Robustly Optimal Reserve Price*

Wei He[†] Jiangtao Li[‡]

April 30, 2019

Abstract

We study a robust version of the single-unit auction problem. The auctioneer has confidence in her estimate of the marginal distribution of a generic bidder's valuation, but does not have reliable information about the joint distribution. In this setting, we analyze the performance of second-price auctions with reserve prices in terms of revenue guarantee, that is, the greatest lower bound of revenue across all joint distributions that are consistent with the marginals. For any finite number of bidders, we solve for the robustly optimal reserve price that generates the highest revenue guarantee. Our analysis has interesting implications in large markets. For any marginal distribution, the robustly optimal reserve price converges to zero as the number of bidders gets large. Furthermore, the second-price auction with no reserve price is asymptotically optimal among all mechanisms.

KEYWORDS: Robust mechanism design, second-price auction, reserve price, correlation, optimal transport, duality approach.

*We are grateful to Dirk Bergemann, Tilman Börgers, Ben Brooks, Gabriel Carroll, Jimmy Chan, Yi-Chun Chen, Eddie Dekel, Songzi Du, Piotr Dworczak, Jeff Ely, Daniel Garrett, Faruk Gul, Marina Halac, Seungjin Han, Bart Lipman, Heng Liu, George Mailath, Meg Meyer, Stephen Morris, Xiaosheng Mu, Wojciech Olszewski, Alessandro Pavan, Harry Di Pei, Eran Shmaya, Marciano Siniscalchi, Bruno Strulovici, Alireza Tahbaz-Salehi, Satoru Takahashi, and participants at various seminars and conferences for helpful discussions. Part of this paper was written while Li was visiting the Northwestern university and the Yale University, and he would like to thank these institutions for their hospitality and support.

†Department of Economics, The Chinese University of Hong Kong, hewei@cuhk.edu.hk

‡School of Economics, Singapore Management University, jtli@smu.edu.sg

Contents

1 Introduction

Suppose that you have an object to sell and you have confidence in your estimate of the distribution of a generic buyer's valuation F. Which mechanism should you use to maximize expected revenue? If there is a single buyer, then the optimal mechanism is to post a price p that maximizes $p(1 - F(p))$; see, for example, Riley and Zeckhauser (1983) and Manelli and Vincent (2007). When there are multiple buyers, you realize that you would be better off, because you can leverage the competition among the buyers to generate a higher expected revenue. But you also realize that you do not want to rely on the competition alone, because there may be cases in which one buyer's valuation is very high and all the other buyers' valuations are very low. Thus, you decide to also impose a reserve price, that is, the lowest price you are willing to sell the object for, to protect yourself from these cases.

Suppose that you do not have reliable information about the joint distribution of the buyers' valuations. How should you set the reserve price? Myerson (1981) shows that if the buyers' valuations are independent, under a regularity condition, the optimal mechanism can be implemented via a second-price auction with a reserve price such that the virtual value function evaluated at the reserve price is zero. But you are skeptical that the valuations are independent. Crémer and McLean (1988) and McAfee and Reny (1992) show that if the buyers' valuations are correlated, one can construct mechanisms that extract the entire surplus. However, such mechanisms entail lotteries and unusual fee schedules that require complex knowledge of the environment. The mere knowledge that the buyers' valuations are correlated (without knowing the exact correlation structure) is not enough for the construction.

The question remains: how should you set the reserve price, when you do not have reliable information about the correlation structure? We take on this question in this paper. The auctioneer in our model has confidence in her estimate of the marginal distribution of a generic bidder's valuation, but has non-Bayesian uncertainty about the correlation structure. We focus on second-price auctions with reserve prices that are both theoretically appealing and widely adopted in practice.[1] Lacking the knowledge of the correlation structure, our auctioneer chooses among reserve prices according to their revenue guarantee, that is, the greatest lower bound of revenue across all joint distributions that are consistent with the marginals.

Traditional models in mechanism design make strong assumptions about detailed knowledge of the auctioneer in the payoff environment. Our model is motivated by the observation that while it is relatively easy to estimate the distribution of a generic bidder's valuation, it seems substantially more difficult to estimate the correlation structure. For example, in wine auctions, it seems plausible that the auction house (such as Christie's or Sotheby's) has an estimate of the distribution of a generic buyer's valuation for a lot of wine, but it is less plausible that the auction house knows the joint distribution of all the bidders' valuations. Besides the statistical aspect that the joint distribution is a much higher-dimensional object, there are many practical reasons. For example, if the auctioneer has never interacted with the same set of bidders in auctions held before, there is no data for the estimation of the correlation structure. This scenario arises in settings in which the bidder pool changes constantly. Furthermore, there are instances in which the auctioneer cannot pin down the identities of the bidders (for example, when bidders bid through proxies). In this case, the auctioneer has no way of estimating the correlation structure.

It's easy to see that the choice of the reserve price depends on the correlation structure. For a simple example, suppose that there are two bidders $I = \{1, 2\}$ and each bidder's valuation v_i is uniformly distributed on the $[0, 1]$ interval. If the bidder's valuations are maximally positively correlated, then the optimal reserve prize is zero, as the competition among the bidders alone extracts the entire surplus. The other extreme case is that the bidders' valuations are maximally negatively correlated. In this case, the auctioneer would want to set a strictly positive reserve price. Setting the reserve price to $\frac{1}{2}$ does not decrease the probability of selling, but generates a strictly higher revenue for almost all valuations profiles. Note that there is an infinite class of correlation structures that are consistent with the

[1] More rigorously, English auctions with reserve prices are widely adopted in practice. English auctions and second-price auctions are strategically equivalent in our setting. To economize on notations, we shall work with second-price auctions.

marginals. Our objective in this paper is to understand the interaction of reserve price versus competition (in the form of second-price auction) when the auctioneer has non-Bayesian uncertainty about the joint distribution.

We first study the case in which the auctioneer is restricted to choosing a deterministic reserve price. Formally, we work with a maxmin optimization problem in which the auctioneer chooses a reserve price to maximize the worst-case expected revenue, where the worst case is taken over all joint distributions that are consistent with the marginals. Our main result, Theorem 1, solves for the robustly optimal reserve price that generates the highest revenue guarantee. The result is general in the sense that it applies to settings with any number of bidders and any marginal distribution. Our analysis has interesting implications in large markets. For any marginal distribution, the robustly optimal reserve price converges to zero as the number of bidders gets large. Furthermore, the second-price auction with no reserve price is asymptotically optimal among all mechanisms.

Our proof is interesting in its own right.[2] Our maxmin optimization problem can be interpreted as a two-player zero-sum game between the auctioneer and Nature. The auctioneer first chooses a reserve price r. Following the choice of the reserve price r, Nature chooses a correlation structure to minimize the expected revenue of the auctioneer subject to the constraint that the correlation structure is consistent with the marginals. This problem is not easy to work with, as the space of such joint distributions is very large. The novelty in our analysis is that we first solve an auxiliary problem in which we put a restriction on what Nature can do. More explicitly, we consider an extreme restriction in the sense that if the auctioneer chooses a reserve price r, Nature has one and exactly one strategy which

[2]In robust mechanism design, there are (at least) two common approaches to solve maxmin optimization problems. In the first approach, the performance of the candidate mechanism is independent of the uncertainty. The key step is to identify one realization of the uncertainty such that the candidate mechanism performs the best among all mechanisms. For example, this is the case in Chung and Ely (2007), Chen and Li (2018), and Yamashita and Zhu (2018) who study the foundations of dominant-strategy mechanisms and in Carroll (2017) who studies separate selling in multi-dimensional screening. The second approach works in settings in which a saddle point solution exists. For example, this is the case in Bergemann, Brooks, and Morris (2017), Du (2018), and Brooks and Du (2019) who study the design of informationally robust optimal auctions. Our approach is different.

we denote by π^r. This auxiliary problem is substantially easier to solve. We show that for the solution to the auxiliary problem r^*, the corresponding correlation structure π^{r^*} is the worst-case correlation structure. Thus, we can conclude that the solution to this auxiliary problem is the robustly optimal reserve price, because for any reserve price r, there exists a correlation structure π^r under which the expected revenue of the reserve price r is weakly lower than the worst-case expected revenue of the reserve price r^*. The construction of the correlation structures will be made clear in the formal analysis, and we will provide the intuition behind our construction.

To show that the corresponding correlation structure π^{r^*} is the worst-case correlation structure for r^*, we adopt a duality approach. This step of our analysis is closely related to the optimal transport theory (see, for example, Villani (2003)). To wit, for any reserve price r, Nature's minimization problem can be interpreted as an optimal transportation problem in which Nature seeks to implement the transportation at minimal cost. A transportation plan is a joint distribution that is consistent with the marginals, and Nature's cost function is the ex post revenue function of the auctioneer. While the literature of optimal transport focuses on two random variables, we work with multiple random variables. To be rigorous and self-contained, we prove an n-dimensional generalization of the Kantorovich duality theorem (see Villani (2003, Theorem 1.1.3)). This generalization is straightforward and follows from a modification of the original proof.

We then extend our analysis to the case in which the auctioneer is allowed to randomize over reserve prices. While a random reserve price may be less practical, this is interesting from a theoretical perspective, as the auctioneer may want to use the randomization in reserve prices to hedge against the uncertainty in the correlation structure. If the auctioneer is allowed to randomize over reserve prices, then the maxmin optimization problem admits a saddle point solution. Section 4 solves for the robustly optimal random reserve price under the condition that the density function does not decrease too fast. The auctioneer randomizes over a large interval of reserve prices, and generates a strictly higher revenue guarantee than the highest revenue guarantee from using a deterministic reserve price.

While we consider two versions of this problem, namely, deterministic reserve prices and random reserve prices, there is a third version of this problem that we do not address in this paper. Indeed, from a more theoretical perspective, one could ask, which mechanism generates the highest revenue guarantee among all mechanisms. We have not yet found a tractable approach to tackle this problem. Nevertheless, it is important to understand the revenue guarantee of standard auction formats[3] such as second-price auctions with reserve prices. This is because, when selecting an auction format, revenue guarantee is one of many criteria that can be used. While second-price auctions with reserve prices studied in this paper may not provide the highest revenue guarantee among all mechanisms, they are nevertheless one of the most common forms of auctioning an object and have many other desirable features aside from revenue guarantee.

The remainder of this introduction discusses related literature. Section 2 presents the model. Section 3 solves for the robustly optimal reserve price if the auctioneer is restricted to choosing a deterministic reserve price, and Section 4 extends our analysis to the case of random reserve prices. Section 5 provides further discussions of our model and results. Section 6 concludes the paper with some open questions. The appendix collects proofs omitted from the main body of the paper.

1.1 Related literature

This paper joins the burgeoning literature of robust mechanism design. The research agenda of this literature is to relax the strong assumptions about the detailed knowledge of the designer in the environment. One strand of this literature focuses on settings in which the designer does not have reliable information about the agents' beliefs; see, for example, Bergemann and Morris (2005), Chung and Ely (2007), Chen and Li (2018), and Yamashita and Zhu (2018). Börgers and Li (2019) propose strategically simple mechanisms in which the outcomes may depend on agents' first-order beliefs, but not on higher-order beliefs. Another strand of the literature focuses on settings in which the designer is uncertain about

[3] In settings in which the auctioneer knows the distribution of the bidders' valuations but is uncertain about additional information that may be received by the bidders, Bergemann, Brooks, and Morris (2017) study the revenue guarantee of first-price auctions and Bergemann, Brooks, and Morris (2018) compare the revenue guarantee of several standard auction formats, including first-price, second-price, and English auctions.

the additional information that may be received by the agents; see, for example, Bergemann, Brooks, and Morris (2017), Bergemann, Brooks, and Morris (2018), Du (2018), Brooks and Du (2019), and Libgober and Mu (2019).

