Correlation-Robust Mechanism Design

Nick Gravin* Pinyan Lu†

Abstract

In this letter, we discuss 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 seller's revenue. In the correlation-robust framework, the seller only knows marginal distribution of each separate item but has no information about 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 given marginal distributions. Carroll's main result states that in multi-item monopoly problem with additive buyer, i.e., buyer's value for any set of items is the sum of values of individual item in the set, the optimal correlation-robust mechanism should sell items separately. We extend the separation result to the case where 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 respective per item budget constraint.

We highlight an alternative approach via 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 worst-case distribution and allowed us to extend the proof to the budget setting.

1 Introduction

In the monopolist setting the 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 case of a single item ($n = 1$) with value drawn from a distribution F the optimal solution [28] is straightforward: the seller offers a fixed take-it-or-leave-it price p chosen to maximize 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 buyer's value for a set of items is additive. This easy-to-state problem despite the simplicity of the single-item case often leads to a very 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 lotteries with probabilistic outcomes, e.g., a 0.6 chance of getting first item and 0.4 chance of getting second item, for some price. In some special cases the optimal mechanism is relatively simple, e.g., in a natural case of values for different goods being independent and uniform $[0, 1]$, the optimal mechanism offers a menu with separate prices for each of the items and a price for the bundle (despite a simple answer the proof

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

of this fact is quite nontrivial [27].) 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 [22, 18]. On the other note, the revenue of the optimal auction may be non-monotone [23] 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 values for the two items are independently distributed.

To avoid the aforementioned complications Carroll [10] has recently proposed a new framework for multidimensional monopolist problem[1] for additive buyer. In this framework the seller knows 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 expected payment with respect to a given *prior* distribution $\mathcal{D}$ over the complete type profiles $\boldsymbol{v} = (v_1, \cdots, v_n)$, in the new framework the seller does not know anything about 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 given marginal distributions $\{F_i\}_{i=1}^n$ of each separate 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 different type $\boldsymbol{v}$ components. Although, Carroll's model is formulated for a buyer with additively separable valuation, the framework easily extends to other more general mechanism design settings, where the buyer does not need to be additive and may potentially have any valuation function for different allocations x of items, e.g., the buyer might be unit-demand (i.e., he does not want more than one item), or have budget constraint.

It is quite remarkable that traditional for computer science worst-case approach was proposed by an economist in an economics journal. There are standard pros and cons of the worst-case versus average-case analysis frameworks in computer science, which also apply to the monopolist setting. However, there are some specific points that we shall discuss below.

1. The underlying assumption of the Bayesian analysis framework is that joint prior distribution is already known to the seller. There is a serious practical concern regarding learning correlated multidimensional distribution: the computational 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 improve his interaction with the seller in the future. In this respect, learning information about separate marginals is much 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 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 quite 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. [25, 29, 2]. 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 to study the unwieldy case of possibly correlated prior distributions.

3. Even with independent prior distribution 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 [24] and obtained a few interesting approximation guarantees.

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

In the case of additive buyer, Babaioff et. al. [1] showed that a simple mechanism, of selling items either separately, or together in one grand bundle gives a constant-factor approximation to the optimal revenue. 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 [1] 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 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 between different type components. Furthermore, it seems to offer 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 mechanism design literature on the monopolist problem and as such deserves more attention from the computer science community. Specifically, it seems quite natural to examine this framework from a computational perspective. A general monopolist problem in the correlation-robust framework can be described with n distributions $\{F_i\}_{i=1}^n$ for each separate item. The goal is to find a truthful mechanism with the best revenue guarantee over all possible joint distributions $\boldsymbol{\mathcal{D}}$ 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 a question of computing the optimal mechanism for any given set of marginals $\{F_i\}_{i=1}^n$. We note that this problem has quite a succinct description. Indeed, the input to this problem can be specified with n one-dimensional distributions $\{F_i\}_{i=1}^n$, each described with $|V_i|$ parameters, where V_i is a support of F_i. This is in contrast with the traditional computational Bayesian framework [6, 5, 7], where the input to a single-buyer monopoly problem (distribution $\mathcal{D}$ of types $\boldsymbol{v} = (v_1, \ldots, v_n)$) might have exponential in the number of items size of $\prod_{i=1}^n |V_i|$ and thus some assumptions about polynomial number of types in the support of $\mathcal{D}$ are necessary.

