

Correlation-Robust Analysis of Single Item Auction

Xiaohui Bei* Nick Gravin† Pinyan Lu† Zhihao Gavin Tang‡

Abstract

We investigate the problem of revenue maximization in single-item auction within the new correlation-robust framework proposed by Carroll [2017] and further developed by Gravin and Lu [2018]. In this framework the auctioneer is assumed to have only partial information about marginal distributions, but does not know the dependency structure of the joint distribution. The auctioneer’s revenue is evaluated in the worst-case over the uncertainty of possible joint distribution.

For the problem of optimal auction design in the correlation robust-framework we observe that in most cases the optimal auction does not admit a simple form like the celebrated Myerson’s auction for independent valuations. We analyze and compare performances of several DSIC mechanisms used in practice. Our main set of results concern the sequential posted-price mechanism (SPM). We show that SPM achieves a constant (4.78) approximation to the optimal correlation-robust mechanism. We also show that in the symmetric (anonymous) case when all bidders have the same marginal distribution, (i) SPM has almost matching worst-correlation revenue as any second price auction with common reserve price, and (ii) when the number of bidders is large, SPM converges to optimum. In addition, we extend some results on approximation and computational tractability for lookahead auctions to the correlation-robust framework.

1 Introduction

The monopolist’s theory of Bayesian revenue maximization for single-item auction is one of the most well studied topics in mechanism design literature. The central result in this field is by Myerson [28] that gives the optimal auction for the case when the buyers’ private values are drawn according to *independent* prior distributions which are known to the auctioneer. As Myerson pointed out, the independence is a strong assumption that is crucial for the optimality of his auction and which does not hold in many scenarios. Naturally, the general problem where the joint prior distribution of bidders valuations may be correlated has attracted significant amount of attention in economics (for a survey see [25]). A notable result from this literature is [18] by Crémer and McLean who showed that the auctioneer can extract full social surplus if the joint prior distribution satisfies certain mild conditions¹. This result implies that for those sufficiently generic joint distributions the auctioneer can always allocate the item to the highest value bidder and at the same time make the expected utility of each bidder to be zero by collecting winner’s value for the good as a payment. This result does not seem to be applicable in any practical setting and, therefore, “casts doubt on the value of the current mechanism design paradigm” according to [27].

*Nanyang Technological University. xhbei@ntu.edu.sg

†Shanghai University of Finance and Economics. nikolai@mail.shufe.edu.cn, lu.pinyan@mail.shufe.edu.cn

‡University of Hong Kong. zhtang@cs.hku.hk

¹It was later shown in [16] that full surplus extraction is possible for generic prior in a certain topological space defined on the set of prior distributions.

Another practical issue concerning the general Bayesian problem with any (correlated) prior distribution was articulated in [23] for a related but different context of the multi-product monopoly problem with a single buyer. Namely, the corresponding learning problem of a multi-dimensional prior distribution has exponential in the dimension (i.e., in the number of goods) representation and sampling complexity. We note that the same issue is pertinent to the Bayesian single-item auction with general multi-dimensional prior distribution of buyers’ private values. In other words, if anyone decides to obtain an accurate statistical estimate of the joint distribution of buyers’ values they would need to observe unrealistically many examples of the joint value profiles. Furthermore, if the seller and the buyers do not have publicly available and statistically accurate estimate of the joint prior, then they will need to agree on the common Bayesian prior, each party having much uncertainty about this distribution. The standard approach in theoretical economics to model such uncertainty is via extremely expressive and rich type spaces of [24]. This route is subject to the criticism from [8], who state that “very large type spaces would be needed and applied work remains highly sensitive to sometimes unexamined modeling choices about types; nowhere is this more true than in mechanism design”.

The robustness analysis of the single-item auction is the topic of our paper. To this end we employ a new *correlation-robust* framework of [13] originally proposed for the monopoly problem with multiple goods and a single buyer, but which could be easily and naturally adapted to Bayesian multiple bidder scenario with possible correlation among the bidders’ types as was mentioned in [23]. In this framework the valuation profile $\mathbf{v} = (v_1, \dots, v_n)$ is drawn from a joint general (correlated) distribution \mathbf{J} of bidder’s values, which is *not completely known* to the auctioneer. It is assumed that the seller knows the marginal distributions of \mathbf{J} for the private values $v_i \sim F_i$ of every individual bidder i , but neither the seller nor the bidders have any knowledge about possible correlation across different bidders. Any single-item auction is evaluated according to the auctioneer’s expected profit derived in *the worst-case*, over all possible joint distributions consistent with the given set of individual distributions F_i of each separate bidder $i \in [n]$. This evaluation gives the seller a robust guarantee on the expected profit of his auction which holds for any distribution with possible dependencies across different bidders. The seller seeks a truthful auction, i.e., dominant strategy incentive compatible (DSIC) and ex-post individually rational (IR)². On the other hand, because of the von Neumann minimax theorem, there exists a prior distribution for which the latter guarantee is tight even in the standard Bayesian framework where the auctioneer besides marginal distributions is given the specific joint prior.

Interestingly, the correlation-robust framework was first used in [13, 23] as a tractable mathematical framework to study the unwieldy and difficult multi-dimensional monopoly problem with a single buyer. On the other hand, in this setting it is often possible for the seller to acquire more information about dependencies, say, between different pairs or even triplets of items by doing more extensive market research. However, in the scenarios with multiple buyers these pair-wise dependencies might be quite hard to observe and learn. For example, a particular pair of buyers could have never participated in the same auction together, or the auctioned item is of the type that is usually sold directly via one-to-one negotiations or by price posting. Another common scenario where the auctioneer may only know marginal distribution for each buyer is the one where identities of the participating bidders cannot be observed. In this case it is reasonable to assume that each buyer has identical prior distribution, which can be deduced from an empirical study of

²Truthfulness is a standard assumption usually adopted in the computer science literature. In economics a more popular assumption for multi-agent environment is a weaker Bayesian incentive compatibility (BIC) and interim IR. We note that BIC and interim IR assumptions are not well compatible with the correlation-robust framework. Indeed, the prior joint distribution \mathbf{J} is not explicitly given. It is hard to model and even harder to predict how the bidders would behave under such uncertainty.

a small random subgroup of the buyers' population.

The correlation-robust framework allows one to directly compare the practical results in Myerson's setting³. One can view both scenarios as a problem where the auctioneer only knows prior distribution of each individual bidder. The difference is that in our case we assume that the buyers' valuations can be correlated and take the worst-case approach while in the Myerson's setting the seller simply assumes that buyers' valuations are independent, i.e., speaking in computer science terms, does the average case analysis.

1.1 Our Results.

Our first group of results concerns the design of optimal correlation-robust auctions, which, turn out to be not as simple as the Myerson's auction. In Appendix B we give a few simple numerical examples with an intricate structure of the optimal auction. Next, we turn to the design of approximately optimal, but much simpler and more practical auctions. To this end, we look at a number of single-item DSIC auctions used in practice in the correlated-robust framework. We show that among them a sequential posted-price mechanism (SPM) can always provide strong correlation-robust revenue guarantees. Our main result is a construction of SPM which is 4.78-approximation to the optimal revenue in this correlation-robust framework. In anonymous case when all buyers share the same marginal distribution, we show that (1) when the number of bidders exceeds the number of types in the support of the marginal prior distribution, SPM is the optimal correlation-robust auction. In other words, as the number n of bidders grows the SPM's revenue approaches optimal correlation-robust revenue. We also show for anonymous case that (2) SPM extracts at least $\frac{n-1}{n}$ -fraction of the revenue of the best second price auction with common reserve price.

Our second group of the results is closely related to the computer science literature on lookahead auctions initiated by Ronen [30]. The practical restriction in the family of lookahead auctions is that only the highest bidder can get the item. We note that this restriction can be quite handy for the designer as together with truthfulness it guarantees that the outcome is envy-free. On the other hand, it is quite easy to optimize within this class of auctions, because the problem reduces to a single-dimensional pricing problem for the highest bidder. It is also known that the revenue of the optimal lookahead auction is at least half of the revenue of the optimal auction. We note, however, that all the previous results assume perfect knowledge of the prior distribution of buyer types. In this paper, we study the worst-correlation performance of lookahead auctions. It is easy to show that the optimal lookahead auction remains 2-approximation to the optimal revenue in the correlation-robust framework. However, it is not that easy to compute one without the access to the joint prior distribution. We solve this computational question in the anonymous case. Specifically, we identify the structure of the corresponding worst-case distribution, which allows us to find a polynomial time algorithm that computes the optimal correlation-robust lookahead auction.

1.2 Related work

Most of the work on Bayesian single-item auction for general (correlated) prior distribution concerns the optimal auction design with respect to a given prior. It was shown in [18] that full social surplus can be extracted if the joint prior distribution is generic in a certain sense. Their result was extended in [27] to continuous type spaces, and by [16] to generic prior distributions in a properly defined topological space over the set of all prior distributions. From a computational perspective,

³Here, by saying practical we refer to the comparison between Myerson's setting to the general (correlated) case. Of course, one may argue that even Myerson's optimal auction is not practical, but in certain important special cases its theoretical solution is close to what is observed in practice.

it was shown in [29] that computing the optimal mechanism under ex-post individual rationality constraint is NP-hard. On the other hand, for randomized mechanisms [19] showed that optimum can be computed in time polynomial in the size of the distribution’s support.

Another line of work is concerned with the design of simple and approximately optimal auctions. In particular, the lookahead auctions introduced by Ronen [30] admits both computational efficiency and good approximation to the optimal revenue: it is a 2-approximation to the optimal mechanism. Ronen also introduced the generalized k -lookahead auction, which allows allocating the item to one of the k highest bidders. Later, [19] showed that the best k -lookahead auction has $\frac{2k-1}{3k-1}$ -approximation and can be computed efficiently in the oracle access model defined in [31]. [15] further improved the ratio to $\frac{e^{1-1/k}}{e^{1-1/k}+1}$, which was recently proved to be asymptotically tight in [20].

When the correlated prior distribution is unknown and needs to be learned, computational and sampling complexity concerns have been raised in several papers [1, 26, 23]. In contrast to the case of independent prior, much less is known about learning problem of a correlated prior or designing a mechanism from given samples. Fu et al. [21] proved that the optimal mechanism can be efficiently learned when there is a finite set of distributions from which the true distribution is drawn. However, there is a fundamental distinction between infinite sets of distributions and finite sets, and their results do not extend to the general value distributions [3].

