

Separation in Correlation-Robust Monopolist Problem with Budget*

Nick Gravin[†]

Pinyan Lu[‡]

Abstract

We consider a monopolist seller that has n heterogeneous items to sell to a single buyer. The seller’s goal is to maximize her revenue. We study this problem in the correlation-robust framework recently proposed by Carroll [Econometrica 2017]. In this framework, the seller only knows marginal distributions for 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 buyer, whose value for a set of items is additive, the optimal correlation-robust mechanism should sell items separately.

We use alternative dual Linear Programming formulation for the optimal correlation-robust mechanism design problem. This LP can be used to compute optimal mechanisms in general settings. We give an alternative proof for the additive monopoly problem without constructing worst-case distribution. As a surprising byproduct of our approach we get that separation result continues to hold even when buyer has a budget constraint on her total payment. Namely, 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.

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 [30] 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 remains one of the primary open challenges in algorithmic mechanism design.

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 of this fact is quite nontrivial [29].) 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 [24, 19]. On the other note, the revenue of the optimal auction may be non-monotone [25] 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 problem¹ 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 $\mathbf{v} = (v_1, \dots, v_n)$, in the new framework the seller does not know anything about correlation of types across dif-

*Gravin is supported by a China Youth 1000-Talent grant.

[†]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

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

ferent 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 \mathbf{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 multi-dimensional 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 respond to the seller’s offer in each single interaction and who also might have incentives 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. The most common case studied in the algorithmic mechanism design literature is the case of known independent prior distribution. 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. There are many scenarios in which it is natural to assume correlation across different items. The studies for correlated priors are not uncommon in the literature, both for the cases of positively or negatively correlated distributions, see e.g. [27, 31, 2]. However, the case of independent priors is usually more tractable and much better

studied in algorithmic mechanism design literature than the case of correlated priors. Indeed, it seems reasonable to resolve first a simpler and more regular case of independent distributions before studying a general and more difficult problem with correlated priors. In this respect, correlation-robust framework offers an alternative tractable model to study the unwieldy case of possibly correlated prior distributions.

3. As was mentioned earlier, 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 in algorithmic game theory studied the monopolist problem in the simple versus optimal framework [26] and obtained a few interesting approximation guarantees. 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 \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 or budget constrained buyer) 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 $\mathbf{v} = (v_1, \dots, v_n)$) 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.

An inquisitive reader may wonder at this point in which other than additive settings the new framework can lead to interesting and tractable auction designs. Carroll's result for additive buyer shows optimality of quite special and simple mechanism that sells items separately. In this setting selling items separately seems to be a perfect idea to the seller who knows only marginal distributions and who then would equalize her expected profit over all possible joint distributions of the buyer. We note however that even this apparently simple and intuitive result requires a highly non-trivial proof [10]. The proof goes by constructing the worst-case joint distribution with fixed marginals. Carroll shows that the worst-case distribution is nothing but simple². Another central concept in the paper are generalized virtual values, which add extra layer of difficulty compared to the analysis with the classic virtual values defined in [30].

1.1 Our Results The central problem in the new worst-case framework is a maxmin problem of finding optimal correlation-robust auction with a given set of marginal distributions. In other words, this problem is a zero sum game played between the auction designer who picks a mechanism and the adversary who chooses the joint distribution with fixed marginals. In this work we give a new *Linear Program formulation* for this *maxmin problem*³. Our LP has intuitive interpretation: in addition to the standard variables representing mechanism's allocation and payment functions $\{x(\mathbf{v}), p(\mathbf{v}), \mathbf{v} \in \mathbf{V}\}$ it has a set of new variables $\{\lambda_i(v_i), v_i \in V_i\}$ for each separate item $i \in [n]$; the LP objective captures the best additive approximation of the payment function $p(\mathbf{v})$

²The joint distribution is neither completely independent, nor perfectly correlated and is constructed as a result of converging Poisson process.

³Carroll used a similar idea for a different problem in [10].

with $\{\lambda_i(v_i)\}_{i=1}^n$. The LP has succinct description, i.e., it has similar number of variables and constraints as the LP describing a truthful auction (Incentive Compatibility and Individual Rationality constraints). In particular, our result implies that one can solve the LP for any given set of marginal distributions, any valuation function, and any feasibility constraint on the allocation and payments in time polynomial in the number of possible types⁴ $\prod_{i=1}^n |V_i|$. Another important feature of our LP is that it can witness optimality of the auction without constructing the worst case distribution. Thus, our proof completely avoids the explicit construction of the highly non trivial worst-case distribution, which was an essential part in the Carroll's proof [10]. We note that generalized virtual values are defined as dual variables to a different LP. In our LP dual and primal formulations, we also avoid explicit construction of the generalized virtual values.

We study next a new setting with a budget constrained buyer. In this setting the buyer still has additive valuation, but has a publicly known budget cap on the maximum amount he can pay to the seller. As the budget constraint applies to all items rather than each individual item, the strategy of selling items together (i.e., bundle items) seems more plausible than in the case without a budget constraint. Surprisingly, it turned out that the optimal correlation-robust auction would still sell items separately. More specifically, we show that the optimal auction should split the total budget across different items (the division depends only on the marginal distributions, but not the bids), and then sell each item separately with a respective per item budget constraint. Our proof of optimality is obtained from our new LP formulation without a construction of the worst case distribution. We further derive a solution for the budget partition problem (i.e., how to split the total budget across different items) and provide an efficient algorithm that computes the optimal auction in $O(\sum_{i=1}^n |V_i|)$ time.

