the Nash strategy of nature is mixed and constitutes a probability distribution on the uncertainty set under which each bidder's values for the items are comonotonic. We also study a restricted auction design problem over deterministic mechanisms. In this setting, we characterize the suboptimality of a separable second price auction with deterministic reserve prices and show that this auction becomes optimal if the bidders are symmetric. The optimal mechanism is derived in closed form and can easily be implemented by the practitioners.

Key words: multi-item auctions, robust auctions, mechanism design, minimax regret, robust optimization

1. Introduction

Consider the problem of designing an auction for selling J items to I bidders. The bidders assign each item a private value, which captures the maximum amount of money they would be willing to pay for this item. The set of values that a bidder assigns to *all* items is referred to as his value profile. We assume that the bidders' preferences are quasilinear and additively separable, that is, the bidders assign any bundle of items a value equal to the sum of the values of its constituents.

In the standard Bayesian setting, the seller's beliefs about the bidders' value profiles are modeled via a commonly known probability distribution, and it is assumed that the seller aims to maximize her expected revenues. If there is only one item ($J = 1$), the optimal mechanism is well-understood under relatively general conditions, see, *e.g.*, Myerson (1981) and Cremer and McLean (1988). If there are multiple items ($J > 1$), on the other hand, computing the optimal mechanism is #P-hard even in unrealistically simple situations (Daskalakis et al. 2014). Even though Daskalakis et al. (2017) and Cai et al. (2021) recently proposed duality schemes for solving multi-item auction design problems, closed-form solutions remain limited to special probabilistic models and/or small numbers of items, see, *e.g.*, Daskalakis et al. (2013) or Giannakopoulos and Koutsoupias (2014).

Assuming that the probability distribution of the bidders' values is commonly known not only renders the mechanism design problem intractable, but it is also difficult to justify in practice. Instead, it is natural to seek mechanisms that are optimal under limited distributional information. When the probability distribution of the bidders' values is ambiguous, the term 'optimal' becomes ambiguous itself. The literature on (distributionally) robust mechanism design regards an auction as optimal if it maximizes the worst-case expected revenues in view of all possible distributions consistent with the information available. The bulk of this literature focuses on single-item auctions, see, *e.g.*, Bose et al. (2006), Bei et al. (2019), Koçyiğit et al. (2020) and Suzdaltsev (2020). As a notable exception, Bandi and Bertsimas (2014) propose a numerical procedure to solve a robust multi-item auction design problem with budget constraints. Chen et al. (2022) investigate robust single-item and bundle pricing assuming that the seller only knows the marginal mean and variance of the value distribution. Carroll (2017) explicitly characterizes the optimal mechanism of a

correlation-robust screening problem, where the marginal distributions of the agent’s multidimensional type are precisely known to the principal, while their joint distribution remains unknown. The multidimensional monopoly pricing problem, which is equivalent to the single-bidder multi-item auction design problem, constitutes a special case of this screening problem. For this special case, Carroll (2017) shows that it is optimal to sell the items separately. Gravin and Lu (2018) then demonstrate that this separation result remains valid even if the bidder is subject to a budget constraint. Koçyiğit et al. (2022) consider a variant of the multidimensional monopoly pricing problem with a minimax regret objective, where the seller has no knowledge of the value distribution apart from its support, and they identify the best randomized as well as the best deterministic mechanism. In both cases, the optimal mechanism sells the items separately via single-item mechanisms that were first characterized by Bergemann and Schlag (2008).

The separation results reviewed above are not easily generalized to multi-bidder auctions. In this paper, we consider the multi-bidder extension of the mechanism design problem studied by Koçyiğit et al. (2022). Specifically, we assume that the seller perceives each bidder’s value profile as an uncertain parameter that is only known to range over a rectangular uncertainty set spanned by the origin and a vector of non-negative upper bounds. When aiming to maximize the worst-case revenue, the seller faces a special case of the robust mechanism design problem studied by Bandi and Bertsimas (2014). Under the box uncertainty set considered here, however, the set of optimal mechanisms is very rich and contains naïve mechanisms that have little practical appeal. For example, it is optimal for the seller to keep all items to herself. This prompts us to adopt a minimax regret objective, that is, we assume in this paper that the seller seeks a mechanism that minimizes her worst-case regret. The regret of a mechanism is defined as the difference between the revenues that could have been achieved under full knowledge of the bidders’ value profiles and the actual revenues generated by the mechanism. The worst-case regret is obtained by maximizing the realized regret across all possible value profiles of the bidders. Caldentey et al. (2017) as well as Poursoltani and Delage (2022) argue that, in a general robust optimization context, minimizing

the worst-case regret results in less conservative decisions than maximizing the worst-case revenue.

The main contributions of this paper can be summarized as follows.

- (i) We propose a mechanism under which each item is sold separately via a randomized second price auction. Using duality and limiting arguments, we prove that this mechanism minimizes the worst-case regret of the seller. If there is only one bidder, the separate single-item mechanisms reduce to randomized posted-price mechanisms that were first described by Bergemann and Schlag (2008) in the context of monopoly pricing.
- (ii) We interpret the robust multi-bidder multi-item mechanism design problem as a zero sum game between the seller, who chooses a mechanism to auction the items, and a fictitious adversary or ‘nature,’ who chooses the bidders’ value profiles from within a box uncertainty set. For symmetric bidders we characterize the Nash equilibrium of this game analytically, *i.e.*, we prove that the seller’s Nash strategy is the separable mechanism identified in (i), whereas nature’s Nash strategy is mixed and thus represents a probability distribution on the uncertainty set. Under this distribution, each bidder’s values for the items are comonotonic, and any bidder’s value profile can be non-zero only if all other bidders’ value profiles vanish.
- (iii) We also study a restricted auction design problem over deterministic mechanisms. In this setting, we characterize the suboptimality of a separable second price auction with deterministic reserve prices and show that this auction becomes optimal if the bidders are symmetric.

The mechanism design model studied in this paper requires no distributional information except for upper bounds on the bidders’ values for each item. This model is relevant if there is no trustworthy distributional information or if any distributional information is costly or time-consuming to acquire. Such a situation could arise, for example, when firms use auctions for initial public offerings. In this case, there is indeed no distributional information available about the bidders’ values for the offered shares. On the other hand, the model studied here may be overly conservative when data is abundant, as is typically the case in online advertisement, where auctions for ad placements are held in real time within fractions of seconds. To our best knowledge, this paper

establishes the first non-trivial robust optimality guarantee for a separable mechanism involving multiple bidders as well as multiple items. We hope that the insights distilled in this paper will pave the way towards more general separation results with a broader range of applications.

This paper also relates to the literature on approximately optimal mechanism design, see, *e.g.*, Dhangwatnotai et al. (2015), Hart and Nisan (2017), Allouah and Besbes (2020) and the references therein. Under this modeling paradigm, the seller aims to identify a mechanism for which some objective function (*e.g.*, the expected revenue) is guaranteed to be close to a full information benchmark value (*e.g.*, the maximum expected revenue achievable) under every probability distribution consistent with the assumptions made. The vast majority of the existing approximation results critically rely on certain independence assumptions (*e.g.*, the values must be independent across bidders or items). In the context of a monopoly pricing problem with a single buyer it has been shown, for example, that if the buyer’s values for the items are independent, then simple mechanisms (such as selling the goods separately or as a single grand bundle at deterministic posted prices) provide constant-factor approximations to the expected revenue of the unknown optimal mechanism (Hart and Nisan 2017). However, if the buyer’s values are correlated, these approximation guarantees cease to hold (Hart and Nisan 2019). The robust approach adopted in this paper does not rely on any independence assumptions. Instead, it is able to provide optimality guarantees for simple mechanisms even when the bidders’ values may be dependent.

Notation. For any closed set $\mathcal{A} \subseteq \mathbb{R}^n$, we denote by $\Delta(\mathcal{A})$ the family of all probability distributions on \mathcal{A} , and for any $\mathbb{P} \in \Delta(\mathcal{A})$, $\text{supp}(\mathbb{P})$ represents the support of \mathbb{P} . The set of all Borel-measurable functions from a Borel set $\mathcal{D} \subseteq \mathbb{R}^n$ to a Borel set $\mathcal{R} \subseteq \mathbb{R}^m$ is denoted by $\mathcal{L}(\mathcal{D}, \mathcal{R})$. Random variables are designated by tildes (*e.g.*, $\tilde{\mathbf{v}}$), and their realizations are denoted by the same symbols without tildes (*e.g.*, \mathbf{v}). For a logical expression \mathcal{E} , we define $\mathbf{1}_{\mathcal{E}} = 1$ if \mathcal{E} is true and $\mathbf{1}_{\mathcal{E}} = 0$ otherwise. Throughout the paper, bidders are indexed by superscripts and items by subscripts.

2. Problem Formulation

We consider the problem of designing a mechanism for selling $J \geq 1$ different items to $I \geq 1$ bidders. The sets of items and bidders are denoted by $\mathcal{J} = \{1, 2, \dots, J\}$ and $\mathcal{I} = \{1, 2, \dots, I\}$, respectively.

Each bidder $i \in \mathcal{I}$ assigns each item $j \in \mathcal{J}$ a value v_j^i that reflects his willingness to pay. In the following we denote by $\mathbf{v}^i = (v_1^i, \dots, v_J^i)$ the row vector of the values that bidder i assigns to all items and by $\mathbf{v}_j = (v_j^1, \dots, v_j^I)^\top$ the column vector of all bidders' values for item j . In addition, we let $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_J)$ be the matrix of all bidders' values for all items. While bidder i has full knowledge of his value profile \mathbf{v}^i , the seller perceives the matrix \mathbf{v} as uncertain. For each $j \in \mathcal{J}$, we assume that the seller only knows an upper bound $\bar{v}_j^i > 0$ on the value v_j^i for all $i \in \mathcal{I}$. We denote by $\bar{v}_j^* = \max_{i \in \mathcal{I}} \bar{v}_j^i$ the largest upper bound across all bidders for item $j \in \mathcal{J}$. We will say that the bidders are *symmetric* if $\bar{v}_j^i = \bar{v}_j^*$ for all $i \in \mathcal{I}$ and $j \in \mathcal{J}$. The seller has no other information about the distribution of \mathbf{v} or suspects that any available information is not trustworthy. For ease of exposition, we assume that the seller incurs no costs for supplying any of the items to any of the bidders. In the following, we denote by $\mathcal{V}^i = \times_{j \in \mathcal{J}} [0, \bar{v}_j^i]$ the uncertainty set of the value profiles of any fixed bidder $i \in \mathcal{I}$ and by $\mathcal{V} = \times_{i \in \mathcal{I}} \mathcal{V}^i$ the uncertainty set of the value profiles of all bidders. We also let \mathcal{W}_j^i , $i \in \mathcal{I}$ and $j \in \mathcal{J}$, be any partition of the set $\mathcal{V}_j = \times_{i \in \mathcal{I}} [0, \bar{v}_j^i]$ such that \mathcal{W}_j^i contains only scenarios \mathbf{v}_j for which bidder i is among the highest bidders for item j . In other words, $\mathbf{v}_j \in \mathcal{W}_j^i$ implies that $i \in \arg \max_{k \in \mathcal{I}} v_j^k$. If there are multiple highest bidders, an arbitrary tie-breaking rule is used (e.g., the lexicographic tie-breaker sets $\mathcal{W}_j^i = \{\mathbf{v}_j \in \mathcal{V}_j : i = \min \arg \max_{k \in \mathcal{I}} v_j^k\}$).

An auction mechanism (\mathbf{q}, \mathbf{m}) consists of an allocation rule $\mathbf{q} \in \mathcal{L}(\mathcal{V}, \mathbb{R}_+^{I \times J})$ and a payment rule $\mathbf{m} \in \mathcal{L}(\mathcal{V}, \mathbb{R}^I)$. Given a matrix $\mathbf{v} \in \mathcal{V}$ of value profiles reported by all bidders, the mechanism (\mathbf{q}, \mathbf{m}) outputs the allocation probabilities of the items to the bidders as well as the payments charged to the bidders. Specifically, in scenario \mathbf{v} , the seller allocates item j to bidder i with probability $q_j^i(\mathbf{v})$ and charges this bidder the amount $m^i(\mathbf{v})$. As a result, the utility of bidder i evaluates to $\sum_{j \in \mathcal{J}} q_j^i(\mathbf{v}) v_j^i - m^i(\mathbf{v})$ and is therefore quasilinear and additively separable across the items.

A (dominant strategy) incentive compatible and (ex-post) individually rational mechanism (\mathbf{q}, \mathbf{m}) satisfies the following constraints.

$$\sum_{j \in \mathcal{J}} q_j^i(\mathbf{v}) v_j^i - m^i(\mathbf{v}) \geq \sum_{j \in \mathcal{J}} q_j^i(\mathbf{w}^i, \mathbf{v}^{-i}) v_j^i - m^i(\mathbf{w}^i, \mathbf{v}^{-i}) \quad \forall i \in \mathcal{I}, \forall \mathbf{v} \in \mathcal{V}, \forall \mathbf{w}^i \in \mathcal{V}^i \quad (\text{IC})$$

$$\sum_{j \in \mathcal{J}} q_j^i(\mathbf{v}) v_j^i - m^i(\mathbf{v}) \geq 0 \quad \forall i \in \mathcal{I}, \forall \mathbf{v} \in \mathcal{V} \quad (\text{IR})$$

$$\sum_{i \in \mathcal{I}} q_j^i(\mathbf{v}) \leq 1 \quad \forall j \in \mathcal{J}, \forall \mathbf{v} \in \mathcal{V} \quad (\text{Inv})$$

The incentive compatibility constraint (IC) ensures that each bidder maximizes his utility by reporting his true value profile irrespective of the values reported by the other bidders. The individual rationality constraint (IR) ensures that the bidders earn non-negative utilities from participation under truthful reporting. Incentive compatibility and individual rationality constraints are routinely used in mechanism design and may be imposed essentially without any loss of generality thanks to the revelation principle (Krishna 2009, Chapter 5). The inventory constraint (Inv) ensures that the seller allocates each item $j \in \mathcal{J}$ with a probability of at most one. The inequality expresses the possibility that the seller may keep any item j to herself with a positive probability.

The seller's ex-post regret is defined as the difference between the maximum profit that could have been realized under complete information about \mathbf{v} and the profit earned with the mechanism (\mathbf{q}, \mathbf{m}) . If the seller was fully aware of the bidders' values \mathbf{v} , she would sell item j at the price $\max_{i \in \mathcal{I}} v_j^i$ to any bidder $i \in \arg \max_{i \in \mathcal{I}} v_j^i$. The maximum profit under complete information can thus be expressed as $\sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} v_j^i)$. The profit earned with mechanism (\mathbf{q}, \mathbf{m}) , on the other hand, amounts to $\sum_{i \in \mathcal{I}} m^i(\mathbf{v})$. In summary, the ex-post regret thus equals $\sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} v_j^i) - \sum_{i \in \mathcal{I}} m^i(\mathbf{v})$, and the worst-case regret is obtained by maximizing the ex-post regret over all value profiles $\mathbf{v} \in \mathcal{V}$.

Throughout this paper we assume that the seller aims to design an incentive compatible and individually rational mechanism that minimizes her worst-case regret. This mechanism design

problem can be formalized as the following robust optimization problem.

$$\begin{aligned}
z^* &= \inf_{\mathbf{q}, \mathbf{m}} \sup_{\mathbf{v} \in \mathcal{V}} \sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} v_j^i) - \sum_{i \in \mathcal{I}} m^i(\mathbf{v}) \\
\text{s.t. } & \mathbf{q} \in \mathcal{L}(\mathcal{V}, \mathbb{R}_+^{I \times J}), \mathbf{m} \in \mathcal{L}(\mathcal{V}, \mathbb{R}^I) \\
& \text{(IC), (IR), (Inv)}
\end{aligned} \tag{MDP}$$

From now on, we use the shorthand \mathcal{X} to denote the set of all mechanisms feasible in (MDP). We also denote by $\text{Reg}(\mathbf{q}, \mathbf{m})$ the worst-case regret of any fixed feasible mechanism (\mathbf{q}, \mathbf{m}) .

REMARK 1. One can show that problem (MDP) is equivalent to

$$\inf_{(\mathbf{q}, \mathbf{m}) \in \mathcal{X}} \sup_{\mathbb{P} \in \Delta(\mathcal{V})} \sup_{(\mathbf{q}', \mathbf{m}') \in \mathcal{X}} \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (m')^i(\tilde{\mathbf{v}}) - \sum_{i \in \mathcal{I}} m^i(\tilde{\mathbf{v}}) \right], \tag{1}$$

where $\tilde{\mathbf{v}}$ constitutes a random vector of bidder values governed by an unknown distribution $\mathbb{P} \in \Delta(\mathcal{V})$. Thus, \mathbb{P} is only known to be supported on \mathcal{V} . Note that problem (1) minimizes an alternative notion of worst-case regret, where the worst case is evaluated with respect to all distributions $\mathbb{P} \in \Delta(\mathcal{V})$. Here, the regret of a given mechanism (\mathbf{q}, \mathbf{m}) under a fixed distribution \mathbb{P} is defined as the difference between the maximum expected revenues achievable under perfect knowledge of \mathbb{P} and the expected revenues generated by (\mathbf{q}, \mathbf{m}) , which is ignorant of \mathbb{P} .