Our work follows a trend in economics literature on robust mechanism design [5], where the goal is to provide performance guarantees even when there is uncertainty in the distribution over bidders’ valuations. A number of works in different settings employ a similar to our correlation-robust approach of searching solution to a parametrized optimization problem in the worst-case over the space of parameters, e.g., [12, 14, 9, 17, 22]. In particular, [4] and [11] study the setting of single-item auction. As in the standard Bayesian framework the bidders’ values for the item are drawn from a commonly known prior, however bidders there may have arbitrary information (certain high-order beliefs) about the prior distribution unknown to the seller. Auction’s performance is measured according to the lowest expected revenue across a class of incomplete information correlated equilibria termed Bayes correlated equilibria (BCE) in [6, 7]. In our work we completely ignore the beliefs of the buyers by focusing on the DSIC mechanisms. The authors in [4] also admit that allowing for model uncertainty only for the seller but not for the buyers is a valid concern, and “the possibility that the buyers face model uncertainty is eminently worthy of study”. When one cannot say much about buyers’ beliefs, BCE or BIC cannot be applied anymore, and our approach with DSIC and ex-post IR requirement seems to be more appropriate.

In AI community [2] also studied mechanisms that are robust to uncertainty in the distribution, assuming that the space of distributions has polynomial number of dimensions. In contrast, the dimension of the space of distributions in our correlation-robust framework is exponential in the number of buyers.

2 Preliminaries

We consider a single-round auction where one item is being sold to n bidders. Each bidder $i \in [n]$ has a privately known value v_i for the item and submits a sealed bid b_i to the auctioneer. Upon receiving submitted bids $\mathbf{b} = (b_1, \dots, b_n)$ from all bidders, the auctioneer decides which bidder i (if any) receives the item and the amount bidder i pays $p_i(\mathbf{b})$. Bidder i utility is the difference between her value v_i and her payment $p_i(\mathbf{b})$ if she gets the item; otherwise, she pays 0 and gets utility of 0. More generally, e.g., for a randomized mechanism, we write allocation probabilities for the bidders as $\mathbf{x}(\mathbf{b}) = (x_1(\mathbf{b}), \dots, x_n(\mathbf{b}))$. In general, throughout the paper we use the convention to denote vectors, or any multidimensional objects, with bold face script. The only exception will be the set

of price function $\{p_i(\mathbf{b})\}_{i=1}^n$ that we simply denote by p .

The type \mathbf{v} is drawn from a joint distribution \mathbf{J} , which is not completely known to the auctioneer and which may admit correlation between different components v_i and v_j of \mathbf{v} . The auctioneer only knows marginal distributions F_i of \mathbf{J} for each separate component i but does not know how these components are correlated with each other. We assume that every distribution F_i has finite support⁴ V_i . We use f_i to denote the probability density function of the distribution F_i . For notational convenience we also use F_i to denote the respective cumulative density function. The joint support of all F_i is $\mathbf{V} = \times_{i=1}^n V_i$. We use Π to denote all possible distributions π supported on \mathbf{V} that are consistent with the marginal distributions F_1, F_2, \dots, F_n , i.e., $\Pi = \{\pi \mid \sum_{\mathbf{v}_{-i}} \pi(v_i, \mathbf{v}_{-i}) = f_i(v_i), \forall i \in [n], v_i \in V_i\}$. Since we want to study the worst-case performance of a mechanism over this uncertainty, we measure performance of a mechanism in the worst-case over $\pi \in \Pi$.

$$\min_{\pi \in \Pi} \sum_{\mathbf{v}} \pi(\mathbf{v}) \cdot \sum_{i=1}^n p_i(\mathbf{v}). \quad (1)$$

Note that (1) is a linear program, since Π is given by a set of linear inequalities. We also write the corresponding dual problem.

$$\begin{aligned} \min \sum_{\mathbf{v}} \pi(\mathbf{v}) \cdot \sum_{i=1}^n p_i(\mathbf{v}) & \qquad \max \sum_{i=1}^n \sum_{v_i} f_i(v_i) \cdot \lambda_i(v_i) & (2) \\ \text{s. t. } \sum_{\mathbf{v}_{-i}} \pi(v_i, \mathbf{v}_{-i}) = f_i(v_i) & \quad \text{dual var. } \lambda_i(v_i) & \text{s. t. } \sum_{i=1}^n \lambda_i(v_i) \leq \sum_{i=1}^n p_i(\mathbf{v}) \quad \forall \mathbf{v} \\ \pi(\mathbf{v}) \geq 0 & & \lambda_i(v_i) \in \mathbb{R} \end{aligned}$$

The (2) LP can be simplified if we consider identical distributions ($F_i = F$ for all $i \in [n]$) and when mechanisms are symmetric, i.e., mechanisms are invariant under permutation of bidder identities. Indeed, we can take the optimal solution π^* to the primal LP and average it out over all permutations of bidders (note that every permutation of bidders again yields an optimal solution to the primal LP). The simplified dual LP looks as follows.

$$\begin{aligned} \max \quad & n \cdot \sum_v f(v) \cdot \lambda(v) \\ \text{s. t. } \quad & \sum_{i=1}^n \lambda(v_i) \leq \sum_{i=1}^n p_i(\mathbf{v}) & \forall \mathbf{v} = (v_1, \dots, v_n) \in V^n \\ & \lambda(v) \in \mathbb{R} & \forall v \in V \end{aligned}$$

The seller's problem. Besides evaluating the performance of a given mechanism (\mathbf{x}, p) , the seller's goal is to find a mechanism with the maximal worst-correlation revenue for a given set of marginal distributions $\{F_i\}_{i=1}^n$. Formally, we want to find a truthful mechanism (\mathbf{x}^*, p^*) such that

$$(\mathbf{x}^*, p^*) \in \operatorname{argmax}_{(\mathbf{x}, p)} \min_{\substack{\pi(\mathbf{x}, p) \\ \pi \in \Pi}} \sum_{\mathbf{v} \in V} \pi(\mathbf{v}) \sum_{i=1}^n p_i(\mathbf{v}). \quad (3)$$

⁴Similar to [13, 23] all our non computational and some computational results extend to the distributions with continuous types and bounded support.

We consider only truthful mechanisms (\mathbf{x}, p) that are dominant strategy incentive compatible (IC) and ex-post individually rational (IR). Formally, for all i , type \mathbf{v} , and v_i' , (\mathbf{x}, p) satisfies:

$$x_i(\mathbf{v}) \cdot v_i - p_i(\mathbf{v}) \geq x_i(v_i', \mathbf{v}_{-i}) \cdot v_i - p_i(v_i', \mathbf{v}_{-i}) \quad (\mathbf{IC})$$

$$x_i(\mathbf{v}) \cdot v_i - p_i(\mathbf{v}) \geq 0 \quad (\mathbf{IR})$$

For problem (3) we can write an LP formulation based on the dual of LP (2).

$$\begin{aligned} \max \quad & \sum_{i=1}^n \sum_{v_i} f_i(v_i) \cdot \lambda_i(v_i) \\ \text{s. t.} \quad & \sum_{i=1}^n \lambda_i(v_i) \leq \sum_{i=1}^n p_i(\mathbf{v}) \quad \forall \mathbf{v} \in \mathbf{V} \\ & \lambda_i(v_i) \in \mathbb{R} \\ & (\mathbf{x}(\mathbf{v}), \{p_i(\mathbf{v})\}_{i \in [n]}) : (\mathbf{IC}), (\mathbf{IR}) \end{aligned} \quad (4)$$

Robust Analysis of specific Auctions. The correlation-robust framework can be used to evaluate any specific auction in the *worst-correlation metric*. This metric serves as another dimension along which one can compare different auction formats. We note however, that it is not a trivial computational task to evaluate worst-case revenue of a given mechanism over all distributions $\pi \in \Pi$. However, for most of the truthful mechanisms employed in practice we can compute in polynomial time worst-correlation revenue (see Appendix A for details): sequential posted price mechanisms, second price auction with arbitrary set of reserves, and Myerson's optimal auction. Our result for sequential posted price mechanism continue to hold even when one can randomize between different set of prices.

Optimal Auctions: Examples. For any set of marginal distributions $\{F_i\}_{i \in [n]}$, we can calculate (in theory) the optimal mechanism in the correlation-robust framework by solving the LP (4). In practice, we can solve for optimal correlation-robust auctions for the instances with a constant number of bidders and prior marginal distributions with reasonably small support sizes. In Appendix B, we illustrate with a few numerical examples how optimal correlation-robust auctions may look like. From these examples, it seems hard to find a reasonable description of such mechanisms other than expressing it in the explicit form. These examples do not completely rule out the possibility of a reasonably simple description for the optimal solution, but they were enough to refute all our conjectures about optimality of already known simple formats for single-item auction. We believe a better understanding of these examples could shed light on the problem of optimal correlation-robust mechanism design.

3 Sequential Posted-Price Auctions

In this section, we study the class of sequential posted price mechanisms (SPM). First, we construct an SPM which is a 4.78-approximation to the optimal correlation-robust mechanism over any distribution $\pi \in \Pi$. Next we focus on the anonymous setting where buyers share an identical marginal distribution. We give an SPM that achieves optimal worst-case revenue when the number of bidders n is large enough; we also compare SPM with the second price auctions with anonymous reserve price: for any such second price auction we construct an SPM which is $\frac{n-1}{n}$ -approximation to the worst-case revenue of this auction.

Sequential Posted Price Mechanism. A sequential posted price mechanism (SPM) \mathcal{M} can be parameterized by an ordering over the buyers σ , and a vector of prices \mathbf{p} . Given a set of marginal distributions $\{F_i\}_{i=1}^n$, the mechanism is implemented by sequentially proposing price $p_{\sigma(i)}$ to buyer $\sigma(i)$ for $i = 1, 2, \dots, n$, and selling the item to the first buyer who accepts his price.

We will sometimes represent the price offered to a buyer in the quantile space with respect to that buyer's distribution. More specifically, given price p_i offered to buyer i , let its distribution quantile be $q_i = 1 - F_i(p_i)$, i.e., q_i is the probability that buyer i accepts price p_i . When the function F_i is continuous and strictly monotone, the mapping $1 - F_i : \text{supp}(F_i) \rightarrow [0, 1]$ is one-to-one. Hence, we can parametrize each distribution F_i in the quantile space as a mapping $V_i : [0, 1] \rightarrow \text{supp}(F_i)$ ⁵ with $V_i(q) = F_i^{-1}(1 - q)$. The use of quantiles is convenient in several places in this section. We slightly abuse the notations and denote the mechanism as both $\mathcal{M} = (\sigma, \mathbf{p})$ and $\mathcal{M} = (\sigma, \mathbf{q})$ interchangeably, where $\mathbf{q} = (q_1, \dots, q_n)$.

3.1 Worst-Case Revenue Guarantee

We first argue that the worst-case distribution for any SPM that posts non-increasing sequence of prices is always the distribution π^* with the maximum positive correlation⁶. To construct π^* one randomly draws a quantile $q \sim U[0, 1]$ and takes $v_i = V_i(q)$ for all $i \in [n]$. Thus $\pi^* \stackrel{\text{def}}{=} \{(V_1(q), \dots, V_n(q)) | q \sim U[0, 1]\}$.

