

Correlation-Robust Mechanism Design

NICK GRAVIN

and

PINIAN LU

ITCS, Shanghai University of Finance and Economics

In this letter, we discuss the correlation-robust framework proposed by Carroll [Econometrica 2017] and our new development [SODA 2018]. Consider a monopolist seller that has n heterogeneous items to sell to a single buyer with the objective of maximizing the seller's revenue. In the correlation-robust framework, the seller only knows marginal distribution of each item but has no information about the correlation across different items in the joint distribution. Any mechanism is then evaluated according to its expected profit in the worst-case over all possible joint distributions with the given marginal distributions. Carroll's main result states that when the buyer's value for any set of her items is the sum of the values of individual items in the set, the optimal correlation-robust mechanism should sell items separately. We extend this result to the case where the buyer has a budget constraint on her total payment. Namely, we show that the optimal robust mechanism splits the total budget in a fixed way across different items independent of the bids, and then sells each item separately with a per item budget constraint.

We highlight an alternative approach via a dual Linear Programming formulation for the optimal correlation-robust mechanism design problem. This LP can be used to compute optimal mechanisms in general (other than additive) settings. It also yields an alternative proof for the additive monopoly problem without constructing the worst-case distribution and allows us to extend the proof to the budget setting.

Categories and Subject Descriptors: F.0 [Theory of Computation]: ANALYSIS OF ALGORITHMS AND PROBLEM COMPLEXITY

1. INTRODUCTION

In the problem of monopolist revenue maximization, a seller has n heterogeneous items to sell to a single buyer. The monopolist has a prior belief about the distribution of buyer's values and wants to sell the goods so as to maximize her expected revenue. In the case of a single item ($n = 1$) with the value drawn from a distribution F the optimal solution [Myerson 1981] is straightforward: the seller offers a fixed take-it-or-leave-it price p chosen to maximize the expected payment $p \cdot (1 - F(p))$. As an example of the multidimensional problem let us consider the most basic and widely studied version, where the buyer's value for a set of items is additive. This easy-to-state problem, despite the simplicity of its solution in the single-item case, often leads to complex and unwieldy solutions.

The problem of finding the right auction format and proving its optimality is quite difficult even in the case of two items ($n = 2$). The monopolist may use quite a few selling strategies: she may sell items independently by posting a separate price for each of the two items, or offer a bundle of both goods, at yet another price. In general, the seller can offer a menu with many options that may involve

Authors' addresses: nikolai@mail.shufe.edu.cn, lu.pinyan@mail.shufe.edu.cn

lotteries with probabilistic outcomes, e.g., a 0.6 chance of getting the first item and 0.4 chance of getting the second, for some price. In some special cases the optimal mechanism is relatively simple, e.g., in the natural case of values for different goods being drawn from $[0, 1]$ independently and uniformly at random, the optimal mechanism offers a menu with separate prices for each of the items and a price for the bundle (the proof of this seemingly simple fact is quite nontrivial [Manelli and Vincent 2007].) For general distributions it has been shown that randomization might be necessary and even that the seller might have to offer an infinite menu of lotteries [Hart and Nisan 2013; Daskalakis et al. 2014]. As another indication of the problem’s complexity, the revenue of the optimal auction may decrease [Hart and Reny 2015] when the buyer’s values in the prior distribution are moved upwards (in the stochastic dominance sense). These issues not only appear when values for two or more items are correlated, but also when they are independently distributed.

To avoid the aforementioned complications Carroll [Carroll 2017] has recently proposed a new framework for the multidimensional monopolist problem¹ for an additive buyer. In this framework the seller knows the prior distribution of types $v_i \sim F_i$ for each individual item $i \in [n]$. However, unlike the traditional approach, in which the seller maximizes the expected payment with respect to a given prior distribution \mathcal{D} over the complete type profiles $\mathbf{v} = (v_1, \dots, v_n)$, in the new framework the seller does not know the correlation of types across different items. Any mechanism then is evaluated according to its expected profit in the worst case, over all possible joint distributions with the given marginal distributions $\{F_i\}_{i=1}^n$ of each item $i \in [n]$. In other words, the seller wants to get a guarantee on the expected profit of a mechanism which is *robust* to any *correlation* across items. Although, Carroll’s model is formulated for a buyer with additively separable valuation, the framework easily extends to other more general mechanism design settings, e.g., settings with multiple buyers, or settings where the buyers are unit-demand (i.e., each buyer does not want more than one item), or have budget constraints.