Conceptual contribution. The contribution of this paper is twofold. First, we propose an alternative approach with intuitive LP formulation that simplifies previous proofs and allows to consider many other computational mechanism design problem in a correlation-robust framework. Second, we identify another class of problems with simple optimal solution, which strengthen the Carroll's counterpoint regarding advantages of bundling in mechanism design literature.

⁴Note that the dependency on the number of possible types $\prod_{i=1}^n |V_i|$ is unavoidable, as the description of a truthful mechanism as a pair of allocation and payment functions would already require $\Omega(\prod_{i=1}^n |V_i|)$ space.

Technical contribution. We propose a simpler alternative to the existing techniques which rely on the relatively new concept of generalized virtual values that simultaneously and independently to Carroll’s work appeared in [8]. The idea of generalized virtual values was first proposed in [22] in the context of verification of action optimality. Our approach is built on the idea of studying dual LP to the adversary’s problem (one who picks the worst-case distribution) and avoids the explicit construction of the joint distribution in the Carroll’s proof. We demonstrate the technical benefit of our LP formulation by extending the separation result of Carroll to a budgeted setting. It is unclear whether the previous techniques of generalized virtual values would yield that result.

1.2 Related work For a given distribution (either with independent or correlated values), the revenue-maximizing auction is usually quite complex. In particular, it may involve bundling and lottery pricing, the menu size could be very large or even infinite [32, 29, 24, 18]. The task of finding the optimal mechanism could be computationally intractable [17, 19].

The computational framework of Cai-Daskalakis-Weinberg addresses the problem of multidimensional mechanism design [6, 5, 7] from a computational complexity perspective. This line of work proposes a computationally tractable Bayesian incentive compatible (BIC) solution for mechanism design problems with multiple buyers. For a few quite general mechanism design settings, they showed that the computational problem can be solved with a black-box access to an algorithmic problem, where the goal is to optimize a perturbed version of the initial objective without any incentives constraints. However, for a single-buyer problem they need assumptions on the support size of joint value distribution.

Another line of work employs a simple versus optimal framework [26] to justify why simple auctions are employed in the real world by showing approximation guarantees of a simple format compared to the optimal mechanism. For the case of unit-demand buyers, it was shown in Chawla et. al. [12, 13, 14] that posted-price mechanisms achieve constant approximation to the optimum. For the case of additive buyers, it was shown in [23, 28, 1, 33] that simple auction format of either selling items separately and a VCG mechanism with per-bidder entry fee yields constant approximation to the optimal revenue. In this work we focus on the exact optimal auction in the correlation-robust worst-case framework. It is possible that our LP formulation can help to obtain approximation guarantees in the correlation-robust auction design framework.

A recent work by Cai et al. [8] provided a unified view on some of the above “simple versus optimal” results by an LP duality based approach of generalized virtual values. The work of Carroll [10]⁵ also heavily relies on the concept of generalized virtual values. Our LP formulation approach is also based on the LP duality. However, we view the mechanism as a parameter and work with the dual problem of designing the worst-case distribution. Combining our techniques with the concept of generalized virtual values might find its applications and can be beneficial in other settings in the correlation-robust framework.

There is a few lines of work in economics literature on robust mechanism design, that aims to explain the usage of intuitive and simple mechanisms by providing performance guarantees in uncertain environment [9, 11, 4, 3, 16, 20]. This body of literature employs a similar to our approach of searching a worst-case solution over the uncertainty in the optimization problem.

2 Preliminaries

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 compact measurable set in $[0, 1]^n$; *type space* $\mathbf{V} = \prod_{i=1}^n V_i$, $V_i \subseteq \mathbb{R}_{\geq 0}$ and a value function $\text{val} : X \times \mathbf{V} \rightarrow \mathbb{R}_{\geq 0}$. In general we will use bold face script to denote vectors, or multidimensional objects like \mathbf{V} , e.g., n -dimensional vector of 1 we denote as $\mathbf{1}$. We consider the set of convex combinations $\Delta(X)$ achievable in expectation by a random feasible allocation X . The agent is risk-neutral, and thus his value naturally extends to $\mathbf{E}[\text{val}] : \Delta(X) \times \mathbf{V} \rightarrow \mathbb{R}$. We will use the variable x to denote either an element of X or of $\Delta(X)$; and function val to denote $\mathbf{E}_x[v]$ for $x \in \Delta(X)$. We use $\mathbf{v} = (v_1, v_2, \dots, v_n) \in \mathbf{V}$ to denote a multidimensional type of the agent. We employ standard formulation of incentive compatible 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 $u(\mathbf{v}, \hat{\mathbf{v}})$.

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

In this formulation, the agent has a true type \mathbf{v} and submits a bid $\hat{\mathbf{v}}$ to the auctioneer. The auctioneer

⁵Carroll proposed the notion of generalized virtual values independently and simultaneously with [8]

sees the agent's bids $\widehat{\mathbf{v}}$ and outputs an allocation vector $x(\widehat{\mathbf{v}}) = (x_1, x_2, \dots, x_n) \in X$ and a payment $p(\widehat{\mathbf{v}})$. The agent's valuation is $\text{val}(\mathbf{v}, x)$ and his utility is $\text{val}(\mathbf{v}, x) - p$ for the true type \mathbf{v} . When buyer has additive valuation, we have $\text{val}(\mathbf{v}, x) = \langle \mathbf{v}, x \rangle = \sum_{i=1}^n v_i \cdot x_i$. A mechanism (x, p) is called truthful if it satisfies IR and IC conditions.