Lemma 3.1. *For any sequential posted price mechanism \mathcal{M} that posts non-increasing sequence of prices, $\min_{\pi \in \Pi} \text{Rev}(\mathcal{M}, \pi) = \text{Rev}(\mathcal{M}, \pi^*)$.*

Proof. We first calculate the revenue extracted by \mathcal{M} for distribution π^* . Observe that buyer i takes the item and pays p_i if and only if $v_i \geq p_i$ and $v_j < p_j$ for all $j < i$, which happens with probability $\max_{j \leq i} (1 - F_j(p_j)) - \max_{j \leq i-1} (1 - F_j(p_j))$. Therefore,

$$\text{Rev}(\mathcal{M}, \pi^*) = \sum_{i=1}^n \left(\max_{j \leq i} (1 - F_j(p_j)) - \max_{j \leq i-1} (1 - F_j(p_j)) \right) \cdot p_i.$$

On the other hand, the probability that the item is sold in the first i rounds for any distribution $\pi \in \Pi$ is $\Pr[\exists j \leq i, v_j \geq p_j] \geq \max_{j \leq i} (1 - F_j(p_j))$. Since the sequence of p_i is non-increasing, the latter probability bound also means that the item is sold at a price equal to, or greater than p_i . This gives us a lower bound on the expected revenue of \mathcal{M} , which matches the revenue $\text{Rev}(\mathcal{M}, \pi^*)$ of \mathcal{M} on π^* . Indeed,

$$\begin{aligned} \text{Rev}(\mathcal{M}, \pi) &= \sum_{i=1}^n p_i \cdot (\Pr[\exists j \leq i, v_j \geq p_j] - \Pr[\exists j \leq i-1, v_j \geq p_j]) \\ &= \sum_{i=1}^n (p_i - p_{i+1}) \cdot \Pr[\exists j \leq i, v_j \geq p_j] \geq \sum_{i=1}^n (p_i - p_{i+1}) \cdot \max_{j \leq i} (1 - F_j(p_j)) \\ &= \sum_{i=1}^n \left(\max_{j \leq i} (1 - F_j(p_j)) - \max_{j \leq i-1} (1 - F_j(p_j)) \right) \cdot p_i = \text{Rev}(\mathcal{M}, \pi^*), \end{aligned} \quad (5)$$

where $p_{n+1} = 0$. □

⁵In general case, when F_i is not a continuous distribution, we let $V_i(q) \stackrel{\text{def}}{=} \inf\{v \mid F_i(v) \geq 1 - q\}$.

⁶From a revenue maximization point of view, the assumption that a SPM always posts non-increasing sequence of prices is without loss of generality. Because otherwise one can always re-order the prices in non-increasing order and this can only improve the revenue.

This characterization of the worst-case distribution provides a simple and convenient way to analyze the worst-case revenue of an SPM. For example, given a set of marginal distributions $\{F_i\}_{i=1}^n$, we can assume without loss of generality that any Pareto-optimal SPM $\mathcal{M} = (\sigma, \mathbf{q})$ has $q_{\sigma(i)} < q_{\sigma(i+1)}$ for every $1 \leq i < n$. This is because if there exists $i < j$ such that $q_{\sigma(i)} > q_{\sigma(j)}$, then when the type \mathbf{v} is drawn from distribution π^* , the item will never be sold to buyer $\sigma(j)$. In this case one might as well put buyer $\sigma(j)$ to the very end of the queue and set $q_{\sigma(j)} = 1$ as it will never hurt the revenue. Furthermore, when the offers are in the increasing order of quantiles, the worst-case revenue of an SPM \mathcal{M} is simply

$$\text{Rev}(\mathcal{M}, \pi^*) = \sum_{i=1}^n (q_{\sigma(i)} - q_{\sigma(i-1)}) p_{\sigma(i)}.$$

Next we focus on a specific class of sequential posted price mechanisms, where the quantiles proposed to different buyers are all “spread-out”, in the sense that any two quantiles are at least factor of 2 apart from each other.

Definition. *Spread-out SPM is a sequential posted pricing $\mathcal{M}_S = (\sigma, \mathbf{q})$ such that for any two $i \neq j$ either $q_i \geq 2q_j$, or $q_j \geq 2q_i$.*

A good feature of SPM with spread-out quantiles is that, as the following lemma shows, the worst-case revenue of an SPM $\mathcal{M} = (\sigma, \mathbf{p})$ with spread-out quantiles is within a constant factor of the revenue of the n independent sales in which the seller has n copies of the items and independently offers one to each buyer i at price p_i . This allows us to approximate the revenue of \mathcal{M} as the sum of individual revenues from price posting and to simplify the comparison with the worst-case revenue of the optimal mechanism.

Lemma 3.2. *Given a set of marginal distributions $\{F_i\}_{i=1}^n$, if an SPM $\mathcal{M}_S = (\sigma, \mathbf{q})$ satisfies $2q_{\sigma(i)} \leq q_{\sigma(i+1)}$ for every $i \in [n-1]$, then for any $\pi \in \Pi$, $\text{Rev}(\mathcal{M}_S, \pi) \geq \frac{1}{2} \sum_i q_i p_i$.*

Proof. By Lemma 3.1, it suffices to show that $\text{Rev}(\mathcal{M}_S, \pi^*) \geq \frac{1}{2} \sum_i q_i p_i$. We have

$$\begin{aligned} \text{Rev}(\mathcal{M}_S, \pi^*) &= \sum_{i=1}^n (q_{\sigma(i)} - q_{\sigma(i-1)}) p_{\sigma(i)} \geq \sum_{i=1}^n (q_{\sigma(i)} - q_{\sigma(i)}/2) p_{\sigma(i)} \\ &= \frac{1}{2} \sum_{i=1}^n q_{\sigma(i)} p_{\sigma(i)} = \frac{1}{2} \sum_{i=1}^n q_i p_i. \quad \square \end{aligned}$$

Now we are ready to prove that SPM is a constant approximation to the mechanism with the optimal worst-case revenue. Given the marginal distributions $\{F_i\}_{i=1}^n$, we already know the worst-case joint distribution for an SPM is π^* , the one with the maximum positive correlation. Our plan is to construct a specific joint distribution $\hat{\pi} \in \Pi$, as well as a revenue upper bound U , such that

- U is an upper bound on the revenue of any IC and IR mechanism on distribution $\hat{\pi}$.
- Find a spread-out SPM \mathcal{M}_S such that $\text{Rev}(\mathcal{M}_S, \pi^*) = \Omega(U)$.

We construct a spread-out SPM \mathcal{M}_S in the following iterative algorithm ALGSEQUENCE (see below). In the output of ALGSEQUENCE, the order σ and quantiles \mathbf{q} define the specific SPM. We note that in ALGSEQUENCE the sets $\mathbf{Q}_{\text{free}}[0] \supset \mathbf{Q}_{\text{free}}[1] \supset \dots \supset \mathbf{Q}_{\text{free}}[n]$. On the other hand, each I_i is indeed an interval, since non of the intervals $[\frac{q_i}{2}, 2q_j]$ is contained in any other such interval. Also, by construction intervals I_i for all $i \in [n]$ are disjoint.

Algorithm 1 ALGSEQUENCE

```

1: Offered  $\leftarrow \emptyset$  ▷ Set of buyers who has been offered a price
2:  $\mathbf{Q}_{\text{free}}[0] \leftarrow [0, 1]$  ▷ Initially available quantiles for spread-out pricing
3: for  $k = 1$  to  $n$  do
4:   Pick  $(i, q_i) \in \operatorname{argmax}_{(i, q_i)} \{q_i \cdot V_i(q_i) \mid i \notin \text{Offered}, q_i \in \mathbf{Q}_{\text{free}}[k-1]\}$ .
5:    $\sigma(k) \leftarrow i$ ;   Offered  $\leftarrow \text{Offered} \cup \{i\}$  ▷ Offer price  $V_i(q_i)$  to  $i$ 
6:    $I_i \stackrel{\text{def}}{=} \mathbf{Q}_{\text{free}}[k-1] \cap [\frac{q_i}{2}, 2q_i]$  ▷ Reserve quantile interval around  $q_i$ 
7:    $\mathbf{Q}_{\text{free}}[k] \stackrel{\text{def}}{=} \mathbf{Q}_{\text{free}}[k-1] \setminus I_i$  ▷ Update available quantiles for spread-out pricing
8: end for
9: return  $(\sigma, \mathbf{q} = (q_1, \dots, q_n))$ 

```

Construction of distribution $\hat{\pi}$. The distribution $\hat{\pi}$ is closely connected to the description of the algorithm ALGSEQUENCE. To construct $\hat{\pi}$ we define $R_{\sigma(k)} \stackrel{\text{def}}{=} \mathbf{Q}_{\text{free}}[k]$ for all $k \in [n-1]$. For notational convenience, we redefine the last interval $I_{\sigma(n)} \stackrel{\text{def}}{=} \mathbf{Q}_{\text{free}}[n-1]$ by including all uncovered quantiles in $[0, 1] \setminus (I_{\sigma(1)} \cup \dots \cup I_{\sigma(n-1)})$ and also define the last $R_{\sigma(n)} \stackrel{\text{def}}{=} \mathbf{Q}_{\text{free}}[n-1]$. Hence, the collection of sets $\mathbf{I} = \{I_1, \dots, I_n\}$ becomes a partition of $[0, 1]$. By construction we have $R_{\sigma(k)} \cap I_{\sigma(j)} = \emptyset$ for all $j \leq k < n$ and $I_{\sigma(1)} \cup \dots \cup I_{\sigma(k)} \cup R_{\sigma(k)} = [0, 1]$ for all $k \in [n]$. The partition \mathbf{I} of the quantile space $[0, 1]$ and \mathbf{R} are used to generate the distribution $\hat{\pi}$.

Construction of distribution $\hat{\pi}$.

```

1: Draw a quantile  $q \sim U[0, 1]$ . Find  $k$  such that  $q \in I_{\sigma(k)}$ .
2: for  $i = \sigma(1) \dots \sigma(k-1)$  do
3:   Independently draw  $q_i \sim U[R_i]$  and let  $v_i = V_i(q_i)$ .
4: end for
5: for  $i = \sigma(k) \dots \sigma(n)$  do
6:   Let  $v_i \leftarrow V_i(q)$ .
7: end for
8: return  $(v_1, \dots, v_n)$ 

```

Figure 1 illustrates the structure of distribution $\hat{\pi}$. It is easy to check that $\hat{\pi}$ is consistent with the marginal distributions $\{F_i\}_{i=1}^n$. Indeed, for each given $i = \sigma(k)$, if the quantile q was drawn from the set $I_{\sigma(j)}$ with $j \leq k$ then $v_i = V_i(q)$. If q was drawn from $I_{\sigma(j)}$ with $j > k$, then q was drawn from $U[R_i]$, since $R_i = R_{\sigma(k)} = \cup_{j>k} I_{\sigma(j)}$. The latter means that $v_i = V_i(q_i)$ where q_i is sampled from $U[R_i]$ has exactly the same distribution as $V_i(q)$. Hence, the marginal distribution of v_i is $V_i(q)$ for $q \sim U[0, 1]$, i.e., $v_i \sim F_i$.