In a moral hazard problem, Carroll (2015) provides a foundation for linear contracts in settings in which the principal has uncertainty about the agent's technology. Carroll and Segal (2018) consider a setting in which the auction's winner may resell to another bidder and the auctioneer has non-Bayesian uncertainty about such resale opportunities. Nakada, Nitzan, and Ui (2018) study the choice of a voting rule on a succession of two alternatives by a group of individuals who are uncertain about their future preferences.

The focus of our paper is on the uncertainty about the payoff environment, that is, the distribution of the bidders' valuations. More explicitly, our auctioneer has an estimate of the distribution of a generic bidder's valuation, but has non-Bayesian uncertainty about the correlation structure. In terms of the source of uncertainty, the closest to our paper is Carroll (2017), who considers a multi-dimensional screening setting in which the seller knows the marginal distribution of the buyer's valuation for each good but does not know the joint distribution. In a screening environment, Carrasco, Luz, Kos, Messner, Monteiro, and Moreira (2018) study the revenue maximization problem of a seller who is partially informed about the distribution of the buyer's valuation, only knowing its first m moments. Relatedly, Suzdaltsev (2018) considers a second-price auction in which the auctioneer knows an upper bound for valuations, the distribution's mean (and possibly variance), and further knows that the bidders' valuations are identically and independently distributed.

2 Preliminaries

2.1 Notations

We first introduce some notations that will be used in the sequel. For any real-valued vector $x \in \mathbb{R}^l$, we write $x(k)$ for the k-th largest element of the vector. For any set S, we denote by $|S|$ its cardinality. If Y is a measurable set, then ΔY is the set of all probability measures on Y. If Y is a metric space, then we treat it as a measurable space with its Borel σ-algebra.

2.2 The auction environment

An auctioneer seeks to sell a single indivisible object. There are $n \geq 2$ risk-neutral bidders competing for the object. We denote by $I = \{1, 2, \ldots, n\}$ the set of bidders and denote by i a typical bidder. Each bidder i holds private information about her valuation for the object, which is modeled as a random variable v_i with cumulative distribution function F_i. We denote by V_i the set of possible valuations of bidder i. The set of possible valuation profiles is $V = \times_{i \in I} V_i$, and we write v for a typical element of V. Apart from their private information, all bidders are identical. That is, $F_i = F_j$ for all $i, j \in I$. Without loss of generality, we assume that $V_i = [0, 1]$ for all $i \in I$, and we denote by F the common cumulative distribution function. We assume that F has a positive density f on $[0, 1]$.

While the auctioneer has confidence in her estimate of the marginal distribution of each bidder's valuation, she does not have reliable information about the joint distribution. In other words, our auctioneer has non-Bayesian uncertainty about the correlation structure. To the auctioneer, any joint distribution is plausible as long as the joint distribution is consistent with the marginals. We denote by

$$\Pi(F) = \Big\{\pi \in \Delta V : \forall i \in I, \forall A_i \subseteq V_i, \pi(A_i \times V_{-i}) = F(A_i)\Big\}$$

the collection of such joint distributions. When there is no confusion, we shall drop the dependence of $\Pi(F)$ on the marginal distribution F.

2.3 Second-price auctions with reserve prices

We focus on second-price auctions with reserve prices. In a second-price auction with a reserve price r, each bidder i submits a bid $m_i \in \mathbb{R}_+$. Conditional on the submitted bids $m = (m_1, m_2, \ldots, m_n)$, bidder i's probability of winning the object $q_i(m)$ and the payment from bidder i to the auctioneer $t_i(m)$ are given as follows:

$$q_i(m) = \begin{cases} \frac{1}{|W(m)|} & \text{if } i \in W(m) \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad t_i(m) = \begin{cases} \frac{\max(m(2),r)}{|W(m)|} & \text{if } i \in W(m) \\ 0 & \text{otherwise} \end{cases}$$

where $W(m) = \{i \in I : m_i = m(1), m_i \geq r\}$.

We are interested in the expected revenue of second price auctions with reserve prices in the dominant-strategy equilibrium in which each bidder submits a bid that is equal to her valuation of the object. For any reserve price r, let

$$REV(r, v) = \begin{cases} 0 & \text{if } v(1) < r; \\ r & \text{if } v(2) < r \leq v(1); \\ v(2) & \text{if } v(2) \geq r, \end{cases}$$

and let

$$REV(r, \pi) = \int_V REV(r, v) \, d\pi(v).$$

That is, we use $REV(r, v)$ to denote the auctioneer's ex post revenue by using the second-price auction with a reserve price r when the realized valuation profile is v, and we use $REV(r, \pi)$ to denote the auctioneer's expected revenue by using the second-price auction with a reserve price r under the joint distribution π.

2.4 Revenue guarantee as a criterion

We say that R is a revenue guarantee of the second-price auction with a reserve price r if for all $\pi \in \Pi$, $REV(r, \pi) \geq R$. We say that R is the revenue guarantee of the second-price auction with a reserve price r if it is a revenue guarantee and there is no higher revenue guarantee.

Our auctioneer chooses among reserve prices according to the revenue guarantee. Formally, the auctioneer solves the following maxmin optimization problem:

$$\sup_{r \in [0,1]} \inf_{\pi \in \Pi} REV(r, \pi).$$

We refer to the solution to this maxmin optimization problem as the robustly optimal reserve price.

Remark 1. An alternative interpretation is that our auctioneer is a maxmin decision maker. Our auctioneer has non-Bayesian uncertainty about the correlation structure, and chooses a reserve price to maximize the worst-case expected revenue, where the worst case is taken over all joint distributions that are consistent with the marginals.

3 Main results

The maxmin optimization problem can be interpreted as a two-player zero-sum game. The two players are the auctioneer and Nature. The auctioneer first chooses a deterministic reserve price $r \in [0, 1]$. After observing the auctioneer's choice of the reserve price, Nature chooses a correlation structure $\pi \in \Pi$. The auctioneer's payoff is $REV(r, \pi)$, and Nature's payoff is $-REV(r, \pi)$.

One may wish to solve this problem directly. That is, we first ask, is there a systematic way of solving for the worst-case correlation structure for any reserve price $r \in [0, 1]$? In principle, if we have a way of identifying the worst-case correlation structure for any reserve price, we could first work out the worst-case expected revenue for any reserve price, and then maximize the worst-case expected revenue (now a function of the reserve price only) by choosing the reserve price. However, it is not clear (at least to us) what would be the worst-case correlation structure for any reserve price.

We take an indirect approach. In the maxmin optimization problem, for each reserve price r, Nature can choose any joint distribution that is consistent with the marginals. This is not easy to work with, as the space of such joint distributions is very large. The novelty in our analysis is that we work with an auxiliary problem which has the interpretation that we impose a restriction on what Nature can do. More explicitly, for each reserve price r, we construct a correlation structure π^r that is consistent with the marginals. The auxiliary problem corresponds to an extreme restriction on Nature's strategies in the sense that if the auctioneer chooses a reserve price r, Nature has no choice but to choose π^r. We show that the solution to the auxiliary problem

$$\max_{r \in [0,1]} REV(r, \pi^r)$$

is also the solution to the maxmin optimization problem.[4]

The key step of our analysis is thus the construction of $\{\pi^r\}_{r \in [0,1]}$. The construction of $\{\pi^r\}_{r \in [0,1]}$ depends on the number of bidders and the marginal distribution, and will be made clear in the formal analysis. But before we move on to the formal analysis, we wish to provide a sketch of our analysis. The sketch

[4]Our construction of $\{\pi^r\}_{r \in [0,1]}$ ensures that a solution to the auxiliary problem exists.

highlights the requirements on $\{\pi^r\}_{r\in[0,1]}$ and should also make our approach more transparent.

In the first step, for each reserve price r, we explicitly construct a joint distribution π^r that is consistent with the marginals. At this stage, we do not know whether the constructed joint distribution π^r is the worst-case correlation structure for the reserve price r. Nevertheless, since π^r is consistent with the marginals, the worst-case expected revenue of a reserve price r is weakly lower than its expected revenue under the correlation structure π^r. Formally, for any r,

$$\inf_{\pi\in\Pi} REV(r,\pi) \le REV(r,\pi^r).$$

In the second step, we solve the following auxiliary maximization problem:

$$\max_{r\in[0,1]} REV(r,\pi^r).$$

Our construction of $\{\pi^r\}_{r\in[0,1]}$ ensures that a solution to the auxiliary problem exists. Let r^* denote a solution to the auxiliary problem. By definition,

$$REV(r^*,\pi^{r^*}) \ge REV(r,\pi^r)$$

for all r.

In the third step, we show that for the reserve price r^*, the correlation structure π^{r^*} is indeed the worst-case correlation structure. Formally, we show that

$$REV(r^*,\pi^{r^*}) = \min_{\pi\in\Pi} REV(r^*,\pi).$$

Our logic can be succinctly summarized below via a series of inequalities and equalities. For any r,

$$\inf_{\pi\in\Pi} REV(r,\pi) \le REV(r,\pi^r) \le REV(r^*,\pi^{r^*}) = \min_{\pi\in\Pi} REV(r^*,\pi).$$

Thus, r^* is a solution to the maxmin optimization problem.

Section 3.1 considers the setting with only two bidders. In this case, the construction of $\{\pi^r\}_{r\in[0,1]}$ is fairly intuitive. We also present a direct proof that shows π^{r^*} is the worst-case correlation structure for the reserve price r^*. Section 3.2 studies the general setting with n bidders. The construction of the correlation structures $\{\pi^r\}_{r\in[0,1]}$ is slightly more complicated, but it is still somewhat intuitive. It is also more difficult to show that π^{r^*} is the worst-case correlation structure for the reserve price r^* directly. We present an indirect proof based on a duality approach.[5]

For ease of exposition, we now introduce one more notation. For any $r \in [0,1]$ and any subset of bidders $S \subseteq I$, let

$$V^{r,S} := \{v \in V : v_i \geq r \text{ if and only if } i \in S\}.$$

In words, for any valuation profile $v \in V^{r,S}$, bidders in S have valuations weakly higher that r and bidders not in S have valuations lower than r. When S consists of a single bidder i, we write $V^{r,i}$ rather than $V^{r,\{i\}}$.

3.1 Two bidders

For the sake of clarity, we first consider the case in which there are only two bidders. We allow for arbitrary distribution F.

To understand the intuition behind our construction of $\{\pi^r\}_{r\in[0,1]}$, it is instructive to study the worst-case correlation structure when there is no reserve price.

Observation 1. *If there is no reserve price, then the worst-case correlation structure is the maximally negative correlation, which is defined by randomly drawing $q \sim U[0,1]$ and taking*

$$v_1 = F^{-1}(q) \text{ and } v_2 = F^{-1}(1-q).$$

An equivalent but indirect way of defining the maximally negative correlation is as follows. The maximally negative correlation is the unique joint distribution such that (1) the probability concentrates on the following curve

$$L_0 : F(1) - F(v_2) = F(v_1) - F(0),\ v_1 \in [0,1];$$

[5]To be clear, our analysis in the case of n bidders can be easily adopted in the setting with only two bidders. We organize our analysis in the current form so as to present the intuition of our analysis in the clearest way possible.

and (2) the joint distribution is consistent with the marginals. While indirect, this alternative definition is somewhat more intuitive. For this reason, in the reminder of the paper, we shall construct joint distributions indirectly.

To see why this is the worst case, note that in the second-price auction with no reserve price, the auctioneer's ex post revenue function is simply $REV(0, v) = v(2) = \min(v_1, v_2)$, which is a supermodular function. Since Nature chooses a joint distribution to minimize the expected value of a supermodular function, the worst-case correlation structure for the auctioneer is indeed the maximally negative correlation.[6]

We now consider an arbitrary reserve price $r \in [0, 1]$. It is less clear what would be the worst-case correlation structure for an arbitrary r. Nevertheless, we have a similar observation as in the case of no reserve price if Nature can only put positive probability in the regions $V^{r,\emptyset}$ and $V^{r,\{1,2\}}$.