By interchanging the two supremum operators in (1) and noting that the regret of any fixed mechanism (\mathbf{q}, \mathbf{m}) is linear in \mathbb{P} , one readily verifies that the regret of (\mathbf{q}, \mathbf{m}) is actually maximized by a Dirac distribution in $\Delta(\mathcal{V})$. Hence, problem (1) is equivalent to

$$\inf_{(\mathbf{q}, \mathbf{m}) \in \mathcal{X}} \sup_{\mathbf{v} \in \mathcal{V}} \sup_{(\mathbf{q}', \mathbf{m}') \in \mathcal{X}} \sum_{i \in \mathcal{I}} (m')^i(\mathbf{v}) - \sum_{i \in \mathcal{I}} m^i(\mathbf{v}) \leq \inf_{(\mathbf{q}, \mathbf{m}) \in \mathcal{X}} \sup_{\mathbf{v} \in \mathcal{V}} \sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} v_j^i) - \sum_{i \in \mathcal{I}} m^i(\mathbf{v}),$$

where the inequality holds because $\sum_{i \in \mathcal{I}} (m')^i(\mathbf{v}) \leq \sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} v_j^i)$ for any mechanism $(\mathbf{q}', \mathbf{m}')$ that satisfies the individual rationality constraint (IR) as well as the inventory constraint (Inv). In fact, the inequality is tight. To see this, fix an arbitrary $\mathbf{v} \in \mathcal{V}$, and consider the posted-price mechanism $(\mathbf{q}', \mathbf{m}')$ that offers any item $j \in \mathcal{J}$ only to bidder $i_j = \min \arg \max_{i \in \mathcal{I}} v_j^i$ at the posted price $v_j^{i_j}$. This mechanism (which depends on \mathbf{v}) satisfies $(m')^i(\mathbf{v}) = v_j^{i_j} = \max_{i \in \mathcal{I}} v_j^i$ for every $i \in \mathcal{I}$. This confirms that the mechanism design problems (MDP) and (1) are indeed equivalent. \square

In the remainder we will sometimes study a deterministic variant of problem (MDP), where the allocation rule \mathbf{q} must be chosen from $\mathcal{L}(\mathcal{V}, \{0, 1\}^{I \times J})$ instead of $\mathcal{L}(\mathcal{V}, \mathbb{R}_+^{I \times J})$. Thus, the allocation probabilities must be set to 0 or 1 in each scenario. Deterministic mechanisms are easier to understand and communicate and may therefore be preferred in practice. In addition, the absence of any randomness in the allocation decisions may increase the acceptance of a mechanism.

One particularly simple policy for the seller would be to auction each of the J items individually. Any such mechanism is separable in the sense of the following definition.

DEFINITION 1 (SEPARABILITY). A mechanism (\mathbf{q}, \mathbf{m}) is called separable if there exists an item-wise allocation rule $\hat{\mathbf{q}}_j \in \mathcal{L}(\mathcal{V}_j, \mathbb{R}_+^I)$ and an item-wise payment rule $\hat{\mathbf{m}}_j \in \mathcal{L}(\mathcal{V}_j, \mathbb{R}^I)$ for all $j \in \mathcal{J}$ such that $\mathbf{q}(\mathbf{v}) = (\hat{\mathbf{q}}_1(\mathbf{v}_1), \dots, \hat{\mathbf{q}}_J(\mathbf{v}_J))$ and $\mathbf{m}(\mathbf{v}) = \sum_{j \in \mathcal{J}} \hat{\mathbf{m}}_j(\mathbf{v}_j)$ for all $\mathbf{v} \in \mathcal{V}$.

In the remainder we will show that problem (MDP) admits an optimal separable mechanism that is available in closed form.

3. Single-Bidder Pricing Problems

The mechanism design problem (MDP) has been studied by Bergemann and Schlag (2008) in the special case when $I = J = 1$ (single-item pricing) and by Koçyiğit et al. (2022) in the special case when $I = 1$ (multi-item pricing). This paper is the first to study problem (MDP) for general $I, J \geq 1$. As some of our results rely on a reduction of the multi-bidder auction design problem to a single-bidder pricing problem, we review here the results for pricing problems by Koçyiğit et al. (2022). Specifically, if $I = 1$, then a randomized separable posted-price mechanism is optimal in (MDP). A separable posted-price mechanism assigns each item j a posted price r_j . If the (single) bidder's value v_j^1 exceeds r_j , then he receives item j at price r_j . Otherwise, there is no transaction. A *randomized* separable posted-price mechanism assigns each item j a random posted price $\tilde{r}_j \sim \mathbb{Q}_j$.

In the following we will focus on the randomized separable posted-price mechanism (\mathbf{q}^*, m^*) induced by the posted-price distributions \mathbb{Q}_j^* , $j \in \mathcal{J}$, defined through

$$\mathbb{Q}_j^*(\tilde{r}_j \leq x) = \begin{cases} 1 + \log\left(\frac{x}{\bar{v}_j^1}\right) & \text{if } \frac{\bar{v}_j^1}{e} \leq x \leq \bar{v}_j^1, \\ 0 & \text{if } 0 \leq x < \frac{\bar{v}_j^1}{e}. \end{cases} \quad (2a)$$

The corresponding allocation and payment rules can be expressed as $\mathbf{q}(\mathbf{v}) = (\hat{q}_1(v_1^1), \dots, \hat{q}_J(v_J^1))$ and $m(\mathbf{v}) = \sum_{j \in \mathcal{J}} \hat{m}_j(v_j^1)$, respectively, where (\hat{q}_j, \hat{m}_j) is a single-item mechanism defined through

$$\hat{q}_j(v_j^1) = \mathbb{Q}_j^*(\tilde{r}_j \leq v_j^1) \quad \text{and} \quad \hat{m}_j(v_j^1) = \mathbb{E}_{\mathbb{Q}_j^*}[\tilde{r}_j \mathbb{1}_{(\tilde{r}_j \leq v_j^1)}] \quad \forall v_j^1 \in [0, \bar{v}_j^1]. \quad (2b)$$

In this case, the single-item allocation rules $\hat{q}_j(v_j^1)$ are piece-wise logarithmic and the single-item payment rules $\hat{m}_j(v_j^1)$ are piece-wise linear. Indeed, a direct calculation reveals that

$$\hat{m}_j(v_j^1) = \begin{cases} v_j^1 - \frac{\bar{v}_j^1}{e} & \text{if } \frac{\bar{v}_j^1}{e} \leq v_j^1 \leq \bar{v}_j^1, \\ 0 & \text{if } 0 \leq v_j^1 < \frac{\bar{v}_j^1}{e}. \end{cases}$$

One can show that if $I = 1$, then the separable mechanism (\mathbf{q}^*, m^*) is optimal in problem (MDP).

THEOREM 1 (Koçyiğit et al. 2022). *If $I = 1$, then the minimum of (MDP) is $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^1}{e}$, and the separable mechanism (\mathbf{q}^*, m^*) corresponding to the single-item mechanisms (2) is optimal.*

In addition, it is easy to see that for $I = 1$ problem (MDP) is equivalent to the zero-sum game

$$z^* = \inf_{(\mathbf{q}, m) \in \mathcal{X}} \sup_{\mathbb{P} \in \Delta(\mathcal{V})} z(m, \mathbb{P}),$$

where $z(m, \mathbb{P}) = \mathbb{E}_{\mathbb{P}}[\sum_{j \in \mathcal{J}} \tilde{v}_j^1 - m(\tilde{\mathbf{v}})]$ represents the expected regret of the mechanism (\mathbf{q}, m) if the bidder's values are governed by the probability distribution \mathbb{P} . In the following, we will focus on a distribution \mathbb{P}^* under which the bidder's values are comonotonic and which is defined through

$$\mathbb{P}^*(\tilde{\mathbf{v}}^1 \leq \mathbf{v}^1) = \begin{cases} \min_{j \in \mathcal{J}} \left(1 - \frac{1}{e} \left(\frac{\bar{v}_j^1}{v_j^1} \right) \right)^+ & \text{if } \mathbf{v}^1 \in \mathcal{V} \setminus \{(\bar{v}_1^1, \dots, \bar{v}_J^1)\}, \\ 1 & \text{if } \mathbf{v}^1 = (\bar{v}_1^1, \dots, \bar{v}_J^1). \end{cases} \quad (3)$$

One can show that if $I = 1$, then the separable mechanism (\mathbf{q}^*, m^*) and the comonotonic probability distribution \mathbb{P}^* form a Nash equilibrium in the above zero-sum game equivalent to (MDP).

THEOREM 2 (Koçyiğit et al. 2022). *If $I = 1$, then the separable mechanism (\mathbf{q}^*, m^*) corresponding to the single-item mechanisms (2) and the distribution \mathbb{P}^* defined in (3) satisfy the saddle point condition $\max_{\mathbb{P} \in \Delta(\mathcal{V})} z(m^*, \mathbb{P}) \leq z(m^*, \mathbb{P}^*) \leq \min_{(\mathbf{q}, m) \in \mathcal{X}} z(m, \mathbb{P}^*)$.*

Finally, we address the deterministic variant of problem (MDP) with $I = 1$, where the allocation rule \mathbf{q} must be chosen from $\mathcal{L}(\mathcal{V}, \{0, 1\}^J)$. In this case, the deterministic separable posted-price mechanism (\mathbf{q}^*, m^*) that sells item j at price $r_j = \bar{v}_j^1/2$ for any $j \in \mathcal{J}$ is optimal. This separable mechanism is induced by the single-item mechanisms (\hat{q}_j, \hat{m}_j) , $j \in \mathcal{J}$, which are defined through

$$(\hat{q}_j(v_j^1), \hat{m}_j(v_j^1)) = \begin{cases} (1, \bar{v}_j^1/2) & \text{if } v_j^1 \geq \bar{v}_j^1/2, \\ (0, 0) & \text{otherwise.} \end{cases} \quad (4)$$

THEOREM 3 (Koçyiğit et al. 2022). *If $I = 1$, then the minimum of the deterministic version of the mechanism design problem (MDP) amounts to $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^1}{2}$, and the separable mechanism (\mathbf{q}^*, m^*) corresponding to the single-item mechanisms (4) is optimal.*

4. Optimality Results

We now return to the general auction design problem (MDP) corresponding to arbitrary $I, J \geq 1$, and we construct a mechanism $(\mathbf{q}', \mathbf{m}')$, under which the seller implements a separate second price auction for each item $j \in \mathcal{J}$ with a random reserve price \tilde{r}_j governed by a probability distribution $\mathbb{Q}_j^* \in \Delta([0, \bar{v}_j^*])$, where $\bar{v}_j^* = \max_{i \in \mathcal{I}} \bar{v}_j^i$. Specifically, \mathbb{Q}_j^* is defined through

$$\mathbb{Q}_j^*(\tilde{r}_j \leq x) = \begin{cases} 1 + \log\left(\frac{x}{\bar{v}_j^*}\right) & \text{if } \frac{\bar{v}_j^*}{e} \leq x \leq \bar{v}_j^* \\ 0 & \text{if } 0 \leq x < \frac{\bar{v}_j^*}{e}. \end{cases}$$

Note that this definition of \mathbb{Q}_j^* naturally extends definition (2a) to an arbitrary number of bidders $I \geq 1$. For each item $j \in \mathcal{J}$, the respective second price auction proceeds as follows. First, the reserve price r_j is sampled from the distribution \mathbb{Q}_j^* . Note that the smallest possible value of r_j under this distribution is $\frac{\bar{v}_j^*}{e}$. The seller then asks the bidders to report their bids for item j . After collecting all bids, the seller allocates item j to the highest bidder provided that his bid exceeds the reserve price r_j , and the winner pays an amount equal to the maximum of the second highest bid and r_j . In the case of ties, item j is given to the unique bidder whose index i satisfies $\mathbf{v}_j \in \mathcal{W}_j^i$.

By construction, the single-item mechanisms (\hat{q}'_j, \hat{m}'_j) , $j \in \mathcal{J}$, corresponding to $(\mathbf{q}', \mathbf{m}')$ satisfy

$$(\hat{q}'_j)_j^i(\mathbf{v}_j, r_j) = \begin{cases} 1 & \text{if } \mathbf{v}_j \in \mathcal{W}_j^i \text{ and } v_j^i \geq r_j \\ 0 & \text{otherwise,} \end{cases} \quad (5a)$$

$$(\hat{m}')_j^i(\mathbf{v}_j, r_j) = \begin{cases} \max\{\max_{k \neq i} v_j^k, r_j\} & \text{if } \mathbf{v}_j \in \mathcal{W}_j^i \text{ and } v_j^i \geq r_j \\ 0 & \text{otherwise} \end{cases} \quad (5b)$$

for all $\mathbf{v}_j \in \mathcal{V}_j$. Note that $(\mathbf{q}', \mathbf{m}')$ is manifestly randomized because it depends on the realizations of the random reserve prices. Next, we introduce a new mechanism $(\mathbf{q}^*, \mathbf{m}^*)$, which is constructed by averaging $(\mathbf{q}', \mathbf{m}')$ across the random reserve prices. This new mechanism inherits separability from $(\mathbf{q}', \mathbf{m}')$, and the corresponding single-item mechanisms (\hat{q}_j, \hat{m}_j) , $j \in \mathcal{J}$, satisfy

$$\hat{q}_j^i(\mathbf{v}_j) = \mathbb{E}_{\mathbb{Q}_j^*}[(\hat{q}')_j^i(\mathbf{v}_j, \tilde{r}_j)] \quad \text{and} \quad \hat{m}_j^i(\mathbf{v}_j) = \mathbb{E}_{\mathbb{Q}_j^*}[(\hat{m}')_j^i(\mathbf{v}_j, \tilde{r}_j)] \quad (6)$$

for all $i \in \mathcal{I}$ and $\mathbf{v}_j \in \mathcal{V}_j$. For any $j \in \mathcal{J}$, the single-item mechanisms (\hat{q}_j, \hat{m}_j) and (\hat{q}'_j, \hat{m}'_j) represent two different implementations of the second price auction with a random reserve price, and they are equivalent in terms of expected payoffs to all agents. The key difference is that the allocation probabilities and payments are deterministic under (\hat{q}_j, \hat{m}_j) , whereas they are random and rely on \tilde{r}_j under (\hat{q}'_j, \hat{m}'_j) . Figure 1 visualizes a single-item mechanism of the form (6) with $I = 2$ bidders, $\bar{v}_1^1 = 1$ and $\bar{v}_1^2 = 2$. Note that bidder 2 is more likely to receive the item in a majority of all scenarios. When the item is allocated to bidder 1, however, the seller may earn considerably less. These phenomena arise whenever $\min_{i \in \mathcal{I}} \bar{v}_1^i \ll \max_{i \in \mathcal{I}} \bar{v}_1^i$. To gain a better understanding, it is instructive to evaluate the expectations in (6) explicitly. A simple calculation yields

$$\hat{q}_j^i(\mathbf{v}_j) = \begin{cases} 1 + \log\left(\frac{v_j^i}{\bar{v}_j^*}\right) & \text{if } \mathbf{v}_j \in \mathcal{W}_j^i \text{ and } v_j^i \geq \frac{\bar{v}_j^*}{e} \\ 0 & \text{otherwise,} \end{cases}$$

$$\hat{m}_j^i(\mathbf{v}_j) = \begin{cases} v_j^i + (\max_{k \neq i} v_j^k) \log\left(\max_{k \neq i} \frac{v_j^k}{\bar{v}_j^*}\right) & \text{if } \mathbf{v}_j \in \mathcal{W}_j^i \text{ and } v_j^i \geq \max_{k \neq i} v_j^k \geq \frac{\bar{v}_j^*}{e} \\ v_j^i - \frac{\bar{v}_j^*}{e} & \text{if } \mathbf{v}_j \in \mathcal{W}_j^i \text{ and } v_j^i \geq \frac{\bar{v}_j^*}{e} > \max_{k \neq i} v_j^k \\ 0 & \text{otherwise.} \end{cases}$$

REMARK 2. All second price auctions with deterministic reserve prices are (dominant strategy) incentive compatible (Krishna 2009, Chapter 2). Thus, $(\mathbf{q}', \mathbf{m}')$ is incentive compatible for any




Figure 1 Visualization of the separable mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ for $I = 2$ bidders, $J = 1$ item, $\bar{v}_1^1 = 1$ and $\bar{v}_1^2 = 2$

fixed realizations of the random reserve prices. As incentive compatibility can be encoded through (infinitely many) linear inequalities, it is preserved by averaging $(\mathbf{q}', \mathbf{m}')$ across the random reserve prices $\tilde{r}_j \sim \mathbb{Q}_j^*$, $j \in \mathcal{J}$. Hence, $(\mathbf{q}^*, \mathbf{m}^*)$ is incentive compatible, and the bidders have a weak preference to report their true values under $(\mathbf{q}^*, \mathbf{m}^*)$. All second price auctions with reserve prices are also (ex-post) individually rational because the bidders' utilities are always non-negative. Indeed, under truthful bidding, a bidder pays at most his own bid. Using similar arguments as above, one can thus show that $(\mathbf{q}^*, \mathbf{m}^*)$ is also individually rational. As the mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ ostensibly satisfies the inventory constraint (Inv), we thus conclude that it is feasible in (MDP). \square

In the following, we investigate the optimality properties of the mechanism $(\mathbf{q}^*, \mathbf{m}^*)$. Specifically, we will show that $(\mathbf{q}^*, \mathbf{m}^*)$ attains a worst-case regret of $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e} = \sum_{j \in \mathcal{J}} \max_{i \in \mathcal{I}} \frac{\bar{v}_j^i}{e}$ and that the optimal value of (MDP) is bounded below by the same value $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$. This implies that $(\mathbf{q}^*, \mathbf{m}^*)$ is optimal in (MDP). We first quantify the worst-case regret of $(\mathbf{q}^*, \mathbf{m}^*)$.