The next lemma sets up the upper bound U for the revenue of any mechanism on distribution $\hat{\pi}$.

Lemma 3.3. $\max_{\mathcal{M}} \operatorname{Rev}(\mathcal{M}, \hat{\pi}) \leq U \stackrel{\text{def}}{=} \sum_{k=1}^{n-1} \int_{I_{\sigma(k)}} \max_{j \geq k} \{V_{\sigma(j)}(q)\} dq + \sum_{i=1}^n \max_{q \in R_i} \{q \cdot V_i(q)\}.$

Proof. Let \mathcal{M} be the optimal mechanism for distribution $\hat{\pi}$. Fix quantile parameter $q \in [0, 1]$ which is used in $\hat{\pi}$ to generate \mathbf{v} . Suppose $q \in I_{\sigma(k)}$ and $i = \sigma(k)$. We consider the following two scenarios based on the buyers who contribute to the revenue of \mathcal{M} .

1. When $k \neq n$, and the payment comes from the buyers in $\{\sigma(k), \dots, \sigma(n)\}$. In this case, the payment is upper bounded by $\max_{j \geq k} \{v_{\sigma(j)}\}$, since it cannot exceed the valuation of

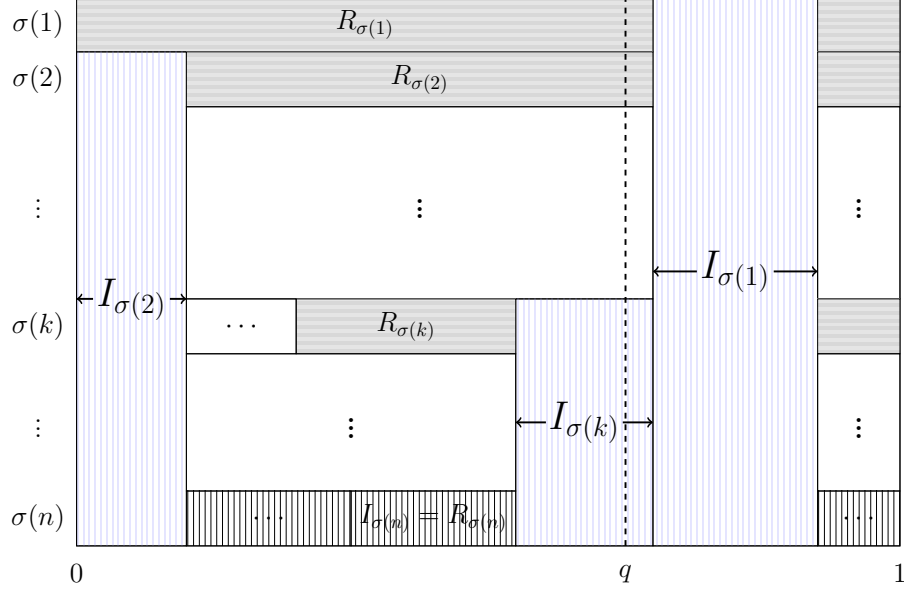


Figure 1: structure of distribution $\hat{\pi}$: after $q \sim U[0, 1]$ is drawn, the quantiles of buyers in $I_{\sigma(k)}$ region all equal q , i.e. $q_i = q$ for $i = \sigma(k) \dots \sigma(n)$; the quantiles of the remaining buyers are drawn independently from their respective R region, i.e. $q_i \sim U[R_i]$ for $i = \sigma(1) \dots \sigma(k-1)$.

the winning agent. When we integrate this over q ranging from 0 to 1, we get the term $\sum_{k=1}^{n-1} \int_{I_{\sigma(k)}} \max_{j \geq k} \{V_{\sigma(j)}(q)\} dq$ in RHS.⁷

2. Contribution to the revenue from agents $\{\sigma(1), \dots, \sigma(k-1)\}$, or when $k = n$ payment of agent $j = \sigma(n)$. In this case, we allow \mathcal{M} to observe the set $I_{\sigma(k)}$ from which quantile q was drawn.

If $k < n$, then we do not need to think about agents $\{\sigma(k), \dots, \sigma(n)\}$, since the payments from these buyers are accounted for in the previous case. On the other hand, according to the construction of $\hat{\pi}$ the values of agents $\{\sigma(1), \dots, \sigma(k-1)\}$ are drawn independently from each other (each $q_{\sigma(j)} \sim U[R_{\sigma(j)}]$ for $j \in [k-1]$). Similarly, if $k = n$, the values of agents $\{\sigma(1), \dots, \sigma(n-1)\}$ are drawn independently from each other and also independently from the initial quantile $q \in I_{\sigma(n)}$ ($v_{\sigma(n)} = V_{\sigma(n)}(q)$, where $q \sim U[I_{\sigma(n)}]$). Either way, the values of the agents $\{\sigma(1), \dots, \sigma(k-1)\}$ are drawn from independent distributions. In this situation the Myerson's auction gives the optimal revenue. Instead we allow the auctioneer to sell n items. Then the best mechanism is to sell items separately by posting the best individual prices. This gives an upper bound on the revenue of Myerson's auction. We integrate this bound over all quantiles $q \in [0, 1]$ and get an upper bound of $\max_{q_i \in R_i} \{q_i \cdot V_i(q_i)\}$ on the revenue extracted from each agent $i \in [n]$ (in fact, the bound is slightly better, because the probability that i accepts price $V_i(q_i)$, where $q_i \in R_i$ is smaller than or equal to q_i). Summing over all agents $i \in [n]$ we get the second term $\sum_{i=1}^n \max_{q_i \in R_i} \{q_i \cdot V_i(q_i)\}$ in RHS.

We conclude the proof by adding these two upper bounds together. \square

⁷Note that this upper bound could be achieved in theory by \mathcal{M} . Because when $q \in I_{\sigma(k)}$ is selected, the valuations of buyers $\sigma(k), \dots, \sigma(n)$ are perfectly correlated. An optimal mechanism will use other bids to infer the quantile parameter q and hence the value $v_i = V_i(q)$. Therefore, it will extract the full social surplus of buyer i .

Now we show that the output of algorithm ALGSEQUENCE the spread-out SPM $\mathcal{M}_{\mathcal{S}} = (\sigma, \mathbf{q})$ has expected revenue that approximates U for its worst-case distribtuion π^* . The upper bound in Lemma 3.3 consists of two terms. In the following (Lemma 3.4 and Lemma 3.5), we relate $\text{Rev}(\mathcal{M}_{\mathcal{S}}, \pi^*)$ to each of them separately.

Lemma 3.4. $2 \ln(4) \cdot \text{Rev}(\mathcal{M}_{\mathcal{S}}, \pi^*) \geq \sum_{k=1}^{n-1} \int_{I_{\sigma(k)}} \max_{j \geq k} \{V_{\sigma(j)}(q)\} dq.$

Proof. We treat I_i separately for each $i = \sigma(k)$, where $k \in [n - 1]$. To simplify notations let us define the function $g(q) \stackrel{\text{def}}{=} \max_{j \geq k} \{V_{\sigma(j)}(q)\}$. By the selection rule of q_i and i in ALGSEQUENCE we have $q \cdot g(q) = q \cdot \max_{j \geq k} \{V_{\sigma(j)}(q)\} \leq q_i \cdot V_i(q_i)$ for every $q \in I_i$. It implies that $g(q) \leq q_i V_i(q_i) \cdot \frac{1}{q}$ for every $q \in I_i$. Therefore, the terms under summation in the RHS of Lemma 3.4 are not more than

$$\int_{I_{\sigma(k)}} \max_{j \geq k} \{V_{\sigma(j)}(q)\} dq \leq \int_{I_i} q_i V_i(q_i) \cdot \frac{1}{q} dq \leq q_i V_i(q_i) \int_{\frac{q_i}{2}}^{2q_i} \frac{1}{q} dq = q_i V_i(q_i) \ln 4, \quad (6)$$

where the second inequality holds as $I_{\sigma(k)} \subseteq [\frac{q_i}{2}, 2q_i]$. The revenue of the SPM $\mathcal{M} = (\sigma, \mathbf{q})$ is at least $\frac{1}{2} \sum_i q_i V_i(q_i)$ according to Lemma 3.2. \square

Lemma 3.5. $2 \cdot \text{Rev}(\mathcal{M}_{\mathcal{S}}, \pi^*) \geq \sum_{i=1}^n \max_{q \in R_i} \{q \cdot V_i(q)\}.$

Proof. Let $i = \sigma(k)$ be the agent chosen at step $k \in [n]$ in ALGSEQUENCE. By definition $R_i \subset \mathbf{Q}_{\text{free}}[k - 1]$, which means that $q_i \cdot p_i = q_i \cdot V_i(q_i) \geq \max_{q \in R_i} \{q \cdot V_i(q)\}$. We sum these inequalities over all $i \in [n]$ and conclude the proof by employing Lemma 3.2. \square

By putting everything together we obtain the final theorem.

Theorem 3.1. *The revenue guarantee of $\mathcal{M}_{\mathcal{S}}$ is a $(2 \ln 4 + 2) (\approx 4.78)$ -approximation to the best worst-case revenue among all mechanisms.*

Proof. We have

$$\begin{aligned} (2 \ln 4 + 2) \cdot \min_{\pi \in \Pi} \text{Rev}(\mathcal{M}_{\mathcal{S}}, \pi) &= (2 \ln 4 + 2) \cdot \text{Rev}(\mathcal{M}_{\mathcal{S}}, \pi^*) \\ &\geq \sum_{k=1}^{n-1} \int_{I_{\sigma(k)}} \max_{j \geq k} \{V_{\sigma(j)}(q)\} dq + \sum_{i=1}^n \max_{q \in R_i} \{q \cdot V_i(q)\} \\ &\geq \max_{\mathcal{M}} \text{Rev}(\mathcal{M}, \hat{\pi}) \geq \max_{\mathcal{M}} \min_{\pi \in \Pi} \text{Rev}(\mathcal{M}, \pi), \end{aligned}$$

where the first line follows from Lemma 3.1, the second line is a combination of Lemma 3.4 and Lemma 3.5, and the third line is by Lemma 3.3. \square

3.2 Large Market Auctions

In this subsection we study a regime where the number of bidders is large and the auctioneer cannot observe bidder identities, e.g., in many auction scenarios on the internet. More concretely, we assume that bidders have identical marginal distributions $F_i = F$ for every $i \in [n]$ and consider the case when support V of F has a small size, not more than n , where n is the number of bidders. We show that in this case a sequential posted-price mechanism can achieve the optimal worst-case revenue. This is because by Lemma 3.1, the worst-case distribution for any SPM is the distribution

π^* with the maximal positive correlation. On the other hand, as the number of possible valuations is no more than the number of bidders, the seller running an SPM has enough attempts to query every possible value in the support of F , therefore extract the full social surplus.