Observation 2. *Fix an arbitrary reserve price $r \in [0, 1]$. In the constrained minimization problem in which Nature can only put positive probability in the regions $V^{r,\emptyset}$ and $V^{r,\{1,2\}}$, the worst-case correlation structure is such that*

1. *in the region $V^{r,\{1,2\}}$, the probability concentrates on the following curve*

$$L_r : F(1) - F(v_2) = F(v_1) - F(r),\ v_1 \in [r, 1];$$

2. *in the region $V^{r,\emptyset}$, the probability concentrates on the following curve*

$$v_2 = v_1,\ v_1 \in [0, r];$$

3. *the joint distribution is consistent with the marginals.*

We denote this joint distribution by π^r (see Figure 1 for a graphical illustration of π^r in the case in which F is the uniform distribution on the $[0, 1]$ interval).

[6] A function $g : V \to \mathbb{R}$ is supermodular if

$$g(v \vee v') + g(v \wedge v') \geq g(v) + g(v')$$

for all $v, v' \in V$, where $\vee$ denotes the component-wise maximum and $\wedge$ denotes the component-wise minimum. For detailed discussions on the ordering of joint distributions based on the integrals of supermodular functions, see for example Meyer and Strulovici (2012).

To see why this is a worst-case correlation structure when Nature is constrained to put positive probability only in the regions $V^{r,\emptyset}$ and $V^{r,\{1,2\}}$, note that we can think of Nature's minimization problem as two sub-problems, as the choice of the joint distribution in the region $V^{r,\emptyset}$ and the choice of the joint distribution in the region $V^{r,\{1,2\}}$ do not interact with each other. In the region $V^{r,\{1,2\}}$, the auctioneer's ex post revenue function is $REV(r, v) = v(2) = \min(v_1, v_2)$, which is a supermodular function. Therefore, our logic in Observation 1 applies here. In the region $V^{r,\emptyset}$, the exact distribution does not matter as long as it is consistent with the marginals, as the ex post revenue for any valuation profile in this region is zero. For concreteness, when constructing π^r, we pick a joint distribution such that the probability in the region $V^{r,\emptyset}$ concentrates on the curve $v_2 = v_1, v_1 \in [0, r]$. This particular choice plays no role in our analysis.

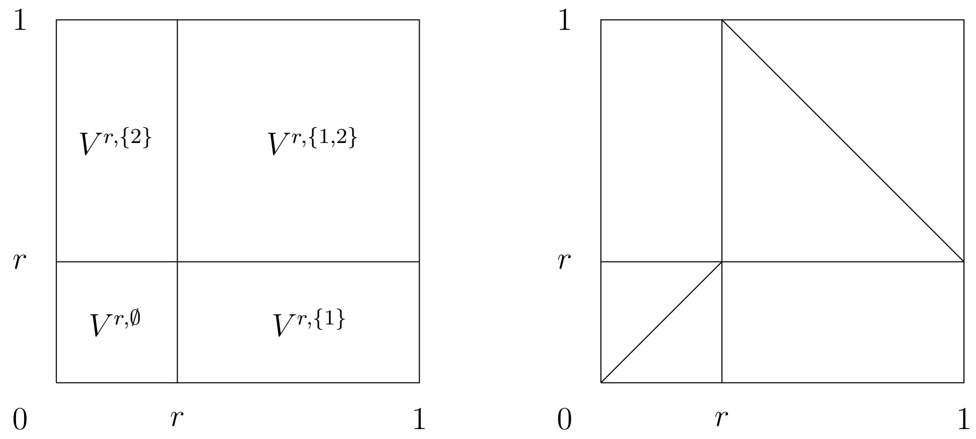

Figure 1: For each $r \in [0, 1]$, the figure on the left depicts the four regions given by $V^{r,\emptyset}$, $V^{r,1}$, $V^{r,2}$, and $V^{r,\{1,2\}}$. The figure on the right depicts the correlation structure π^r that we construct in the case in which F is the uniform distribution on the $[0, 1]$ interval.

Our logic so far is incomplete to pin down the worst-case correlation structure for an arbitrary r, as Nature may want to allocate some probability to the regions $V^{r,1}$ and $V^{r,2}$.

Nevertheless, our observations lead us to consider an auxiliary maximization problem that we formulate below. Now that we have constructed the correlation structure π^r for each $r \in [0, 1]$, the expected revenue of the reserve price r under π^r can be calculated as follows:

$$REV(r, \pi^r) = \int_{[r,1]^2} \min(v_1, v_2)\, d\pi^r(v) = 2 \int_r^{c(r)} x\, dF(x),$$

where $c(r) := F^{-1}(\frac{1+F(r)}{2})$. Consider the following auxiliary maximization problem:

$$\max_{r\in[0,1]} REV(r, \pi^r) = 2\int_r^{c(r)} x\, dF(x).$$

Proposition 1 below shows that the solution to this auxiliary maximization problem is the robustly optimal reserve price and generates the highest revenue guarantee of $REV(r^*, \pi^{r^*})$.

Proposition 1. *Suppose that there are two bidders and each bidder's valuation is distributed according to F. Let r^* denote a solution to the following maximization problem:*

$$\max_{r\in[0,1]} REV(r, \pi^r) = 2\int_r^{c(r)} x\, dF(x).$$

Then, r^* is the robustly optimal reserve price and generates the highest revenue guarantee of $REV(r^*, \pi^{r^*})$.

It suffices to show that for the reserve price r^*, the correlation structure π^{r^*} that we construct is the worst-case correlation structure. This is because the worst-case expected revenue of any reserve price r is weakly lower than its expected revenue under π^r, which is weakly lower than the expected revenue of the reserve price r^* under π^{r^*}.

The details of the proof can be found in the Appendix. Here, we provide a sketch of the proof. We first show that r^* necessarily satisfies that

$$F(2r^*) = \frac{1 + F(r^*)}{2}.$$

This is derived from the first-order condition of the auxiliary problem. We proceed to show that for any r such that $F(2r) = \frac{1+F(r)}{2}$, the correlation structure π^r that we construct is the worst-case correlation structure. To prove this, we use a direct approach, which highlights the role played by the condition that $F(2r) = \frac{1+F(r)}{2}$. More explicitly, we show that for any correlation structure π that is consistent with the marginals, there exists a correlation structure $\hat{\pi}$ such that

1. $\hat{\pi}$ is consistent with the marginals;
2. $\hat{\pi}$ only puts positive probability in the regions $V^{r,\emptyset}$ and $V^{r,\{1,2\}}$; and
3. the auctioneer's expected revenue is weakly lower under $\hat{\pi}$ than under π.

Thus, to solve for the worst-case correlation structure, we only have to consider joint distributions that are consistent with the marginals and only put positive probability in the regions $V^{r,\emptyset}$ and $V^{r,\{1,2\}}$. This, combined with Observation 2, implies that π^r is indeed the worst-case correlation structure.

Remark 2. Note that we have not imposed any condition on the marginal distribution F. Thus, there is nothing much we can say about r^* except that $F(2r^*) = \frac{1+F(r^*)}{2}$. Nevertheless, for any distribution F, it is easy to solve for the solution to the auxiliary problem.

Example 1 below applies our analysis to the case in which each bidder's valuation is uniformly distributed on the $[0, 1]$ interval.

Example 1 (Two bidders and uniform distribution)**.** Suppose that there are two bidders and each bidder's valuation is uniformly distributed on the $[0, 1]$ interval. From our analysis above, r^* necessarily satisfies that $2r^* = \frac{1+r^*}{2}$. The robustly optimal reserve price is $\frac{1}{3}$, and generates the highest revenue guarantee of $\frac{1}{3}$.

3.2 n bidders

In the case of two bidders, our analysis relies on the observation that when both bidders' valuations are weakly higher than r, the ex post revenue is $REV(r, v) = \min(v_1, v_2)$, which is a supermodular function. This simple observation leads us to construct the correlation structures $\{\pi^r\}_{r\in[0,1]}$ and to work with an auxiliary maximization problem in which the objective function is $REV(r, \pi^r)$ and the choice variable is r. By first-order condition, the solution of the auxiliary problem necessarily satisfies $F(2r) = \frac{1+F(r)}{2}$. For any such r, we can directly verify that π^r is the worst-case correlation structure.

When there are more than two bidders, we need to deal with two difficulties that are absent in the case of two bidders. The first difficulty is that, when there are more than two bidders, even in the case in which the reserve price is set to zero, the ex post revenue function $REV(0, v) = v(2)$ is not a supermodular function. This means that the construction of the correlation structures does not generalize in a straightforward manner.

The second difficulty is that, if we somehow manage to select the appropriate correlation structures, we still have to show that for the solution to the auxiliary problem, the corresponding correlation structure that we construct is the worst-case correlation structure for the auctioneer. This is somewhat straightforward in the case of two bidders, because there are only three types of regions, namely, regions in which both bidders' valuations are weakly higher than r, exactly one bidder's valuation is weakly higher than r, and none of the bidders' valuations is weakly higher than r. We exploit the property of the solution to the auxiliary problem and show that the worst-case correlation structure puts zero probability in the regions in which exactly one bidder's valuation is weakly higher than r. The proof is done by shifting probability from these regions to the other regions in a way that respects the marginals and weakly decreases the auctioneer's expected revenue. This procedure is less hopeful when there are more bidders, because there are many more types of regions.

What would be the appropriate correlation structures to work with? As in the case of two bidders, we get some intuition by first working with the case of no reserve price.

Observation 3. *Suppose that there are $n \geq 3$ bidders. In the second-price auction with no reserve price, Nature's objective is to minimize the expectation of $v(2)$ by choosing a joint distribution that is consistent with the marginals. The key observation is that if we fix $v(1)$ and $v(2)$, Nature would want to choose $v(3), v(4), \ldots, v(n)$ such that $v(2) = v(3) = \ldots = v(n)$. This is because for any values of $v(3), v(4), \ldots, v(n)$, the ex post revenue is always $v(2)$ and Nature's objective is to minimize the expectation of $v(2)$ by choosing a joint distribution that is consistent with the marginals.*

This observation leads us to work with the correlation structures $\{\pi^r\}_{r\in[0,1]}$ as follows. For each reserve price $r \in [0,1]$, the correlation structure π^r is such that

1. it only puts positive probability in the regions $V^{r,\emptyset}$ and $V^{r,I}$;
2. in the region $V^{r,I}$, the probability concentrates on n curves $L_r^1, L_r^2, \ldots, V_r^n$ where

$$L_r^i := \Big\{ v \in V^{r,I} : F(v_j) - F(r) = (n-1)(1 - F(v_i)), \, \forall j \neq i,$$
$$v_i \in [F^{-1}(\frac{(n+1) - F(r)}{n}), 1] \Big\};$$

3. in the region $V^{r,\emptyset}$, the probability concentrates on the following curve

$$v_i = v_1,\ \forall i \in I,\ v_1 \in [0, r];$$

4. the joint distribution is consistent with the marginals.

The interpretation of the curve L_r^i is that in the region in which every bidder's valuation is weakly higher than r and bidder i has the highest valuation, Nature puts probability in a way such that bidders other than i have the same valuation, and bidder i's valuation is maximally negatively correlated with the other bidders' valuation.

We consider the auxiliary problem as follows:

$$\max_{r\in[0,1]} REV(r, \pi^r) = \int_{[0,1]^n} v(2)\, d\pi^r(v) = \frac{n}{n-1}\int_r^{c_n(r)} x\, dF(x)$$

where $F(c_n(r)) = \frac{(n-1)+F(r)}{n}$. Let r_n^* be the solution of the auxiliary problem. Theorem 1 below shows that r_n^* is the robustly optimal reserve price and generates the highest revenue guarantee of $REV(r_n^*, \pi^{r_n^*})$.