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 agent is risk-neutral and his value extends to any convex combination of feasible allocations $x \in X$); *type space* $\boldsymbol{V} = \prod_{i=1}^n V_i$, $V_i \subseteq \mathbb{R}_{\geq 0}$ and a value function $\mathsf{val} : X \times \boldsymbol{V} \to \mathbb{R}_{\geq 0}$. We use $\boldsymbol{v} = (v_1, v_2, \cdots, v_n) \in \boldsymbol{V}$ to denote a multidimensional type of the agent. When buyer has additive valuation, we have $\mathsf{val}(\boldsymbol{v}, x) = \langle \boldsymbol{v}, x \rangle = \sum_{i=1}^n v_i \cdot x_i$. We employ standard formulation of incentive compatible (a.k.a. truthful) mechanism as a pair of allocation $x : \boldsymbol{V} \to X$ and payment $p : \boldsymbol{V} \to \mathbb{R}_{\geq 0}$ functions satisfying incentive compatibility (IC) and individual rationality (IR) constraints for quasi-linear utility $u(\boldsymbol{v}, \widehat{\boldsymbol{v}})$.

$$u(\boldsymbol{v}, \widehat{\boldsymbol{v}}) \stackrel{\text{def}}{=} \mathsf{val}(\boldsymbol{v}, x(\widehat{\boldsymbol{v}})) - p(\widehat{\boldsymbol{v}}) \leq u(\boldsymbol{v}, \boldsymbol{v}) = \mathsf{val}(\boldsymbol{v}, x(\boldsymbol{v})) - p(\boldsymbol{v}) \qquad \text{for all } \boldsymbol{v}, \widehat{\boldsymbol{v}} \in V \quad \text{(IC)}.$$

$$u(\boldsymbol{v}, \boldsymbol{v}) = \mathsf{val}(\boldsymbol{v}, x(\boldsymbol{v})) - p(\boldsymbol{v}) \geq 0 \qquad \text{for all } \boldsymbol{v} \in V \quad \text{(IR)}.$$

A mechanism is *budget feasible* if the agent’s payment to the seller is bounded by a budget B, i.e., $p(\boldsymbol{v}) \leq B$ for all $\boldsymbol{v} \in V$. The agent derives utility of $-\infty$ when $p(\boldsymbol{v}) > B$ and the same quasi-linear utility of $\mathsf{val}(\boldsymbol{v}, x(\boldsymbol{v})) - p(\boldsymbol{v})$, when $p(\boldsymbol{v}) \leq B$. We assume that agent's budget B is *public*, i.e., the budget B is known to the auctioneer[2].

The type $\boldsymbol{v}$ is drawn from a joint distribution $\mathcal{D}$, which is not completely known to the auctioneer and which may admit correlation between different components v_i and v_j of $\boldsymbol{v}$. 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 support[3] 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 $\boldsymbol{V} = \times_{i=1}^{n} V_i$. We use Π to denote all possible distributions π supported on $\boldsymbol{V}$ that are consistent with the marginal distributions $F_1, F_2, \cdots, F_n$, i.e., $\Pi = \left\{\pi \mid \sum_{\boldsymbol{v}_{\text{-}i}} \pi(v_i, \boldsymbol{v}_{\text{-}i}) = f_i(v_i), \quad \forall i \in [n], v_i \in V_i\right\}$. The goal is to design a truthful mechanism that maximizes 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

$$(x^*, p^*) \in \underset{(x,p)}{\operatorname{argmax}} \min_{\substack{\pi(x,p) \\ \pi \in \Pi}} \sum_{\boldsymbol{v} \in V} \pi(\boldsymbol{v}) p(\boldsymbol{v}). \tag{1}$$

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 $F_1, \cdots, F_n$ and whose objective is to minimize the auctioneer's revenue. We note that the strategy space of the auctioneer, i.e., the set of truthful mechanisms given by $x : \boldsymbol{V} \to X$ and $p : \boldsymbol{V} \to \mathbb{R}_{\geq 0}$, is convex (because a random mixture of truthful mechanisms is a truthful mechanism itself) and is compact[4]. 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 equilibrium[5], 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 $\mathcal{M} = (x, p)$:

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

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

[2]We 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.

[3]Similar to [10] our results extend to the distributions with continuous type distributions.

[4]Indeed, as there are only finite number of types, one can think of a pair of allocation x and payment p functions as $|\boldsymbol{V}|$ vectors in X and $|\boldsymbol{V}|$ 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.

[5]by Glicksberg Theorem for continues games

corresponding dual problem.

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

The value of the primal LP 3 is worst-case revenue $\mathsf{Rev}(\mathcal{M})$ of the mechanism $\mathcal{M} = (x, p)$. Intuitively, the dual LP (3) captures the best additive approximation of the payment function $p(\boldsymbol{v})$ 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:

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

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 for buyer's valuations such as additive, unit-demand, budget additive and many other tractable settings which allow succinct LP description of X. However, the optimal solution to these problems would normally require description proportional to the size of the type domain $|\boldsymbol{V}| = \prod_{i=1}^n |V_i|$, which makes it not efficient for problems with 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 $\mathsf{Rev}(F_i, B_i)$ the optimal revenue of a single-item single-bidder auction that can be extracted from a single agent with a value distribution $v_i \sim F_i$ and a public budget B_i. To simplify notations we use $\mathsf{Rev}_i(B_i) \stackrel{\text{def}}{=} \mathsf{Rev}(F_i, B_i)$. 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 $\mathsf{Rev}_i(B_i)$. 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 $\mathsf{Rev}(\{F_i\}_{i=1}^n, B)$ to denote

$$\max \sum_{i=1}^{n} \mathsf{Rev}_i(B_i), \quad s.t. \quad \sum_{i=1}^{n} B_i \le B.$$