There are standard pros and cons of worst-case versus average-case analysis frameworks in computer science, which also apply here. Beyond those, there are specific points that we shall discuss below.

- (1) The underlying assumption of the Bayesian framework is that the joint prior distribution is already known to the seller. There is serious practical concern regarding learning correlated multidimensional distributions: the representation and sampling complexity of this problem is exponential in the dimension (i.e., number of items). Another challenge in learning the prior distribution arises as a result of strategic behavior of the buyer, who does not usually report his type but responds to the seller’s offer in each single interaction and might want to conceal data in order to gain from his interaction with the seller in the future. In this respect, learning information about separate marginals is much a simpler econometrics task that does not suffer from the curse of dimensionality.
- (2) It is standard in the literature to assume that the prior distribution is independent across items. In this case it is expected that one can get better

¹Carroll considered a more general setting of multidimensional screening with additively separable payoff structure.

revenue guarantees than in the worst-case framework. However, in practice, the independence assumption does not always hold and even verifying it (in the property testing sense) is a non trivial statistical task. The studies of correlated priors are scarce but not uncommon in the literature, both for the cases of positively or negatively correlated distributions, see e.g. [Levin 1997; Tang and Wang 2016; Bateni et al. 2015]. The case of correlated distribution is significantly more challenging than the case of independent priors. In this respect, correlation-robust framework offers an alternative tractable model of studying the unwieldy case of possibly correlated prior distributions.

- (3) Even with independent prior distributions the optimal mechanism can be very complex and as such is not employed in practice. A recent line of work studies the monopolist problem in the simple versus optimal framework [Hartline and Roughgarden 2009] and obtained a few interesting approximation guarantees. In the case of additive buyer, Babaioff et al. [Babaioff et al. 2014] showed that simple mechanism of selling items either separately or together in one grand bundle gives a constant-factor approximation to the optimal revenue. A recent work by Cai et al. [Cai et al. 2016] provided a unified view on some of the above “simple versus optimal” results by an LP duality based approach of generalized virtual values. In the worst-case framework, Carroll has shown that the optimal *correlation-robust* mechanism is to sell items separately, without any bundling. His result compliments the result of [Babaioff et al. 2014] by adding a valuable counterpoint to the algorithmic mechanism design literature, as Carroll puts it, “If you don’t know enough to see how to bundle, then don’t.”
- (4) The prior distribution usually represents a belief of the seller about the buyer’s types, but not the exact distribution. As such the prior might not accurately capture the actual distribution and thus some robustness guarantees and insensitivity to the precise data can be useful. The new framework addresses the issue of possible correlation among different type components. Furthermore, it seems to offer a more tractable way to analyze other robustness issues, such as mistakes in the beliefs about marginal distributions.

To conclude, the new framework complements and adds a few valuable points to the literature on the monopolist problem. Specifically, it seems quite natural to examine this framework from a computational perspective.

The computational problem in the correlation-robust framework can be described with n distributions $\{F_i\}_{i=1}^n$, each F_i given by $|V_i|$ parameters, where V_i is the support of F_i . The goal is to find a truthful mechanism with the best revenue guarantee over all possible joint distributions \mathcal{D} with specified marginals $\{F_i\}_{i=1}^n$. We know from Carroll’s work what the optimal solution is for the case of additive buyer. However, for other versions of the problem (e.g., for unit-demand) the structure of the optimal mechanism is unclear and it is natural to ask the question of computing the optimal mechanism. This problem has a succinct description in contrast with the traditional computational Bayesian framework [Cai et al. 2012b; 2012a; 2013], where the input (distribution \mathcal{D} of types $\mathbf{v} = (v_1, \dots, v_n)$) may be exponential in the number of items².