Another property for mechanism is budget feasibility: the agent's payment to the seller is bounded by a budget B . 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 agent's budget B is *public*, i.e., the budget B is known to the auctioneer⁶. In this setting, the auctioneer must use budget feasible mechanism, i.e., a mechanism such that $p(\mathbf{v}) \leq B$ for all $\mathbf{v} \in V$.

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 between different components v_i and v_j 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 support⁷ 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., $\Pi = \{\pi \mid \sum_{\mathbf{v}_i} \pi(v_i, \mathbf{v}_i) = f_i(v_i), \forall i \in [n], v_i \in V_i\}$. 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

$$(2.1) \quad (x^*, p^*) \in \operatorname{argmax}_{(x,p)} \min_{\substack{\pi \in \Pi \\ \mathbf{v} \in V}} \sum \pi(\mathbf{v}) p(\mathbf{v}).$$

3 LP formulation of Maxmin

We begin by looking at equation (2.1) 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 that the strat-

⁶We 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.

⁷Similar to [10] our results extend to the distributions with continuous type distributions.

egy space of the auctioneer, i.e., the set of truthful mechanisms given by $x : \mathbf{V} \rightarrow \Delta(X)$ 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 compact⁸. 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 equilibrium⁹, 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 (2.1).

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

$$(3.2) \quad \min_{\pi \in \Pi} \sum_{\mathbf{v}} p(\mathbf{v}) \cdot \pi(\mathbf{v}).$$

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

$$(3.3) \quad \begin{aligned} \min \quad & \sum_{\mathbf{v}} p(\mathbf{v}) \cdot \pi(\mathbf{v}) \\ \text{s. t.} \quad & \sum_{\mathbf{v}_i} \pi(v_i, \mathbf{v}_i) = f_i(v_i) \quad \text{dual var. } \lambda_i(v_i) \\ & \pi(\mathbf{v}) \geq 0 \\ \max \quad & \sum_{i=1}^n \sum_{v_i} f_i(v_i) \cdot \lambda_i(v_i) \\ \text{s. t.} \quad & \sum_{i=1}^n \lambda_i(v_i) \leq p(\mathbf{v}) \quad \forall \mathbf{v} \\ & \lambda_i(v_i) \in \mathbb{R} \end{aligned}$$

The value of the primal LP (3.3) is worst-case revenue $\text{Rev}(\mathcal{M})$ of the mechanism $\mathcal{M} = (x, p)$. Intuitively, the dual LP (3.3) 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

⁸Indeed, 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.

⁹by Glicksberg Theorem for continues games.

problems (3.3) are equal for any fixed truthful mechanism $\mathcal{M} = (x, p)$. This allows us to convert the maximin problem (2.1) to a maximization LP problem:

$$(3.4) \quad \begin{aligned} & \max \sum_{i=1}^n \sum_{v_i} f_i(v_i) \cdot \lambda_i(v_i) \\ & \text{s. t. } \sum_{i=1}^n \lambda_i(v_i) \leq p(v) \quad \forall v; \\ & (x, p) : (\text{IC}), (\text{IR}); \quad x(v) \in \Delta(X). \end{aligned}$$

One can solve LP (3.4) 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 $\Delta(X)$. However, the optimal solution to these problems would normally require description proportional to the size of the type domain $|\mathbf{V}| = \prod_{i=1}^n |V_i|$, 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. One such problem is monopolist setting for additive buyer, for which Carroll [10] showed that simple auction of selling items separately achieves the optimal revenue in the worst-case.

4 Additive Separation with Budget

Let us denote by $\text{Rev}(F_i, B_i)$ 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 . We will use $\text{Rev}(F_i)$ to denote the revenue of the optimal posted-price auction without budget constraint. To simplify notations we use $\text{Rev}_i(B_i) \stackrel{\text{def}}{=} \text{Rev}(F_i, B_i)$. 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 $\text{Rev}_i(B_i)$. 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 $\text{Rev}(\{F_i\}_{i=1}^n, B)$ to denote

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

The solution to this problem gives us the expected revenue of the the optimal item-budgets mechanism. We will discuss how to find an optimal mechanism for

each individual item i and budget B_i (i.e., those with revenue $\text{Rev}(F_i, B_i)$) and how to split the budget in the optimal way over different items in the next Section 5). In the following theorem we show 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 (3.4), 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} .

Proof. Let us assume to the contrary that \mathcal{M} is a truthful and budget-feasible mechanism, such that its revenue $\text{Rev}(\mathcal{M}) > \sum_{i=1}^n \text{Rev}(F_i, B_i)$ for any partition $\{B_i\}_{i=1}^n$ of the budget B . Now let $\{\lambda_i(v_i)\}_{i \in [n], v_i \in V_i}$ be a solution to the dual LP 3.3 for the mechanism $\mathcal{M} = (x, p)$. Notice that the solution to the dual LP 3.3 is not unique, since we can do the following linear transformation with $\{\lambda_i(v_i)\}$ for any fixed pair of items $i, j \in [n]$ and $\delta \in \mathbb{R}$:

$$(4.5) \quad \begin{cases} \lambda_i(v_i) \leftarrow \lambda_i(v_i) + \delta & \forall v_i \in V_i \\ \lambda_j(v_j) \leftarrow \lambda_j(v_j) - \delta & \forall v_j \in V_j, \end{cases}$$

and get another feasible solution with the same value $\sum_{i=1}^n \sum_{v_i} f_i(v_i) \cdot \lambda_i(v_i)$. In particular, by applying a few of these transformation we can get the following set of $\{\lambda_i(v_i)\}_{i \in [n], v_i \in V_i}$

CLAIM 1. $\exists \{\lambda_i(v_i)\}_{i \in [n], v_i \in V_i}$ s.t. $\sum_{v_i \in V_i} \lambda_i(v_i) \cdot f_i(v_i) > \text{Rev}_i(B_i^o)$ with $B_i^o = \max_{v_i \in V_i} \{0, \lambda_i(v_i)\} \quad \forall i \in [n]$.