Theorem 1. *Suppose that there are n bidders and each bidder's valuation is distributed according to F. Let r_n^* denote a solution to the following maximization problem:*

$$\max_{r\in[0,1]} REV(r, \pi^r) = \frac{n}{n-1}\int_r^{c_n(r)} x\, dF(x).$$

Then, r_n^* is the robustly optimal reserve price and generates the highest revenue guarantee of $REV(r_n^*, \pi^{r_n^*})$.

The auxiliary problem is easy to solve. By the first-order condition, r_n^* necessarily satisfies that

$$c_n(r) = nr.$$

Before we present the proof of Theorem 1, we provide an example to illustrate how to apply the theorem to solve for the robustly optimal reserve price.

Example 2 (n bidders and uniform distribution)**.** Suppose that there are n bidders and each bidder's valuation is uniformly distributed on the $[0, 1]$ interval. From our analysis above, r^* necessarily satisfies that $nr_n^* = \frac{(n-1)+r_n^*}{n}$. The robustly optimal reserve price is $r_n^* = \frac{1}{n+1}$, and generates the highest revenue guarantee of $\frac{n}{2(n+1)}$.

3.3 Proof of Theorem 1

It suffices to show that for the reserve price r_n^*, $\pi^{r_n^*}$ is the worst-case correlation structure for the auctioneer. Since r_n^* is a solution to the auxiliary maximization problem, for any reserve price r, there exists a correlation structure π^r such that the expected revenue of the reserve price r under π^r is weakly lower than the worst-case expected revenue of the reserve price r_n^*. This establishes that r_n^* is the robustly optimal reserve price and generates the highest revenue guarantee of $REV(r_n^*, \pi^{r_n^*})$. The proof proceeds as follows. We first show that the r_n^* satisfies

$$F(nr_n^*) = \frac{(n-1) + F(r_n^*)}{n}.$$

We then show that for any r such that $F(nr) = \frac{(n-1)+F(r)}{n}$, π^r is the worst-case correlation structure.

The first step is straightforward. This requirement on r_n^* is an immediate implication of the first-order condition. Consider the auxiliary maximization problem:

$$\max_{r\in[0,1]} REV(r, \pi^r) = \frac{n}{n-1} \int_r^{c_n(r)} x\, dF(x).$$

By the first-order condition, we have $\frac{dREV(r,\pi^r)}{dr} = \frac{n}{n-1} f(r)(\frac{c_n(r)}{n} - r)$. Let

$$\begin{aligned} R_n :=& \{r \in [0,1] : \frac{n}{n-1} f(r)(\frac{c_n(r)}{n} - r) = 0\} \\ =& \{r \in [0,1] : F(nr) = \frac{(n-1) + F(r)}{n}\} \end{aligned}$$

denote the set of stationary points. Note that the auxiliary problem has an interior solution, since the first-order derivative has a positive value at $r = 0$ and has a negative value at $r = 1$. Thus, it must be that $r_n^* \in R_n$.

In what follows, we show that for any reserve price $r \in R_n$, π^r is the worst-case correlation structure. That is, π^r is a solution to the following minimization problem,

$$\min_{\pi\in\Pi} REV(r, \pi). \tag{Primal}$$

The minimization problem is hard to solve directly. In particular, the direct approach that we use to prove the analogous statement in the case of two bidders does not easily generalize to the general setting. In what follows, we adopt a duality approach. We shall refer to this minimization problem as the primal minimization problem.

We construct the dual maximization problem of the primal minimization problem and show that the optimal value of the maximization problem is weakly less than the optimal value of the minimization problem. That is, we establish a weak duality property. The weak duality property that we establish can be viewed as the n-dimensional generalization of the weak duality property in the Kantorovich duality theorem. We then proceed to construct the primal variables and dual variables such that the value of the objective function of the minimization problem under the constructed primal variables and the value of the objective function of the maximization problem under the constructed dual variables are the same. This implies that the constructed primal variables is a solution to the primal minimization problem.

Consider the following dual maximization problem of the primal minimization problem:

$$\max_{\mu_1,\mu_2,\ldots,\mu_n} \quad \mathbb{J}(\mu_1,\mu_2,\ldots,\mu_n) = \sum_{i\in I}\int_{V_i} \mu_i(v_i)\,dF(v_i) \tag{Dual}$$

$$\text{subject to} \quad \text{for all } v \in V, \sum_{i\in I} \mu_i(v_i) \le REV(r,v).$$

Lemma 1 (n-dimensional generalization of the Kantorovich duality theorem)**.** *The optimal value of the dual maximization problem is weakly less than the optimal value of the primal minimization problem.*

Remark 3. The Kantorovich duality theorem shows a strong duality result in the case of two random variables. For our results, it suffices to prove the weak duality result. The extension to the case of n random variables is a straightforward. To be self-contained, we present the short proof here.

Proof of Lemma 1. It suffices to show that, for any feasible variables $\mu = (\mu_1,\mu_2,\ldots,\mu_N)$ of the dual maximization problem and any feasible variables π of the primal minimization problem, the value of the maximization problem under μ is weakly less than the value of the minimization problem under π. As we shall see below, this follows immediately from the respective feasibility constraints of the primal minimization problem and the dual maximization problem.

Let π be feasible variables of the primal minimization problem. That is, for all $i \in I$ and for all measurable sets $A_i \in V_i$,

$$\pi(A_i \times V_{-i}) = F(A_i). \tag{1}$$

Let $\mu = (\mu_1, \mu_2, \ldots, \mu_N)$ be feasible variables of the dual maximization problem. That is, for all $v \in V$,

$$\sum_{i \in I} \mu_i(v_i) \leq REV(r, v). \tag{2}$$

Thus, we have

$$\begin{aligned}
\mathbb{J}(\mu) &= \sum_{i \in I} \int_{V_i} \mu_i(v_i)\, dF(v_i) \\
&= \sum_{i \in I} \int_{V} \mu_i(v_i)\, d\pi(v) \\
&= \int_{V} \sum_{i \in I} \mu_i(v_i)\, d\pi(v) \\
&\leq \int_{V} REV(r, v)\, d\pi(v) \\
&= REV(r, \pi),
\end{aligned}$$

where the second line follows from (1) and the fourth line follows from (2). □

We are now ready to show that for any reserve price $r \in R_n$, π^r is the worst-case correlation structure. The proof proceeds as follows. Step (1) calculates the value of the objective function of the primal minimization problem under π^r. Step (2) constructs dual variables and calculates the value of the objective function of the dual maximization problem under the constructed dual variables. Step (3) verifies that these two values are the same for any $r \in R_n$.

Step (1). The value of the objective function of the primal minimization problem under π^r is

$$\frac{n}{n-1} \int_{r}^{c_n(r)} x\, dF(x)$$

where $c_n(r) = F^{-1}(\frac{(n-1)+F(r)}{n})$.

Step (2). For each $i \in I$, let

$$\mu_i(v_i) = \begin{cases} 0, & \text{if } v_i < r; \\ \frac{1}{n-1}(v_i - r), & \text{if } r \leq v_i < nr; \\ r, & \text{if } v_i \geq nr. \end{cases}$$

It is easy to verify that these dual variables satisfy the constraints of the dual maximization problem. Indeed, since $\mu_i(v_i)$ is an increasing function of v_i,

1. if $v(1) \geq v(2) \geq nr$, then $\sum_{i\in I} \mu_i(v_i) \leq nr \leq v(2) = REV(r, v)$;
2. if $v(1) \geq nr > v(2) \geq r$, then
$$\sum_{i\in I} \mu_i(v_i) \leq r + (n-1)\frac{1}{n-1}(v(2) - r) = v(2) = REV(r, v);$$
3. if $v(1) \geq nr$ and $r > v(2)$, then $\sum_{i\in I} \mu_i(v_i) = r = REV(r, v)$;
4. if $nr > v(1) \geq v(2) \geq r$, then
$$\sum_{i\in I} \mu_i(v_i) \leq \frac{1}{n-1}(nr - r) + (n-1)\frac{1}{n-1}(v(2) - r) = v(2) = REV(r, v);$$
5. if $nr > v(1) \geq r > v(2)$, $\sum_{i\in I} \mu_i(v_i) = \frac{1}{n-1}(v(1) - r) < r = REV(r, v)$; and
6. if $r > v(1)$, then $\sum_{i\in I} \mu_i(v_i) = 0 = REV(r, v)$.

We now calculate the value of the objective function of the dual maximization problem under the constructed dual variables as follows:

$$\begin{aligned}\mathbb{J}(\mu_1, \mu_2, \ldots, \mu_n) &= \sum_{i\in I} \int_{V_i} \mu_i(v_i)\, dF(v_i) \\ &= n \int_{V_1} \mu_1(v_1)\, dF(v_1) \\ &= n \int_r^{nr} \frac{1}{n-1}(v_1 - r)\, dF(v_1) + n \int_{nr}^1 r\, dF(v_1) \\ &= \frac{n}{n-1} \int_r^{nr} v_1\, dF(v_1) - \frac{n}{n-1} \int_r^{nr} r\, dF(v_1) + n \int_{nr}^1 r\, dF(v_1).\end{aligned} \tag{3}$$

Step (3). Recall that $R_n = \{r \in [0, 1] : F(nr) = \frac{(n-1)+F(r)}{n}\}$. Thus, for any $r \in R_n$, $c_n(r) = nr$, and the value of the objective function of the primal minimization problem under π^r is simply

$$\frac{n}{n-1} \int_r^{nr} x\, dF(x).$$

The value of the objective function of the dual maximization problem under the dual variables constructed in Step (2) is also

$$\frac{n}{n-1} \int_r^{nr} x\, dF(x),$$

since the last two terms in (3) cancel off. This completes the proof.

3.4 Large number of bidders

Our analysis so far solves for the robustly optimal reserve price for any number of bidders and any marginal distribution. In this subsection, we consider the setting with a large number of bidders.

Example 3 (large number of bidders and uniform distribution)**.** *Suppose that there are n bidders and the marginal distribution is the uniform distribution on* $[0, 1]$*. As n gets large, the robustly optimal reserve price $r_n^* = \frac{1}{n+1}$ converges to zero, and the highest revenue guarantee $\frac{n}{2(n+1)}$ converges to $\frac{1}{2}$ which is the expectation of a generic bidder's valuation.*

These features generalize to any marginal distribution. Indeed, Theorem 1 has immediate implications as follows.

Corollary 1. *For any marginal distribution F,*

1. $r_n^* < \frac{1}{n}$ *for any n;*
2. $\lim_{n\to\infty} r_n^* \to 0$*;*
3. $\lim_{n\to\infty} REV(r_n^*, \pi^{r_n^*}) \to \int_0^1 x\, dF(x)$*; and*
4. *The second-price auction with no reserve price is asymptotically optimal among all mechanisms.*

In words, Corollary 1 says that for any marginal distribution F, the robustly optimal reserve price converges to zero as the number of bidders gets large. Thus, our analysis lends extra support to the observation that in reality the reserve price is typically set at a very low level. Furthermore, our analysis gives a special role to the second-price auction with no reserve price even if the auctioneer can choose any mechanism. That is, the second-price auction with no reserve price is asymptotically optimal among all mechanisms.

The first three statements In Corollary 1 are immediate implications of our analysis in the case of n bidders. By Theorem 1, r_n^* necessarily satisfies that

$$r_n^* \in R_n = \{r \in [0, 1] : F(nr) = \frac{(n-1) + F(r)}{n}\}.$$

Thus, it must be that $r_n^* < \frac{1}{n}$ for any n. Subsequently, we have $\lim_{n\to\infty} r_n^* \to 0$. We also note that $\lim_{n\to\infty} c_n(r_n^*) = \lim_{n\to\infty} F^{-1}(\frac{n-1+F(r_n^*)}{n}) = 1$. Again by Theorem

1, the highest revenue guarantee with n bidders is

$$REV(r_n^*, \pi^{r_n^*}) = \frac{n}{n-1}\int_{r_n^*}^{c_n(r_n^*)} x\,dF(x) \to \int_0^1 x\,dF(x)$$

as $n \to \infty$.