²To make the computational problem reasonable some assumptions about polynomial number of

2. THE MODEL

We consider a canonical multidimensional auction environment where one agent is selling n heterogeneous items to a single buyer. This environment can be specified by an *allocation space* X , which is assumed to be a convex set in $[0, 1]^n$ (we assume that the agent is risk-neutral and his value extends to any convex combination of feasible allocations $x \in X$); *type space* $\mathbf{V} = \prod_{i=1}^n V_i$, $V_i \subseteq \mathbb{R}_{\geq 0}$ and a value function $\text{val} : X \times \mathbf{V} \rightarrow \mathbb{R}_{\geq 0}$. We use $\mathbf{v} = (v_1, v_2, \dots, v_n) \in \mathbf{V}$ to denote a multidimensional type of the agent. When the buyer has additive valuation, we have $\text{val}(\mathbf{v}, x) = \langle \mathbf{v}, x \rangle = \sum_{i=1}^n v_i \cdot x_i$. We employ standard formulation of incentive compatible (a.k.a. truthful) mechanism as a pair of allocation $x : \mathbf{V} \rightarrow X$ and payment $p : \mathbf{V} \rightarrow \mathbb{R}_{\geq 0}$ functions satisfying incentive compatibility (IC) and individual rationality (IR) constraints for quasi-linear utility $u(\mathbf{v}, \hat{\mathbf{v}})$.

$$\begin{aligned} u(\mathbf{v}, \hat{\mathbf{v}}) &\stackrel{\text{def}}{=} \text{val}(\mathbf{v}, x(\hat{\mathbf{v}})) - p(\hat{\mathbf{v}}) \leq u(\mathbf{v}, \mathbf{v}) = \text{val}(\mathbf{v}, x(\mathbf{v})) - p(\mathbf{v}) \quad \text{for all } \mathbf{v}, \hat{\mathbf{v}} \in V \quad (\text{IC}). \\ u(\mathbf{v}, \mathbf{v}) &= \text{val}(\mathbf{v}, x(\mathbf{v})) - p(\mathbf{v}) \geq 0 \quad \text{for all } \mathbf{v} \in V \quad (\text{IR}). \end{aligned}$$

A mechanism is *budget feasible* if the agent's payment to the seller is bounded by a budget B , i.e., $p(\mathbf{v}) \leq B$ for all $\mathbf{v} \in V$. The agent derives utility of $-\infty$ when $p(\mathbf{v}) > B$ and the same quasi-linear utility of $\text{val}(\mathbf{v}, x(\mathbf{v})) - p(\mathbf{v})$, when $p(\mathbf{v}) \leq B$. We assume that the agent's budget B is *public*, i.e., the budget B is known to the auctioneer³.

The type \mathbf{v} is drawn from a joint distribution \mathcal{D} , which is not completely known to the auctioneer and which may admit correlation among different components of \mathbf{v} . The auctioneer only knows marginal distributions F_i of \mathcal{D} for each separate component i but does not know how these components are correlated with each other. We assume that every distribution F_i is discrete and has finite support⁴ V_i . We use f_i to denote the probability density function of the distribution F_i . We also slightly abuse notations and 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., $\Pi = \{\pi \mid \sum_{\mathbf{v}_i} \pi(v_i, \mathbf{v}_i) = f_i(v_i), \quad \forall i \in [n], v_i \in V_i\}$. The goal is to design a truthful mechanism that maximizes the auctioneer's expected revenue in the worst case with respect to the unknown joint distribution \mathcal{D} . Formally, we want to find a truthful (budget feasible) mechanism (x^*, p^*) such that

$$(x^*, p^*) \in \underset{(x,p)}{\text{argmax}} \min_{\substack{\pi(x,p) \\ \pi \in \Pi}} \sum_{\mathbf{v} \in V} \pi(\mathbf{v}) p(\mathbf{v}). \quad (1)$$

3. LP FORMULATION

We begin by looking at equation (1) as a zero-sum game played between the auction designer and an adversary, who gets to pick a distribution π with given marginals F_1, \dots, F_n and whose objective is to minimize the auctioneer's revenue. We note

types in the support of \mathcal{D} usually are made.