Proof. First, note that by applying transform (4.5) to $\{\lambda_i(v_i)\}$ we increase $\max_{v_i \in V_i} \lambda_i(v_i)$ by δ and decrease $\max_{v_j \in V_j} \lambda_j(v_j)$ by δ . Thus we can make $\max_{v_i \in V_i} \lambda_i(v_i) \geq$

0 for all $i \in [n]$. If we keep $\{\lambda_i(v_i)\}$ within the region $\max_{v_i \in V_i} \lambda_i(v_i) \geq 0$ for all $i \in [n]$, then $B_i^o \stackrel{\text{def}}{=} \max_{v_i \in V_i} \{0, \lambda_i(v_i)\} = \max_{v_i \in V_i} \lambda_i(v_i)$.

Second, let $v_i^* \stackrel{\text{def}}{=} \operatorname{argmax}_{v_i \in V_i} \lambda_i(v_i)$ for each $i \in [n]$, then from the dual LP 3.3 we have

$$\sum_{i \in [n]} B_i^o = \sum_{i \in [n]} \lambda_i(v_i^*) \leq p(\mathbf{v}^*) \leq B,$$

I.e., the set $\{B_i^o\}_{i \in [n]}$ is budget feasible. Now, since $\sum_{i \in [n]} \sum_{v_i \in V_i} \lambda_i(v_i) \cdot f_i(v_i) > \sum_{i \in [n]} \operatorname{Rev}_i(B_i^o)$, we can always find $j \in [n]$ s.t. $\sum_{v_j \in V_j} \lambda_j(v_j) \cdot f_j(v_j) > \operatorname{Rev}_j(B_j^o) + \varepsilon$ for some fixed small $\varepsilon > 0$.

Finally, we show in Section 5 that $\operatorname{Rev}_i(B_i)$ is concave as a function of B_i . Also note that the revenue can never exceed the budget $\operatorname{Rev}_i(B_i) \leq B_i$, $\operatorname{Rev}_i(B)$ is increasing function, $\operatorname{Rev}_i(0) = 0$, and $\operatorname{Rev}_i(\infty) = \operatorname{Rev}_i$. I.e., the derivative of $\operatorname{Rev}_i(x)$ is equal to or smaller than 1 for any $x \geq 0$. If for some $i \in [n]$ $\sum_{v_i \in V_i} \lambda_i(v_i) \cdot f_i(v_i) \leq \operatorname{Rev}_i(B_i^o)$, then we find $j \in [n]$ with $\sum_{v_j \in V_j} \lambda_j(v_j) \cdot f_j(v_j) > \operatorname{Rev}_j(B_j^o) + \varepsilon$ and do transform (4.5) with some small fixed $\delta > 0$ s.t. $\sum_{v_j \in V_j} \lambda_j(v_j) \cdot f_j(v_j) > \operatorname{Rev}_j(B_j^o)$. As $\operatorname{Rev}_i(B)$ is a bounded function $\sum_{v_i \in V_i} \lambda_i(v_i) \cdot f_i(v_i)$ will become greater than $\operatorname{Rev}_i(B_i^o)$ after finitely many steps. This way we can make $\sum_{v_i \in V_i} \lambda_i(v_i) \cdot f_i(v_i) > \operatorname{Rev}_i(B_i^o)$ for all $i \in [n]$.

Now we fix \mathcal{M} and $\{\lambda_i(v_i)\}_{v_i \in V_i}$ for each $i \in [n]$ and try to find marginal distribution with the smallest support that still has a gap between $\operatorname{Rev}_i(B_i^o)$ and corresponding revenue term $\sum_{v_i \in V_i} \lambda_i(v_i) \cdot f_i(v_i)$ of \mathcal{M} :

$$(4.6) \quad f_i^* = \operatorname{argmin}_{f_i} \left\{ \left| \operatorname{supp}(f_i) \right| \left| \Delta_i^{\operatorname{rev}}(f_i) \stackrel{\text{def}}{=} \sum_{v_i \in V_i} \lambda_i(v_i) \cdot f_i(v_i) - \operatorname{Rev}(f_i, B_i^o) > 0 \right. \right\}$$

Notice that either $0 \notin \operatorname{supp}(f_i^*)$, or $\operatorname{supp}(f_i^*) = \{0\}$. Indeed, if $\lambda_i(0) \leq 0$ one can either remove 0 from the support of f_i^* and proportionally increase the remaining $f_i^*(v_j)$, or when $\lambda_i(0) > 0$ simply set $\operatorname{supp}(F_i) = \{0\}$ with $f_i^*(0) = 1$ (in this case $\operatorname{Rev}(f_i^*, B_i^o) = 0$). Notice, that in the latter case we have a counter-example to the original problem where one of the items has a deterministic value 0 to the buyer. Then we can have a counter-example with a smaller number of items

$n - 1$, as we can use the same \mathcal{M} for the setting where item i is excluded with the revenue larger than $\operatorname{Rev}(\{F_j\}_{j \in [n], j \neq i}, B)$. Note that we cannot continue reducing the counter-example indefinitely, as there is no counter-example for $n = 1$. Thus we assume that $0 \notin \operatorname{supp}(f_i^*)$. Furthermore, we observe the following properties of $\lambda_i(v_i)$ for $v_i \in \operatorname{supp}(f_i^*)$.