The solution to this problem gives us the expected revenue of the the optimal item-budgets mechanism. Our main result from [20] 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* $\mathsf{Rev}(\{F_i\}_{i=1}^n, B)$*.*

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 a problem that variables $\{\lambda_i(v_i)\}$ can be negative. To this end, we regularize the problem by restricting the domain $v_i \in V_i$ and ensure that $\{\lambda_i(v_i)\}_{i\in[n],v_i\in V_i}$ are non-negative and monotonically increasing for each $i \in [n]$. We construct an item-pricing mechanism such that its payment function is point-wise dominated (upper bounded) by $\sum_{i\in[n]} \lambda_i(v_i)$. Finally, we get a contradiction by combining certain tight IC and IR constraints for the item-pricing 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 our and Carroll's papers 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 buyer to have additive valuations. It is 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 important research direction to extend the correlation-robust framework to the case of multiple buyers. Two possible extensions include (i) a model where 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 describe the optimal auction succinctly, i.e., find a polynomial in the input size representation? We know that for additive buyer, and also for 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 unit-demand buyer.

Approximation. In this work, we focused on studying exact optimality of 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 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

[1] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. 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*, pages 21–30, 2014.

[2] MohammadHossein Bateni, Sina Dehghani, MohammadTaghi Hajiaghayi, and Saeed Seddighin. Revenue maximization for selling multiple correlated items. In *Algorithms-ESA 2015*, pages 95–105. Springer, 2015.

[3] Dirk Bergemann and Stephen Morris. Robust mechanism design. *Econometrica*, 73(6):1771–1813, 2005.

[4] Dirk Bergemann and Karl Schlag. Robust monopoly pricing. *Journal of Economic Theory*, 146(6):2527–2543, 2011.

[5] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. 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*, pages 459–478, 2012.

[6] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. 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*, pages 130–139, 2012.

[7] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. 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*, pages 618–627, 2013.

[8] Yang Cai, Nikhil R. Devanur, and S. Matthew Weinberg. 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*, pages 926–939, 2016.

[9] Gabriel Carroll. Robustness and linear contracts. *The American Economic Review*, 105(2):536–563, 2015.

[10] Gabriel Carroll. Robustness and separation in multidimensional screening. *Econometrica*, 85(2):453–488, 2017.

[11] Gabriel Carroll and Delong Meng. Robust contracting with additive noise. *Journal of Economic Theory*, 166:586–604, 2016.

[12] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings 8th ACM Conference on Electronic Commerce (EC-2007), San Diego, California, USA, June 11-15, 2007*, pages 243–251, 2007.

[13] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 311–320, 2010.

[14] Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior*, 91:297–317, 2015.

[15] Kim-Sau Chung and Jeffrey C Ely. Foundations of dominant-strategy mechanisms. *The Review of Economic Studies*, 74(2):447–476, 2007.

[16] Vincent Conitzer and Tuomas Sandholm. Complexity of mechanism design. In *Proceedings of the Eighteenth conference on Uncertainty in artificial intelligence*, pages 103–110. Morgan Kaufmann Publishers Inc., 2002.

[17] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Mechanism design via optimal transport. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 269–286. ACM, 2013.

[18] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. The complexity of optimal mechanism design. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1302–1318. SIAM, 2014.

[19] Daniel F Garrett. Robustness of simple menus of contracts in cost-based procurement. *Games and Economic Behavior*, 87:631–641, 2014.

[20] Nick Gravin and Pinyan Lu. Separation in correlation-robust monopolist problem with budget. In *Proceedings of SODA 2018*.

[21] Sergiu Hart and Noam Nisan. Approximate revenue maximization with multiple items. In *ACM Conference on Electronic Commerce, EC '12, Valencia, Spain, June 4-8, 2012*, page 656, 2012.

[22] Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. In *ACM Conference on Electronic Commerce, EC '13, Philadelphia, PA, USA, June 16-20, 2013*, pages 565–566, 2013.

[23] Sergiu Hart and Philip J Reny. Maximal revenue with multiple goods: Nonmonotonicity and other observations. *Theoretical Economics*, 10(3):893–922, 2015.

[24] Jason D Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 225–234. ACM, 2009.

[25] Jonathan Levin. An optimal auction for complements. *Games and Economic Behavior*, 18(2):176–192, 1997.

[26] Xinye Li and Andrew Chi-Chih Yao. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences*, 110(28):11232–11237, 2013.

[27] Alejandro M Manelli and Daniel R Vincent. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic theory*, 137(1):153–185, 2007.

[28] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[29] Pingzhong Tang and Zihe Wang. Optimal auctions for negatively correlated items. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 103–120. ACM, 2016.

[30] John Thanassoulis. Haggling over substitutes. *Journal of Economic theory*, 117(2):217–245, 2004.

[31] Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 92–109, 2015.