³We note that optimal auction problem in a *private* budget setting is quite complex even in the single-item case. Thus the public budget assumption is indeed necessary if our goal is to find settings with simple optimal auctions.

⁴Similar to [Carroll 2017] our results extend to the distributions with continuous type distributions.

that the strategy space of the auctioneer, i.e., the set of truthful mechanisms given by $x : \mathbf{V} \rightarrow X$ and $p : \mathbf{V} \rightarrow \mathbb{R}_{\geq 0}$, is convex (because a random mixture of truthful mechanisms is a truthful mechanism itself) and is compact⁵. Similarly the strategy space Π of the adversary (distribution player) is also a compact convex set. Thus the sets of both players' mixed strategies coincide with their respective sets of pure strategies. Now, our two-player game admits at least one mixed Nash equilibrium⁶, which is also a pure Nash equilibrium: $\mathcal{M}^* = (x^*, p^*)$ for the auctioneer player and π^* for the adversary. This Nash equilibrium defines a unique value of a zero sum game and, therefore, yields a solution to minmax problem (1).

We restrict our attention to the minimization problem of the distribution player for any fixed truthful mechanism $\mathcal{M} = (x, p)$:

$$\min_{\pi \in \Pi} \sum_{\mathbf{v}} p(\mathbf{v}) \cdot \pi(\mathbf{v}). \quad (2)$$

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

$$\begin{aligned} \min \sum_{\mathbf{v}} p(\mathbf{v}) \cdot \pi(\mathbf{v}) & & \max \sum_{i=1}^n \sum_{v_i} f_i(v_i) \cdot \lambda_i(v_i) & (3) \\ \text{s. t. } \sum_{\mathbf{v}_i} \pi(v_i, \mathbf{v}_{-i}) = f_i(v_i) & \quad \text{dual var. } \lambda_i(v_i) & \text{s. t. } \sum_{i=1}^n \lambda_i(v_i) \leq p(\mathbf{v}) & \quad \forall \mathbf{v} \\ \pi(\mathbf{v}) \geq 0 & & \lambda_i(v_i) \in \mathbb{R} & \end{aligned}$$

The value of the primal LP 3 is worst-case revenue $\text{Rev}(\mathcal{M})$ of the mechanism $\mathcal{M} = (x, p)$. Intuitively, the dual LP (3) captures the best additive approximation of the payment function $p(\mathbf{v})$ of \mathcal{M} with $\{\lambda_i(v_i), v_i \in V_i\}_{i=1}^n$. The values of the primal and dual problems (3) are equal for any fixed truthful mechanism $\mathcal{M} = (x, p)$. This allows us to convert the maxmin problem (1) to a maximization LP problem:

$$\begin{aligned} \max \sum_{i=1}^n \sum_{v_i} f_i(v_i) \cdot \lambda_i(v_i) & (4) \\ \text{s. t. } \sum_{i=1}^n \lambda_i(v_i) \leq p(\mathbf{v}) \quad \forall \mathbf{v}; & \quad (x, p) : (\text{IC}), (\text{IR}); \quad x(\mathbf{v}) \in X. \end{aligned}$$

One can solve LP (4) with standard polynomial time techniques to get an optimal auction in a variety of settings. For example we can compute optimal auctions when the buyer has additive, unit-demand, budget additive, or other valuations which allows succinct LP description of X . However, the optimal solution to these

⁵Indeed, as there are only finite number of types, one can think of a pair of allocation x and payment p functions as $|\mathbf{V}|$ vectors in X and $|\mathbf{V}|$ real numbers in $\mathbb{R}_{\geq 0}$. Thus we get a natural notion of convergence and distance for the mechanisms. As the set of truthful mechanisms is described by a finite set of not strict IC and IR inequalities, we conclude that truthful mechanisms form a closed set. Note that allocation domain is compact and payment function of a truthful mechanism is bounded by a constant, which makes the set of truthful mechanisms to be bounded as well. Therefore, it is compact.