Descending Sequential Posted Price Mechanism DSPM.

1. Let the support of F be $V = \{s_1, \dots, s_m\}$ ($m \leq n$). W.l.o.g. let $s_1 > s_2 > \dots > s_m$.
2. For $i = 1$ through m do:
 - Offer take-it-or-leave price s_i to bidder i

Theorem 3.2. *When n bidders have identical marginal distributions $F_i = F$ with support size $|V| \leq n$, sequential posted price mechanism DSPM has optimal worst-case revenue.*

Proof. Since DSPM is an SPM with monotonically decreasing sequence of prices, the worst-case distribution is π^* by Lemma 3.1. We repeat here the definition of π^* in the case when distributions F_i are identical: randomly draw $v \sim F$ and make $v_i = v$ for all $i \in [n]$. The revenue of DSPM achieved on any profile with $v_i = v$ for all $i \in [n]$ is exactly v , since the support of F is equal to or smaller than n and DSPM will have a chance to offer the price of v to a bidder with value v . Therefore, the worst-case revenue achieved by this mechanism is no less than $\sum_{v \in V} f(v) \cdot v$.

On the other hand, the revenue of any IR mechanism with respect to distribution π^* cannot be larger than the social welfare, which in this case is exactly $\sum_{v \in V} f(v) \cdot v$. \square

The above result can be extended to instances where the support of F is larger than the number of bidders, at a small multiplicative approximation error. To work with the continuous spaces we simply group the values into the ranges of $[(1 + \epsilon)^i, (1 + \epsilon)^{i+1})$ and can get the following corollary.

Corollary 3.1. *When n bidders have identical marginal distributions $F_i = F$ supported on $[1, h]$, an SPM achieves $(1 + \epsilon)$ -approximation to the worst-correlation revenue of the optimal mechanism, provided that $n \geq \frac{\ln h}{\ln(1+\epsilon)} (\approx \frac{\ln h}{\epsilon})$.*

3.3 Comparison with Second Price Auction with Common Reserve Price

In this subsection we investigate the relation of the sequential posted-price mechanisms and second price auctions with reserve prices in our correlation-robust framework. These two families of mechanisms are arguably the most important and well-studied DSIC mechanisms in the mechanism design area. In particular, in the classic setting when buyers' valuations are drawn i.i.d, Myerson's celebrated auction falls into the family of second price auction with common reserve price, while SPM is shown to approach optimal revenue asymptotically [10].

Interestingly, in our correlation-robust framework, under the assumption that buyers share an identical marginal distribution, the best SPM has almost matching or even better worst-correlation revenue than the best second price auction with uniform (anonymous) reserve price. We formalize the statement below and provide the proof in Appendix C.

Lemma 3.6. *When buyers' valuations follow identical marginal distributions. The revenue guarantee of the best sequential posted price auction is always a $\frac{n}{n-1}$ -approximation to the best worst-case revenue among all second price auctions with common reserve price.*

4 Lookahead Auctions

In this section, we study mechanisms with an additional restriction that the auctioneer must sell the item to the highest bidder. This family of mechanisms is also known as the lookahead auctions in the literature. We first give the formal definition of lookahead auctions, which is slightly different from the literature when there are multiple buyers sharing the same highest bid.

Lookahead Auction Family \mathcal{L} . A mechanism $\mathcal{M} = (\mathbf{x}, \mathbf{p})$ is a lookahead auction if for any $\mathbf{v} \in \mathbf{V}$ and $i \in [n]$ such that $v_i \neq \max_j v_j$, the allocation to the i -th buyer $x_i(\mathbf{v}) = 0$. Let $\mathcal{M}_{\mathcal{L}}$ be the lookahead auction with the best robust-correlation revenue.

Theorem 4.1. *The revenue guarantee of $\mathcal{M}_{\mathcal{L}}$ is a 2-approximation to the best worst-case revenue among all mechanisms.*

Proof. According to Ronen's result [30], we know that for any joint distribution π , $\max_{\mathcal{M}} \text{Rev}(\mathcal{M}, \pi) \leq 2 \cdot \max_{\mathcal{M} \in \mathcal{L}} \text{Rev}(\mathcal{M}, \pi)$. Therefore,

$$\begin{aligned} \max_{\mathcal{M}} \min_{\pi(\mathcal{M}) \in \Pi} \text{Rev}(\mathcal{M}, \pi) &= \min_{\pi \in \Pi} \max_{\mathcal{M}(\pi)} \text{Rev}(\mathcal{M}, \pi) \\ &\leq 2 \cdot \min_{\pi \in \Pi} \max_{\mathcal{M}(\pi) \in \mathcal{L}} \text{Rev}(\mathcal{M}, \pi) = 2 \cdot \max_{\mathcal{M} \in \mathcal{L}} \min_{\pi(\mathcal{M}) \in \Pi} \text{Rev}(\mathcal{M}, \pi). \end{aligned}$$

Where the first and the last equality follow from the von Neumann minimax theorem and the convexity of \mathcal{L} and the set of all truthful auctions. \square

For known correlated distribution and under appropriate oracle assumption of how to access the distribution (since the distribution itself may have exponential in n size), the best lookahead auction can be computed in polynomial time.

Is the optimal correlation-robust lookahead auction computable in polytime?

Note that this question does not require any additional assumptions on the oracle access model, as the input to the problem is given by n single dimensional distributions $\{F_i\}_{i \in [n]}$ that can be explicitly described as the input to an algorithm. It is an interesting open question, which we do not know how to solve. In the remaining part of this section, we show how to solve this problem for an important special case of the identical marginal distributions.

For each permutation $\sigma \in S_n$, let \mathbf{v}^σ be the value profile that $v_i^\sigma = v_{\sigma(i)}$ for all i . As mentioned in Section 2, for each distribution and mechanism, by taking the average over all permutations of the buyer identities, we may restrict our study to symmetric joint distributions and mechanisms. Let $\Pi_{\text{sym}} \stackrel{\text{def}}{=} \{\pi \in \Pi : \pi(\mathbf{v}) = \pi(\mathbf{v}^\sigma), \forall \mathbf{v}, \sigma\}$ and $\mathcal{L}_{\text{sym}} \stackrel{\text{def}}{=} \{\mathcal{M} \in \mathcal{L} : x_i(\mathbf{v}) = x_{\sigma(i)}(\mathbf{v}^\sigma), p_i(\mathbf{v}) = p_{\sigma(i)}(\mathbf{v}^\sigma), \forall \mathbf{v}, \sigma, i\}$. We have that

$$\max_{\mathcal{M} \in \mathcal{L}} \min_{\pi \in \Pi} \text{Rev}(\mathcal{M}, \pi) = \max_{\mathcal{M} \in \mathcal{L}_{\text{sym}}} \min_{\pi \in \Pi_{\text{sym}}} \text{Rev}(\mathcal{M}, \pi). \quad (7)$$

Consider the following simple family of value profiles, which we refer as *high-low* types:

$$\mathbf{V}_{\text{HL}} \stackrel{\text{def}}{=} \{\mathbf{v} : \exists i \in [n], v_{\text{H}} \geq v_{\text{L}} \text{ such that } v_i = v_{\text{H}} \text{ and } v_j = v_{\text{L}} \text{ for } j \neq i\}.$$

In other words, there exists a high-type buyer i with value v_{H} while all other buyers have the same value v_{L} . Let Π_{HL} be the family of distributions that are supported on \mathbf{V}_{HL} , i.e. $\Pi_{\text{HL}} \stackrel{\text{def}}{=} \{\pi \in \Pi_{\text{sym}} : \pi(\mathbf{v}) = 0, \forall \mathbf{v} \notin \mathbf{V}_{\text{HL}}\}$. Our next goal will be to show that the worst-correlation distribution with

respect to (7) belongs to Π_{HL} . Note that all lookahead mechanisms can be described as posting a randomized price to the highest bidder based on the other bidders' values. Intuitively, the worst distribution should reveal the least possible information about the highest bid by showing the remaining bids. We show that following this logic the most efficient way to hide information is by using only high-low types.

Given a distribution $\pi \in \Pi_{\text{sym}}$, we consider the following transformation into a distribution $\pi_{\text{HL}} \in \Pi_{\text{HL}}$. For each $\mathbf{v} = (v_1, \dots, v_n)$, We distribute the probability mass $\pi(\mathbf{v})$ uniformly among the following high-low types $\{(v_i, v_j, \dots, v_j)\}_{j \neq i}$, where i is sampled uniformly among highest bidders and then j is sampled uniformly among all bidders except i . We denote this partial distribution on \mathbf{V}_{HL} transformed from \mathbf{v} as \mathbf{v}_{HL} . Applying this process to all \mathbf{v} , we obtain a distribution π_{HL} supported on \mathbf{V}_{HL} . Moreover, by construction π_{HL} must be consistent with the marginal f , i.e., $\pi_{\text{HL}} \in \Pi$. Indeed, our transformation per any $\mathbf{v} \in \mathbf{V}$ does not change the average (over the set $[n]$ of bidders) marginal distribution⁸. We claim that $\max_{\mathcal{M}} \text{Rev}(\mathcal{M}, \pi) \geq \max_{\mathcal{M}} \text{Rev}(\mathcal{M}, \pi_{\text{HL}})$.

Before going to the details of our proof, we introduce extra notations to describe a partial lookahead mechanism defined on high-low types. For $v_{\text{H}} > v_{\text{L}}$, the mechanism will look at the value v_{L} and offer a possibly randomized price $p \geq v_{\text{L}}$ to the high type bidder. We denote $\delta(v_{\text{L}}, p)$, where $p \geq v_{\text{L}}$, the probability of posting price p . IC and IR conditions are satisfied automatically in this form, and the feasibility of any symmetric single-item lookahead auction can be described by the two following conditions: $\sum_{p \geq v_{\text{L}}} \delta(v_{\text{L}}, p) \leq 1$ and $\delta(v_{\text{L}}, v_{\text{L}}) \leq \frac{1}{n}$. For any truthful symmetric lookahead auction \mathcal{M} , we define \mathcal{M}_{HL} as its restriction on high-low type and we know that it can be represented by the above $\delta(\cdot, \cdot)$.