PROPOSITION 1. *The worst-case regret of the separable mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ corresponding to the single-item mechanisms (6) is given by $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$.*

Proof. We will first show that the regret of $(\mathbf{q}^*, \mathbf{m}^*)$ is bounded above by $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$ in any scenario $\mathbf{v} \in \mathcal{V}$ and then identify a scenario for which this upper bound is attained.

To this end, we introduce an artificial bidder whose value profile \mathbf{v}^a belongs to the box uncertainty set $\mathcal{V}^a = \times_{j \in \mathcal{J}} [0, \bar{v}_j^*]$. For any $j \in \mathcal{J}$, we denote by (f_j^a, g_j^a) a single-buyer mechanism for selling item j to the artificial bidder, where $f_j^a(v_j^a) = \mathbb{Q}_j^*(\tilde{r}_j \leq v_j^a)$ and $g_j^a(v_j^a) = \mathbb{E}_{\mathbb{Q}_j^*}[\tilde{r}_j \mathbb{1}_{(v_j^a \geq \tilde{r}_j)}]$ for all values $v_j^a \in [0, \bar{v}_j^*]$ of the artificial bidder. Note that f_j^a and g_j^a can thus be interpreted as the expected allocation and payment rules of the randomized posted-price mechanism with random price $\tilde{r}_j \sim \mathbb{Q}_j^*$, respectively; see also (2b). In addition, by Theorem 1 the separable single-buyer multi-item mechanism corresponding to the single-item single-buyer mechanisms $\{(f_j^a, g_j^a)\}_{j \in \mathcal{J}}$ is optimal in problem (MDP) if there is only the artificial bidder.

Fix now any scenario $\mathbf{v} \in \mathcal{V}$, and define $i^*(j)$ as the unique $i \in \mathcal{I}$ with $\mathbf{v}_j \in \mathcal{W}_j^i$. By construction, $i^*(j) \in \mathcal{I}$ identifies one of the highest bidders for item j in scenario \mathbf{v} . As the mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ sells each item j separately to bidder $i^*(j)$, its regret in scenario \mathbf{v} can thus be expressed as

$$\sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} v_j^i) - \sum_{i \in \mathcal{I}} \hat{m}^i(\mathbf{v}) = \sum_{j \in \mathcal{J}} (v_j^{i^*(j)} - \hat{m}_j^{i^*(j)}(\mathbf{v}_j)) \leq \sum_{j \in \mathcal{J}} (v_j^{i^*(j)} - g_j^a(v_j^{i^*(j)})) \leq \sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}, \quad (7)$$

where the first inequality holds because

$$\hat{m}_j^{i^*(j)}(\mathbf{v}_j) = \mathbb{E}_{\mathbb{Q}_j^*} \left[\max \left\{ \max_{i \neq i^*(j)} v_j^i, \tilde{r}_j \right\} \mathbb{1}_{(v_j^{i^*(j)} \geq \tilde{r}_j)} \right] \geq \mathbb{E}_{\mathbb{Q}_j^*} \left[\tilde{r}_j \mathbb{1}_{(v_j^{i^*(j)} \geq \tilde{r}_j)} \right] = g_j^a(v_j^{i^*(j)}).$$

Indeed, $\hat{m}_j^{i^*(j)}(\mathbf{v}_j)$ is the expected payment of bidder $i^*(j)$ under the second price auction with random reserve price $\tilde{r}_j \sim \mathbb{Q}_j^*$, whereas $g_j^a(v_j^{i^*(j)})$ is the expected payment under the posted-price mechanism with random price $\tilde{r}_j \sim \mathbb{Q}_j^*$. The second inequality in (7) holds because the separable mechanism corresponding to the single-item mechanisms $\{(f_j^a, g_j^a)\}_{j \in \mathcal{J}}$ is optimal in problem (MDP) if there is only one bidder whose value for item j is bounded above by \bar{v}_j^* for any $j \in \mathcal{J}$. In fact, we know from Theorem 1 that the worst-case regret of this mechanism is given by $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$. As $\mathbf{v} \in \mathcal{V}$ was chosen arbitrarily, the regret of $(\mathbf{q}^*, \mathbf{m}^*)$ is indeed bounded above by $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$.

Consider now the scenario $\hat{\mathbf{v}}$ defined through $\hat{v}_j^i = \frac{\bar{v}_j^*}{e}$ for all $i \in \mathcal{I}$ and $j \in \mathcal{J}$. In this scenario, we have $\sum_{j \in \mathcal{J}} \max_{i \in \mathcal{I}} \hat{v}_j^i = \sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$. Moreover, no bidder receives the good with a positive probability

and makes a non-zero payment because $\mathbb{Q}_j^*(\tilde{r}_j \leq \frac{\bar{v}_j^*}{e}) = 0$ for all $j \in \mathcal{J}$. In other words, the bidders' values for any item j remain almost surely below the reserve price \tilde{r}_j . The seller's regret in this scenario therefore attains the upper bound $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$. \square

Next, we show that $(\mathbf{q}^*, \mathbf{m}^*)$ is optimal in (MDP). The proof of this result establishes a lower bound on the optimal value of (MDP) that matches the worst-case regret of $(\mathbf{q}^*, \mathbf{m}^*)$.

THEOREM 4. *The separable mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ corresponding to the single-item mechanisms (6) is optimal in (MDP). The optimal value z^* of problem (MDP) amounts to $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$.*

In order to prove Theorem 4, we first introduce some useful notation, that is, we denote by

$$\mathcal{J}^i = \left\{ j \in \mathcal{J} : i = \min \arg \max_{i' \in \mathcal{I}} \bar{v}_j^{i'} \right\} \quad \text{and} \quad \hat{\mathcal{I}} = \{i \in \mathcal{I} : \mathcal{J}^i \neq \emptyset\}$$

the set of items for which bidder i has the smallest index among those who have the highest upper bound and the subset of bidders for whom \mathcal{J}^i is nonempty, respectively. By the construction of \mathcal{J}^i , we have $\bar{v}_j^i = \bar{v}_j^*$ for all $j \in \mathcal{J}^i$. Note that $\mathcal{J}^i, i \in \hat{\mathcal{I}}$, constitutes a partition of \mathcal{J} . We also introduce an extremal value profile $\hat{\mathbf{v}} \in \mathbb{R}^{I \times J}$ defined through

$$\hat{v}_j^i = \begin{cases} \bar{v}_j^i & \text{if } j \in \mathcal{J}^i, \\ 0 & \text{otherwise.} \end{cases}$$

Consider now the following discrete approximation of the mechanism design problem (MDP).

$$z_n^* = \inf_{\mathbf{q}, \mathbf{m}} \sup_{\mathbf{v} \in \mathcal{V}_n} \sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} v_j^i) - \sum_{i \in \mathcal{I}} m^i(\mathbf{v})$$

$$\text{s.t. } \mathbf{q} \in \mathcal{L}(\mathcal{V}_n, \mathbb{R}_+^{I \times J}), \mathbf{m} \in \mathcal{L}(\mathcal{V}_n, \mathbb{R}^I)$$

$$\sum_{j \in \mathcal{J}} q_j^i(\mathbf{v}) v_j^i - m^i(\mathbf{v}) \geq \sum_{j \in \mathcal{J}} q_j^i(\mathbf{w}^i, \mathbf{v}^{-i}) v_j^i - m^i(\mathbf{w}^i, \mathbf{v}^{-i}) \quad \forall i \in \mathcal{I}, \forall \mathbf{v} \in \mathcal{V}_n, \forall \mathbf{w}^i \in \mathcal{V}_n^i \quad (\text{IC}_n)$$

$$\sum_{j \in \mathcal{J}} q_j^i(\mathbf{v}) v_j^i - m^i(\mathbf{v}) \geq 0 \quad \forall i \in \mathcal{I}, \forall \mathbf{v} \in \mathcal{V}_n \quad (\text{IR}_n)$$

$$\sum_{i \in \mathcal{I}} q_j^i(\mathbf{v}) \leq 1 \quad \forall j \in \mathcal{J}, \forall \mathbf{v} \in \mathcal{V}_n \quad (\text{Inv}_n)$$

(8)

Problem (8) is parametrized in $n \in \mathbb{N}$ and differs from (MDP) only in that it involves a discrete uncertainty set $\mathcal{V}_n = \times_{i \in \mathcal{I}} \mathcal{V}_n^i$, where $\mathcal{V}_n^i = \times_{j \in \mathcal{J}} \mathcal{V}_{n,j}^i$ and $\mathcal{V}_{n,j}^i = \{0, \frac{1}{n} \hat{v}_j^i, \frac{2}{n} \hat{v}_j^i, \dots, \hat{v}_j^i\}$. For each $i \in \mathcal{I}$, $j \in \mathcal{J}$ and $n \in \mathbb{N}$, $\mathcal{V}_{n,j}^i$ represents a uniform one-dimensional grid within the interval $[0, \hat{v}_j^i]$. Particularly, we have $|\mathcal{V}_{n,j}^i| = n + 1$ if $j \in \mathcal{J}^i$, whereas $\mathcal{V}_{n,j}^i = \{0\}$ and $|\mathcal{V}_{n,j}^i| = 1$ if $j \notin \mathcal{J}^i$. Note also that any (scalar) function defined on \mathcal{V}_n corresponds to an $(n + 1)^J$ -dimensional vector.

LEMMA 1. *For any $n \in \mathbb{N}$, we have $z_n^* \leq z^*$.*

Proof. By construction we have that $\hat{v}_j^i \leq \bar{v}_j^i$, for all $i \in \mathcal{I}$, $j \in \mathcal{J}$, and as a result, $\mathcal{V}_n \subseteq \mathcal{V}$. The objective function of (8) is thus smaller than or equal to that of (MDP) uniformly across all \mathbf{q} and \mathbf{m} , and the feasible set of (8) contains that of (MDP) as it relaxes all constraints associated with value profiles $\mathbf{v} \in \mathcal{V} \setminus \mathcal{V}_n$. Thus, the optimal value of (8) cannot be larger than that of (MDP). \square

As its objective function is convex and piecewise linear, problem (8) can be reformulated as the following equivalent finite linear program, where r represents an auxiliary epigraphical variable.

$$\begin{aligned}
z_n^* &= \inf_{\mathbf{q}, \mathbf{m}, r} r \\
\text{s.t. } & \mathbf{q} \in \mathcal{L}(\mathcal{V}_n, \mathbb{R}_+^{I \times J}), \mathbf{m} \in \mathcal{L}(\mathcal{V}_n, \mathbb{R}^I), r \in \mathbb{R} \\
& r \geq \sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} v_j^i) - \sum_{i \in \mathcal{I}} m^i(\mathbf{v}) \quad \forall \mathbf{v} \in \mathcal{V}_n \\
& (\text{IC}_n), (\text{IR}_n), (\text{Inv}_n)
\end{aligned} \tag{9}$$

The linear program dual to (9) is given by

$$\begin{aligned}
z_n^* &= \sup_{\alpha, \gamma, \beta, \lambda} \sum_{j \in \mathcal{J}} \sum_{\mathbf{v} \in \mathcal{V}_n} \alpha(\mathbf{v}) (\max_{i \in \mathcal{I}} v_j^i) - \sum_{j \in \mathcal{J}} \sum_{\mathbf{v} \in \mathcal{V}_n} \lambda_j(\mathbf{v}) \\
\text{s.t. } & \alpha \in \mathcal{L}(\mathcal{V}_n, \mathbb{R}_+), \gamma^i \in \mathcal{L}(\mathcal{V}_n \times \mathcal{V}_n^i, \mathbb{R}_+) \quad \forall i \in \mathcal{I}, \beta \in \mathcal{L}(\mathcal{V}_n, \mathbb{R}_+^I), \lambda \in \mathcal{L}(\mathcal{V}_n, \mathbb{R}_+^J) \\
& \sum_{\mathbf{v} \in \mathcal{V}_n} \alpha(\mathbf{v}) = 1 \\
& \beta^i(\mathbf{v}) + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma^i(\mathbf{v}, \mathbf{w}^i) - \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i) = \alpha(\mathbf{v}) \quad \forall i \in \mathcal{I}, \forall \mathbf{v} \in \mathcal{V}_n \\
& \lambda_j(\mathbf{v}) \geq \beta^i(\mathbf{v}) v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma^i(\mathbf{v}, \mathbf{w}^i) v_j^i \\
& \quad - \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i) w_j^i \quad \forall i \in \mathcal{I}, \forall j \in \mathcal{J}, \forall \mathbf{v} \in \mathcal{V}_n,
\end{aligned} \tag{10}$$

where α represents the dual variable of the epigraphical constraint, γ and β are the dual variables of the incentive compatibility and individual rationality constraints, respectively, and λ collects the dual variables of the upper probability bounds in (9). Strong duality holds because the trivial mechanism that sets $\mathbf{q}(\mathbf{v}) = \mathbf{0}$ and $\mathbf{m}(\mathbf{v}) = \mathbf{0}$ for all $\mathbf{v} \in \mathcal{V}_n$ is feasible in (9) if $r \geq \sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} \hat{v}_j^i)$.

Since the linear program (10) seeks to make the decision variables $\lambda_j(\mathbf{v})$ as small as possible while ensuring that they remain non-negative and satisfy the last constraint of (10), it is clear that

$$\begin{aligned} \lambda_j(\mathbf{v}) &= \max_{i \in \mathcal{I}} \left(\beta^i(\mathbf{v})v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma^i(\mathbf{v}, \mathbf{w}^i)v_j^i - \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i)w_j^i \right)^+ \\ &= \max_{i \in \mathcal{I}} \left(\alpha(\mathbf{v})v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i)(v_j^i - w_j^i) \right)^+ \quad \forall j \in \mathcal{J}, \forall \mathbf{v} \in \mathcal{V}_n \end{aligned}$$

at optimality, where the second equality exploits the second equality constraint in (10) to eliminate $\beta^i(\mathbf{v})$. By substituting the above expression for $\lambda_j(\mathbf{v})$ into the objective of problem (10), we then obtain the following equivalent non-linear program in the decision variables α and γ .

$$\begin{aligned} z_n^* &= \sup_{\alpha, \gamma} \sum_{j \in \mathcal{J}} \sum_{\mathbf{v} \in \mathcal{V}_n} \alpha(\mathbf{v}) (\max_{i \in \mathcal{I}} v_j^i) - \sum_{j \in \mathcal{J}} \sum_{\mathbf{v} \in \mathcal{V}_n} \max_{i \in \mathcal{I}} \left(\alpha(\mathbf{v})v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i)(v_j^i - w_j^i) \right)^+ \\ \text{s.t. } &\alpha \in \mathcal{L}(\mathcal{V}_n, \mathbb{R}_+), \gamma^i \in \mathcal{L}(\mathcal{V}_n \times \mathcal{V}_n^i, \mathbb{R}_+) \quad \forall i \in \mathcal{I} \\ &\sum_{\mathbf{v} \in \mathcal{V}_n} \alpha(\mathbf{v}) = 1 \\ &\alpha(\mathbf{v}) + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i) - \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma^i(\mathbf{v}, \mathbf{w}^i) \geq 0 \quad \forall i \in \mathcal{I} \quad \forall \mathbf{v} \in \mathcal{V}_n \end{aligned} \tag{11}$$

The last constraint group in (11) is equivalent to the requirement that $\beta^i(\mathbf{v}) \geq 0$ for all $i \in \mathcal{I}$ and $\mathbf{v} \in \mathcal{V}_n$ in (10). Note that, by construction, the optimal value of (11) is still equal to z_n^* .

LEMMA 2. We have $\liminf_{n \rightarrow \infty} z_n^* \geq \sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$.

Proof. To prove the claim, we will construct a feasible solution for problem (11), whose objective function value necessarily provides a lower bound on z_n^* . We will then show that the objective function value of this solution converges to $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$ as n tends to infinity.