CLAIM 2. $\lambda_i(v_i) < \lambda_i(\hat{v}_i)$ for any $v_i < \hat{v}_i$ in the support of f_i^* .

Proof. If this is not the case, then modification $f_i^*(v_i) \leftarrow f_i^*(v_i) + f_i^*(\hat{v}_i)$, $f_i^*(\hat{v}_i) \leftarrow 0$ does not decrease $\Delta_i^{\operatorname{rev}}(f_i^*)$. The new distribution has a smaller support – a contradiction with the definition of f_i^* .

CLAIM 3. $\lambda_i(v_i) > 0$ for any v_i in the support of f_i^* .

Proof. Let us assume towards a contradiction that $\lambda_i(v_i^0) \leq 0$ for some $v_i^0 \in \operatorname{supp}(f_i^*)$. We note that v_i^0 is not unique point in the support of f_i^* , since $\Delta_i^{\operatorname{rev}}(f_i^*) = \sum_{v_i \in V_i} \lambda_i(v_i) \cdot f_i^*(v_i) - \operatorname{Rev}(f_i^*, B_i^o) > 0$. We define $\hat{f}_i(v_i) \leftarrow \frac{f_i^*(v_i)}{1 - f_i^*(v_i^0)}$ for $v_i \neq v_i^0$, and $\hat{f}_i(v_i^0) \leftarrow 0$. For the new distribution \hat{f}_i we have $\operatorname{Rev}(\hat{f}_i, B_i^o) \leq \frac{\operatorname{Rev}(f_i^*, B_i^o)}{1 - f_i^*(v_i^0)}$. Indeed, the probability of any particular type increases only by at most factor of $\frac{1}{1 - f_i^*(v_i^0)}$. On the other hand, $\sum_{v_i \in V_i} \lambda_i(v_i) \cdot \hat{f}_i(v_i) = \sum_{v_i \in V_i} \lambda_i(v_i) \cdot \frac{f_i^*(v_i)}{1 - f_i^*(v_i^0)} - \lambda_i(v_i^0) \cdot \frac{f_i^*(v_i^0)}{1 - f_i^*(v_i^0)} \geq \frac{\sum_{v_i \in V_i} \lambda_i(v_i) \cdot f_i^*(v_i)}{1 - f_i^*(v_i^0)}$. Thus we have $\Delta_i^{\operatorname{rev}}(\hat{f}_i) = \sum_{v_i \in V_i} \lambda_i(v_i) \cdot \hat{f}_i(v_i) - \operatorname{Rev}(\hat{f}_i, B_i^o) \geq \frac{1}{1 - f_i^*(v_i^0)} \left(\sum_{v_i \in V_i} \lambda_i(v_i) \cdot f_i^*(v_i) - \operatorname{Rev}(f_i^*, B_i^o) \right) > \Delta_i^{\operatorname{rev}}(f_i^*)$ and \hat{f}_i has smaller support than f_i^* – a contradiction.

Given that $\{\lambda_i(v_i)\}_{v_i \in \operatorname{supp}(f_i^*)}$ is non negative and monotonically increasing according to Claims 2-3, we try to construct a mechanism \mathcal{M}_i selling individual item i with the payment $p_{\mathcal{M}_i}(v_i)$ matching $\lambda_i(v_i)$ for $v_i \in \operatorname{supp}(f_i^*)$. Let $\operatorname{supp}(f_i^*) = \{v_j\}_{j=1}^k$, $0 < v_1 < v_2 < \dots < v_k$ for some $k \in \mathbb{N}$. We set $q(v_j) \stackrel{\text{def}}{=} \frac{\lambda_i(v_j) - \lambda_i(v_{j-1})}{v_j}$ for $k \geq j \geq 2$ and $q(v_1) \stackrel{\text{def}}{=} \frac{\lambda_i(v_1)}{v_1}$. If $\sum_{j \in [k]} q(v_j) \leq 1$, then we can define \mathcal{M}_i as the mechanism that posts a random price v_j with probability $q(v_j)$ for each $j \in [k]$ and price 0 with the remaining probability. Then $p_{\mathcal{M}_i}(v_i) = \lambda_i(v_i) \leq B_i^o$ and

$$\begin{aligned} \operatorname{Rev}(f_i^*, B_i^o) &\geq \operatorname{Rev}_{\mathcal{M}_i}(f_i^*) = \sum_{v_i \in V_i} f_i^*(v_i) \cdot p_{\mathcal{M}_i}(v_i) \\ &= \sum_{v_i \in V_i} f_i^*(v_i) \cdot \lambda_i(v_i). \end{aligned}$$

I.e., $\Delta_i^{\text{rev}}(f_i^*) = \sum_{v_i \in V_i} \lambda_i(v_i) \cdot f_i^*(v_i) - \text{Rev}(f_i^*, B_i^o) \leq 0$ – a contradiction. Therefore, $\sum_{j \in [k]} q(v_j) > 1$. Then we define \mathcal{M}_i as the mechanism that posts a random price v_j with probability $\frac{q(v_j)}{\sum_{j \in [k]} q(v_j)}$ for each $j \in [k]$. The expected payment $p_{\mathcal{M}_i}(v_i) = \frac{\lambda_i(v_i)}{\sum_{j \in [k]} q(v_j)} < \lambda_i(v_i)$ of \mathcal{M}_i for each $v_i \in \text{supp}(f_i^*)$. We denote these payments by $\tilde{\lambda}_i(v_i) \stackrel{\text{def}}{=} p_{\mathcal{M}_i}(v_i)$. Let $\tilde{\mathcal{M}} \stackrel{\text{def}}{=} (\tilde{\mathbf{x}}, \tilde{p}) = \bigoplus_{i=1}^n \mathcal{M}_i$ be the mechanism that sells each item $i \in [n]$ separately using \mathcal{M}_i . Let us restrict the domain $v_i \in V_i$ to $\tilde{V}_i \stackrel{\text{def}}{=} \text{supp}(f_i^*)$ for each $i \in [n]$ and respectively $\mathbf{V} = \prod_{i=1}^n V_i$ to $\tilde{\mathbf{V}} = \prod_{i=1}^n \tilde{V}_i$. By construction we have