⁶by Glicksberg Theorem for continues games

problems would normally require descriptions of length proportional to the size of the type domain $|\mathbf{V}| = \prod_{i=1}^n |V_i|$, which makes it not efficient for problems with a large number of items. Thus a next most natural question is to find special classes of problems that admit succinct and simple auctions in the correlation-robust framework.

4. RESULTS AND PROOF OUTLINE

Let us denote by $\text{Rev}(F_i, B_i)$ the optimal revenue of a single-item auction that can be extracted from a single agent with value distribution F_i and a public budget B_i . We propose the following straightforward format of budget feasible mechanisms: split the budget $B = \sum_{i=1}^n B_i$ across all items $\{B_i\}_{i=1}^n$; independently for each item i run an optimal single-item auction with the revenue $\text{Rev}(F_i, B_i)$. We call this class of budget feasible mechanisms *item-budgets* mechanisms. We note that this is fairly large class of mechanisms, as there are many ways in which the budget B can be split over the different items. We use $\text{Rev}(\{F_i\}_{i=1}^n, B)$ to denote

$$\max \sum_{i=1}^n \text{Rev}(F_i, B_i), \quad \text{s.t.} \quad \sum_{i=1}^n B_i \leq B.$$

The solution to this problem gives us the expected revenue of the the optimal item-budgets mechanism. Our main result from [Gravin and Lu] says that the optimal correlation-robust mechanism is in fact an item-budgets mechanism.

THEOREM 4.1. *The optimal correlation-robust mechanism has the revenue of $\text{Rev}(\{F_i\}_{i=1}^n, B)$.*

Proof Outline. We assume towards a contradiction that there is a mechanism \mathcal{M} with higher revenue. Then we fix \mathcal{M} and consider the variables $\{\lambda_i(v_i)\}_{i \in [n], v_i \in V_i}$ in the dual LP (4), which give an additive approximation (lower bound) on the payment function of \mathcal{M} . It is natural to interpret $\{\lambda_i(v_i)\}_{v_i \in V_i}$ as prices for each separate item $i \in [n]$. However, we need to deal with the problem that variables $\{\lambda_i(v_i)\}$ can be negative. To this end, we can find a smaller counter example (a mechanism \mathcal{M}' and set of variables $\{\lambda'_i(v_i)\}_{i \in [n], v_i \in V'_i}$ with a smaller domain $\mathbf{V}' \subset \mathbf{V}$), such that $\lambda'_i(v_i)$ are non-negative and monotonically increasing for each $i \in [n]$. We construct an item-budgets mechanism such that its payment function is point-wise dominated (strictly upper bounded) by $\sum_{i \in [n]} \lambda'_i(v_i)$ and by the constraints of the dual LP (3) is also point-wise dominated by the payment function of \mathcal{M}' . Finally, we get a contradiction by combining certain tight IC and IR constraints for the item-budgets mechanism that together yield an upper bound on a weighted sum of the payments of \mathcal{M}' .

5. OPEN PROBLEMS

Correlation-robust approach offers a new optimization framework for design and analysis of mechanisms. It addresses some reasonable practical concerns and also brings closer Bayesian and worst-case frameworks in algorithmic mechanism design literature. The results in Carroll's paper and ours seem to be only initial steps in this framework and there are multiple open avenues for future work. Here, we list a

few interesting directions. We believe that the LP formulation approach developed in this paper may find its applications as a useful initial step in the future work on this topic.

Beyond additive valuations. All current work on the topic has assumed the buyer to have additive valuations. It is an intriguing research direction to investigate other types of valuations. It is particularly interesting to understand optimal correlation-robust auctions for another class of simple unit-demand valuations. It is not clear if the optimal mechanism will have to use lotteries as sometimes is required in the Bayesian framework with independent values. Another natural simple class of valuations to study is the class of budget additive buyer's valuations.