The third statement is of particular interest. We interpret $\int_0^1 x\,dF(x)$ as the "full surplus" in our setting. This is because our auctioneer can never rule out the possibility that all the bidders' valuations are the same. Thus, the expectation of a generic bidder's valuation is the best revenue guarantee that the auctioneer can hope for. Thus, for whatever mechanism that the auctioneer might use, be it a second-price auction or a more complex mechanism, the expectation of a generic bidder's valuation is always an upper bound of the revenue guarantee of the mechanism.

The fourth statement in Corollary 1 follows from the above arguments and an additional observation. Fix any joint distribution that is consistent with the marginals. The expected revenue of the second-price auction with a sufficiently small reserve price and the revenue of the second-price auction with no reserve price are sufficiently close, since the ex post revenue function of the two mechanisms are the same except in regions in which at most one bidder's valuation is weakly larger than the reserve price. When the reserve price is sufficiently small, the probability of these regions is sufficiently close to zero.

4 Extension to random reserve price

In this section, we extend our analysis to the case in which the auctioneer is allowed to randomize over reserve prices.

By allowing the auctioneer to randomize over reserve prices, we are enlarging the auctioneer's strategy space. Let $\mathcal{G}$ denote the set of all cumulative distribution functions on the $[0, 1]$ interval. The auctioneer now chooses a distribution $G \in \mathcal{G}$ rather than a deterministic reserve price $r \in [0, 1]$. For any random reserve price G, let

$$REV(G, v) = \int_0^1 REV(r, v)\,dG(r)$$

and let

$$REV(G, \pi) = \int_V REV(G, v)\,d\pi(v).$$

That is, we use $REV(G, v)$ to denote the auctioneer's ex post revenue by using a random reserve price G when the realized valuation profile is v, and we use $REV(G, \pi)$ to denote the auctioneer's expected revenue by using a random reserve price G under the joint distribution π.

The auctioneer solves the following maxmin optimization problem:

$$\sup_{G \in \mathcal{G}} \inf_{\pi \in \Pi} REV(G, \pi).$$

We refer to the solution of this maxmin optimization problem as the robustly optimal random reserve price.

Note that both $\mathcal{G}$ and Π are convex, and $REV(G, \pi)$ is linear in both G and π. Thus, $\max_{G \in \mathcal{G}} \min_{\pi \in \Pi} REV(G, \pi) = \min_{\pi \in \Pi} \max_{G \in \mathcal{G}} REV(G, \pi)$. Our approach is thus to identify a saddle point (G^*, π^*) such that

$$REV(G^*, \pi) \geq REV(G^*, \pi^*) \geq REV(G, \pi^*)$$

for all $G \in \mathcal{G}$ and $\pi \in \Pi$. Since

$$\begin{aligned}
\max_{G \in \mathcal{G}} \min_{\pi \in \Pi} REV(G, \pi) &\geq \min_{\pi \in \Pi} REV(G^*, \pi) \\
&= REV(G^*, \pi^*) \\
&= \max_{G \in \mathcal{G}} REV(G, \pi^*) \\
&\geq \min_{\pi \in \Pi} \max_{G \in \mathcal{G}} REV(G, \pi) \\
&= \max_{G \in \mathcal{G}} \min_{\pi \in \Pi} REV(G, \pi),
\end{aligned}$$

we can conclude that G^* is the robustly optimal random reserve price, and $REV(G^*, \pi^*)$ is the highest revenue guarantee.

Section 4.1 studies the setting in which there are n bidders and each bidder's valuation is uniformly distributed on the $[0, 1]$ interval. Section 4.2 extends our analysis to the class of distributions that satisfy the condition that $xf(x)$ is weakly increasing in x. For example, this holds if the distribution is convex.

4.1 Uniform distribution

For the sake of clarity, we first study the case in which there are $n \geq 2$ bidders and each bidder's valuation is uniformly distributed on the $[0, 1]$ interval. We construct a particular random reserve price G^* and a correlation structure π^* such that (G^*, π^*) is a saddle point.

We first construct the correlation structure π^*. Some intuition behind the construction is as follows: π^* is such that given that the auctioneer knows the joint distribution is π^*, the auctioneer is indifferent among a range of reserve prices.

Construction of π^*. Let $\bar{b} = \frac{2n-1}{2n}$. The correlation structure π^* is such that

1. it only puts positive probability in the regions $V^{\bar{b},\emptyset}$ and $V^{\bar{b},i}$ for each $i \in I$;
2. in the region $V^{\bar{b},i}$, π^* is uniformly distributed on the following region

$$D_i := \Big\{(v_1, v_2, \ldots, v_n) : \bar{b} \le v_i \le 1, \\ 0 \le v_1 = v_2 = \ldots = v_{i-1} = v_{i+1} = \ldots = v_n < \bar{b}\Big\}$$

   with total measure $1 - \bar{b}$;
3. in the region $V^{\bar{b},\emptyset}$, π^* is uniformly distributed on the following line

$$D_0 := \{(v_1, v_2, \ldots, v_n) : 0 \le v_1 = \ldots = v_n < \bar{b}\}$$

   with total measure $\frac{1}{2}$.

In words, in the region $V^{\bar{b},i}$, bidders other than bidder i have the same valuation which is independent of v_i. In the region $V^{\bar{b},\emptyset}$, all the bidders have the same valuation. It is straightforward to verify that π^* is consistent with the marginals. Figure 2 illustrates π^* in the case in which there are two bidders.

Before we proceed, let us discuss some intuition behind the construction of the correlation structure π^*. One intuition that we have already mentioned is that π^* creates a lot of indifferences for the choice of the reserve price. The other intuition is described as follows. We focus on the region in which only bidder i's valuation is weakly larger than $\bar{b}$. The ex post revenue for each v in this region and each realization of the random reserve price is then the maximum of $v_{-i}(1)$ and the realization of the random reserve price, which is a submodular function. To minimize the expectation of a submodular function, in this region, bidders other than bidder i have the same valuation.

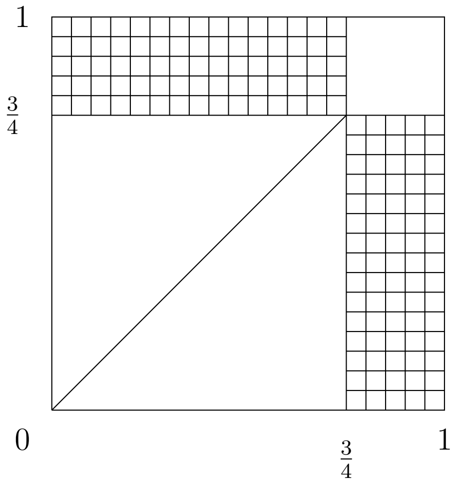

Figure 2: The figure depicts the correlation structure π^* that we construct in the case in which there are two bidders and F is the uniform distribution on $[0, 1]$. In this case, $\bar{b} = \frac{3}{4}$.

We now calculate the expected revenue for each reserve price $r \in [0, 1]$ against π^*. It is straightforward to calculate that

$$REV(r, \pi^*) = \begin{cases} \frac{2n-1}{4n}, & \text{if } r \in [0, \bar{b}]; \\ nr(1-r), & \text{if } r \in (\bar{b}, 1]. \end{cases}$$

Since $\bar{b} = \frac{2n-1}{2n} \geq \frac{1}{2}$, $nr(1-r) < n\bar{b}(1-\bar{b}) = \frac{2n-1}{4n}$ whenever $r > \bar{b}$. Thus,

$$\arg\max_{r \in [0,1]} REV(r, \pi^*) = [0, \bar{b}].$$

Construction of G^*. Let

$$G^*(r) = \bar{b}^{-\frac{1}{n-1}} r^{\frac{1}{n-1}}$$

with support $[0, \bar{b}]$. Since every reserve price in the support of G^* maximizes the auctioneer's expected revenue against π^*,

$$G^* \in \arg\max_{G \in \mathcal{G}} REV(G, \pi^*).$$

Thus, $REV(G^*, \pi^*)$ is an upper bound of the revenue guarantee.

Proposition 2. *Suppose that there are n bidders and each bidder's valuation is uniformly distributed on the $[0, 1]$ interval. Then, G^* is the robustly optimal random reserve price, and generates the highest revenue guarantee of $\frac{2n-1}{4n}$.*

It remains to show that

$$\pi^* \in \arg\min_{\pi \in \Pi} REV(G^*, \pi).$$

Here, as in the case of a deterministic reserve price, we adopt the duality approach. In words, the revenue guarantee of G^* is $REV(G^*, \pi^*)$. Since we have established that $REV(G^*, \pi^*)$ is an upper bound of the revenue guarantee, G^* is the robustly optimal random reserve price, and achieves the highest revenue guarantee $REV(G^*, \pi^*)$. The details of the proof can be found in the Appendix.

4.2 General distribution

We now extend our analysis to a large class of marginal distributions. Throughout this section, we make the following assumption: $xf(x)$ is weakly increasing in x. In words, this assumption says that the density function does not decrease too fast. Our analysis here parallels that in the case of uniform distribution. In the case of uniform distribution, it is relatively easy to construct π^* such that the auctioneer is indifferent among a range of reserve prices against π^*. In the general case, the construction of such a correlation structure is more complicated.

We first present the following technical lemma which is used in the construction of the saddle point. The lemma is a consequence of the assumption that $xf(x)$ is weakly increasing in x, and the detailed proof can be found in the appendix.

Lemma 2. *Fix a marginal distribution F such that $xf(x)$ is weakly increasing in x. Let $\psi(x) := x - \frac{1-F(x)}{f(x)}$, and let*

$$\gamma(x) := 1 - F(x) - \frac{1}{n-1} x^{-\frac{n}{n-1}} \int_0^x y^{\frac{n}{n-1}} f(y)\, dy.$$

Then,

1. $\lim_{x \to 0} xf(x) = 0$;
2. *there exists a unique $b^* \in (0, 1)$ such that $\psi(b^*) = 0$; and*
3. *there exists $x \in [b^*, 1]$ such that $\gamma(x) = 0$.*

Let $\bar{b}_F$ be such that $\bar{b}_F \in [b^*, 1]$ and $\gamma(\bar{b}_F) = 0$. We are now ready to construct a particular random reserve price G_F^* and a correlation structure π_F^* such that (G_F^*, π_F^*) is a saddle point.

Construction of π_F^*. As in the case of uniform distribution, π_F^* only puts positive probability in the regions $V^{\bar{b}_F,\emptyset}$ and $V^{\bar{b}_F,i}$ for each $i \in I$.