$$\begin{aligned} \forall \mathbf{v} \in \tilde{\mathbf{V}} \quad \tilde{p}(\mathbf{v}) &= \sum_{i \in [n]} \tilde{\lambda}_i(\mathbf{v}) \\ \forall i \in [n], v_i \in \tilde{V}_i \quad \tilde{\lambda}_i(v_i) &< \lambda_i(v_i). \end{aligned}$$

We denote by v_i^{\max}, v_i^{\min} the largest and the smallest types in \tilde{V}_i , and by v_i^+ denote the next type in \tilde{V}_i larger than v_i for each $v_i \in \tilde{V}_i \setminus \{v_i^{\max}\}$. We use $\mathbf{v}^{\min} = (v_1^{\min}, \dots, v_n^{\min})$ to denote the vector of minimal types. For any vector $\mathbf{v} \in \tilde{\mathbf{V}}$ we define a set $I_{\max}(\mathbf{v}) \stackrel{\text{def}}{=} \{i \mid v_i = v_i^{\max}\}$ and decompose $\mathbf{v} = \mathbf{v}^m + \mathbf{w}$ into

$$\begin{aligned} \mathbf{v}^m &\stackrel{\text{def}}{=} \begin{cases} v_i^m = v_i^{\max} & \forall i \in I_{\max}(\mathbf{v}) \\ v_i^m = 0 & \forall i \notin I_{\max}(\mathbf{v}) \end{cases} \\ \mathbf{w} &\stackrel{\text{def}}{=} \begin{cases} w_i = 0 & \forall i \in I_{\max}(\mathbf{v}) \\ w_i = v_i & \forall i \notin I_{\max}(\mathbf{v}) \end{cases} \end{aligned}$$

For \mathbf{v} (or \mathbf{w}) and $i \in [n]$ s.t. $v_i \neq v_i^{\max}$ ($w_i \neq v_i^{\max}$), we define $\mathbf{v}^{i+} \stackrel{\text{def}}{=} (v_{-i}, v_i^+)$, $\mathbf{w}^{i+} \stackrel{\text{def}}{=} (w_{-i}, w_i^+)$.

CLAIM 4. $\mathbf{w} = \sum_{i \notin I_{\max}(\mathbf{v})} \alpha_i \cdot \mathbf{w}^{i+}$, where $\alpha_i \geq 0$ and $\sum_{i \notin I_{\max}(\mathbf{v})} \alpha_i < 1$.

Proof. It is easy to verify that $\alpha_j \stackrel{\text{def}}{=} \frac{w_j}{w_j^+ - w_j} \left(1 + \sum_{i \notin I_{\max}(\mathbf{v})} \frac{w_i}{w_i^+ - w_i} \right)^{-1}$, where $j \notin I_{\max}(\mathbf{v})$ satisfy the condition.

We summarize below some important properties¹⁰ of $\tilde{\mathcal{M}} = (\tilde{\mathbf{x}}, \tilde{p})$:

1. $\tilde{x}_i(v_i^{\max}, \mathbf{v}_{-i}) = 1$ tight alloc. constraint at v_i^{\max}
2. $u(\mathbf{v}^{\min}, \mathbf{v}^{\min}) = 0$ tight IR at \mathbf{v}^{\min}

¹⁰See the description of the optimal single-item mechanism in Section 5.

3. $u(\mathbf{v}^{i+}, \mathbf{v}^{i+}) = u(\mathbf{v}^{i+}, \mathbf{v}) = \langle \mathbf{v}^{i+}, \tilde{\mathbf{x}}(\mathbf{v}) \rangle - \tilde{p}(\mathbf{v})$ tight IC at $\mathbf{v}^{i+} \rightarrow \mathbf{v}$.

We can write down the following inequality for our mechanism $\mathcal{M} = (\mathbf{x}, p)$ and corresponding set of $\{\lambda_i(v_i)\}_{i \in [n], v_i \in \tilde{V}_i}$ for any $\mathbf{v} \in \tilde{\mathbf{V}}$

$$\begin{aligned} (4.7) \quad u(\mathbf{v}, \mathbf{v}) &= \langle \mathbf{v}^m + \mathbf{w}, \mathbf{x}(\mathbf{v}) \rangle - p(\mathbf{v}) \\ &= \sum_{i \notin I_{\max}(\mathbf{v})} \alpha_i \cdot (\langle \mathbf{w}^{i+} + \mathbf{v}^m, \mathbf{x}(\mathbf{v}) \rangle - p(\mathbf{v})) + \\ &\quad + \left(1 - \sum_{i \notin I_{\max}(\mathbf{v})} \alpha_i \right) (\langle \mathbf{v}^m, \mathbf{x}(\mathbf{v}) \rangle - p(\mathbf{v})) \\ &= \sum \alpha_i \cdot u(\mathbf{v}^{i+}, \mathbf{v}) + \left(1 - \sum \alpha_i \right) (\langle \mathbf{v}^m, \mathbf{x}(\mathbf{v}) \rangle - p(\mathbf{v})) \\ &\leq \sum \alpha_i \cdot u(\mathbf{v}^{i+}, \mathbf{v}^{i+}) + \left(1 - \sum \alpha_i \right) \langle \mathbf{v}^m, \mathbf{1} \rangle \\ &\quad - \left(1 - \sum \alpha_i \right) \sum_{i=1}^n \lambda_i(v_i). \end{aligned}$$