For $n > e(1+e)$, our construction of a feasible (α_n, γ_n) for problem (11) relies on bidder-specific values $\hat{\alpha}_n^i \in \mathcal{L}(\mathcal{V}_n^i, \mathbb{R}_+)$ and $\hat{\gamma}^i \in \mathcal{L}(\mathcal{V}_n^i \times \mathcal{V}_n^i, \mathbb{R}_+)$, $i \in \hat{\mathcal{I}}$, which are defined through

$$\hat{\alpha}_n^i(\mathbf{v}^i) = \begin{cases} \frac{n}{ek(k+1)} & \text{if } \exists k \in \{\lfloor \frac{n}{e} \rfloor, \dots, n-1\} \text{ with } \mathbf{v}^i = \frac{k}{n} \hat{\mathbf{v}}^i, \\ 1 - \sum_{k=\lfloor \frac{n}{e} \rfloor}^{n-1} \frac{n}{ek(k+1)} & \text{if } \mathbf{v}^i = \hat{\mathbf{v}}^i, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$\hat{\gamma}_n^i(\mathbf{w}^i, \mathbf{v}^i) = \begin{cases} \frac{(n-e-e^2)n}{e(n-e)(k+1)} & \text{if } \exists k \in \{\lfloor \frac{n}{e} \rfloor, \dots, n-1\} \text{ with } \mathbf{v}^i = \frac{k}{n} \hat{\mathbf{v}}^i \text{ and } \mathbf{w}^i = \mathbf{v}^i + \frac{1}{n} \hat{\mathbf{v}}^i, \\ 0 & \text{otherwise.} \end{cases}$$

We will first derive some useful properties of $(\hat{\alpha}_n^i, \hat{\gamma}_n^i)_{\{i \in \hat{\mathcal{I}}\}}$. Next, we will use $(\hat{\alpha}_n^i, \hat{\gamma}_n^i)_{\{i \in \hat{\mathcal{I}}\}}$ to construct a feasible solution (α_n, γ_n) for problem (11). To this end, note first that $(\hat{\alpha}_n^i, \hat{\gamma}_n^i)$ is feasible in

$$\begin{aligned} & \sup_{\hat{\alpha}^i, \hat{\gamma}^i} \sum_{j \in \mathcal{J}^i} \sum_{\mathbf{v}^i \in \mathcal{V}_n^i} \hat{\alpha}^i(\mathbf{v}^i) v_j^i - \sum_{j \in \mathcal{J}^i} \sum_{\mathbf{v}^i \in \mathcal{V}_n^i} \left(\hat{\alpha}^i(\mathbf{v}^i) v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \hat{\gamma}^i(\mathbf{w}^i, \mathbf{v}^i) (v_j^i - w_j^i) \right)^+ \\ & \text{s.t. } \hat{\alpha}^i \in \mathcal{L}(\mathcal{V}_n^i, \mathbb{R}_+), \hat{\gamma}^i \in \mathcal{L}(\mathcal{V}_n^i \times \mathcal{V}_n^i, \mathbb{R}_+) \\ & \sum_{\mathbf{v}^i \in \mathcal{V}_n^i} \hat{\alpha}^i(\mathbf{v}^i) = 1 \\ & \hat{\alpha}^i(\mathbf{v}^i) + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \hat{\gamma}^i(\mathbf{w}^i, \mathbf{v}^i) - \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \hat{\gamma}^i(\mathbf{v}^i, \mathbf{w}^i) \geq 0 \quad \forall \mathbf{v}^i \in \mathcal{V}_n^i. \end{aligned} \tag{12}$$

Problem (12) is equivalent to problem (7) in Koçyiğit et al. (2022), that is, it is equivalent to the dual of a single-buyer multi-item mechanism design problem for selling the items in \mathcal{J}^i to bidder $i \in \hat{\mathcal{I}}$. To see this, recall that $v_j^i = 0$ for all $j \notin \mathcal{J}^i$ and $\mathbf{v}^i \in \mathcal{V}_n^i$. Thus, we can assume without loss of generality that $\hat{\alpha}^i$ is only a function of $(v_j^i)_{\{j \in \mathcal{J}^i\}}$ and that $\hat{\gamma}^i$ is only a function of $((v_j^i)_{\{j \in \mathcal{J}^i\}}, (w_j^i)_{\{j \in \mathcal{J}^i\}})$. By (Koçyiğit et al. 2022, Lemma 3), we know that $(\hat{\alpha}_n^i, \hat{\gamma}_n^i)$ is feasible in (12) and that its objective function value converges to $\sum_{j \in \mathcal{J}^i} \frac{\hat{v}_j^*}{e} = \sum_{j \in \mathcal{J}^i} \frac{\bar{v}_j^*}{e}$ as n tends to infinity.

We now construct a feasible solution (α_n, γ_n) for problem (11), which is defined through

$$\begin{aligned} \alpha_n(\mathbf{v}) &= \prod_{i \in \hat{\mathcal{I}}} \hat{\alpha}_n^i(\mathbf{v}^i) \quad \forall \mathbf{v} \in \mathcal{V}_n \\ \gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i) &= \begin{cases} \prod_{i' \in \hat{\mathcal{I}} \setminus \{i\}} \hat{\alpha}_n^{i'}(\mathbf{v}^{i'}) \hat{\gamma}_n^i(\mathbf{w}^i, \mathbf{v}^i) & \forall i \in \hat{\mathcal{I}} \\ 0 & \forall i \in \mathcal{I} \setminus \hat{\mathcal{I}} \end{cases} \quad \forall \mathbf{v} \in \mathcal{V}_n, \forall \mathbf{w}^i \in \mathcal{V}_n^i. \end{aligned}$$

Next, we will show that (α_n, γ_n) is feasible in (11). Both α_n and γ_n are non-negative because $\hat{\alpha}_n^i$ and $\hat{\gamma}_n^i$ are non-negative for all $i \in \hat{\mathcal{I}}$. Since $\sum_{\mathbf{v}^i \in \mathcal{V}_n^i} \hat{\alpha}_n^i(\mathbf{v}^i) = 1$ for all $i \in \hat{\mathcal{I}}$ and $\mathcal{V}_n^i = \{\mathbf{0}\}$ for all $i \in \mathcal{I} \setminus \hat{\mathcal{I}}$, one can verify that $\sum_{\mathbf{v} \in \mathcal{V}_n} \alpha_n(\mathbf{v}) = 1$. It remains to be shown that $\alpha_n(\mathbf{v}) + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i) - \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma_n^i(\mathbf{v}, \mathbf{w}^i) \geq 0$ for all $i \in \mathcal{I}$ and $\mathbf{v} \in \mathcal{V}_n$. By the definitions of α_n and γ_n , it suffices to show that this inequality holds for any $i \in \hat{\mathcal{I}}$ when there exists an integer $k \in \{\lfloor \frac{n}{e} \rfloor, \dots, n\}$ with $\mathbf{v}^i = \frac{k}{n} \hat{\mathbf{v}}^i$. In fact, for any $i \in \mathcal{I} \setminus \hat{\mathcal{I}}$, the inequality trivially holds as $\alpha_n(\mathbf{v}) \geq 0$ and $\gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i) = \gamma_n^i(\mathbf{v}, \mathbf{w}^i) = 0$ for all $\mathbf{v} \in \mathcal{V}_n$ and $\mathbf{w}^i \in \mathcal{V}_n^i$. For any $i \in \hat{\mathcal{I}}$, on the other hand, if there does not exist a $k \in \{\lfloor \frac{n}{e} \rfloor, \dots, n\}$ with $\mathbf{v}^i = \frac{k}{n} \hat{\mathbf{v}}^i$, then the left-hand side of the inequality trivially evaluates to zero. Consider now any $i \in \hat{\mathcal{I}}$, and assume that there exists such k . When $k = \lfloor \frac{n}{e} \rfloor$, we have $\gamma_n^i(\mathbf{v}, \mathbf{w}^i) = 0$ for all $\mathbf{w}^i \in \mathcal{V}_n^i$ and $\mathbf{v}^{-i} \in \mathcal{V}_n^{-i}$, and the inequality holds because α_n and γ_n are non-negative. On the other hand, when $k = n$ or, equivalently, when $\mathbf{v}^i = \hat{\mathbf{v}}^i$ we have $\gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \hat{\mathbf{v}}^i) = 0$ for all $\mathbf{w}^i \in \mathcal{V}_n^i$ and $\mathbf{v}^{-i} \in \mathcal{V}_n^{-i}$ and $\gamma_n^i((\hat{\mathbf{v}}^i, \mathbf{v}^{-i}), \mathbf{w}^i) = 0$ for all $\mathbf{w}^i \in \mathcal{V}_n^i \setminus \{\hat{\mathbf{v}}^i - \hat{\mathbf{v}}^i/n\}$ and $\mathbf{v}^{-i} \in \mathcal{V}_n^{-i}$. Thus, we have

$$\begin{aligned}
& \alpha_n((\hat{\mathbf{v}}^i, \mathbf{v}^{-i})) + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \hat{\mathbf{v}}^i) - \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma_n^i((\hat{\mathbf{v}}^i, \mathbf{v}^{-i}), \mathbf{w}^i) \\
&= \alpha_n((\hat{\mathbf{v}}^i, \mathbf{v}^{-i})) - \gamma_n^i((\hat{\mathbf{v}}^i, \mathbf{v}^{-i}), \hat{\mathbf{v}}^i - \hat{\mathbf{v}}^i/n) \\
&= \alpha_n((\hat{\mathbf{v}}^i, \mathbf{v}^{-i})) - \prod_{i' \in \hat{\mathcal{I}} \setminus \{i\}} \hat{\alpha}_n^{i'}(\mathbf{v}^{i'}) \hat{\gamma}_n^i(\hat{\mathbf{v}}^i, \hat{\mathbf{v}}^i - \hat{\mathbf{v}}^i/n) \\
&= \left(\prod_{i' \in \hat{\mathcal{I}} \setminus \{i\}} \hat{\alpha}_n^{i'}(\mathbf{v}^{i'}) \right) (\hat{\alpha}_n^i(\hat{\mathbf{v}}^i) - \hat{\gamma}_n^i(\hat{\mathbf{v}}^i, \hat{\mathbf{v}}^i - \hat{\mathbf{v}}^i/n)) \\
&= \left(\prod_{i' \in \hat{\mathcal{I}} \setminus \{i\}} \hat{\alpha}_n^{i'}(\mathbf{v}^{i'}) \right) \left(\hat{\alpha}_n^i(\hat{\mathbf{v}}^i) + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \hat{\gamma}_n^i(\mathbf{w}^i, \hat{\mathbf{v}}^i) - \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \hat{\gamma}_n^i(\hat{\mathbf{v}}^i, \mathbf{w}^i) \right) \geq 0,
\end{aligned}$$

where the equalities follow from the definitions of α_n , γ_n^i , and $\hat{\gamma}_n^i$, and the inequality holds because $\hat{\alpha}_n$ is non-negative and because $(\hat{\alpha}_n^i, \hat{\gamma}_n^i)$ satisfies the last constraint of (12) whenever $i \in \hat{\mathcal{I}}$. Finally,

when $\lfloor \frac{n}{e} \rfloor + 1 \leq k \leq n-1$, we have

$$\begin{aligned}
& \alpha_n(\mathbf{v}) + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i) - \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma_n^i(\mathbf{v}, \mathbf{w}^i) \\
&= \alpha_n(\mathbf{v}) + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \left(\prod_{i' \in \hat{\mathcal{I}} \setminus \{i\}} \hat{\alpha}_n^{i'}(\mathbf{v}^{i'}) \right) \hat{\gamma}_n^i(\mathbf{w}^i, \mathbf{v}^i), \mathbf{v}^i) - \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \left(\prod_{i' \in \hat{\mathcal{I}} \setminus \{i\}} \hat{\alpha}_n^{i'}(\mathbf{v}^{i'}) \right) \hat{\gamma}_n^i(\mathbf{v}^i, \mathbf{w}^i) \\
&= \alpha_n(\mathbf{v}) + \prod_{i' \in \hat{\mathcal{I}} \setminus \{i\}} \hat{\alpha}_n^{i'}(\mathbf{v}^{i'}) \hat{\gamma}_n^i((\mathbf{v}^i + \hat{\mathbf{v}}^i/n, \mathbf{v}^i), \mathbf{v}^i) - \prod_{i' \in \hat{\mathcal{I}} \setminus \{i\}} \hat{\alpha}_n^{i'}(\mathbf{v}^{i'}) \hat{\gamma}_n^i(\mathbf{v}^i, \mathbf{v}^i - \hat{\mathbf{v}}^i/n) \\
&= \left(\prod_{i' \in \hat{\mathcal{I}} \setminus \{i\}} \hat{\alpha}_n^{i'}(\mathbf{v}^{i'}) \right) (\hat{\alpha}_n^i(\mathbf{v}^i) + \hat{\gamma}_n^i(\mathbf{v}^i + \hat{\mathbf{v}}^i/n, \mathbf{v}^i), \mathbf{v}^i) - \hat{\gamma}_n^i(\mathbf{v}^i, \mathbf{v}^i - \hat{\mathbf{v}}^i/n) \\
&= \left(\prod_{i' \in \hat{\mathcal{I}} \setminus \{i\}} \hat{\alpha}_n^{i'}(\mathbf{v}^{i'}) \right) \left(\hat{\alpha}_n^i(\mathbf{v}^i) + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \hat{\gamma}_n^i(\mathbf{w}^i, \mathbf{v}^i) - \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \hat{\gamma}_n^i(\mathbf{v}^i, \mathbf{w}^i) \right) \geq 0,
\end{aligned}$$

where the equalities follow again from the definitions of α_n , γ_n^i , and $\hat{\gamma}_n^i$, and the inequality holds similarly because $\hat{\alpha}_n$ is non-negative and because $(\hat{\alpha}_n^i, \hat{\gamma}_n^i)$ satisfies the last constraint of (12) whenever $i \in \hat{\mathcal{I}}$. In summary, we have thus shown that (α_n, γ_n) is feasible in (11).

We will now reformulate the objective function value of (α_n, γ_n) in terms of $(\hat{\alpha}_n^i, \hat{\gamma}_n^i)_{\{i \in \hat{\mathcal{I}}\}}$. This value can be expressed as $z^+ - z^-$, where

$$z^+ = \sum_{j \in \mathcal{J}} \sum_{\mathbf{v} \in \mathcal{V}_n} \alpha_n(\mathbf{v}) (\max_{i \in \mathcal{I}} v_j^i), \quad z^- = \sum_{j \in \mathcal{J}} \sum_{\mathbf{v} \in \mathcal{V}_n} \max_{i \in \mathcal{I}} \left(\alpha_n(\mathbf{v}) v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i) (v_j^i - w_j^i) \right)^+.$$

To reformulate z^+ in terms of $(\hat{\alpha}_n^i, \hat{\gamma}_n^i)_{\{i \in \hat{\mathcal{I}}\}}$, we observe that

$$z^+ = \sum_{i \in \hat{\mathcal{I}}} \sum_{j \in \mathcal{J}^i} \sum_{\mathbf{v} \in \mathcal{V}_n} \alpha_n(\mathbf{v}) v_j^i = \sum_{i \in \hat{\mathcal{I}}} \sum_{j \in \mathcal{J}^i} \sum_{\mathbf{v}^i \in \mathcal{V}_n^i} \left(\sum_{\mathbf{v}^{-i} \in \mathcal{V}_n^{-i}} \prod_{i' \in \hat{\mathcal{I}}} \hat{\alpha}_n^{i'}(\mathbf{v}^{i'}) \right) v_j^i = \sum_{i \in \hat{\mathcal{I}}} \sum_{j \in \mathcal{J}^i} \sum_{\mathbf{v}^i \in \mathcal{V}_n^i} \hat{\alpha}_n^i(\mathbf{v}^i) v_j^i,$$

where the first equality holds because \mathcal{J}^i , $i \in \hat{\mathcal{I}}$, is a partition of \mathcal{J} and because $v_j^i = 0$ for all $i \in \mathcal{I}$ such that $j \notin \mathcal{J}^i$, and the second equality follows from the definition of α_n .