Next, we provide a reversal transformation to extend a partial mechanism \mathcal{M}_{HL} to a symmetric truthful mechanism $\overline{\mathcal{M}_{\text{HL}}}$ on the full support. For each value profile \mathbf{v} , let v_{second} be the second highest bid, which is equal to the highest bid when there are multiple highest bidders. Recall that \mathbf{v}_{HL} is a partial distribution on \mathbf{V}_{HL} with total probability mass $\pi(\mathbf{v})$. The mechanism $\overline{\mathcal{M}_{\text{HL}}}$ on \mathbf{v} simply mimics the mechanism \mathcal{M}_{HL} on \mathbf{v}_{HL} with the following small modification: when the mechanism \mathcal{M}_{HL} offers a price which is less than v_{second} , the mechanism $\overline{\mathcal{M}_{\text{HL}}}$ will offer the price v_{second} instead. By this definition, it is easy to verify that $\overline{\mathcal{M}_{\text{HL}}}$ is indeed a symmetric truthful lookahead auction on the whole support. Another conclusion from this transformation is that the payment of $\overline{\mathcal{M}_{\text{HL}}}$ on \mathbf{v} is at least the expected payment of \mathcal{M}_{HL} on \mathbf{v}_{HL} since $\overline{\mathcal{M}_{\text{HL}}}$ can always sell the item when \mathcal{M}_{HL} can sell it and with a same or larger price. Sum over all \mathbf{v} , we get the conclusion that

$$\text{Rev}(\overline{\mathcal{M}_{\text{HL}}}, \pi) \geq \text{Rev}(\mathcal{M}_{\text{HL}}, \pi_{\text{HL}}).$$

Now, we are ready to prove $\max_{\mathcal{M}} \text{Rev}(\mathcal{M}, \pi) \geq \max_{\mathcal{M}} \text{Rev}(\mathcal{M}, \pi_{\text{HL}})$ by the following chain of inequalities:

$$\max_{\mathcal{M}} \text{Rev}(\mathcal{M}, \pi_{\text{HL}}) = \max_{\mathcal{M}_{\text{HL}}} \text{Rev}(\mathcal{M}_{\text{HL}}, \pi_{\text{HL}}) \leq \max_{\mathcal{M}_{\text{HL}}} \text{Rev}(\overline{\mathcal{M}_{\text{HL}}}, \pi) \leq \max_{\mathcal{M}} \text{Rev}(\mathcal{M}, \pi).$$

Taking the minimum over all $\pi \in \Pi_{\text{sym}}$, we have

$$\min_{\pi \in \Pi_{\text{sym}}} \max_{\mathcal{M}} \text{Rev}(\mathcal{M}, \pi) = \min_{\pi \in \Pi_{\text{HL}}} \max_{\mathcal{M}_{\text{HL}}} \text{Rev}(\mathcal{M}_{\text{HL}}, \pi).$$

Therefore, we can calculate the optimal partial mechanism by solving the following LP and

⁸Speaking in algebraic terms, our transformation does not change bidder marginal distributions when applied to all types in the orbit of any particular $\mathbf{v} \in \mathbf{V}$ under the action of permutation group S_n .

extend it into the whole support by the previous procedure.

$$\begin{aligned}
& \max \quad n \cdot \sum_v f(v) \cdot \lambda(v) \\
& \text{s. t.} \quad \lambda(v) \leq v\delta(v, v) \quad \forall v \\
& \quad \lambda(v_H) + (n-1)\lambda(v_L) \leq \sum_{v_H \geq p \geq v_L} p\delta(v_L, p) \quad \forall v_H > v_L \\
& \quad \sum_{p \geq v_L} \delta(v_L, p) \leq 1 \quad \forall v_L \\
& \quad \delta(v, v) \leq \frac{1}{n} \quad \forall v
\end{aligned}$$

Note that this LP involves only quadratic number of variables and polynomial number of constraints, hence, is computable in polynomial time.

For notational simplicity, we presented our result for the original lookahead auction, which only sells item to the highest value bidder. Both previous results can be extended to more general k -lookahead auction, which allows allocating the item to one of the top k bidders.

Theorem 4.2. *The revenue guarantee of k -lookahead auction is a $\frac{e^{1-1/k}}{e^{1-1/k}+1}$ -approximation to the best worst-case revenue among all mechanisms. For constant k and symmetric distributions, the optimal k -lookahead auction can be computed in polynomial time.*

5 Conclusion and Open problems

In this work we adapted the correlation-robust framework developed for the additive monopoly problem [13] to the well studied setting of single-item auction. Our approach can be considered as a worst-case counterpart to the extensively studied Myerson's setting with independent prior. In this sense the Myerson's auction can be interpreted as the optimal average-case mechanism. This framework provides meaningful guarantees for the difficult problem with general (correlated) prior, and also is much easier to apply in practice than the standard Bayesian approach, since it only requires information about single dimensional distributions from each agent (which can be efficiently learned from data) and also avoids unrealistically good conclusions of [18].

In this work we showed the power of sequential posted price mechanisms and proved that SPMs achieve approximately optimal revenue in the correlation-robust framework. We also studied lookahead auctions and gave some computational results. Our work leaves a lot of open questions. Perhaps the most important (as our numerical experiments indicate also quite difficult) is the question of describing/computing the optimal correlation-robust auction. On the other hand, we found that, when the number of bidders is large, the format of the optimal auction does not matter much and a simple SPM extracts the optimal worst-correlation revenue. A simpler but still very interesting question is to find other regimes where the optimal correlation-robust mechanisms are easy to describe or to compute.

We also showed how to evaluate the worst-case revenue of most of the existing deterministic DSIC mechanisms such as SPM, second price auctions with reserves, and Myerson's auction. The question to which extent these computational results extend to the randomized mechanisms remains widely open. We only know the answer for the class of randomized SPM (and that one only for the prices offered in the decreasing order).

Correlation-robust framework offers some interesting computational questions for lookahead auctions. In the classic Bayesian framework the computation of the best lookahead auction is quite

trivial and simply requires careful description of the oracle access to the prior distribution. In the correlation-robust framework the input to the problem has a succinct representation, and for special case of anonymous bidders we already know that optimal lookahead auction can be computed in polynomial time. It is intriguing question to extend this result to an arbitrary set of distributions. On the other hand, it is unclear if the correlation-robust approximation guarantee of lookahead auctions can be improved for the case of identical marginals. Consequently, it is also an interesting question whether the approximation ratio of lookahead auctions can be improved in the standard Bayesian setting if the prior distribution is symmetric, i.e., invariant under any permutation of bidders' identities.

References

- [1] Michael Albert, Vincent Conitzer, and Giuseppe Lopomo. Assessing the robustness of cremer-mclean with automated mechanism design. In *AAAI*, pages 763–769, 2015.
- [2] Michael Albert, Vincent Conitzer, and Peter Stone. Automated design of robust mechanisms. In *AAAI*, pages 298–304, 2017.
- [3] Michael Albert, Vincent Conitzer, and Peter Stone. Mechanism design with unknown correlated distributions: Can we learn optimal mechanisms? In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, pages 69–77, 2017.
- [4] Dirk Bergemann, Benjamin A Brooks, and Stephen Morris. Informationally robust optimal auction design. 2016.
- [5] Dirk Bergemann and Stephen Morris. Robust mechanism design. *Econometrica*, 73(6):1771–1813, 2005.
- [6] Dirk Bergemann and Stephen Morris. Robust predictions in games with incomplete information. *Econometrica*, 81(4):1251–1308, 2013.
- [7] Dirk Bergemann and Stephen Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522, 2016.
- [8] Dirk Bergemann, Stephen Morris, and Satoru Takahashi. Efficient auctions and interdependent types. *The American Economic Review*, 102(3):319–324, 2012.
- [9] Dirk Bergemann and Karl Schlag. Robust monopoly pricing. *Journal of Economic Theory*, 146(6):2527–2543, 2011.
- [10] Liad Blumrosen and Thomas Holenstein. Posted prices vs. negotiations: an asymptotic analysis. In *EC*, page 49. ACM, 2008.
- [11] Benjamin A Brooks and Songzi Du. Optimal auction design with common values: An informationally-robust approach. 2018.
- [12] Gabriel Carroll. Robustness and linear contracts. *The American Economic Review*, 105(2):536–563, 2015.
- [13] Gabriel Carroll. Robustness and separation in multidimensional screening. *Econometrica*, 85(2):453–488, 2017.

- [14] Gabriel Carroll and Delong Meng. Robust contracting with additive noise. *Journal of Economic Theory*, 166:586–604, 2016.
- [15] Xue Chen, Guangda Hu, Pinyan Lu, and Lei Wang. On the approximation ratio of k-lookahead auction. In *International Workshop on Internet and Network Economics*, pages 61–71. Springer, 2011.
- [16] Yi-Chun Chen and Siyang Xiong. Genericity and robustness of full surplus extraction. *Econometrica*, 81(2):825–847, 2013.
- [17] Kim-Sau Chung and Jeffrey C Ely. Foundations of dominant-strategy mechanisms. *The Review of Economic Studies*, 74(2):447–476, 2007.
- [18] Jacques Crémer and Richard P. McLean. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53(2):345–361, 1985.
- [19] Shahar Dobzinski, Hu Fu, and Robert D Kleinberg. Optimal auctions with correlated bidders are easy. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 129–138. ACM, 2011.
- [20] Shahar Dobzinski and Nitzan Uziely. Revenue loss in shrinking markets. *arXiv preprint arXiv:1706.08148*, 2017.
- [21] Hu Fu, Nima Haghpanah, Jason Hartline, and Robert Kleinberg. Optimal auctions for correlated buyers with sampling. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 23–36. ACM, 2014.
- [22] Daniel F Garrett. Robustness of simple menus of contracts in cost-based procurement. *Games and Economic Behavior*, 87:631–641, 2014.
- [23] Nick Gravin and Pinyan Lu. Separation in correlation-robust monopolist problem with budget. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2069–2080. SIAM, 2018.
- [24] John C. Harsanyi. Games with incomplete information played by "bayesian" players, i-iii. *Management Science*, 14(5):320–334, 1968.
- [25] Paul Klemperer. Auction theory: A guide to the literature. *Journal of Economic Surveys*, 13(3):227–286, 1999.
- [26] Giuseppe Lopomo, Luca Rigotti, and Chris Shannon. Uncertainty in mechanism design. *Revise and resubmit at Review of Economic Studies*, 2009.
- [27] Randolph McAfee and Philip Reny. Correlated information and mechanism design. *Econometrica*, 60(2):395–421, 1992.
- [28] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- [29] Christos H Papadimitriou and George Pierrakos. On optimal single-item auctions. In *Proceedings of the 43rd annual ACM Symposium on Theory of Computing*, pages 119–128. ACM, 2011.

- [30] Amir Ronen. On approximating optimal auctions. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 11–17. ACM, 2001.
- [31] Amir Ronen and Amin Saberi. On the hardness of optimal auctions. In *Proceedings of the 43rd Symposium on Foundations of Computer Science*, pages 396–405, 2002.

A Robustness Analysis of Existing Auctions

We emphasize that the correlation-robust framework can be used to evaluate any specific auction in the *worst-correlation metric*. These evaluations can help to compare and decide between different auction formats. Motivated by this consideration we study the problem of computing minimal revenue of a given auction over all distributions Π with feasible marginals. Specifically, we show in this section how to calculate the revenue of sequential posted price mechanisms, second price auction with individual reserves, and Myerson auction in the correlation-robust framework. Our result for the sequential posted price mechanisms (SPM) holds even for randomized mechanisms, provided that randomization is only used over Pareto-efficient deterministic SPMs, i.e., the sequence of posted prices that are offered to the buyers is non-increasing.