Observe that the inequality (4.7) is tight for $\tilde{\mathcal{M}}$ and $\{\tilde{\lambda}_i(v_i)\}_{i \in [n], v_i \in \tilde{V}_i}$. In the following we are going to apply inequality (4.7) multiple times to arrive at a contradiction. We start with an inequality $u(\mathbf{v}^{\min}, \mathbf{v}^{\min}) \geq 0$, which is also tight for $\tilde{\mathcal{M}}$ and apply (4.7) to write an upper bound on $u(\mathbf{v}^{\min}, \mathbf{v}^{\min})$. We obtain inequality of the form:

$$(4.8) \quad C_1 + \sum_{\mathbf{v} \in \tilde{\mathbf{V}}} \beta(\mathbf{v}) \cdot u(\mathbf{v}, \mathbf{v}) - \sum_{i \in [n]} \sum_{v_i \in \tilde{V}_i} \gamma_i(v_i) \cdot \lambda_i(v_i) \geq 0,$$

where $\beta(\mathbf{v}) \geq 0, \gamma_i(v_i) \geq 0, C_1 = \text{Const} > 0$.

We can apply (4.7) again to substitute a term $u(\mathbf{v}, \mathbf{v})$ in LHS of (4.8) with positive coefficient $\beta(\mathbf{v}) > 0$ and derive another inequality of the form (4.8). Notice that at each iteration of this process, the types \mathbf{v} in $u(\mathbf{v}, \mathbf{v})$ in LHS of (4.8) are increasing. Therefore, after finite number of steps we can transform the inequality (4.8) into one of the form

$$(4.9) \quad C \geq \sum_{i \in [n]} \sum_{v_i \in \tilde{V}_i} \delta_i(v_i) \cdot \lambda_i(v_i)$$

where $\delta_i(v_i) \geq 0, C = \text{Const} > 0$

Note that since each step in the derivation of inequality (4.9) was tight for $\tilde{\mathcal{M}}$, we also have (4.9) to be tight for $\tilde{\mathcal{M}}$ meaning that $C = \sum_{i \in [n]} \sum_{v_i \in \tilde{V}_i} \delta_i(v_i) \cdot \tilde{\lambda}_i(v_i)$. I.e., at least one of $\delta_i(v_i)$ is strictly greater than 0. Since $\tilde{\lambda}_i(v_i) < \lambda_i(v_i)$ for any $i \in [n]$ and $v_i \in \tilde{V}_i$ we

get

$$C = \sum_{i \in [n]} \sum_{v_i \in \tilde{V}_i} \gamma_i(v_i) \cdot \tilde{\lambda}_i(v_i)$$

$$< \sum_{i \in [n]} \sum_{v_i \in \tilde{V}_i} \gamma_i(v_i) \cdot \lambda_i(v_i) \leq C.$$

A contradiction.

5 Optimal Item-Budgets Mechanism

In this section, we show how to find the optimal item-budgets auction and prove some properties of it, which are used in the previous section.

Optimal single-item auction. We note that the solution to this problem is known for more general setting [15], where the buyer has private valuation and private budget and the seller only knows the joint distribution of possibly correlated agent's budget and value. Here, however, we give an independent and simple solution to this problem in the special case of public budget. It will help us to solve the next problem of designing optimal item-budgets mechanism (i.e., finding optimal budget partition) for the multi-item problem. We begin by recalling a well-known equivalence between truthful single-item single-buyer auctions [21] and randomized posted-price mechanisms. I.e., any truthful mechanism for single-item single-buyer auction can be expressed as a convex combination of posted-price mechanisms (in other words, it is a random mechanism with a randomly selected take-it-or-leave price offer). Indeed, one can express any monotone allocation curve x of a given truthful mechanism $\mathcal{M} = (x, p)$ as a convex combination of step-functions each corresponding to a posted-price mechanism. The resulting random posted-price mechanism has the same payment function p as \mathcal{M} by revenue equivalence principle.

We emphasize the convenience of such representation of truthful mechanisms as convex combinations of single posted-price auctions. Indeed, the budget constraint can be stated as a single linear inequality $p(v_i^{\max}) \leq B_i$ for the payment of the highest type v_i^{\max} . Furthermore, both the revenue objective and the payment $p(v_i^{\max})$ are linear functions for the operation of taking convex combinations of mechanisms. With these observations at hand, we can plot a simple graph representing every deterministic posted-price mechanism on the plane with the axes corresponding to the payment of the highest type $p(v_i^{\max})$ (x -axis) and expected revenue Rev (y -axis) and get the set of all truthful mechanisms as a convex hull of these points (see Figure 1).

Given the convex polygon of truthful mechanisms we can easily find the revenue maximizing mechanism




Figure 1: Bold dots represent single price mechanisms with posted prices in $\{0, V_i\} \subset [0, v_i^{\max}]$. Shaded area represents set of all truthful mechanisms.

satisfying budget constraint B_i : we just need to consider intersection of the polygon with half-space $x \leq B_i$ and take the maximal y -axis point in this set. In particular, the optimal solution will be either a single vertex (i.e., deterministic posted price mechanism), or a convex combination $q_1 \cdot \mathcal{M}_1 + q_2 \cdot \mathcal{M}_2$ of two deterministic posted-price mechanisms $\mathcal{M}_1, \mathcal{M}_2$ with respective prices r_1 and r_2 and probabilities $q_1 + q_2 = 1$. In the latter case the expected payment of the highest type $p(v_i^{\max}) = q_1 \cdot r_1 + q_2 \cdot r_2 = B_i$.¹¹ We now return to the general case of multi-item auction and describe optimal budget partition across different items.