To reformulate z^- in terms of $(\hat{\alpha}_n^i, \hat{\gamma}_n^i)_{\{i \in \hat{\mathcal{I}}\}}$, we first note that

$$\begin{aligned}
& \alpha_n(\mathbf{v}) v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i) (v_j^i - w_j^i) \\
&= \begin{cases} \alpha_n(\mathbf{v}) v_j^i - \gamma_n^i((\mathbf{v}^i + \hat{\mathbf{v}}^i/n, \mathbf{v}^{-i}), \mathbf{v}^i) \hat{v}_j^i/n & \text{if } j \in \mathcal{J}^i, \exists k \in \{\lfloor \frac{n}{e} \rfloor, \dots, n-1\} \text{ with } \mathbf{v}^i = \frac{k}{n} \hat{\mathbf{v}}^i, \\ \alpha_n(\mathbf{v}) v_j^i & \text{if } j \in \mathcal{J}^i, \mathbf{v}^i = \hat{\mathbf{v}}^i, \\ 0 & \text{otherwise,} \end{cases} \quad (13)
\end{aligned}$$

for all $i \in \hat{\mathcal{I}}$, whereas $\alpha_n(\mathbf{v})v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i)(v_j^i - w_j^i)$ evaluates to zero for all $i \in \mathcal{I} \setminus \hat{\mathcal{I}}$, $j \in \mathcal{J}$ and $\mathbf{v} \in \mathcal{V}_n$. We will next show that $\alpha_n(\mathbf{v})v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i)(v_j^i - w_j^i)$ is always non-negative. It suffices to show this for all $i \in \hat{\mathcal{I}}$, $j \in \mathcal{J}^i$ and $\mathbf{v} \in \mathcal{V}_n$ such that there exists $k \in \{\lfloor \frac{n}{e} \rfloor, \dots, n-1\}$ with $\mathbf{v}^i = \frac{k}{n} \hat{\mathbf{v}}^i$. In fact, in all other cases, $\alpha_n(\mathbf{v})v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i)(v_j^i - w_j^i)$ is non-negative because it is either equal to $\alpha_n(\mathbf{v})v_j^i$ or to zero and because α_n is non-negative. For any $i \in \hat{\mathcal{I}}$, $j \in \mathcal{J}^i$ and $\mathbf{v} \in \mathcal{V}_n$ such that there exists $k \in \{\lfloor \frac{n}{e} \rfloor, \dots, n-1\}$ with $\mathbf{v}^i = \frac{k}{n} \hat{\mathbf{v}}^i$, we have

$$\begin{aligned} \alpha_n(\mathbf{v})v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i)(v_j^i - w_j^i) &= \prod_{i' \in \hat{\mathcal{I}}} \hat{\alpha}_n^{i'}(\mathbf{v}^{i'})v_j^{i'} - \prod_{i' \in \hat{\mathcal{I}} \setminus \{i\}} \hat{\alpha}_n^{i'}(\mathbf{v}^{i'})\hat{\gamma}_n^i(\mathbf{v}^i + \hat{\mathbf{v}}^i/n, \mathbf{v}^i)\hat{v}_j^i/n \\ &= \prod_{i' \in \hat{\mathcal{I}} \setminus \{i\}} \hat{\alpha}_n^{i'}(\mathbf{v}^{i'}) (\hat{\alpha}_n^i(\mathbf{v}^i)v_j^i - \hat{\gamma}_n^i(\mathbf{v}^i + \hat{\mathbf{v}}^i/n, \mathbf{v}^i)\hat{v}_j^i/n), \end{aligned}$$

where the first and second equalities exploit the definitions of α_n , γ_n^i and $\hat{\gamma}_n^i$. It remains to be shown that $\hat{\alpha}_n^i(\mathbf{v}^i) - \hat{\gamma}_n^i(\mathbf{v}^i + \hat{\mathbf{v}}^i/n, \mathbf{v}^i)\hat{v}_j^i/n$ is non-negative, which will establish our claim. Using the definitions of $\hat{\alpha}_n^i$ and $\hat{\gamma}_n^i$, we obtain

$$\begin{aligned} \hat{\alpha}_n^i(\mathbf{v}^i)v_j^i - \hat{\gamma}_n^i(\mathbf{v}^i + \hat{\mathbf{v}}^i/n, \mathbf{v}^i)\hat{v}_j^i/n &= \frac{n}{ek(k+1)} \frac{k}{n} \hat{v}_j^i - \frac{(n-e-e^2)n}{e(n-e)(k+1)} \frac{1}{n} \hat{v}_j^i \\ &= \frac{1}{e(k+1)} \left(1 - \frac{(n-e-e^2)}{(n-e)} \right) \hat{v}_j^i = \frac{e}{(n-e)(k+1)} \hat{v}_j^i. \end{aligned}$$

As we selected $n > e(1+e) > e$, $\hat{\alpha}_n^i(\mathbf{v}^i)v_j^i - \hat{\gamma}_n^i(\mathbf{v}^i + \hat{\mathbf{v}}^i/n, \mathbf{v}^i)\hat{v}_j^i/n$ is necessarily non-negative. Using this observation and the fact that \mathcal{J}^i , $i \in \hat{\mathcal{I}}$, is a partition of \mathcal{J} , we can now reformulate z^- as

$$\begin{aligned} z^- &= \sum_{i \in \hat{\mathcal{I}}} \sum_{j \in \mathcal{J}^i} \sum_{\mathbf{v} \in \mathcal{V}_n} \max_{i' \in \mathcal{I}} \left(\alpha_n(\mathbf{v})v_j^{i'} + \sum_{\mathbf{w}^{i'} \in \mathcal{V}_n^{i'}} \gamma_n^{i'}((\mathbf{w}^{i'}, \mathbf{v}^{-i'}), \mathbf{v}^{i'}) (v_j^{i'} - w_j^{i'}) \right) \\ &= \sum_{i \in \hat{\mathcal{I}}} \sum_{j \in \mathcal{J}^i} \sum_{\mathbf{v}^i \in \mathcal{V}_n^i} \sum_{\mathbf{v}^{-i} \in \mathcal{V}_n^{-i}} \left(\alpha_n(\mathbf{v})v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i)(v_j^i - w_j^i) \right), \end{aligned}$$

where the second equality holds because $\alpha_n(\mathbf{v})v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i)(v_j^i - w_j^i)$ vanishes unless $j \in \mathcal{J}^i$. We next show that

$$\sum_{\mathbf{v}^{-i} \in \mathcal{V}_n^{-i}} \left(\alpha_n(\mathbf{v})v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i)(v_j^i - w_j^i) \right) = \hat{\alpha}^i(\mathbf{v}^i)v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \hat{\gamma}^i(\mathbf{w}^i, \mathbf{v}^i)(v_j^i - w_j^i) \quad (14)$$

for any $i \in \hat{\mathcal{I}}$, $j \in \mathcal{J}^i$ and $\mathbf{v}^i \in \mathcal{V}_n^i$. By definition of α_n , it is clear that $\sum_{\mathbf{v}^{-i} \in \mathcal{V}_n^{-i}} \alpha_n(\mathbf{v}) = \hat{\alpha}^i(\mathbf{v}^i)$. If there does not exist $k \in \{\lfloor \frac{n}{e} \rfloor, \dots, n-1\}$ with $\mathbf{v}^i = \frac{k}{n} \hat{\mathbf{v}}^i$, then $\gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i) = 0$ for all $\mathbf{w}^i \in \mathcal{V}_n^i$

and $\mathbf{v}^{-i} \in \mathcal{V}_n^{-i}$. Under the same condition, we similarly have that $\hat{\gamma}^i(\mathbf{w}^i, \mathbf{v}^i) = 0$ for all $\mathbf{w}^i \in \mathcal{V}_n^i$. Thus, we find $\sum_{\mathbf{v}^{-i} \in \mathcal{V}_n^{-i}} \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i)(v_j^i - w_j^i) = \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \hat{\gamma}^i(\mathbf{w}^i, \mathbf{v}^i)(v_j^i - w_j^i) = 0$. If there exists $k \in \{\lfloor \frac{n}{e} \rfloor, \dots, n-1\}$ with $\mathbf{v}^i = \frac{k}{n} \hat{\mathbf{v}}^i$, on the other hand, then

$$\begin{aligned} \sum_{\mathbf{v}^{-i} \in \mathcal{V}_n^{-i}} \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \gamma_n^i((\mathbf{w}^i, \mathbf{v}^{-i}), \mathbf{v}^i)(v_j^i - w_j^i) &= \sum_{\mathbf{v}^{-i} \in \mathcal{V}_n^{-i}} \left(\prod_{i' \in \hat{\mathcal{I}} \setminus \{i\}} \hat{\alpha}_n^{i'}(\mathbf{v}^{i'}) \right) \hat{\gamma}_n^i(\mathbf{v}^i + \hat{\mathbf{v}}^i/n, \mathbf{v}^i)(-\hat{v}_j^i/n) \\ &= \hat{\gamma}_n^i(\mathbf{v}^i + \hat{\mathbf{v}}^i/n, \mathbf{v}^i)(-\hat{v}_j^i/n) = \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \hat{\gamma}_n^i(\mathbf{w}^i, \mathbf{v}^i)(v_j^i - w_j^i), \end{aligned}$$

where the first equality follows from the definitions of γ_n^i and $\hat{\gamma}_n^i$, the second equality holds because $\hat{\alpha}_n^{i'}$ is non-negative and satisfies $\sum_{\mathbf{v}^{i'} \in \mathcal{V}_n^{i'}} \hat{\alpha}_n^{i'}(\mathbf{v}^{i'}) = 1$ for all $i' \in \hat{\mathcal{I}}$, and the last equality follows again from the definition of $\hat{\gamma}_n^i$. Thus, (14) follows. In summary, we have shown that

$$\begin{aligned} z^+ - z^- &= \sum_{i \in \hat{\mathcal{I}}} \sum_{j \in \mathcal{J}^i} \sum_{\mathbf{v}^i \in \mathcal{V}_n^i} \hat{\alpha}_n^i(\mathbf{v}^i) v_j^i - \sum_{i \in \hat{\mathcal{I}}} \sum_{j \in \mathcal{J}^i} \sum_{\mathbf{v}^i \in \mathcal{V}_n^i} \left(\hat{\alpha}_n^i(\mathbf{v}^i) v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \hat{\gamma}_n^i(\mathbf{w}^i, \mathbf{v}^i)(v_j^i - w_j^i) \right) \\ &= \sum_{i \in \hat{\mathcal{I}}} \sum_{j \in \mathcal{J}^i} \sum_{\mathbf{v}^i \in \mathcal{V}_n^i} \hat{\alpha}_n^i(\mathbf{v}^i) v_j^i - \sum_{i \in \hat{\mathcal{I}}} \sum_{j \in \mathcal{J}^i} \sum_{\mathbf{v}^i \in \mathcal{V}_n^i} \left(\hat{\alpha}_n^i(\mathbf{v}^i) v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \hat{\gamma}_n^i(\mathbf{w}^i, \mathbf{v}^i)(v_j^i - w_j^i) \right)^+, \end{aligned}$$

where the last equality follows from our insight that $\hat{\alpha}_n^i(\mathbf{v}^i) v_j^i + \sum_{\mathbf{w}^i \in \mathcal{V}_n^i} \hat{\gamma}_n^i(\mathbf{w}^i, \mathbf{v}^i)(v_j^i - w_j^i)$ is always non-negative. Thus, the objective function value of (α_n, γ_n) in (11) equals the sum of the objective function values of $(\hat{\alpha}_n^i, \hat{\gamma}_n^i)_{\{i \in \hat{\mathcal{I}}\}}$ in (12) over all $i \in \hat{\mathcal{I}}$. As the objective function value of $(\hat{\alpha}_n^i, \hat{\gamma}_n^i)$ in (12) converges to $\sum_{j \in \mathcal{J}^i} \frac{\hat{v}_j^*}{e} = \sum_{j \in \mathcal{J}^i} \frac{\bar{v}_j^*}{e}$, the objective function value of (α_n, γ_n) in (11) converges to $\sum_{i \in \hat{\mathcal{I}}} \sum_{j \in \mathcal{J}^i} \frac{\hat{v}_j^*}{e} = \sum_{j \in \mathcal{J}} \frac{\hat{v}_j^*}{e} = \sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$ as n tends to infinity. \square

Proof of Theorem 4. By Remark 2 and Proposition 1, the mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ is feasible in (MDP), and its objective function value is given by $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$. By Lemma 2, the optimal value z^* of (MDP) is bounded below by $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$. As the objective function value $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$ of $(\mathbf{q}^*, \mathbf{m}^*)$ is equal to this lower bound, we have $z^* = \sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$, and thus $(\mathbf{q}^*, \mathbf{m}^*)$ is optimal in (MDP). \square

REMARK 3. When the bidders are symmetric in the sense that $\bar{v}_j^i = \bar{v}_j^*$ for all $i \in \mathcal{I}$ and $j \in \mathcal{J}$, there is a simple way to establish the lower bound $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$, which matches the worst-case regret of $(\mathbf{q}^*, \mathbf{m}^*)$, for the optimal value of (MDP). Fix an arbitrary feasible mechanism $(\mathbf{q}, \mathbf{m}) \in \mathcal{X}$ and an arbitrary bidder $k \in \mathcal{I}$. In addition, define $\mathcal{V}(k) = \{\mathbf{v} \in \mathcal{V} : \mathbf{v}^i = \mathbf{0} \ \forall i \in \mathcal{I} \setminus \{k\}\}$ as the set of all

value profiles under which the bidders $i \in \mathcal{I} \setminus \{k\}$ assign no value to any item. The worst-case regret of (\mathbf{q}, \mathbf{m}) thus satisfies

$$\sup_{\mathbf{v} \in \mathcal{V}} \sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} v_j^i) - \sum_{i \in \mathcal{I}} m^i(\mathbf{v}) \geq \sup_{\mathbf{v} \in \mathcal{V}(k)} \sum_{j \in \mathcal{J}} v_j^k - \sum_{i \in \mathcal{I}} m^i(\mathbf{v}) \geq \sup_{\mathbf{v}^k \in \mathcal{V}^k} \sum_{j \in \mathcal{J}} v_j^k - m^k(\mathbf{v}^k, \mathbf{0}), \quad (15)$$

where the first inequality holds because $\mathcal{V}(k) \subseteq \mathcal{V}$ and $\max_{i \in \mathcal{I}} v_j^i \geq v_j^k$ for every $j \in \mathcal{J}$, while the second inequality holds due to (IR), which implies that $m^i(\mathbf{v}) \leq 0$ for all $i \neq k$ and $\mathbf{v} \in \mathcal{V}(k)$. Note that $(\mathbf{q}^k(\mathbf{v}^k, \mathbf{0}), m^k(\mathbf{v}^k, \mathbf{0}))$ can be interpreted as an incentive compatible and individually rational single-bidder mechanism for selling the items to bidder k . From Theorem 1, we know that the worst-case regret of this single-buyer mechanism is bounded below by $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^k}{e}$. In view of (15), the worst-case regret of (\mathbf{q}, \mathbf{m}) is therefore also bounded below by $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^k}{e}$. In fact, this is true for every $k \in \mathcal{I}$, and thus $\max_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \frac{\bar{v}_j^i}{e}$ constitutes a lower bound on the worst-case regret of *any* mechanism. By Proposition 1, $(\mathbf{q}^*, \mathbf{m}^*)$ attains a worst-case regret of $\sum_{j \in \mathcal{J}} \max_{i \in \mathcal{I}} \frac{\bar{v}_j^i}{e}$, which matches the simple lower bound if the bidders are symmetric or if there is only $J = 1$ item. If the bidders are asymmetric and there are $J \geq 2$ items, however, the lower bound is loose. Hence, this simple reasoning cannot be used to prove the optimality of $(\mathbf{q}^*, \mathbf{m}^*)$ in general. The complexity thus originates from the asymmetry of the bidders provided that there are multiple items. \square

Note that problem (MDP) can be interpreted as a zero-sum game between the seller, who chooses the mechanism (\mathbf{q}, \mathbf{m}) , and some fictitious adversary or ‘nature,’ who chooses the bidders’ value profiles \mathbf{v} with the goal to inflict maximum damage to the seller. In the remainder of this section, we assume that the bidders are symmetric and show that (MDP) admits a Nash equilibrium in mixed strategies, which can be computed analytically. To this end, note first that the mechanisms in the seller’s convex feasible set \mathcal{X} can be seen as mixed strategies because the allocation rules are probabilistic. Conversely, the value profiles in nature’s convex box uncertainty set $\mathcal{V} = (\mathcal{V}^1)^I$ with $\mathcal{V}^1 = \times_{j \in \mathcal{J}} [0, \bar{v}_j^*]$ represent pure strategies. While the objective function of problem (MDP) is affine and thus convex in the payment rule \mathbf{m} for a fixed scenario $\mathbf{v} \in \mathcal{V}$, it is generically non-concave in \mathbf{v} for a fixed mechanism $(\mathbf{q}, \mathbf{m}) \in \mathcal{X}$. Thus, nature’s decision problem is non-convex. To convert the zero-sum game (MDP) to an equivalent convex-concave saddle point problem, we

allow nature to play mixed strategies corresponding to distributions $\mathbb{P} \in \Delta(\mathcal{V})$. With this standard trick, problem (MDP) can be reformulated as

$$\inf_{(\mathbf{q}, \mathbf{m}) \in \mathcal{X}} \sup_{\mathbb{P} \in \Delta(\mathcal{V})} \mathbb{E}_{\mathbb{P}} \left[\sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} \tilde{v}_j^i) - \sum_{i \in \mathcal{I}} m^i(\tilde{\mathbf{v}}) \right]. \quad (16)$$

Next, we prove that the seller's Nash strategy is given by the separable mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ corresponding to the single-item mechanisms (6). In order to describe nature's Nash strategy, we first introduce a marginal distribution $\hat{\mathbb{P}} \in \Delta(\mathcal{V}^1)$ of the value profile $\tilde{\mathbf{v}}^1$, which we define through

$$\hat{\mathbb{P}}(\tilde{\mathbf{v}}^1 \leq \mathbf{v}^1) = \begin{cases} \min_{j \in \mathcal{J}} \left(1 - \frac{1}{e} \left(\frac{\bar{v}_j^*}{v_j^1} \right) \right)^+ & \text{if } \mathbf{v}^1 \in \mathcal{V}^1 \setminus \{(\bar{v}_1^*, \dots, \bar{v}_J^*)\} \\ 1 & \text{if } \mathbf{v}^1 = (\bar{v}_1^*, \dots, \bar{v}_J^*). \end{cases}$$

By Theorem 2, the distribution $\hat{\mathbb{P}}$ represents nature's Nash strategy in problem (16) if there is only $I = 1$ bidder, that is, in the special case where the auction design problem collapses to a monopoly pricing problem. Note also that the values $(\tilde{v}_1^1, \dots, \tilde{v}_J^1)$ of bidder 1 for the different items are comonotonic under $\hat{\mathbb{P}}$. In the following, we will use $\hat{\mathbb{P}}$ to construct an explicit Nash strategy for nature when $I > 1$. To this end, we introduce for each $i \in \mathcal{I}$ a probability distribution $\hat{\mathbb{P}}^i \in \Delta(\mathcal{V})$ of the random matrix $\tilde{\mathbf{v}}$, which is uniquely determined through the relations

$$\hat{\mathbb{P}}^i(\tilde{\mathbf{v}}^i \leq \mathbf{w}) = \hat{\mathbb{P}}(\tilde{\mathbf{v}}^1 \leq \mathbf{w}) \quad \forall \mathbf{w} \in \mathcal{V}^i \quad \text{and} \quad \hat{\mathbb{P}}^i(\tilde{\mathbf{v}}^k = \mathbf{0}) = 1 \quad \forall k \neq i.$$