Theorem A.1. *There are polynomial time algorithms to calculate the worst-correlation revenue for*

1. *second price auction with individual reserves*⁹,
2. *Myerson auction*,
3. *sequential posted price mechanisms*,
4. *randomized sequential posted price mechanisms*.

For the last result we assume that every deterministic pricing mechanism in the support of the randomized mechanism offers prices in non-increasing order.

To calculate the revenue in the correlation-robust framework we shall consider the dual LP (2). We observe that it has polynomially many variables, namely, $\{\lambda_i(v_i)\}_{i \in [n], v_i \in V_i}$, but exponential in n number of constraints. Therefore, a naïve approach of solving this LP directly would take exponential in n number of steps. Instead we adopt three different strategies for the above families of mechanisms. First, for the second price auction with individual reserves, we show that $\{\lambda_i(v_i)\}_{v_i \in V_i}$ are monotone in the optimal solution of the dual LP (2) for any $i \in [n]$. To this end, we prove a more general fact that as long as $\sum_{i=1}^n p_i(\mathbf{v})$ is a monotone function of \mathbf{v} the set of $\{\lambda_i(v_i)\}_{v_i \in V_i}$ must also be monotone for each $i \in [n]$. The monotonicity of $\{\lambda_i(v_i)\}_{v_i \in V_i}$ allows us to prune redundant constraints from the dual LP (2) and efficiently reduce it to a manageable size. We note that this approach fails to solve the case of the Myerson auction, since the total payment in the Myerson auction may exhibit non-monotone behavior. Instead, we introduce extra variables and develop a polynomial size LP that is equivalent to LP (2). Our construction builds on the observation that for the Myerson auction, there are only polynomially many possible outcomes, i.e., the winner of the auction and her payment, and each outcome defines a simple-to-describe set

⁹The auction is defined as follows: Each bidder is associated with a reserve price. Consider the bidders who bid at least their individual reserves. If there are at least two such bidders, the highest bidder wins and pays the larger between his reserve and the second highest bid. If there is only one such bidder, that bidder wins and pays the reserve price.

of buyers' types. Finally, we explicitly construct the worst case distribution for all sequential posted price mechanisms with non-increasing prices (SPM) and give closed-form expression for the worst-correlation revenue. Moreover, since all mechanisms within this family share the same worst-case distribution, it is easy to compute the worst-correlation revenue of randomized sequential posted price mechanisms.

Proof of Theorem A.1:

(1) *Second price auction with individual reserves.*

First, we consider second price auction with individual reserves. To this end, we show that $\{\lambda_i\}_{v_i \in V_i}$ must be monotone for all $i \in [n]$ for a large family of mechanisms.

Lemma A.1. *For any \mathcal{M} with monotone payments (i.e., $\forall \check{\mathbf{v}} \preceq \mathbf{v}, \sum_i p_i(\check{\mathbf{v}}) \leq \sum_i p_i(\mathbf{v})$), there exists an optimal $\lambda = \{\lambda_i(v_i)\}_{i \in [n], v_i \in V_i}$ such that $\{\lambda_i(v_i)\}_{v_i \in V_i}$ are non decreasing for all i .*

Proof. Let λ be the optimal solution to the LP. We argue that $\bar{\lambda}_i(v_i) = \max_{v \leq v_i} \lambda_i(v)$ is also an optimal solution. Note that $\bar{\lambda}_i$ is monotonically increasing for all i and $\bar{\lambda}_i(v_i) \geq \lambda_i(v_i)$ for all i and v_i . Hence, the objective function $\sum_{i=1}^n \sum_{v_i \in V_i} f_i(v_i) \cdot \bar{\lambda}_i(v_i)$ is equal to or greater than the optimal. It suffices to show that $\bar{\lambda} = \{\bar{\lambda}_i\}_{i \in [n], v_i \in V_i}$ is feasible, i.e., $\sum_{i=1}^n \bar{\lambda}_i(v_i) \leq \sum_{i=1}^n p_i(\mathbf{v})$ for all \mathbf{v} .

Fix a value profile \mathbf{v} . Suppose $\lambda_i(v_i) = \lambda_i(\check{v}_i)$ for all i , where $\check{v}_i \leq v_i$. Let $\check{\mathbf{v}} = (\check{v}_1, \dots, \check{v}_n)$. Note that $\check{\mathbf{v}} \preceq \mathbf{v}$. By the original constraint corresponding to $\check{\mathbf{v}}$, we have

$$\sum_{i=1}^n \bar{\lambda}_i(v_i) = \sum_{i=1}^n \lambda_i(\check{v}_i) \leq \sum_{i=1}^n p_i(\check{\mathbf{v}}) \leq \sum_{i=1}^n p_i(\mathbf{v}). \quad \square$$

We note that any second price auction with individual reserves has monotone payments. By Lemma A.1 the solution to the dual LP (2) does not change if we add more linear constraints saying that λ is monotone. Now we are going to point out only a few (polynomially many) important constraints in the latter LP, so that all the remaining constraints do not affect the optimal solution. Let r_i be the reserve price for each buyer $i \in [n]$. We partition the space of all value profiles \mathbf{V} into $\{\mathbf{V}_{i,z_i}\}_{i \in [n], z_i \in V_i}$ according to the winner of the auction and her payment z_i and \mathbf{V}_0 when no buyer gets the item. Specifically, in \mathbf{V}_{i,z_i} we consider all value profiles \mathbf{v} such that buyer i wins the auction and pays exactly z_i . In this case, we know that buyer i has the highest bid and every other buyer $j \neq i$ has a value either smaller than her individual reserve price r_j , or smaller than the price z_i . That is,

$$\mathbf{V}_{i,z_i} \subseteq \left\{ \mathbf{v} \mid v_i \geq z_i, \quad \forall j \neq i, v_j \leq \max\{z_i, r_j^-\} \right\}, \quad \text{where } r_j^- \stackrel{\text{def}}{=} \max_{v_j < r_j} v_j.$$

Then any constraint in the dual LP (2) for $\mathbf{v} \in \mathbf{V}_{i,z_i}$ is implied by the monotonicity of λ and the following inequality

$$\sum_{j \neq i} \lambda_j(\max\{z_i, r_j^-\}) + \lambda_i(v_i^{\max}) \leq z_i,$$

where v_i^{\max} is the maximum value in the support of F_i . Moreover, the latter inequality is weaker than the constraint of dual LP (2) corresponding to the value profile $\check{\mathbf{v}}$, where $\check{v}_i = v_i^{\max}$ and $\check{v}_j = \max\{z_i, r_j^-\}$ for all $j \neq i$. For \mathbf{V}_0 , we know that all buyers bid smaller than their individual reserves, i.e. $\mathbf{V}_0 = \{\mathbf{v} \mid \forall i, v_i < r_i\}$. Then, due to monotonicity of λ it suffices to use only $\sum_i \lambda_i(r_i^-) \leq 0$ constraint. We observe that there are only polynomially many different categories $\{\mathbf{V}_{i,z_i}\}_{i \in [n], z_i \in V_i}$ and the monotonicity of λ can be described with only $\sum_{i \in [n]} |V_i|$ linear constraints. Thus we can prune the remaining constraints and solve the LP efficiently.

(2) *Myerson auction.*

Next, we study the worst-correlation revenue of the Myerson's auction. Let r_i be the reserve price and ϕ_i be the ironed virtual valuation for each buyer i . Following the proof for second price auction with individual reserves, we partition \mathbf{V} into $\{\mathbf{V}_{i,z_i}\}_{i \in [n], z_i \in V_i}$ according to the winner of the auction and her payment $z_i \geq r_i$ and \mathbf{V}_0 when no buyer wins the item. Specifically, we consider all value profiles \mathbf{v} so that buyer i wins the auction and pays exactly z_i . We know that the buyer i has the highest virtual value $\phi_i(z_i) \geq 0$ and every other buyer $j \neq i$ has a virtual value smaller than $\phi_i(z_i)$ ¹⁰. That is

$$\mathbf{V}_{i,z_i} \subseteq \{\mathbf{v} \mid v_i \geq z_i, \quad \forall j \neq i, \phi_j(v_j) < \phi_i(z_i)\}.$$

Then any constraint in dual LP (2) for $\mathbf{v} \in \mathbf{V}_{i,z_i}$ is implied by the following inequality

$$\sum_{j \neq i} \max_{v_j: \phi_j(v_j) < \phi_i(z_i)} \lambda_j(v_j) + \max_{v_i \geq z_i} \lambda_i(v_i) \leq z_i. \quad (8)$$

Moreover, as in the previous case this inequality is implied by the constraint of the dual LP (2) corresponding to the value profile $\check{\mathbf{v}}$, where $\check{v}_i = \operatorname{argmax}_{v_i \geq z_i} \lambda_i(v_i)$ and $\check{v}_j = \operatorname{argmax}_{v_j: \phi_j(v_j) < \phi_i(z_i)} \lambda_j(v_j)$ for all $j \neq i$. Also $\mathbf{V}_0 = \{\mathbf{v} \mid \forall i, v_i < r_i\}$. Hence, the constraints for all $\mathbf{v} \in \mathbf{V}_0$ are implied by

$$\sum_i \max_{v_i \leq r_i^-} \lambda_i(v_i) \leq 0. \quad (9)$$

We observe that there are only polynomially many different categories $\{\mathbf{V}_{i,z_i}\}_{i \in [n], z_i \in V_i}$. To succinctly describe the max operators in (8) and (9), we introduce extra variables $\{\bar{\lambda}_i(v_i)\}_{i \in [n], z_i \in V_i}$ with a few extra constraints:

$$\bar{\lambda}_i(\cdot) \text{ is monotone} \quad \forall i \in [n] \quad \text{and} \quad \bar{\lambda}_i(v_i) \geq \lambda_i(v_i), \quad \forall i, v_i \in V_i.$$

Then constraints (8) and (9) can be expressed as

$$\sum_{j \neq i} \bar{\lambda}_j(\phi_j^{-1}(\phi_i(z_i))) + \bar{\lambda}_i(v_i^{\max}) \leq z_i \quad \text{and} \quad \sum_i \bar{\lambda}_i(r_i^-) \leq 0,$$

where $\phi_j^{-1}(\phi_i(z_i)) \stackrel{\text{def}}{=} \max\{v_j \mid \phi_j(v_j) < \phi_i(z_i)\}$. Thus, we obtain a polynomial size LP that is equivalent to the dual LP (2) and can solve this new LP efficiently.