Optimal partition of item budgets. We recall that the optimal multi-unit auction for a single budget constrained buyer is the one where items are auctioned independently, each item i being sold in the optimal auction with a specified budget constraint B_i . The set of budgets $\{B_i\}_{i=1}^n$ is such that $\sum_{i=1}^n B_i \leq B$. We already know how to find the optimal revenue $\text{Rev}(F_i, B_i) = \text{Rev}_i(B_i)$ in a single-item auction with given budget B_i . We only left to find the optimal partition of budgets $\{B_i\}_{i=1}^n$ that solves the following problem

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

¹¹Note that in this case the mechanism \mathcal{M}_2 (one with the higher reserve $r_2 > r_1$) is not budget feasible, since for the equality $q_1 \cdot r_1 + q_2 \cdot r_2 = B_i$ to hold we must have $r_2 > B_i > r_1$. However, the convex combination of \mathcal{M}_1 and \mathcal{M}_2 is budget feasible and admits the following simple description with a 2-option menu: offer (i) entire item $x_i = 1$ for the price $B_i = q_1 \cdot r_1 + q_2 \cdot r_2$, or (ii) $x_i = q_1$ for the price $q_1 \cdot r_1$

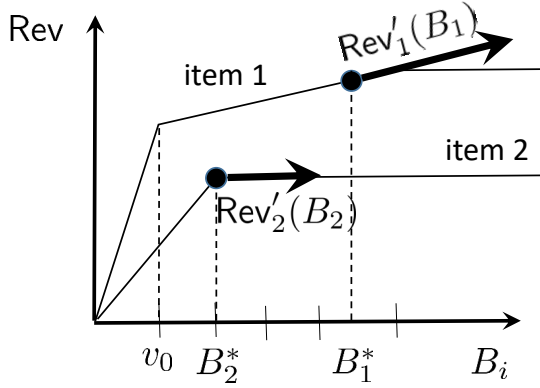


Figure 2: $\text{Rev}_1(B_1)$ and $\text{Rev}_2(B_2)$ concave curves. Bold dots represent the final budget partition. The greedy algorithm starts with $B_1 \leftarrow \theta_0$, then $B_2 \leftarrow B_2^*$ and finally $B_1 \leftarrow B_2^* = B - B_2$.

From the previous paragraph we already know that each $\text{Rev}_i(B_i)$ is a concave, monotone, and piece-wise linear function of B_i . Indeed, it is just an upper envelop of the truthful mechanisms polygon up to the point with the highest y -axis (i.e., posted price mechanism with the highest revenue), after that point the function $\text{Rev}_i(B_i)$ is just a constant. One can recognize here a simple concave maximization problem with budget constraint. We know that this problem can be solved with a simple greedy algorithm that starts at the point of $B_i = 0$ for all $i \in [n]$ and proceeds by greedily increasing one B_i at a time with the current highest derivative (slope) of $\text{Rev}_i(B_i)$ until derivative of $\text{Rev}_i(B_i)$ changes, or when algorithm reaches the point $\sum_{i=1}^n B_i = B$, or when all functions $\text{Rev}_i(B_i)$ reach their respective maximums. We note that such an algorithm will need to make at most $\sum_{i=1}^n |V_i|$ changes of the derivative value and can be implemented in $O(\sum_{i=1}^n |V_i|)$ time. An interesting observation about the optimal solution is that for all but at most one item the optimal mechanism posts deterministic prices, and for only one item the mechanism may use randomized outcomes (see Figure 2).

6 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.

Acknowledgments

We sincerely thank Gabriel Carroll for helpful comments and for pointing out a gap in the proof in an early version of this paper.

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 multi-dimensional 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] Yeon-Koo Che and Ian Gale. The optimal mechanism for selling to a budget-constrained buyer. *Journal of Economic theory*, 92(2):198–233, 2000.
- [16] Kim-Sau Chung and Jeffrey C Ely. Foundations of dominant-strategy mechanisms. *The Review of Economic Studies*, 74(2):447–476, 2007.
- [17] 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.
- [18] 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.
- [19] 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.
- [20] Daniel F Garrett. Robustness of simple menus of contracts in cost-based procurement. *Games and Economic Behavior*, 87:631–641, 2014.
- [21] Andrew V Goldberg, Jason D Hartline, and Andrew Wright. Competitive auctions and digital goods. In *Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pages 735–744. Society for Industrial and Applied Mathematics, 2001.
- [22] Nima Haghpanah and Jason D. Hartline. Reverse mechanism design. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pages 757–758, 2015.
- [23] 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.
- [24] 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.
- [25] Sergiu Hart and Philip J Reny. Maximal revenue with multiple goods: Nonmonotonicity and other observa-

- tions. *Theoretical Economics*, 10(3):893–922, 2015.
- [26] 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.
- [27] Jonathan Levin. An optimal auction for complements. *Games and Economic Behavior*, 18(2):176–192, 1997.
- [28] 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.
- [29] 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.
- [30] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- [31] Pingzhong Tang and Ziheng Wang. Optimal auctions for negatively correlated items. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 103–120. ACM, 2016.
- [32] John Thanassoulis. Haggling over substitutes. *Journal of Economic theory*, 117(2):217–245, 2004.
- [33] 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.