In $V^{\bar{b}_F,i}$, π_F^* concentrates on the following region

$$\begin{aligned} D_i = \Big\{ &(v_1, v_2, \ldots, v_n) : \, \bar{b}_F \le v_i \le 1, \\ &0 \le v_1 = v_2 = \ldots = v_{i-1} = v_{i+1} = \ldots = v_n < \bar{b}_F \Big\}. \end{aligned}$$

The marginal of π_F^* coincides with the restriction of F on $[\bar{b}_F, 1] \subseteq V_i$. Restricted in $V^{\bar{b}_F,i}$, all the $\{v_j\}_{j \neq i}$ are maximally positively correlated with the marginal of π_F^* being H on $[0, \bar{b}_F) \subseteq V_j$ for each $j \neq i$, where

$$H(x) = \frac{1}{n-1} x^{-\frac{n}{n-1}} \int_0^x y^{\frac{n}{n-1}} f(y) \, \mathrm{d}y.$$

Then the restriction of π_F^* on $V^{\bar{b}_F,i}$ is the product measure on V_i and $\prod_{j \neq i} V_j$; that is, v_i and $(v_1, v_2, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n)$ are independently distributed. Note that H is a feasible measure on $[0, \bar{b}_F)$, because

1. $\lim_{x \to 0} H(x) = 0$ and $H(\bar{b}_F) = 1 - F(\bar{b}_F)$;
2. H is continuous;
3. H is weakly increasing since the derivative of H is

$$\begin{aligned} h(x) &= \frac{1}{n-1} \left[-\frac{n}{n-1} \cdot x^{-\frac{2n-1}{n-1}} \cdot \int_0^x y^{\frac{n}{n-1}} f(y) \, \mathrm{d}y + f(x) \right] \\ &\ge \frac{1}{n-1} \left[-\frac{n}{n-1} \cdot x^{-\frac{2n-1}{n-1}} \cdot x f(x) \cdot \int_0^x y^{\frac{1}{n-1}} \, \mathrm{d}y + f(x) \right] \\ &= 0. \end{aligned}$$

In the region $V^{\bar{b}_F,\emptyset}$, π_F^* concentrates on the following line

$$D_0 = \{(v_1, \ldots, v_n) \colon 0 \le v_1 = \ldots = v_n < \bar{b}_F\},$$

with the density on any dimension V_j being $f(v_i) - (n-1)h(v_j)$. The density is well defined since by the construction of H, we have that $f(v_i) - (n-1)h(v_j) \ge 0$ for $v_j \in [0, \bar{b}_F]$.

In words, in the region $V^{\bar{b}_F,i}$, bidders other than bidder i have the same valuation which is independent of v_i. In the region $V^{\bar{b}_F,\emptyset}$, all the bidders have the same valuation. It is straightforward to verify that π^*_F is consistent with the marginal distribution F. The intuition behind the construction here is similar to the case of uniform distribution, and we shall not repeat the arguments.

Now that we have constructed the correlation structure π^*_F, we can calculate $REV(r, \pi^*_F)$ for all $r \in [0, 1]$. If $r \in (\bar{b}_F, 1]$, then $REV(r, \pi^*_F) = nr(1 - F(r))$. If $r \in [0, \bar{b}_F]$, then

$$REV(r, \pi^*_F) = \int_r^{\bar{b}_F} x\Big[f(x) - (n-1)h(x)\Big]\, dx + n\Big[\int_r^{\bar{b}_F} xh(x)\, dx + rH(r)\Big],$$

which is the sum of the expected revenue in the region $V^{\bar{b}_F,\emptyset}$ and the expected revenue in the n symmetric regions $\{V^{\bar{b}_F,i}\}_{i\in I}$. The expected revenue in each of the n symmetric regions is the sum of $\int_r^{\bar{b}_F} xh(x)\, dx$ and $rH(r)$, where $\int_r^{\bar{b}_F} xh(x)\, dx$ (resp. $rH(r)$) is the expected revenue from valuations profiles such that the second highest valuation is weakly higher than (resp. lower than) the reserve price r. This simplifies to

$$\begin{aligned}
REV(r, \pi^*_F) &= \int_r^{\bar{b}_F} x\Big[f(x) + h(x)\Big]\, dx + nrH(r) \\
&= n\int_r^{\bar{b}_F} \Big[xh(x) + H(x)\Big]\, dx + nrH(r) \\
&= n\int_r^{\bar{b}_F} xh(x)\, dx + n\Big[\bar{b}_F H(\bar{b}_F) - rH(r) - \int_r^{\bar{b}_F} x\, dH(x)\Big] + nrH(r) \\
&= n\bar{b}_F H(\bar{b}_F) \\
&= n\bar{b}_F(1 - F(\bar{b}_F)).
\end{aligned}$$

The second equality holds since by the construction of H, $nH(x) + (n-1)xh(x) = xf(x)$ for any $x \in [0, \bar{b}_F]$. The third equality uses integration by parts, and the last equality follows from that $H(\bar{b}_F) = (1 - F(\bar{b}_F))$.

Note that the derivative of $x(1-F(x))$ is $1 - F(x) - xf(x)$, which is negative for $x > b^*$. Since $\bar{b}_F \geq b^*$, for any $r > \bar{b}_F$, $nr(1 - F(r)) < n\bar{b}_F(1 - F(\bar{b}_F))$. Thus,

$$\arg\max_{r\in[0,1]} REV(r, \pi^*_F) = [0, \bar{b}_F].$$

Construction of G^*_F. Let

$$G^*_F(r) = \bar{b}_F^{-\frac{1}{n-1}} r^{\frac{1}{n-1}}$$

with support $[0, \bar{b}_F]$. Since every reserve price in the support of G_F^* maximizes the auctioneer's expected revenue against π_F^*,

$$G_F^* \in \arg\max_{G \in \mathcal{G}} REV(G, \pi_F^*).$$

Thus, $REV(G_F^*, \pi_F^*)$ is an upper bound of the revenue guarantee.

Theorem 2. *Suppose that there are n bidders and each bidder's valuation is distributed according to F. Then, G_F^* is the robustly optimal random reserve price, and generates the highest revenue guarantee of*

$$REV(G_F^*, \pi_F^*) = n\bar{b}_F(1 - F(\bar{b}_F)).$$

It remains to show that

$$\pi_F^* \in \arg\min_{\pi \in \Pi} REV(G_F^*, \pi).$$

In words, the revenue guarantee of G_F^* is $REV(G_F^*, \pi_F^*)$. Since we have established that $REV(G_F^*, \pi_F^*)$ is an upper bound of the revenue guarantee, G_F^* is the robustly optimal random reserve price, and achieves the highest revenue guarantee $REV(G_F^*, \pi_F^*)$. The details of the proof can be found in the Appendix.

Remark 4. For any fixed marginal distribution F such that $xf(x)$ is weakly increasing in x, when the number of bidders gets large, $\bar{b}_F$ converges to 1, the robustly optimal random reserve price is $G_F^*(r) = \bar{b}_F^{-\frac{1}{n-1}} r^{\frac{1}{n-1}}$ which converges to the dirac measure on zero, and the highest revenue guarantee is

$$REV(G_F^*, \pi_F^*) = n\bar{b}_F(1 - F(\bar{b}_F)) = n\bar{b}_F \frac{1}{n-1} \bar{b}_F^{-\frac{n}{n-1}} \int_0^{\bar{b}_F} y^{\frac{n}{n-1}} f(y)\, dy$$

which converges to $\int_0^1 x\, dF(x)$.

5 Discussions

This section provides further discussions of our model and results. Section 5.1 discusses an alternative model in which the auctioneer has local uncertainty about the marginal distribution. Section 5.2 compares our robustly optimal reserve price with the reserve price in Myerson (1981). Section 5.3 revisits our construction of the correlation structures $\{\pi^r\}_{r \in [0,1]}$.

5.1 Local uncertainty about the marginal distribution

Our modeling of the auctioneer's knowledge of the joint distribution and the marginals is somewhat extreme. To capture the idea that the auctioneer has no reliable information about the joint distribution, we impose non-Bayesian uncertainty over a large set of joint distributions. But on the other hand, we assume that the auctioneer has confidence in her estimate of the marginals.

Here, we briefly discuss an alternative model in which the auctioneer has local uncertainty about the marginal distribution. By local uncertainty, we mean that the auctioneer has uncertainty about the true marginal distribution, but is confident that the true marginal distribution is sufficiently close to F. We denote by $\tilde{F}$ the true marginal distribution. We denote by r_F^* the robustly optimal reserve price calculated under F, and denote by $r_{\tilde{F}}^*$ the robustly optimal reserve price calculated under $\tilde{F}$.

We claim that the revenue guarantee of r_F^* is close to the revenue guarantee of $r_{\tilde{F}}^*$ under the true marginal distribution $\tilde{F}$. We provide the intuition below without presenting the formal proof. The key observation is that, for any joint distribution π that is consistent with F, there exists a joint distribution $\tilde{\pi}$ that is consistent with $\tilde{F}$ such that for any reserve price r, $REV(r, \pi)$ and $REV(r, \tilde{\pi})$ are sufficiently close, and vice versa. This implies that the revenue guarantee of the reserve price r_F^* when Nature chooses $\pi \in \Pi(F)$ and when Nature chooses $\pi \in \Pi(\tilde{F})$ cannot be far apart. The same is true for the revenue guarantee of the reserve price $r_{\tilde{F}}^*$. The claim then follows because by the definition of r_F^* (resp. $r_{\tilde{F}}^*$), the revenue guarantee of r_F^* (resp. $r_{\tilde{F}}^*$) is weakly higher than the revenue guarantee of $r_{\tilde{F}}^*$ (resp. r_F^*) under F (resp. $\tilde{F}$).

5.2 Comparison with Myerson (1981)

In a seminal paper, Myerson (1981) studies the design of optimal auction in settings with independent private value. We now compare our robustly optimal reserve price and the optimal reserve price in the setting of Myerson (1981). For ease of comparison, we assume that there are n bidders, the marginal distribution of each bidder's valuation is F with support $[0, 1]$, and that F satisfies the regularity condition that $x - \frac{1-F(x)}{f(x)}$ is weakly increasing in x.

In Myerson (1981), the bidders' valuations are independent, and the optimal mechanism can be implemented via a second-price auction with a reserve price r_M such that

$$r_M - \frac{1 - F(r_M)}{f(r_M)} = 0.$$

Note that r_M is bounded away from zero. Furthermore, r_M is independent of the number of bidders. In the special case in which $F(x) = x$, $r_M = \frac{1}{2}$. In our model, we relax the auctioneer's knowledge of the joint distribution of the bidders' valuations. The robustly optimal reserve price typically varies with the number of the bidders. Corollary 1 shows that for any n, $r_n^* < \frac{1}{n}$, and the robustly optimal reserve price converges to zero as the number of bidders gets large. This implies that there exists N such that when there are more than N bidders, our robustly optimal reserve price is strictly lower than r_M.[7] In the special case in which $F(x) = x$, $r_n^* = \frac{1}{n+1}$.

5.3 The correlation structures $\{\pi^r\}_{r\in[0,1]}$

The key step in our analysis in the case of a deterministic reserve price is the construction of the correlation structures $\{\pi^r\}_{r\in[0,1]}$. When there are n bidders and the marginal distribution is F, we show in Section 3 that for any r such that $F(r) = \frac{(n-1)+F(r)}{n}$, the correlation structure π^r is the worst-case correlation structure. This suffices for our purpose of solving for the robustly optimal reserve price, sparing us the need to solve for the worst-case correlation structure for any $r \in [0, 1]$.

While the correlation structure π^r that we construct is indeed the worst case for any r such that $F(r) = \frac{(n-1)+F(r)}{n}$, this is not the case for any r. The easiest way to see this is to consider a reserve price that is sufficiently small. Note that the ex post revenue for each valuation profile in the regions in which exactly one bidder's valuation is weakly higher than the reserve price is simply the reserve price. Thus, if the reserve price is sufficiently small, Nature would want to allocate positive probability to these regions. We provide such an example below.

[7] There exists F that satisfies the regularity condition such that our robustly optimal reserve price is larger than r_M when n is small.

Example 4. Suppose that $n = 2$ and $F(x) = x$. Let $r = \frac{1}{6}$. By the construction of π^r, $REV(r, \pi^r) = 2\int_r^{\frac{1+r}{2}} x\,dx = \frac{5}{16}$. We now consider a correlation structure $\pi' \in \Pi$ such that (1) $\pi'(V^{r,\emptyset}) = 0$, $\pi'(V^{r,1}) = \pi'(V^{r,2}) = \frac{1}{6}$, and $\pi'(V^{r,\{1,2\}}) = \frac{2}{3}$; and (2) in the region $V^{r,\{1,2\}}$, the probability is uniformly distributed on the line $v_2 = \frac{7}{6} - v_1$, $v_1 \in [\frac{1}{6}, 1]$. It is straightforward to calculate that $REV(r, \pi') = \frac{11}{36} < REV(r, \pi^r) = \frac{5}{16}$.