Note that under $\hat{\mathbb{P}}^i$, the marginal distribution of the value profile $\tilde{\mathbf{v}}^i$ coincides with $\hat{\mathbb{P}}$, while the marginal distributions of the value profiles $\tilde{\mathbf{v}}^k$, $k \neq i$, are all equal to the Dirac distribution that concentrates unit mass at $\mathbf{0}$. As $\hat{\mathbb{P}}$ constitutes a comonotonic distribution, the values of bidder i for the items are comonotonic under $\hat{\mathbb{P}}^i$. The support of the distribution $\hat{\mathbb{P}}^i$ is given by

$$\text{supp}(\hat{\mathbb{P}}^i) = \left\{ \mathbf{v} \in \mathcal{V} : \mathbf{v}^i = s \bar{\mathbf{v}}^* \text{ for some } s \in \left[\frac{1}{e}, 1 \right] \text{ and } \mathbf{v}^k = \mathbf{0} \quad \forall k \neq i \right\},$$

where $\bar{\mathbf{v}}^* = (\bar{v}_1^*, \dots, \bar{v}_J^*)$. Finally, define $\mathbb{P}^* \in \Delta(\mathcal{V})$ as the average of $\hat{\mathbb{P}}^i$, $i \in \mathcal{I}$, that is, set

$$\mathbb{P}^* = \frac{1}{I} \sum_{i \in \mathcal{I}} \hat{\mathbb{P}}^i. \quad (17)$$

Note that this definition of \mathbb{P}^* naturally extends definition (3) from $I = 1$ to $I \geq 1$. By construction, it is easy to verify that the support of \mathbb{P}^* can be represented as $\text{supp}(\mathbb{P}^*) = \cup_{i \in \mathcal{I}} \text{supp}(\hat{\mathbb{P}}^i)$. In addition, under \mathbb{P}^* the highest bidder's value profile exceeds the positive threshold $\frac{1}{e}\bar{\mathbf{v}}^*$ almost surely, while all other bidders' value profiles are almost surely equal to $\mathbf{0}$. The following theorem asserts that the separable mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ corresponding to the single-item mechanisms (6) and the probability distribution \mathbb{P}^* defined in (17) represent the Nash strategies of the seller and of nature in the zero-sum game (16), respectively. To simplify the subsequent discussion, we denote by

$$z(\mathbf{m}, \mathbb{P}) = \mathbb{E}_{\mathbb{P}} \left[\sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} \tilde{v}_j^i) - \sum_{i \in \mathcal{I}} m^i(\tilde{\mathbf{v}}) \right]$$

the expected regret of the feasible mechanism (\mathbf{q}, \mathbf{m}) under the probability distribution $\mathbb{P} \in \Delta(\mathcal{V})$.

THEOREM 5 (Nash Equilibrium). *If the bidders are symmetric, then the separable mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ corresponding to the single-item mechanisms (6) and the distribution \mathbb{P}^* defined in (17) satisfy the saddle point condition*

$$\max_{\mathbb{P} \in \Delta(\mathcal{V})} z(\mathbf{m}^*, \mathbb{P}) \leq z(\mathbf{m}^*, \mathbb{P}^*) \leq \min_{(\mathbf{q}, \mathbf{m}) \in \mathcal{X}} z(\mathbf{m}, \mathbb{P}^*). \quad (18)$$

To prove Theorem 5, we first show that the maximization problem on the left-hand side of (18) is solved by \mathbb{P}^* and attains an optimal value of $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$. Next, we relax the minimization problem on the right-hand side of (18) to a single-buyer multi-item pricing problem with the objective of minimizing the regret under the distribution $\hat{\mathbb{P}}$. This relaxation is facilitated by the symmetric construction of the distribution \mathbb{P}^* . We know from Theorem 2 that the optimal value of the resulting pricing problem amounts to $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$. The observation that $z(\mathbf{m}^*, \mathbb{P}^*) = \sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$ then completes the proof. All technical details are relegated to the online appendix.

5. Deterministic Mechanisms

We now address the deterministic mechanism design problem obtained from (MDP) by restricting attention to deterministic mechanisms with $\mathbf{q} \in \mathcal{L}(\mathcal{V}, \{0, 1\}^{I \times J})$. To this end, we study the separable mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ consisting of the single-item mechanisms $(\hat{q}_j, \hat{\mathbf{m}}_j)$, $j \in \mathcal{J}$, defined through

$$\hat{q}_j^i(\mathbf{v}_j) = \begin{cases} 1 & \text{if } \mathbf{v}_j \in \mathcal{W}_j^i \text{ and } v_j^i \geq \frac{\bar{v}_j^*}{2} \\ 0 & \text{otherwise} \end{cases} \quad (19a)$$

and

$$\hat{m}_j^i(\mathbf{v}_j) = \begin{cases} \max \left\{ \max_{k \neq i} v_j^k, \frac{\bar{v}_j^*}{2} \right\} & \text{if } \mathbf{v}_j \in \mathcal{W}_j^i \text{ and } v_j^i \geq \frac{\bar{v}_j^*}{2} \\ 0 & \text{otherwise} \end{cases} \quad (19b)$$

for all $i \in \mathcal{I}$ and $\mathbf{v}_j \in \mathcal{V}_j$. Note that the single-item mechanism $(\hat{\mathbf{q}}_j, \hat{\mathbf{m}}_j)$ represents a second price auction with a deterministic reserve price $\frac{\bar{v}_j^*}{2}$ and is therefore both incentive compatible and individually rational (Krishna 2009, Chapter 2). This readily implies that the separable mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ is incentive compatible and individually rational. As $(\mathbf{q}^*, \mathbf{m}^*)$ trivially satisfies the inventory constraint, it is thus feasible in the deterministic mechanism design problem at hand.

PROPOSITION 2. *The worst-case regret of the separable mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ corresponding to the single-item mechanisms (19) is given by $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{2}$.*

Proof. As in the proof of Proposition 1, we will first show that the regret of $(\mathbf{q}^*, \mathbf{m}^*)$ is bounded above by $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{2}$ in any scenario $\mathbf{v} \in \mathcal{V}$ and then identify a scenario that attains this bound.

To this end, we introduce again an artificial bidder with value profile $\mathbf{v}^a \in \mathcal{V}^a = \times_{j \in \mathcal{J}} [0, \bar{v}_j^*]$, and for any $j \in \mathcal{J}$, we denote by (f_j^a, g_j^a) the posted-price mechanism with deterministic posted price $\frac{\bar{v}_j^*}{2}$ for selling item j to the artificial bidder. By Theorem 3, the separable single-buyer multi-item mechanism corresponding to the single-item single-buyer mechanisms $\{(f_j^a, g_j^a)\}_{j \in \mathcal{J}}$ is known to be optimal in the deterministic version of problem (MDP) if there is only the artificial bidder.

Fix now any scenario $\mathbf{v} \in \mathcal{V}$, and define $i^*(j)$ as the unique $i \in \mathcal{I}$ with $\mathbf{v}_j \in \mathcal{W}_j^i$. By construction, $i^*(j) \in \mathcal{I}$ identifies one of the highest bidders for item j in scenario \mathbf{v} . As the mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ sells each item j separately to bidder $i^*(j)$, its regret in scenario \mathbf{v} can thus be expressed as

$$\sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} v_j^i) - \sum_{i \in \mathcal{I}} \hat{m}^i(\mathbf{v}) = \sum_{j \in \mathcal{J}} (v^{i^*(j)} - \hat{m}_j^{i^*(j)}(\mathbf{v}_j)) \leq \sum_{j \in \mathcal{J}} (v^{i^*(j)} - g_j^a(v_j^{i^*(j)})) \leq \sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{2},$$

where the second inequality exploits Theorem 3. Further details are omitted because the arguments used here widely parallel those used in the proof of Proposition 1.

For any sufficiently small $\epsilon > 0$ we now define a scenario $\hat{\mathbf{v}}(\epsilon)$ through $\hat{v}_j^i(\epsilon) = \frac{\bar{v}_j^*}{2} - \epsilon$ for all $i \in \mathcal{I}$ and $j \in \mathcal{J}$. In this scenario, we have $\sum_{j \in \mathcal{J}} \max_{i \in \mathcal{I}} \hat{v}_j^i(\epsilon) = \sum_{j \in \mathcal{J}} (\frac{\bar{v}_j^*}{2} - \epsilon)$. Moreover, in this scenario

the seller keeps all items to herself, and no bidder makes a payment because the reserve price $\frac{\bar{v}_j^*}{2}$ for item j exceeds all bidders' values for item j irrespective of $j \in \mathcal{J}$. The seller's regret in scenario $\hat{v}(\epsilon)$ thus evaluates to $\sum_{j \in \mathcal{J}} (\frac{\bar{v}_j^*}{2} - \epsilon)$. Note that all of these arguments hold for all sufficiently small $\epsilon > 0$. The seller's regret thus attains the upper bound $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{2}$ asymptotically as ϵ tends to 0. \square

PROPOSITION 3. *The optimal value z^* of the deterministic version of the mechanism design problem (MDP) is bounded below by $\max_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \frac{\bar{v}_j^i}{2}$.*

Proof. The proof widely parallels to the arguments in Remark 3. Specifically, one can show that the worst-case regret of *any* mechanism for auctioning the items to I bidders is bounded below by the worst-case regret of the *best* mechanism for selling the items to only one bidder. The claim then follows from Theorem 3. Details are omitted for brevity. \square

Propositions 2 and 3 imply that $(\mathbf{q}^*, \mathbf{m}^*)$ attains a worst-case regret of $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{2} = \sum_{j \in \mathcal{J}} \max_{i \in \mathcal{I}} \frac{\bar{v}_j^i}{2}$ and that no mechanism can attain a worst-case regret lower than $\max_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \frac{\bar{v}_j^i}{2}$, which implies that

$$\text{Reg}(\mathbf{q}^*, \mathbf{m}^*) - z_d^* \leq \sum_{j \in \mathcal{J}} \max_{i \in \mathcal{I}} \frac{\bar{v}_j^i}{2} - \max_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \frac{\bar{v}_j^i}{2}, \quad (20)$$

where z_d^* denotes the optimal value of the deterministic version of the mechanism design problem (MDP). The right-hand side of the inequality (20) can be evaluated ex ante and provides an upper bound on the suboptimality of $(\mathbf{q}^*, \mathbf{m}^*)$. Note that this upper bound (as well as the suboptimality gap it majorizes) collapse to zero if the bidders are symmetric, in which case the sums and the maxima in (20) can be interchanged. Propositions 2 and 3 thus imply the following main result.

COROLLARY 1. *If the bidders are symmetric, then the minimum of the deterministic version of the mechanism design problem (MDP) amounts to $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{2}$, and the separable mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ corresponding to the single-item mechanisms (19) is optimal.*

6. Numerical Results

We now design numerical experiments to compare the separable mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ corresponding to the single-item mechanisms (6) against several popular benchmark mechanisms. We know from Theorem 4 that $(\mathbf{q}^*, \mathbf{m}^*)$ is optimal in terms of worst-case regret. Throughout this section we assume, however, that the bidders' values constitute a random vector $\tilde{\mathbf{v}}$ that follows some (unknown) distribution in $\Delta(\mathcal{V})$. Under this assumption, both the revenue as well as the regret generated by any mechanism constitute random variables. We can thus compare different mechanisms in terms of their expected revenues and their regret quantiles, for example.

Throughout this section we assume that the bidders are symmetric and that there is only one item. This allows us to (approximately) compute the expected revenue-maximizing mechanism if I is small. Indeed, recall that computing the expected revenue-maximizing mechanism for multiple items is $\#P$ -hard even under unrealistically simple value distributions (Daskalakis et al. 2014). We also assume that the distribution of $\tilde{\mathbf{v}}$ is representable as $\mathbb{P}_\epsilon = (1 - \epsilon)\mathbb{P}^N + \epsilon\mathbb{P}^*$, where $\mathbb{P}^N \in \Delta(\mathcal{V})$ represents a nominal distribution that can be thought of as the seller's best guess for \mathbb{P}_ϵ , while \mathbb{P}^* is defined as in (17), and $\epsilon \in [0, 1]$. Recall from Theorem 5 that since the bidders are symmetric, \mathbb{P}^* represents nature's Nash strategy. It is useful to interpret \mathbb{P}_ϵ as an ϵ -contamination of \mathbb{P}^N . Hence, the larger ϵ , the more the data-generating distribution \mathbb{P}_ϵ differs from the seller's best guess \mathbb{P}^N . We assume that the bidders' values are independent and identically distributed (i.i.d.) under \mathbb{P}^N .

We compare the mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ against the following benchmarks: (1) the nominal mechanism, which maximizes the expected revenue under \mathbb{P}^N , (2) the single-sample mechanism by Dhangwatnotai et al. (2015), and (3) the classical second price auction without reserve price. The single-sample mechanism is a second price auction with a random reserve price sampled from the true value distribution of any bidder (note that these distributions are identical because the bidders are symmetric). If the bidders' values are i.i.d. and a mild regularity condition holds, then the single-sample mechanism is known to generate at least half of the maximum expected revenue that can be generated by *any* feasible mechanism (Dhangwatnotai et al. 2015). The second

price auction without reserve price is also a competitive benchmark if the number I of bidders is large and the bidders' values are i.i.d. because it is known to extract at least as much revenue as the expected revenue-maximizing mechanism for $I - 1$ bidders (Bulow and Klemperer 1996). As the bidders' values are i.i.d. under \mathbb{P}^N , the single-sample mechanism and the second price auction without reserve price are thus at an advantage for small contamination levels ϵ .

In all experiments we assume that \mathbb{P}^N represents a truncated normal distribution on $\mathcal{V} = [0, 1]^I$. As normal distributions are uniquely determined by their first two moments, they are relatively easy to estimate from data. We emphasize, however, that our mechanism design problem (MDP) targets situations in which there is limited distributional information. Consequently, the seller may not have access to any statistical data and may thus be unable to evaluate the first two moments of \mathbb{P}^N altogether. In such situations, the nominal mechanism is fundamentally unavailable. In addition, even if \mathbb{P}^N is known, the nominal mechanism cannot be computed exactly. To compute the nominal mechanism approximately, we approximate the set \mathcal{V} of all value profiles by an equidistant grid $\hat{\mathcal{V}} \subseteq \mathcal{V}$ and approximate \mathbb{P}^N by a discrete distribution $\hat{\mathbb{P}}^N$ supported on $\hat{\mathcal{V}}$. The nominal mechanism is then approximated by the solution of the linear program that maximizes the expected revenues under $\hat{\mathbb{P}}^N$ across all mechanisms that satisfy (IC), (IR) and (Inv) only on $\hat{\mathcal{V}}$ instead of \mathcal{V} . Thus, the approximate nominal mechanism is at an advantage *vis-à-vis* $(\mathbf{q}^*, \mathbf{m}^*)$, which satisfies (IC), (IR) and (Inv) throughout the entire set \mathcal{V} of value profiles. More details on the setup of our experiments and the construction of the nominal mechanism are provided in the online appendix.

The first experiment focuses on $I = 2$ bidders. Under \mathbb{P}^N the bidders' values follow independent normal distributions with mean 0.5 and variance 0.1, truncated on the interval $[0, 1]$. In this case, the worst-case regrets of the three benchmark mechanisms (approximately computed with respect to the grid $\hat{\mathcal{V}}$) amount to 0.65, 0.65, and 1, respectively. As expected, they exceed the worst-case regret of $(\mathbf{q}^*, \mathbf{m}^*)$, which amounts to 0.37 (for consistency also computed with respect to $\hat{\mathcal{V}}$). The left plot of Figure 2 shows the 75th percentile of the random regret of all mechanisms under $\hat{\mathbb{P}}_\epsilon$ for different contamination levels ϵ . Here, $\hat{\mathbb{P}}_\epsilon$ represents a discrete approximation of \mathbb{P}_ϵ supported on $\hat{\mathcal{V}}$.

We observe that $(\mathbf{q}^*, \mathbf{m}^*)$ dominates all benchmarks across all $\epsilon \geq 0.25$. Furthermore, the 75th regret percentile of $(\mathbf{q}^*, \mathbf{m}^*)$ is immune to contaminations of the nominal distribution, while those of the three benchmark mechanisms quickly deteriorate as ϵ increases. The right plot of Figure 2 shows the expected revenues of all mechanisms relative to the maximum expected revenues achievable under $\hat{\mathbb{P}}_\epsilon$ for different contamination levels ϵ . Although maximizing expected revenues is not the objective of problem (MDP), the optimal mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ compares favorably against the three benchmark mechanisms in terms of expected revenues. It is even competitive with the expected revenue-maximizing mechanism under the unknown data-generating distribution $\hat{\mathbb{P}}_\epsilon$. Specifically, $(\mathbf{q}^*, \mathbf{m}^*)$ extracts more than 70% of the optimal expected revenues when $\epsilon = 0$, and this percentage further grows as ϵ increases. Most importantly, $(\mathbf{q}^*, \mathbf{m}^*)$ dominates the single-sample mechanism as well as the second price auction without reserve price for all $\epsilon \geq 0.25$, which are the only two benchmark mechanisms that are realistically implementable.