(3) & (4) *Deterministic and randomized sequential posted price mechanisms.*

For this case, we first consider SPM that posts non-increasing sequence of prices. It has already been shown in Lemma 3.1 that the worst-case distribution for SPM is the distribution π^* with the maximal positive correlation. Therefore, to calculate the worst revenue guarantee, it suffices to calculate the sum in (5) in Section 3. Furthermore, observe that the worst distribution π^* does not depend on the mechanism. Hence, for any randomized mechanism with only non increasing sequences of posted prices the worst-case distribution is again π^* . Then it is easy to compute the worst-correlation revenue of the mechanism by summing the expression (5) over all random price sequences¹¹.

Finally, the approach used in Myerson auction of partitioning the value space according to the winner of the auction and her payment, can also be used to get a polynomial time algorithm for sequential posted price mechanisms with arbitrary sequence of prices (e.g., which is not necessarily monotone). \square

¹⁰For simplicity of the presentation we assume that there are no ties for virtual values.

¹¹If there are exponentially many price sequences in the support of the randomized mechanism, one can simply do Monte-carlo simulation by sampling valuations from π^* and running \mathcal{M} .

B Optimal Auctions: Examples

This section provides numerical examples of several optimal correlation-robust auctions. In order to better understand the framework, we run experiments for different marginal distributions to explore the structure of the optimal mechanisms.

In our first example there are two buyers with identical marginal distributions $U[0, 1]$. In this case, we could analytically prove that the optimal mechanism has a reasonably simple description: it is the second price auction with a random anonymous reserve, where the reserve price is drawn from $U[0, \frac{3}{4}]$. The optimal $\lambda(v) = \frac{2v^2}{3}$ for $v \in [0, \frac{3}{4}]$ and $\lambda(v) = \frac{3}{8}$ for $v \in (\frac{3}{4}, 1]$ and the worst-correlation revenue equals to $\frac{3}{8}$. It is worthwhile to mention that second price auction and Myerson auction achieve worst-correlation revenue $\frac{1}{4}$ and $\frac{5}{16}$ respectively. Given the simple format of the optimal mechanism in the latter setting, it is natural to ask:

Does the result generalize to (1) more buyers and (2) other distributions?

Our further numerical experiments give a strong evidence that the answer to both of these questions is negative. Consider a minimal extension of the previous example to the case of $n = 3$ buyers with identical marginal distribution $U[0, 1]$. We observe complicated structure of the optimal mechanism¹²: the mechanism not only allocates the item to the highest bidder, but also *allocates to the second highest bidder with significant probability*. The latter fact is quite disconcerting, as the optimal auction exhibits such unruly behavior even in this simple symmetric setting. Furthermore, this fact also rules out the possibility that any lookahead auction can be optimal in general. We remark that we only do not know how to describe the optimal mechanism in a simple language, but still admit that there could be interesting special cases where the optimal auctions are well behaved. We warn our readers about possible difficulties and leave the two latter questions as an interesting open research direction.

Next, we consider the case of two symmetric buyers with marginal uniform distribution $U(1, 2, 3)$, i.e., the uniform distribution over the discrete values 1, 2, and 3. It might seem that at least in this extremely simple setting the optimal mechanism must have a clean description. Surprisingly, it is not the case. We present the optimal mechanism in the following matrices described in a symmetric way, i.e., independent of the buyer identities, where $x(v_1, v_2)$ and $p(v_1, v_2)$ are the allocation and payment to the buyer who bids v_1 while the other one bids v_2 :

$$x : \begin{array}{ccc} & 1 & 2 & 3 \\ \begin{array}{c} 1 \\ 2 \\ 3 \end{array} & \begin{pmatrix} \frac{1}{3} & 0 & 0 \\ \frac{5}{6} & \frac{1}{2} & 0 \\ 1 & 1 & \frac{1}{2} \end{pmatrix} & \text{and } p : & \begin{array}{ccc} & 1 & 2 & 3 \\ \begin{array}{c} 1 \\ 2 \\ 3 \end{array} & \begin{pmatrix} \frac{1}{3} & 0 & 0 \\ \frac{4}{3} & 1 & 0 \\ \frac{11}{6} & \frac{5}{2} & \frac{3}{2} \end{pmatrix} \end{array}$$

The optimal $\lambda(1) = \frac{1}{3}$, $\lambda(2) = 1$, $\lambda(3) = \frac{3}{2}$ and the worst-correlation revenue is $\frac{17}{9}$. One possible description of this mechanism is as a second price auction with random individual reserves. However, the reserve prices are much trickier as we illustrate below. First, one chooses a pair of reserve prices uniformly at random from the multiset $\{(1, 2), (2, 1), (2, 3), (3, 2)\}$. Then we run second price auction with tie broken in favor of the bidder with the larger reserve. For instance, when the buyers have values (3, 2) and reserve prices are (1, 2), the first buyer wins and pays 3, since if she bids 2, the second buyer would win the auction by the higher reserve tie-breaking rule. Some part of the auction's complexity could be explained as a result of the tricky tie-breaking

¹²We do not have an analytical expression for the optimal mechanism in this case. Instead, we discretize the uniform distribution and observe the numerical solution to the LP (4).

issues for the discrete types. Nonetheless, the complicated structure we already observe in such small instances precludes us from making any good conjecture about the optimal mechanism.

Finally, we consider an asymmetric setting to avoid any tie-breaking complications from the previous example. In this setting there are $n = 2$ buyers with marginal distributions $U(1, 3, 5)$ and $U(2, 4, 6)$ respectively. We give the optimal mechanism by the following matrices.

$$x_1 : \begin{array}{c} 2 \quad 4 \quad 6 \\ 1 \begin{pmatrix} 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 \\ 1 & 1 & 0 \end{pmatrix}, \quad x_2 : \begin{array}{c} 2 \quad 4 \quad 6 \\ 1 \begin{pmatrix} \frac{1}{2} & \frac{5}{6} & 1 \\ 0 & \frac{5}{6} & 1 \\ 0 & 0 & 1 \end{pmatrix} \end{array} \quad \text{and} \quad p_1 : \begin{array}{c} 2 \quad 4 \quad 6 \\ 1 \begin{pmatrix} 0 & 0 & 0 \\ 2 & 0 & 0 \\ \frac{11}{3} & 5 & 0 \end{pmatrix}, \quad p_2 : \begin{array}{c} 2 \quad 4 \quad 6 \\ 1 \begin{pmatrix} 1 & \frac{7}{3} & \frac{10}{3} \\ 0 & \frac{10}{3} & \frac{13}{3} \\ 0 & 0 & 6 \end{pmatrix} \end{array}$$

C Proof of Lemma 3.6

Lemma. *When buyers' valuations follow identical marginal distributions. The revenue guarantee of the best sequential posted price auction is always a $\frac{n}{n-1}$ -approximation to the best worst-case revenue among all second price auctions with common reserve price.*

Proof. Let SPA_r be the second price auction with reserve price r . By Lemma 3.1, it suffices to show that for all r , there exists a SPM $\mathcal{M} \in \mathcal{S}$ such that

$$\frac{n}{n-1} \text{Rev}(\mathcal{M}, \pi^*) \geq \min_{\pi \in \Pi} \text{Rev}(\text{SPA}_r, \pi).$$

Again, it would be easy for us to work in the quantile space. We first construct distribution $\hat{\pi}$ to establish an upper bound on $\min_{\pi \in \Pi} \text{Rev}(\text{SPA}_r, \pi)$.

Construction of distribution $\hat{\pi}$.

- 1: Draw a quantile $q \sim U[0, 1]$
 - 2: **if** $q \leq q^* \stackrel{\text{def}}{=} 1 - F(r)$ **then** ▷ All bids are at least r
 - 3: Draw a buyer $h \sim U[n]$ ▷ Pick the winning buyer
 - 4: Let $v_h \leftarrow V(\frac{q}{n})$
 - 5: Let $v_\ell \leftarrow V(\frac{(n-1)q}{n} + \frac{q^*}{n})$ for all $\ell \neq h$.
 - 6: **end if**
 - 7: **if** $q > q^*$ **then** ▷ All bids are less than r and the item remains unsold
 - 8: Let $v_i \leftarrow V(q)$ for all i .
 - 9: **end if**
-

We first explain why $\hat{\pi}$ is consistent with the marginal distribution F . Fix any buyer i .

- When $q \leq q^*$ and i was selected as the winning buyer h , $v_i \leftarrow V(\frac{q}{n})$ where $\frac{q}{n}$ was drawn from $U[0, \frac{q^*}{n}]$. This event happens with probability $\frac{q^*}{n}$.
- When $q \leq q^*$ and i was not selected as the winning buyer, $v_i \leftarrow V(q')$ where $q' = \frac{(n-1)q}{n} + \frac{q^*}{n}$ was drawn from $U[\frac{q^*}{n}, q^*]$. This event happens with probability $q^* - \frac{q^*}{n}$.
- When $q > q^*$, $v_i = V(q)$ where q was drawn from $U[q^*, 1]$. This event happens with probability $1 - q^*$.

Combining these three cases together, $v_i = V(q)$ where $q \sim U[0, 1]$, i.e., $v_i \sim F$.

When SPA_r is run on $\hat{\pi}$, if $q > q^*$, all bids are less than r and the item remains unsold. On the other hand, if $q \leq q^*$, we always have $\frac{q}{n} \leq \frac{(n-1)q}{n} + \frac{q^*}{n}$, which implies $v_h \geq v_\ell$. In this case bidder h wins the auction and pays the price of v_ℓ . Therefore we have $\text{Rev}(\text{SPA}_r, \hat{\pi}) = \int_0^{q^*} v_\ell dq = \int_0^{q^*} V\left(\frac{(n-1)q}{n} + \frac{q^*}{n}\right) dq$. By letting $q' = \frac{(n-1)q}{n} + \frac{q^*}{n}$, we have

$$\begin{aligned} \text{Rev}(\text{SPA}_r, \hat{\pi}) &= \int_{\frac{q^*}{n}}^{q^*} V(q') d\left(\frac{nq' - q^*}{n-1}\right) = \frac{n}{n-1} \int_{\frac{q^*}{n}}^{q^*} V(q) dq \\ &= \frac{n}{n-1} \sum_{i=1}^{n-1} \int_{\frac{iq^*}{n}}^{\frac{(i+1)q^*}{n}} V(q) dq \\ &\leq \frac{n}{n-1} \sum_{i=1}^{n-1} V\left(\frac{iq^*}{n}\right) \cdot \frac{q^*}{n} = \frac{q^*}{n-1} \sum_{i=1}^{n-1} V\left(\frac{iq^*}{n}\right). \end{aligned}$$

Now consider a SPM \mathcal{M} that posts price $V\left(\frac{iq^*}{n}\right)$ to the i -th buyer in sequence for $1 \leq i \leq n-1$ ¹³. By Lemma 3.1, we know the worst-correlation revenue of \mathcal{M} is

$$\text{Rev}(\mathcal{M}, \pi^*) = \frac{q^*}{n} \sum_{i=1}^{n-1} V\left(\frac{iq^*}{n}\right).$$

We conclude the lemma by combining the above two bounds. □

¹³We only post prices to the first $n-1$ buyers in this SPM.