6 Conclusion

We consider a robust version of the single-unit auction problem. In particular, we relax the assumption about the detailed knowledge of the auctioneer about the payoff environment and solve for the robustly optimal reserve price that generates the highest revenue guarantee.

This paper focuses on second-price auctions with reserve prices. These auction formats are both theoretically appealing and widely adopted in practice. However, from a more theoretical perspective, our analysis is not entirely complete. Indeed, an intriguing question that we have not addressed in this paper is the following: which mechanism generates the highest revenue guarantee among all mechanisms? Our analysis relies on the structure of the ex post revenue function in second-price auctions with reserve prices, and it is not yet clear to us how to extend our technique to address this more demanding research question. We leave this question for future research.

Further research might also consider additional restrictions on the set of joint distributions that the auctioneer perceives plausible. While classical papers such as Myerson (1981) consider one extreme formulation of the single-unit auction problem in the sense that the auctioneer knows the exact correlation structure, we consider the other extreme formulation in the sense that the auctioneer has no additional information besides the marginal distribution. It might be fruitful to investigate settings in which the auctioneer has some additional information besides the marginals, such as the knowledge that the bidders' valuations are positively correlated.

References

Bergemann, D., B. A. Brooks, and S. Morris (2017): "Informationally Robust Optimal Auction Design," mimeo.

——— (2018): "Revenue Guarantee Equivalence," American Economic Review, forthcoming.

Bergemann, D., and S. Morris (2005): "Robust Mechanism Design," *Econometrica*, 73, 1771–1813.

Börgers, T., and J. Li (2019): "Strategically Simple Mechanisms," Econometrica, forthcoming.

Brooks, B. A., and S. Du (2019): "Optimal Auction Design with Common Values: An Informationally-Robust Approach," mimeo.

Carrasco, V., V. F. Luz, N. Kos, M. Messner, P. Monteiro, and H. Moreira (2018): "Optimal Selling Mechanisms Under Moment Conditions," *Journal of Economic Theory*, 177, 245–279.

Carroll, G. (2015): "Robustness and Linear Contracts," *American Economic Review*, 105(2), 536–63.

——— (2017): "Robustness and Separation in Multidimensional Screening," *Econometrica*, 85(2), 453–488.

Carroll, G. D., and I. R. Segal (2018): "Robustly Optimal Auctions with Unknown Resale Opportunities," Review of Economic Studies, forthcoming.

Chen, Y.-C., and J. Li (2018): "Revisiting the Foundations of Dominant-Strategy Mechanisms," *Journal of Economic Theory*, 178, 294–317.

Chung, K.-S., and J. C. Ely (2007): "Foundations of Dominant-Strategy Mechanisms," *Review of Economic Studies*, 74(2), 447–476.

Crémer, J., and R. P. McLean (1988): "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56(6), 1247–1257.

Du, S. (2018): "Robust Mechanisms under Common Valuation," *Econometrica*, 86(5), 1569–1588.

Libgober, J., and X. Mu (2019): "Informational Robustness in Intertemporal Pricing," mimeo.

Manelli, A., and D. Vincent (2007): "Multidimensional Mechanism Design: Revenue Maximization and the Multiple-Good Monopoly," *Journal of Economic Theory*, 137(1), 153–185.

McAfee, R. P., and P. J. Reny (1992): "Correlated Information and Mechanism Design," *Econometrica*, 60(2), 395–421.

Meyer, M., and B. Strulovici (2012): "Increasing Interdependence of Multivariate Distributions," *Journal of Economic Theory*, 147(4), 1460–1489.

Myerson, R. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6(1), 58–71.

Nakada, S., S. Nitzan, and T. Ui (2018): "Robust Voting under Knightian Uncertainty," mimeo.

Riley, J., and R. Zeckhauser (1983): "Optimal Selling Strategies: When to Haggle, When to Hold Firm," *Quarterly Journal of Economics*, 98(2), 267–289.

Suzdaltsev, A. (2018): "Distributionally Robust Pricing In Auctions," mimeo.

Villani, C. (2003): *Topics in Optimal Transportation*, no. 58. American Mathematical Society.

Yamashita, T., and S. Zhu (2018): "On the Foundations of Ex Post Incentive Compatible Mechanisms," mimeo.

A Appendix

A.1 Proof of Proposition 1

We first show that r^* necessarily satisfies that $F(2r^*) = \frac{1+F(r^*)}{2}$. Consider the auxiliary maximization problem:

$$\max_{r\in[0,1]} REV(r,\pi^r) = 2\int_r^{c(r)} x\,dF(x).$$

By the first-order condition, we have that $\frac{dREV(r,\pi^r)}{dr} = 2f(r)(\frac{c(r)}{2} - r)$. Let

$$\begin{aligned} R :=& \{r \in [0,1] : 2f(r)(\frac{c(r)}{2} - r) = 0\} \\ =& \{r \in [0,1] : c(r) = 2r\} \\ =& \{r \in [0,1] : F^{-1}(\frac{1+F(r)}{2}) = 2r\} \\ =& \{r \in [0,1] : F(2r) = \frac{1+F(r)}{2}\} \end{aligned}$$

denote the set of stationary points. Note that the first-order derivative takes a positive value at $r = 0$, and takes a negative value at $r = 1$. Thus, the auxiliary problem has an interior solution. That is, $r^* \in R$.

We proceed to show that for any $r \in R$, π^r is the worst-case correlation structure. That is,

$$\pi^r \in \arg\min_{\pi\in\Pi} REV(r,\pi).$$

Without loss of generality, we consider only symmetric joint distributions. We show that for any $\pi \in \Pi$ that is symmetric, there exists π' such that

1. $\pi' \in \Pi$;
2. π' puts zero probability in the regions $V^{r,1}$ and $V^{r,2}$; and
3. $REV(r,\pi') \leq REV(r,\pi)$.

Thus, to solve for the worst-case correlation structure, we only have to consider joint distributions that are consistent with the marginals and only put positive probability in the regions $V^{r,\emptyset}$ and $V^{r,\{1,2\}}$. This, combined with Observation 2, implies that π^r is indeed the worst-case correlation structure.

The idea behind the construction of π' for any symmetric π is intuitive. Unfortunately, the construction does require quite a bit of notations. For ease of reference, we define nine segments as follows (see Figure 3):

$$A_1 = [0, r] \times [0, r]; \qquad A_2 = [r, c(r)] \times [0, r]; \qquad A_3 = [c(r), 1] \times [0, r];$$
$$A_4 = [0, r] \times [r, c(r)]; \qquad A_5 = [0, r] \times [c(r), 1]; \qquad A_6 = [r, c(r)] \times [r, c(r)];$$
$$A_7 = [c(r), 1] \times [r, c(r)]; \qquad A_8 = [r, c(r)] \times [r, 1]; \qquad A_9 = [r, 1] \times [r, 1].$$

For $1 \leq j \leq 9$, we also write $A_j = [\underline{x}^j, \bar{x}^j] \times [\underline{y}^j, \bar{y}^j]$. For example, $\underline{x}^2 = r$, $\bar{x}^2 = c(r)$, $\underline{y}^2 = 0$, and $\bar{y}^2 = r$.

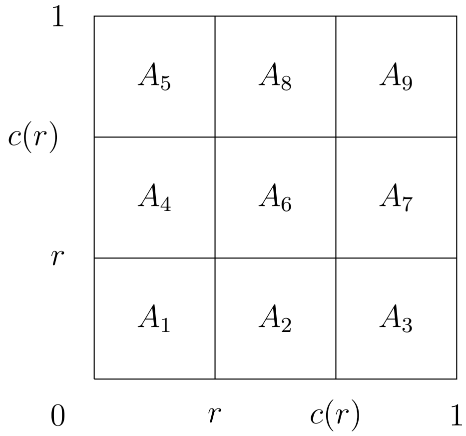

Figure 3: The nine segments that we define on the basis of r and $c(r)$.

Fix any $\pi \in \Pi$ that is symmetric. For $1 \leq j \leq 9$, let $a_j := \pi(A_j)$ denote the total measure of π on A_j. For any $[c_1, c_2] \subseteq [0, 1]$, $[d_1, d_2] \subseteq [0, 1]$ and any $1 \leq j \leq 9$, let

$$\pi_x^j([c_1, c_2]) = \pi([c_1, c_2] \times [\underline{y}^j, \bar{y}^j]) \text{ and } \pi_y^j([d_1, d_2]) = \pi([\underline{x}^j, \bar{x}^j] \times [d_1, d_2]).$$

We consider two cases. In the first case, $a_2 \geq a_3$. In the second case, $a_2 < a_3$. Suppose that $a_2 \geq a_3$. Since π is symmetric, $a_4 \geq a_3$. If $a_3 \neq 0$, we construct a correlation structure π' from π by shifting all the measure from A_3 to A_7 and shifting the same measure from A_4 to A_1 in a way that respects the marginals. Otherwise, we skip this step. This weakly decreases the auctioneer's expected revenue, since the ex post revenue is r for any $v \in A_3 \cup A_4$ and the ex post revenue is capped at $c(r) = 2r$ for any $v \in A_7$. Formally, π' is such that

1. π' coincides with π on A_2, A_5, A_6, A_8, and A_9;

2. $\pi'(A_3) = 0$;

3. for any $[c_1, c_2] \times [d_1, d_2] \subseteq A_7$,

$$\pi'([c_1, c_2] \times [d_1, d_2]) = \pi([c_1, c_2] \times [d_1, d_2]) + \frac{\pi_x^3([c_1, c_2]) \cdot \pi_y^4([d_1, d_2])}{a_4};$$

4. for any $[c_1, c_2] \times [d_1, d_2] \subseteq A_4$,

$$\pi'([c_1, c_2] \times [d_1, d_2]) = \pi([c_1, c_2] \times [d_1, d_2]) - \frac{a_3}{a_4} \cdot \pi([c_1, c_2] \times [d_1, d_2]);$$

5. for any $[c_1, c_2] \times [d_1, d_2] \subseteq A_1$,

$$\pi'([c_1, c_2] \times [d_1, d_2]) = \pi([c_1, c_2] \times [d_1, d_2]) + \frac{\pi_x^4([c_1, c_2]) \cdot \pi_y^3([d_1, d_2])}{a_4}.$$

Analogously, one can construct a correlation structure π'' from π' by shifting all the measure from A_5 to A_8 and shifting the same measure from A_2 to A_1 in a way that respects the marginals and weakly decreases the auctioneer's expected revenue. Note that

$$\pi''(A_3) = \pi''(A_5) = 0 \text{ and } \pi''(A_2) = \pi''(A_4) = a_2 - a_3.$$

If $a_2 = a_3$, then we have proved the desired result. If $a_2 > a_3$, then the last step in this case is to construct a correlation structure $\hat{\pi}$ from π'' by shifting all the measure from A_2 to A_6 and shifting the same measure from A_4 to A_1 in a way that respects the marginals. This weakly decreases the expected revenue, since the ex post revenue is r for any $v \in A_2 \cup A_4$ and the ex post revenue is capped at $c(r) = 2r$ for any $v \in A_6$. Formally, $\hat{\pi}$ is such that

1. $\hat{\pi}$ coincides with π'' on A_3, A_5, A_7, A_8, and A_9;

2. $\hat{\pi}(A_2) = \hat{\pi}(A_4) = 0$;

3. for any $[c_1, c_2] \times [d_1, d_2] \subseteq A_1$,

$$\hat{\pi}([c_1, c_2] \times [d_1, d_2]) = \pi''([c_1, c_2] \times [d_1, d_2]) + \frac{\pi_x''^4([c_1, c_2]) \cdot \pi_y''^2([d_1, d_2])}{a_2 - a_3};$$

4. for any $[c_1, c_2] \times [d_1, d_2] \subseteq A_6$,

$$\hat{\pi}([c_1, c_2] \times [d_1, d_2]) = \pi''([c_1, c_2] \times [d_1, d_2]) + \frac{\pi_x''^2([c_1, c_2]) \cdot \pi_y''^4([d_1, d_2])}{a_2 - a_3}.$$