Figure 3 shows how the results of the above experiment change if the means of the bidder values are set to 0.1 instead of 0.5, while all other parameters are kept unchanged. The left plot reveals that $(\mathbf{q}^*, \mathbf{m}^*)$ dominates the three benchmarks in terms of the 75th regret percentile across all contamination levels $\epsilon \geq 0.3$. The outperformance becomes more pronounced for large values of ϵ . The right plot shows that $(\mathbf{q}^*, \mathbf{m}^*)$ dominates the practically implementable single-sample mechanism and the second price auction without reserve price for all $\epsilon \geq 0.3$ and that it surpasses the nominal mechanism for all $\epsilon \geq 0.55$ in terms of expected revenues. Results for an analogous third experiment, where the bidders' values have means 0.9, are relegated to the online appendix.

The last experiment focuses on $I = 10$ and $I = 100$ bidders. We use the same procedure as in the experiment underlying Figure 2 to construct the nominal distribution \mathbb{P}^N , the data-generating distribution \mathbb{P}_ϵ as well as the respective grid-based approximations $\hat{\mathbb{P}}^N$ and $\hat{\mathbb{P}}_\epsilon$. In this case the nominal mechanisms with respect to $\hat{\mathbb{P}}^N$ and $\hat{\mathbb{P}}_\epsilon$ can no longer be computed because the underlying linear programs suffer from a curse of dimensionality. Figure 4 shows that $(\mathbf{q}^*, \mathbf{m}^*)$ dominates all computable benchmarks in terms of expected revenues for virtually every contamination level ϵ .




Figure 2 Comparison of the minimax regret mechanism (q^*, m^*) against the three benchmarks for $I = 2$ bidders, where \mathbb{P}^N is constructed from independent normal distributions with mean values 0.5 and variances 0.1.

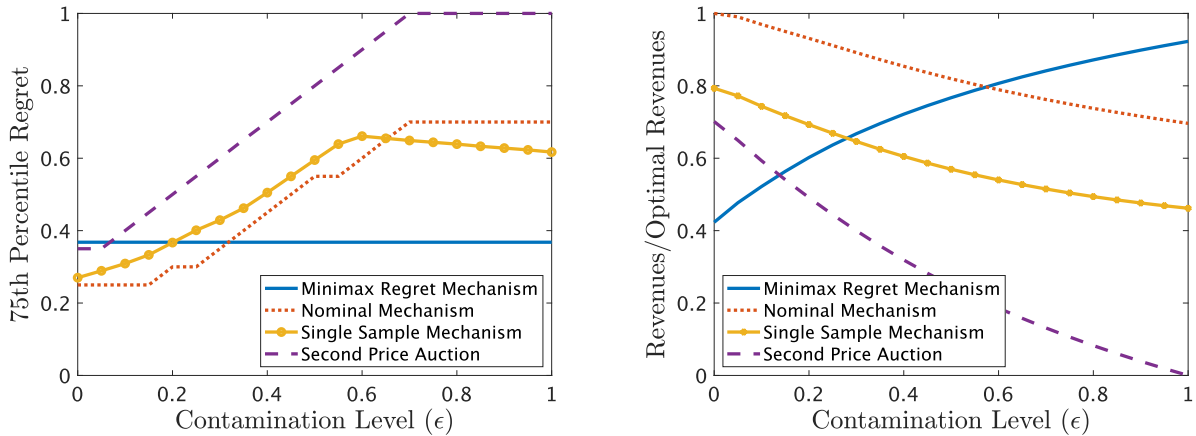


Figure 3 Comparison of the minimax regret mechanism (q^*, m^*) against the three benchmarks for $I = 2$ bidders, where \mathbb{P}^N is constructed from independent normal distributions with mean values 0.1 and variances 0.1.

In summary, our numerical results demonstrate that the minimax regret mechanism (q^*, m^*) not only attains the least worst-case regret but also compares favorably against practical benchmark mechanisms in terms of regret percentiles and expected revenues. The online appendix reports on additional experiments without the assumption that the bidders' values are independent under \mathbb{P}^N .

7. Concluding Remarks

We studied a robust auction design problem with minimax regret objective involving multiple items as well as multiple bidders, where the seller only knows that the values of the bidders range over a




Figure 4 Comparison of the minimax regret mechanism (q^*, m^*) against the two computable benchmarks for $I = 10$ bidders (left) and $I = 100$ bidders (right), where \mathbb{P}^N is constructed from independent normal distributions with mean values 0.5 and variances 0.1.

box uncertainty set. By using duality techniques, we show that a separable mechanism, which sells the items through separate second price auctions with random reserve prices, is optimal. Reinterpreting the auction design problem as a zero-sum game between the seller and nature allows us to prove the existence of a Nash equilibrium in mixed strategies, which we can characterize in close form when the bidders are symmetric. When restricting attention to deterministic mechanisms, we further prove that a separable posted-price mechanism is optimal provided that the bidders are symmetric. We conjecture, however, that separation remains optimal in much more general situations, and we hope that this paper will spur further research on robust auction design problems under different informational assumptions.

Acknowledgments

This research was funded by the SNSF grant BSCGI0.157733 and by the Ministry of Education, Singapore, under its Academic Research Fund Tier 2 grant MOE-T2EP20222-0003.

Appendix

A. Proofs

Proof of Theorem 5. We first show that \mathbb{P}^* solves the problem on the left-hand side of (17) (Step 1), and then we prove that (q^*, m^*) solves the problem on the right-hand side of (17) (Step 2).

Step 1. By Proposition 1, the worst-case regret of the mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ is given by

$$\sup_{\mathbf{v} \in \mathcal{V}} \sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} v_j^i) - \sum_{i \in \mathcal{I}} (m^*)^i(\mathbf{v}) = \sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}. \quad (\text{A1})$$

We will strengthen this equality by showing that the worst-case regret $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$ is in fact attained by any $\mathbf{v} \in \text{supp}(\mathbb{P}^*)$. This will allow us to solve the minimization problem in (17) analytically.

By the construction of \mathbb{P}^* , for any fixed value profile $\mathbf{v} \in \text{supp}(\mathbb{P}^*)$, there exists a unique bidder i^* who assigns a higher value to each item than all other bidders, that is, $v_j^{i^*} \geq \frac{\bar{v}_j^*}{e} \geq v_j^i = 0$ for all $i \in \mathcal{I} \setminus \{i^*\}$ and $j \in \mathcal{J}$. This insight implies that

$$\begin{aligned} \sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} v_j^i) - \sum_{i \in \mathcal{I}} (m^*)^i(\mathbf{v}) &= \sum_{j \in \mathcal{J}} (v_j^{i^*} - \hat{m}_j^{i^*}(\mathbf{v}_j)) = \sum_{j \in \mathcal{J}} \left(v_j^{i^*} - \mathbb{E}_{\mathbb{Q}_j^*} [(\hat{m}')_j^{i^*}(\mathbf{v}_j, \tilde{r}_j)] \right) \\ &= \sum_{j \in \mathcal{J}} \left(v_j^{i^*} - \int_{\frac{\bar{v}_j^*}{e}}^{v_j^{i^*}} x \mathbb{Q}_j^*(dx) \right) = \sum_{j \in \mathcal{J}} \left(v_j^{i^*} - \int_{\frac{\bar{v}_j^*}{e}}^{v_j^{i^*}} dx \right) = \sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}, \end{aligned} \quad (\text{A2})$$

where the first equality holds because i^* represents the highest bidder for any item $j \in \mathcal{J}$ and because only the highest bidders for any of the items are charged under the mechanism $(\mathbf{q}^*, \mathbf{m}^*)$. The third equality follows from the observation that the (random) reserve price \tilde{r}_j almost surely exceeds the second-highest bid $\max_{i \neq i^*} v_j^i = 0$ with respect to the distribution \mathbb{Q}_j^* of \tilde{r}_j .

Fix now an arbitrary distribution $\mathbb{P} \in \Delta(\mathcal{V})$. Then, the expected regret of \mathbf{m}^* under \mathbb{P} satisfies

$$z(\mathbf{m}^*, \mathbb{P}) = \mathbb{E}_{\mathbb{P}} \left[\sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} \tilde{v}_j^i) - \sum_{i \in \mathcal{I}} m^i(\tilde{\mathbf{v}}) \right] \leq \sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e},$$

where the inequality follows from (A1). By (A2), the inequality is tight for $\mathbb{P} = \mathbb{P}^*$. Hence, \mathbb{P}^* solves the maximization problem on the left-hand side of (17).

Step 2. Consider now the expected regret minimization problem on the right-hand side of (17), which can be expressed more explicitly as

$$\begin{aligned} \inf_{\mathbf{q}, \mathbf{m}} \quad & \mathbb{E}_{\mathbb{P}^*} \left[\sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} \tilde{v}_j^i) - \sum_{i \in \mathcal{I}} m^i(\tilde{\mathbf{v}}) \right] \\ \text{s.t.} \quad & \mathbf{q} \in \mathcal{L}(\mathcal{V}, \mathbb{R}_+^{I \times J}), \mathbf{m} \in \mathcal{L}(\mathcal{V}, \mathbb{R}^I) \end{aligned} \quad (\text{A3})$$

(IC), (IR), (Inv).

To prove that $(\mathbf{q}^*, \mathbf{m}^*)$ solves problem (A3), we first relax this problem by replacing its objective function with a lower bound and by reducing the uncertainty sets of its robust constraints (Step 2.a). We then aggregate the constraints of the resulting problem across the bidders to obtain an even looser relaxation of (A3), which turns out to be equivalent to a multi-item pricing problem involving a single bidder (Step 2.b).

By leveraging Theorem 2, we then show that this problem's optimal value matches the objective value of $(\mathbf{q}^*, \mathbf{m}^*)$ in (A3).

Step 2.a. To construct a relaxation of problem (A3), we first establish a lower bound on its objective function. Indeed, for any fixed mechanism (\mathbf{q}, \mathbf{m}) we have

$$\begin{aligned} \mathbb{E}_{\mathbb{P}^*} \left[\sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} \tilde{v}_j^i) - \sum_{i \in \mathcal{I}} m^i(\tilde{\mathbf{v}}) \right] &= \frac{1}{I} \sum_{k \in \mathcal{I}} \mathbb{E}_{\hat{\mathbb{P}}^k} \left[\sum_{j \in \mathcal{J}} (\max_{i \in \mathcal{I}} \tilde{v}_j^i) - \sum_{i \in \mathcal{I}} m^i(\tilde{\mathbf{v}}) \right] \\ &= \frac{1}{I} \sum_{k \in \mathcal{I}} \mathbb{E}_{\hat{\mathbb{P}}^k} \left[\sum_{j \in \mathcal{J}} \tilde{v}_j^k - \sum_{i \in \mathcal{I}} m^i(\tilde{\mathbf{v}}) \right] \geq \frac{1}{I} \sum_{k \in \mathcal{I}} \mathbb{E}_{\hat{\mathbb{P}}^k} \left[\sum_{j \in \mathcal{J}} \tilde{v}_j^k - m^k(\tilde{\mathbf{v}}) \right], \end{aligned}$$

where the first equality follows from the definition of \mathbb{P}^* , while the second equality holds because $\tilde{\mathbf{v}}^i = \mathbf{0}$ for all $i \neq k$ and $\tilde{\mathbf{v}}^k \geq \frac{1}{e} \bar{\mathbf{v}}^* > \mathbf{0}$ almost surely under $\hat{\mathbb{P}}^k$. The inequality exploits the individual rationality constraint (IR), which implies that $m^i(\tilde{\mathbf{v}}) \leq 0$ almost surely under $\hat{\mathbb{P}}^k$ for all $i \neq k$.

For each bidder $i \in \mathcal{I}$, define $\mathcal{S}^i = \{\mathbf{v} \in \mathcal{V} : \mathbf{v}^k = \mathbf{0} \ \forall k \neq i\}$. Next, we relax the incentive compatibility constraint (IC) and the individual rationality constraint (IR) for any bidder $i \in \mathcal{I}$ by enforcing them only for scenarios $\mathbf{v} \in \mathcal{S}^i \subseteq \mathcal{V}$. The resulting relaxations are thus representable as

$$\sum_{j \in \mathcal{J}} q_j^i(\mathbf{v}^i, \mathbf{v}^{-i}) v_j^i - m^i(\mathbf{v}^i, \mathbf{v}^{-i}) \geq \sum_{j \in \mathcal{J}} q_j^i(\mathbf{w}^i, \mathbf{v}^{-i}) v_j^i - m^i(\mathbf{w}^i, \mathbf{v}^{-i}) \quad \forall i \in \mathcal{I}, \forall \mathbf{v} \in \mathcal{S}^i, \forall \mathbf{w}^i \in \mathcal{V}^i \quad (\widehat{\text{IC}})$$

and

$$\sum_{j \in \mathcal{J}} q_j^i(\mathbf{v}^i, \mathbf{v}^{-i}) v_j^i - m^i(\mathbf{v}^i, \mathbf{v}^{-i}) \geq 0 \quad \forall i \in \mathcal{I}, \forall \mathbf{v} \in \mathcal{S}^i. \quad (\widehat{\text{IR}})$$

Similarly, we note that the original inventory constraint (Inv) implies the relaxation

$$\begin{aligned} \sum_{i \in \mathcal{I}} q_j^i(\mathbf{v}) &\leq 1 \quad \forall j \in \mathcal{J}, \forall k \in \mathcal{I}, \forall \mathbf{v} \in \mathcal{S}^k \\ \implies q_j^k(\mathbf{v}) &\leq 1 \quad \forall j \in \mathcal{J}, \forall k \in \mathcal{I}, \forall \mathbf{v} \in \mathcal{S}^k, \end{aligned} \quad (\widehat{\text{Inv}})$$

where the second implication holds because the allocation probabilities are non-negative on \mathcal{V} .

In summary, we obtain the following relaxation of problem (A3).

$$\begin{aligned} \inf_{\mathbf{q}, \mathbf{m}} \quad & \frac{1}{I} \sum_{k \in \mathcal{I}} \mathbb{E}_{\hat{\mathbb{P}}^k} \left[\sum_{j \in \mathcal{J}} \tilde{v}_j^k - m^k(\tilde{\mathbf{v}}) \right] \\ \text{s.t.} \quad & \mathbf{q} \in \mathcal{L}(\mathcal{V}, \mathbb{R}_+^{I \times J}), \mathbf{m} \in \mathcal{L}(\mathcal{V}, \mathbb{R}^I) \\ & (\widehat{\text{IC}}), (\widehat{\text{IR}}), (\widehat{\text{Inv}}) \end{aligned} \quad (\text{A4})$$

Step 2.b. We now use constraint aggregation to construct a relaxation of problem (A4), which constitutes another—even looser—relaxation of problem (A3). To this end, define for any $i \in \mathcal{I}$ the linear embedding $E^i \in \mathcal{L}(\mathcal{V}^i, \mathcal{V}) (= \mathcal{L}(\mathcal{V}^1, \mathcal{V}))$ via

$$E^i(\mathbf{v}) = \underbrace{(\mathbf{0}, \dots, \mathbf{0})}_{i-1}, \mathbf{v}^\top, \underbrace{(\mathbf{0}, \dots, \mathbf{0})}_{I-i}^\top,$$

where, by slight abuse of notation, $\mathbf{v} \in \mathcal{V}^1$ denotes any value profile of a fixed bidder (recall also that the elements of \mathcal{V}^1 constitute row vectors). The proposed aggregation averages all constraints of (A4) across the bidders and expresses the resulting optimization problem in terms of the new auxiliary variables $\mathbf{f} \in \mathcal{L}(\mathcal{V}^1, \mathbb{R}_+^J)$ and $g \in \mathcal{L}(\mathcal{V}^1, \mathbb{R})$ defined via

$$f_j(\mathbf{v}) = \frac{1}{I} \sum_{i \in \mathcal{I}} q_j^i(E^i(\mathbf{v})) \quad \forall j \in \mathcal{J} \quad \text{and} \quad g(\mathbf{v}) = \frac{1}{I} \sum_{i \in \mathcal{I}} m^i(E^i(\mathbf{v})),$$

respectively, where $\mathbf{v} \in \mathcal{V}^1$ again denotes any value profile of a fixed bidder.