Multiple buyers. In the monopolist problem we have only one buyer. It is an important research direction to extend the correlation-robust framework to the case of multiple buyers. Two possible extensions include (i) a model where the worst-case distributions for different buyers are independent (ii) the distributions for different buyers can be correlated and the performance of a mechanism is measured in the worst-case over this correlation. We believe that both extensions are reasonable and deserve further investigation.

Computational complexity. Our LP formulation for the optimal correlation-robust auction has $\Omega(\prod_{i=1}^n |V_i|)$ variables, which has exponential dependency on the input size $\sum_{i=1}^n |V_i|$. When can we describe⁷ the optimal auction succinctly, i.e., find a polynomial in the input size representation? We know that for an additive buyer, and also for an additive buyer with budget constraint the optimal mechanism has a simple form and can be described and computed in polynomial time. But the problem remains open for other settings, such as, e.g., the monopolist problem for a unit-demand buyer.

Approximation. In this work, we focused on studying exactly optimal mechanisms. Similar to the case of independent prior distribution in the Bayesian model, it is reasonable to look at approximately optimal mechanisms in the correlation-robust framework, especially in the case when the exact optimum is too complex to implement in practice. Considering all the complications of the optimal mechanisms in the Bayesian framework, it seems that we are lucky to have simple optimal mechanism for the case of an additive buyer. It is quite likely that this is not going to be the case in many other settings. In this situation a reasonable next step would be to search for simple auctions that are approximately optimal in the correlation-robust framework.

REFERENCES

- BABAIOFF, M., IMMORLICA, N., LUCIER, B., AND WEINBERG, S. M. 2014. A simple and approximately optimal mechanism for an additive buyer. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*. 21–30.
- BATANI, M., DEGHANI, S., HAJIAGHAYI, M., AND SEDDIGHIN, S. 2015. Revenue maximization for selling multiple correlated items. In *Algorithms-ESA 2015*. Springer, 95–105.

⁷Preferably, we want to compute the optimal mechanism in polynomial time.

- CAI, Y., DASKALAKIS, C., AND WEINBERG, S. M. 2012a. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*. 459–478.
- CAI, Y., DASKALAKIS, C., AND WEINBERG, S. M. 2012b. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*. 130–139.
- CAI, Y., DASKALAKIS, C., AND WEINBERG, S. M. 2013. Understanding incentives: Mechanism design becomes algorithm design. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*. 618–627.
- CAI, Y., DEVANUR, N. R., AND WEINBERG, S. M. 2016. A duality based unified approach to bayesian mechanism design. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*. 926–939.
- CARROLL, G. 2017. Robustness and separation in multidimensional screening. *Econometrica* 85, 2, 453–488.
- DASKALAKIS, C., DECKELBAUM, A., AND TZAMOS, C. 2014. The complexity of optimal mechanism design. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 1302–1318.
- GRAVIN, N. AND LU, P. Separation in correlation-robust monopolist problem with budget. In *Proceedings of SODA 2018*.
- HART, S. AND NISAN, N. 2013. The menu-size complexity of auctions. In *ACM Conference on Electronic Commerce, EC '13, Philadelphia, PA, USA, June 16-20, 2013*. 565–566.
- HART, S. AND RENY, P. J. 2015. Maximal revenue with multiple goods: Nonmonotonicity and other observations. *Theoretical Economics* 10, 3, 893–922.
- HARTLINE, J. D. AND ROUGHGARDEN, T. 2009. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*. ACM, 225–234.
- LEVIN, J. 1997. An optimal auction for complements. *Games and Economic Behavior* 18, 2, 176–192.
- MANELLI, A. M. AND VINCENT, D. R. 2007. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic theory* 137, 1, 153–185.
- MYERSON, R. B. 1981. Optimal auction design. *Mathematics of operations research* 6, 1, 58–73.
- TANG, P. AND WANG, Z. 2016. Optimal auctions for negatively correlated items. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, 103–120.