This completes the proof for the first case since

$$\hat{\pi}(A_2) = \hat{\pi}(A_3) = \hat{\pi}(A_4) = \hat{\pi}(A_5) = 0.$$

Next, we study the case in which $a_2 < a_3$. Since π is consistent with the marginals and $F(2r) = \frac{1+F(r)}{2}$, $a_2 + a_6 + a_8 = F(c(r)) - F(r) = 1 - F(c(r)) = a_3 + a_7 + a_9$. Since π is symmetric, $a_7 = a_8$. Thus, $a_2 + a_6 = a_3 + a_9 \geq a_3$, which implies $a_6 \geq a_3 - a_2$. We further divide A_6 by the 45-degree line into three regions: $A_6^u = \{v \in A_6 : v_1 < v_2\}$ (above the 45-degree line), $A_6^m = \{v \in A_6 : v_1 = v_2\}$ (the 45-degree line), and $A_6^d = \{v \in A_6 : v_1 > v_2\}$ (below the 45-degree line). Without loss of generality, we can work with π such that $\pi(A_6^m) = 0$.[8] Thus, $\pi(A_6^u) = \pi(A_6^d) = \frac{a_6}{2}$. For any $[c_1, c_2] \times [d_1, d_2] \subseteq A_6$, let

$$\pi_x^d([c_1, c_2]) = \pi\left(([c_1, c_2] \times [\underline{y}^6, \bar{y}^6]) \cap A_6^d\right)$$

and

$$\pi_y^d([d_1, d_2]) = \pi\left(([\underline{x}^6, \bar{x}^6] \times [d_1, d_2]) \cap A_6^d\right).$$

We construct a correlation structure π' from π by shifting measure $\frac{a_3 - a_2}{2}$ from A_3 to A_2 and shifting the same measure from A_6^d to A_7 in a way that respects the marginals and does not change the expected revenue. Formally, π' is such that

1. π' coincides with π on A_1, A_4, A_5, A_6^u, A_6^m, A_8, and A_9;
2. for any $[c_1, c_2] \times [d_1, d_2] \subseteq A_2$,

$$\pi'([c_1, c_2] \times [d_1, d_2]) = \pi([c_1, c_2] \times [d_1, d_2]) + (a_3 - a_2) \cdot \frac{\pi_x^d([c_1, c_2]) \cdot \pi_y^3([d_1, d_2])}{a_3 \cdot a_6};$$

3. for any $[c_1, c_2] \times [d_1, d_2] \subseteq A_3$,

$$\pi'([c_1, c_2] \times [d_1, d_2]) = \pi([c_1, c_2] \times [d_1, d_2]) - \frac{a_3 - a_2}{2} \cdot \frac{\pi([c_1, c_2] \times [d_1, d_2])}{a_3};$$

[8]Otherwise, let π_6^m be the restriction of π on A_6^m, and π_x^m and π_y^m be the marginal of π_6^m on V_1 and V_2, respectively. Then we can construct another finite measure $\bar{\pi}_6^m$ having the same marginals π_x^m and π_y^m as follows: $\bar{\pi}_6^m$ concentrates on the curve with the maximally negative correlation on A_6: $\pi_x^m[\underline{x}^6, v_1] = \pi_y^m[v_2, \bar{y}^6]$ for $(v_1, v_2) \in A_6$. Let $\bar{\pi}'$ be the finite measure by restricting π on $V \setminus A_6^m$, and $\bar{\pi} = \bar{\pi}' + \bar{\pi}_6^m$. Then $\bar{\pi}$ respects the marginals, $\bar{\pi}(A_6^m) = 0$, and $REV(r, \bar{\pi}) \leq REV(r, \pi)$.

4. for any $[c_1, c_2] \times [d_1, d_2] \subseteq A_6^d$,

$$\pi'([c_1, c_2] \times [d_1, d_2]) = \pi([c_1, c_2] \times [d_1, d_2]) - (a_3 - a_2) \cdot \frac{\pi([c_1, c_2] \times [d_1, d_2])}{a_6};$$

5. for any $[c_1, c_2] \times [d_1, d_2] \subseteq A_7$,

$$\pi'([c_1, c_2] \times [d_1, d_2]) = \pi([c_1, c_2] \times [d_1, d_2]) + (a_3 - a_2) \cdot \frac{\pi_x^3([c_1, c_2]) \cdot \pi_y^d([d_1, d_2])}{a_3 \cdot a_6}.$$

Analogously, one can construct a correlation structure π'' from π' by shifting measure from A_5 to A_4 and shifting the same measure from A_6^u to A_8 in a way that respects the marginals and does not change the expected revenue. Note that $\pi''(A_2) = \pi''(A_3) = \pi''(A_4) = \pi''(A_5) = \frac{a_2+a_3}{2}$. We can then adopt our approach in the first case. This completes the proof.

A.2 Proof of Proposition 2

In what follows, we show that

$$\pi^* \in \arg\min_{\pi \in \Pi} REV(G^*, \pi).$$

We first calculate the ex post revenue of the auctioneer as follows:

$$REV(G^*, v) = \begin{cases} \bar{b}^{-\frac{1}{n-1}} \left[\frac{1}{n} v(1)^{\frac{n}{n-1}} + \frac{n-1}{n} v(2)^{\frac{n}{n-1}}\right], & \text{if } v(1) \leq \bar{b}; \\ \frac{\bar{b}}{n} + \frac{n-1}{n} \bar{b}^{-\frac{1}{n-1}} v(2)^{\frac{n}{n-1}}, & \text{if } v(2) \leq \bar{b} < v(1); \\ v(2), & \text{if } v(2) > \bar{b}. \end{cases}$$

Let

$$u(x) = \begin{cases} \frac{1}{n} \bar{b}^{-\frac{1}{n-1}} x^{\frac{n}{n-1}}, & \text{if } x \leq \bar{b}; \\ \frac{\bar{b}}{n}, & \text{if } x > \bar{b}. \end{cases}$$

One can easily verify that

$$REV(G^*, v) \geq \sum_{i \in I} u(v_i)$$

for all $v \in V$. It follows that for any $\pi \in \Pi$,

$$\begin{aligned} REV(G^*, \pi) &= \int_V REV(G, v)\, d\pi(v) \\ &\geq \int_V \sum_{i \in I} u(v_i)\, d\pi(v) \end{aligned}$$

$$
\begin{aligned}
&= \sum_{i\in I} \int_V u(v_i)\, d\pi(v) \\
&= n\int_0^1 u(x)\, dx \\
&= \frac{2n-1}{4n}.
\end{aligned}
$$

Since $REV(G^*, \pi^*) = \frac{2n-1}{4n}$, we conclude that

$$\pi^* \in \arg\min_{\pi\in\Pi} REV(G^*, \pi).$$

This completes the proof.

A.3 Proof of Lemma 2

1. Suppose that $\lim_{x\to 0} xf(x) = c > 0$. Since $xf(x)$ is weakly increasing in x, for any $x > 0$, we have that $xf(x) \geq c$ and $f(x) \geq \frac{c}{x}$. But then $F(x) \geq \int_0^x \frac{c}{y}\, dy = \infty$ for any $x > 0$. We have a contradiction.

2. Let $\eta(x) := xf(x) - (1 - F(x))$. Since $xf(x)$ is weakly increasing in x, $\eta(x)$ is increasing in x. Since $\lim_{x\to 0} \eta(x) < 0$ and $\eta(1) > 0$, there exists a unique b^* such that

$$\eta(x) \begin{cases} < 0, x < b^*; \\ = 0, x = b^*; \\ > 0, x > b^*. \end{cases}$$

Since $\psi(x) = \frac{\eta(x)}{f(x)}$, we have that

$$\psi(x) \begin{cases} < 0, x < b^*; \\ = 0, x = b^*; \\ > 0, x > b^*. \end{cases}$$

3. We show that (1) $\lim_{x\to 0} \gamma(x) > 0$; and (2) for any $x \leq b^*$ such that $\gamma(x) \leq 0$, we have that $\gamma'(x) \geq 0$. It then follows that $\gamma(b^*) \geq 0$. Since $\gamma(1) < 0$, there exists $x \in [b^*, 1]$ such that $\gamma(x) = 0$.

For (1),

$$\lim_{x\to 0} \gamma(x) = 1 - \frac{1}{n-1} \lim_{x\to 0} \frac{\int_0^x y^{\frac{n}{n-1}} f(y)\, \mathrm{d}y}{x^{\frac{n}{n-1}}}$$

$$
\begin{aligned}
&= 1 - \frac{1}{n-1}\lim_{x\to 0}\frac{x^{\frac{n}{n-1}}f(x)}{\frac{n}{n-1}x^{\frac{1}{n-1}}}\\
&= 1 - \frac{1}{n}\lim_{x\to 0} xf(x)\\
&= 1.
\end{aligned}
$$

For (2), for any $x \leq b^*$ and $\gamma(x) \leq 0$,

$$
\begin{aligned}
\gamma'(x) &= -\frac{n}{n-1}f(x) + \frac{n}{(n-1)^2}x^{-\frac{2n-1}{n-1}}\int_0^x y^{\frac{n}{n-1}}f(y)\,\mathrm{d}y\\
&\geq -\frac{n}{n-1}f(x) + \frac{n}{(n-1)}\frac{1-F(x)}{x}\\
&\geq 0,
\end{aligned}
$$

where the first inequality follows from the definition of the function γ and the assumption that $\gamma(x) \leq 0$, and the second inequality is due to the fact that $\psi(x) \leq 0$ for $x \leq b^*$.

A.4 Proof of Theorem 2

In what follows, we show that

$$
\pi_F^* \in \arg\min_{\pi\in\Pi} REV(G_F^*, \pi).
$$

We calculate the ex post revenue of the auctioneer as follows:

$$
REV(G_F^*, v) = \begin{cases}
\bar{b}_F^{-\frac{1}{n-1}}\left[\frac{1}{n}v(1)^{\frac{n}{n-1}} + \frac{n-1}{n}v(2)^{\frac{n}{n-1}}\right], & \text{if } v(1) \leq \bar{b}_F;\\
\frac{\bar{b}_F}{n} + \frac{n-1}{n}\bar{b}_F^{-\frac{1}{n-1}}v(2)^{\frac{n}{n-1}}, & \text{if } v(2) \leq \bar{b}_F < v(1);\\
v(2), & \text{if } v(2) > \bar{b}_F.
\end{cases}
$$

Let

$$
u(x) = \begin{cases}
\frac{1}{n}\bar{b}_F^{-\frac{1}{n-1}}x^{\frac{n}{n-1}}, & \text{if } x \leq \bar{b}_F;\\
\frac{\bar{b}_F}{n}, & \text{if } x > \bar{b}_F.
\end{cases}
$$

One can easily verify that

$$
REV(G_F^*, v) \geq \sum_{i\in I} u(v_i)
$$

for all $v \in V$. It follows that for any $\pi \in \Pi$,

$$
REV(G_F^*, \pi) \geq \sum_{i\in I}\int_V u(v_i)\,d\pi(v)
$$

$$
\begin{aligned}
&= n\int_0^1 u(x)\,dF(x) \\
&= n\left[\int_0^{\bar{b}_F} \frac{1}{n}\bar{b}_F^{-\frac{1}{n-1}} x^{\frac{n}{n-1}}\,dF(x) + \frac{\bar{b}_F}{n}(1-F(\bar{b}_F))\right] \\
&= n\bar{b}_F(1-F(\bar{b}_F)),
\end{aligned}
$$

where the last equality follows from the construction of $\bar{b}_F$.

Since $REV(G_F^*, \pi_F^*) = n\bar{b}_F(1-F(\bar{b}_F))$, we conclude that

$$
\pi_F^* \in \arg\min_{\pi\in\Pi} REV(G_F^*, \pi).
$$

This completes the proof.