Thanks to the definition of the set \mathcal{S}^i introduced in Step 2.a, the relaxed incentive compatibility constraint $(\widehat{\text{IC}})$ can be expressed as

$$\sum_{j \in \mathcal{J}} q_j^i(E^i(\mathbf{v}))v_j - m^i(E^i(\mathbf{v})) \geq \sum_{j \in \mathcal{J}} q_j^i(E^i(\mathbf{w}))v_j - m^i(E^i(\mathbf{w})) \quad \forall i \in \mathcal{I}, \forall \mathbf{v}, \mathbf{w} \in \mathcal{V}^1,$$

where we use again \mathbf{v} and \mathbf{w} to denote arbitrary value profiles of a fixed bidder. By averaging the above inequality across all bidders $i \in \mathcal{I}$, we obtain the following aggregate constraint, which can be reformulated in terms of the new decision variables \mathbf{f} and g .

$$\begin{aligned} \frac{1}{I} \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} q_j^i(E^i(\mathbf{v}))v_j - \frac{1}{I} \sum_{i \in \mathcal{I}} m^i(E^i(\mathbf{v})) &\geq \frac{1}{I} \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} q_j^i(E^i(\mathbf{w}))v_j - \frac{1}{I} \sum_{i \in \mathcal{I}} m^i(E^i(\mathbf{w})) \quad \forall \mathbf{v}, \mathbf{w} \in \mathcal{V}^1 \\ \iff \sum_{j \in \mathcal{J}} f_j(\mathbf{v})v_j - g(\mathbf{v}) &\geq \sum_{j \in \mathcal{J}} f_j(\mathbf{w})v_j - g(\mathbf{w}) \quad \forall \mathbf{v}, \mathbf{w} \in \mathcal{V}^1 \end{aligned} \quad (\widehat{\text{IC}}')$$

Similarly, we can reformulate the relaxed individual rationality constraint $(\widehat{\text{IR}})$ as

$$\sum_{j \in \mathcal{J}} q_j^i(E^i(\mathbf{v}))v_j - m^i(E^i(\mathbf{v})) \geq 0 \quad \forall i \in \mathcal{I}, \forall \mathbf{v} \in \mathcal{V}^1.$$

By averaging the resulting inequality across all bidders $i \in \mathcal{I}$, we obtain the following aggregate constraint and its reformulation in terms of \mathbf{f} and g .

$$\begin{aligned} \frac{1}{I} \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} q_j^i(E^i(\mathbf{v}))v_j - \frac{1}{I} \sum_{i \in \mathcal{I}} m^i(E^i(\mathbf{v})) &\geq 0 \quad \forall \mathbf{v} \in \mathcal{V}^1 \\ \iff \sum_{j \in \mathcal{J}} f_j(\mathbf{v})v_j - g(\mathbf{v}) &\geq 0 \quad \forall \mathbf{v} \in \mathcal{V}^1 \end{aligned} \quad (\widehat{\text{IR}}')$$

Finally, the relaxed inventory constraint $(\widehat{\text{Inv}})$ can be formulated as

$$q_j^i(E^i(\mathbf{v})) \leq 1 \quad \forall j \in \mathcal{J}, \forall i \in \mathcal{I}, \forall \mathbf{v} \in \mathcal{V}^1.$$

Averaging the above inequality across all bidders $i \in \mathcal{I}$, we obtain

$$\begin{aligned} \frac{1}{I} \sum_{i \in \mathcal{I}} q_j^i(E^i(\mathbf{v})) &\leq 1 \quad \forall j \in \mathcal{J}, \forall \mathbf{v} \in \mathcal{V}^1 \\ \iff f_j(\mathbf{v}) &\leq 1 \quad \forall j \in \mathcal{J}, \forall \mathbf{v} \in \mathcal{V}^1. \end{aligned} \quad (\widehat{\text{Inv}}')$$

We can also re-express the decision-dependent part of the objective function of problem (A4) in terms of the new variables as

$$\frac{1}{I} \sum_{k \in \mathcal{I}} \mathbb{E}_{\hat{\mathbb{P}}^k} [m^k(\tilde{\mathbf{v}})] = \frac{1}{I} \sum_{k \in \mathcal{I}} \mathbb{E}_{\hat{\mathbb{P}}^k} [m^k(E^k(\tilde{\mathbf{v}}^k))] = \frac{1}{I} \sum_{k \in \mathcal{I}} \mathbb{E}_{\hat{\mathbb{P}}} [m^k(E^k(\tilde{\mathbf{v}}^1))] = \mathbb{E}_{\hat{\mathbb{P}}} [g(\tilde{\mathbf{v}}^1)],$$

where the first equality holds because $\tilde{\mathbf{v}}^i = \mathbf{0}$ for all $i \neq k$ almost surely under $\hat{\mathbb{P}}^k$, while the second equality holds because the marginal distribution of $\tilde{\mathbf{v}}^k$ under $\hat{\mathbb{P}}^k$ is given by $\hat{\mathbb{P}}$.

The resulting aggregation of problem (A4) can now be represented as

$$\begin{aligned} \inf_{\mathbf{f}, g} \quad & \mathbb{E}_{\hat{\mathbb{P}}} \left[\sum_{j \in \mathcal{J}} \tilde{v}_j^1 - g(\tilde{\mathbf{v}}^1) \right] \\ \text{s.t.} \quad & \mathbf{f} \in \mathcal{L}(\mathcal{V}^1, \mathbb{R}_+^J), g \in \mathcal{L}(\mathcal{V}^1, \mathbb{R}) \\ & (\widehat{\text{IC}}'), (\widehat{\text{IR}}'), (\widehat{\text{Inv}}'). \end{aligned} \quad (\text{A5})$$

By construction, the problems (A4) and (A5) constitute two increasingly loose relaxations of problem (A3). Moreover, problem (A5) constitutes a multi-item pricing problem involving a single bidder ($I = 1$) that minimizes the expected regret under the distribution $\hat{\mathbb{P}}$. The decision variables \mathbf{f} and g can be interpreted as the allocation and payment rules of the sales mechanism, respectively. By Theorem 2, the optimal value of this problem amounts to $\sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$. As problem (A5) constitutes a relaxation of problem (A3) and as $z(\mathbf{m}^*, \mathbb{P}^*) = \sum_{j \in \mathcal{J}} \frac{\bar{v}_j^*}{e}$ by Step 1, we can conclude that the mechanism $(\mathbf{q}^*, \mathbf{m}^*)$ solves problem (A3). This observation completes the proof. \square

B. Background Information on the Numerical Experiments

Section B.1 first provides additional details on the numerical experiments omitted from the main text.

Section B.2 then reports on the results of additional experiments.

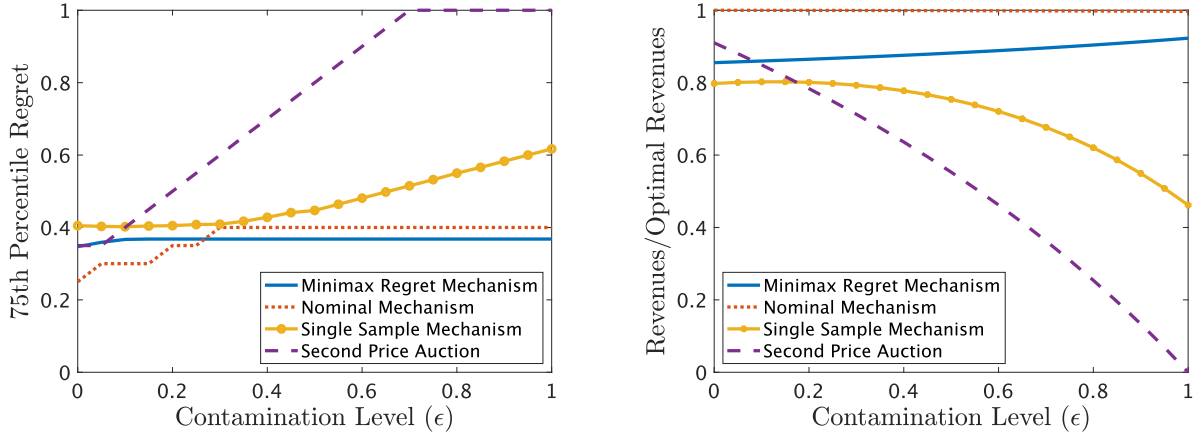


Figure 5 Comparison of the minimax regret mechanism (q^*, m^*) against the three benchmarks for $I = 2$ bidders, where \mathbb{P}^N is constructed from independent normal distributions with mean values 0.9 and variances 0.1.

B.1. Implementation Details

We define $\hat{\mathcal{V}} = \mathcal{V} \cap \delta \cdot \mathbb{Z}^I$ as the discrete grid approximating \mathcal{V} , and we set the grid spacing constant δ to 0.05 in all experiments with $I = 2$ bidders and to 0.001 in all experiments with $I = 10$ or $I = 100$ bidders. The discrete approximation $\hat{\mathbb{P}}^N$ of the nominal distribution \mathbb{P}^N assigns each point in $\hat{\mathcal{V}}$ a probability that is proportional to the probability density function of \mathbb{P}^N at that point. The discrete approximation $\hat{\mathbb{P}}_\epsilon$ of the data-generating distribution \mathbb{P}_ϵ is constructed analogously. Computing an expected revenue-maximizing mechanism requires solving a linear program whose numbers of variables and constraints scales with the cardinality of $\hat{\mathcal{V}}$. Any such linear program is implemented in MATLAB R2022a via the YALMIP interface and solved using GUROBI.

B.2. Additional Results

Figure 5 shows the 75th regret quantile and the (normalized) expected revenues of the different mechanisms under $\hat{\mathbb{P}}_\epsilon$ as a function of the contamination level ϵ . Here, the nominal distribution \mathbb{P}^N is calibrated as in the main text, but the means of the bidder values are set to 0.9. The results are consistent with Figures 2 and 3 in the main text. Indeed, the minimax regret mechanism (q^*, m^*) outperforms the two practically relevant benchmarks in terms of the 75th regret percentile as well as the expected revenues even for small values of ϵ , and it generates more than 80% of the highest achievable expected revenues for every contamination level ϵ .

Figure 6 visualizes the performance of the different mechanisms in view of an alternative nominal distribution \mathbb{P}^N , under which the bidders' values are positively correlated. Specifically, we assume here that,

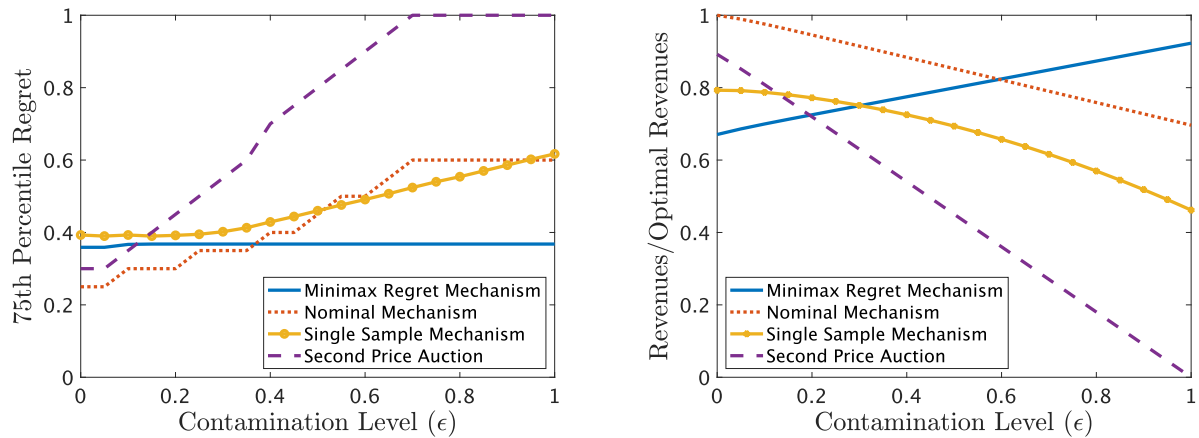


Figure 6 Comparison of the minimax regret mechanism (q^*, m^*) against the three benchmarks for $I = 2$ bidders, where \mathbb{P}^N is constructed from correlated normal distributions with mean values 0.5, variances 0.1 and covariances 0.05.

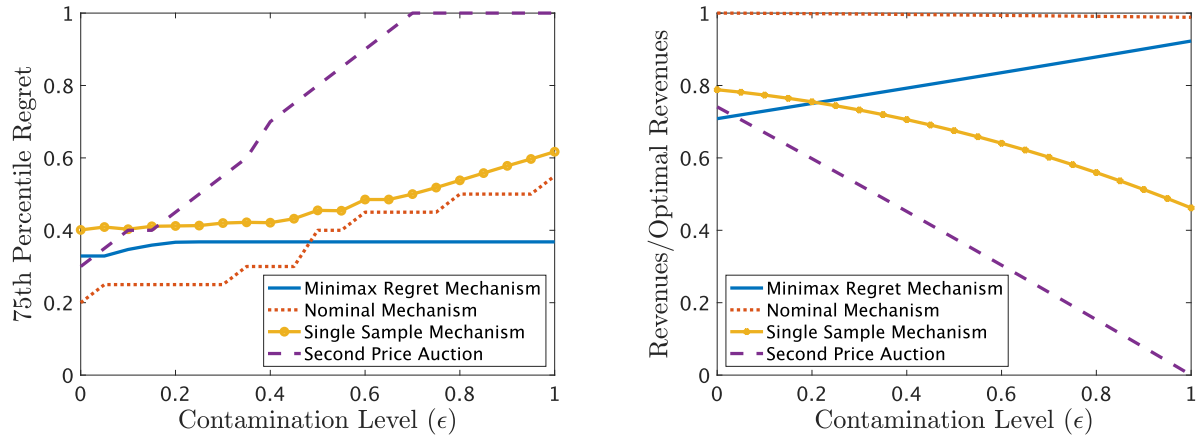


Figure 7 Comparison of the minimax regret mechanism (q^*, m^*) against the three benchmarks for $I = 2$ bidders, where \mathbb{P}^N is constructed from correlated normal distributions with mean values 0.5, variances 0.1 and covariances -0.05 .

under \mathbb{P}^N , the bidders' values follow a joint normal distributions with means 0.5, variances 0.1 and covariances 0.05, truncated on the unit square $\mathcal{V} = [0, 1]^2$. Analogously, Figure 7 shows the results of the same experiment with the only difference that the nominal covariances are set to -0.05 . We observe that the presence of correlations between the bidders' values does not fundamentally change our main conclusions. In particular, the minimax regret mechanism (q^*, m^*) outperforms the two practically relevant benchmarks even for small contamination levels ϵ . In addition, it displays a comparable or better performance than the nominal mechanism.

References

- Allouah A, Besbes O (2020) Prior-independent optimal auctions. *Management Science* 66(10):4417–4432.
- Bandi C, Bertsimas D (2014) Optimal design for multi-item auctions: A robust optimization approach. *Mathematics of Operations Research* 39(4):1012–1038.
- Bei X, Gravin N, Lu P, Tang ZG (2019) Correlation-robust analysis of single item auction. *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, 193–208.
- Bergemann D, Schlag K (2008) Pricing without priors. *Journal of the European Economic Association* 6(2-3):560–569.
- Bose S, Ozdenoren E, Pape A (2006) Optimal auctions with ambiguity. *Theoretical Economics* 1(4):411–438.
- Bulow J, Klemperer P (1996) Auctions versus negotiations. *American Economic Review* 86(1):180–194.
- Cai Y, Devanur NR, Weinberg SM (2021) A duality-based unified approach to bayesian mechanism design. *SIAM Journal on Computing* 50(3):STOC16–160–200.
- Caldentey R, Liu Y, Lobel I (2017) Intertemporal pricing under minimax regret. *Operations Research* 65(1):104–129.
- Carroll G (2017) Robustness and separation in multidimensional screening. *Econometrica* 85(2):453–488.
- Chen H, Hu M, Perakis G (2022) Distribution-free pricing. *Manufacturing & Service Operations Management* 24(4):1939–1958.
- Cremer J, McLean RP (1988) Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica* 56(6): 1247–1257.
- Daskalakis C, Deckelbaum A, Tzamos C (2013) Mechanism design via optimal transport. *Proceedings of the 14th ACM Conference on Electronic Commerce*, 269–286.
- Daskalakis C, Deckelbaum A, Tzamos C (2014) The complexity of optimal mechanism design. *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, 1302–1318.
- Daskalakis C, Deckelbaum A, Tzamos C (2017) Strong duality for a multiple-good monopolist. *Econometrica* 85(3):735–767.

- Dhangwatnotai P, Roughgarden T, Yan Q (2015) Revenue maximization with a single sample. *Games and Economic Behavior* 91:318–333.
- Giannakopoulos Y, Koutsoupias E (2014) Duality and optimality of auctions for uniform distributions. *Proceedings of the 15th ACM Conference on Economics and Computation*, 259–276.
- Gravin N, Lu P (2018) Separation in correlation-robust monopolist problem with budget. *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2069–2080.
- Hart S, Nisan N (2017) Approximate revenue maximization with multiple items. *Journal of Economic Theory* 172:313–347.
- Hart S, Nisan N (2019) Selling multiple correlated goods: Revenue maximization and menu-size complexity. *Journal of Economic Theory* 183:991–1029.
- Koçyiğit Ç, Iyengar G, Kuhn D, Wiesemann W (2020) Distributionally robust mechanism design. *Management Science* 66(1):159–189.
- Koçyiğit Ç, Rujeeapaiboon N, Kuhn D (2022) Robust multidimensional pricing: Separation without regret. *Mathematical Programming* 196:841–874.
- Krishna V (2009) *Auction Theory* (Academic Press).
- Myerson RB (1981) Optimal auction design. *Mathematics of Operations Research* 6(1):58–73.
- Poursoltani M, Delage E (2022) Adjustable robust optimization reformulations of two-stage worst-case regret minimization problems. *Operations Research* 70(5):2906–2930.
- Suzdaltsev A (2020) An optimal distributionally robust auction. *arXiv preprint arXiv:2006.05192*.