

The power of randomness in Bayesian optimal mechanism design

Shuchi Chawla*

David Malec†

Balasubramanian Sivan‡

Abstract

We investigate the power of randomness in the context of a fundamental Bayesian optimal mechanism design problem—a single seller aims to maximize expected revenue by allocating multiple kinds of resources to “unit-demand” agents with preferences drawn from a known distribution. When the agents’ preferences are single-dimensional Myerson’s seminal work [14] shows that randomness offers no benefit—the optimal mechanism is always deterministic. In the multi-dimensional case, where each agent’s preferences are given by different values for each of the available services, Briest et al. [7] recently showed that the gap between the expected revenue obtained by an optimal randomized mechanism and an optimal deterministic mechanism can be unbounded even when a single agent is offered only 4 services. However, this large gap is attained through unnatural instances where values of the agent for different services are correlated in a specific way. We show that when the agent’s values involve no correlation or a specific kind of positive correlation, the benefit of randomness is only a small constant factor (4 and 8 respectively). Our model of positively correlated values (that we call additive values) is a natural model for unit-demand agents and items that are substitutes. Our results extend to multiple agent settings as well.

1 Introduction

A fundamental objective in the design of mechanisms is to maximize the seller’s revenue. In the absence of any information about buyers’ preferences, i.e. in prior-free settings, randomization is a frequently used algorithmic technique (see, e.g., [11] and references therein); In a spirit similar to randomization in online algorithm design, it allows the seller to hedge against adversarial values. While randomization unsurprisingly turns out to be essential for any guarantees on revenue in certain prior-free settings, it appears to be not so in Bayesian settings where the designer has distributional information about the agents’ types and the goal is to maximize revenue in expectation over the distribution. For example, for a single item auction in the Bayesian setting, Myerson’s seminal work [14] shows that the optimal mechanism is always a deterministic one.

In this work we investigate the power of randomness in the context of the following archetypical multi-parameter optimal mechanism design problem — a single seller offers multiple kinds of service, and a number of “unit-demand” agents are each interested in buying any one of the services. Whereas in Myerson’s work each agent has a single-dimensional type (namely a value for the item under sale), in our setting each agent has a multi-dimensional type characterized by a (different) value for each of the services offered by the seller. An example of such a setting is an online travel agency selling airline tickets, hotel rooms, etc.

*Computer Sciences Dept., University of Wisconsin - Madison. shuchi@cs.wisc.edu.

†Computer Sciences Dept., University of Wisconsin - Madison. dmalec@cs.wisc.edu.

‡Computer Sciences Dept., University of Wisconsin - Madison. balu2901@cs.wisc.edu.




Figure 1: An example from [18] contrasting the optimal item pricing and the optimal lottery pricing. The regions R_1 , R_2 , and R_{Lot} denote the sets of valuations at which the agent buys item 1, item 2, and the $(1/2, 1/2)$ lottery respectively.

Customers have different preferences over different available services, but are only interested in buying one. We study the Bayesian version of this problem: the distribution from which the buyers' preferences are drawn is known to the seller. Given Myerson's observation about single-dimensional settings, one might expect that in the multi-dimensional case the optimal mechanism (ignoring computational issues) is once again deterministic. Thanassoulis [18] and Manelli and Vincent [12] independently discovered that this is not the case. This raises the following natural question: *what quantitative benefit do randomized mechanisms offer over deterministic ones in Bayesian optimal mechanism design?*

To answer this question we must first understand the structure of randomized mechanisms in multi-dimensional settings. In the context of a single unit-demand agent and a seller offering multiple items, any deterministic mechanism is simply a pricing for each of the items with the agent picking the one that maximizes her utility (her value for the item minus its price). Likewise, randomized mechanisms can be thought of as pricings for distributions or convex combinations over items. These convex combinations are called *lotteries*. A risk-neutral buyer with a quasiconcave utility function buys the lottery that maximizes his expected value minus the price of the lottery.

The following example due to Thanassoulis explains how lotteries work. Suppose that a seller offers two items for sale to a single buyer, and that the buyer's value for each of the items is independently uniformly distributed in the interval $[5, 6]$. The optimal deterministic mechanism for the seller is to simply price each of the items at $p^* = \$5.097$ (see Figure 1). In a randomized mechanism, the seller may in addition price a $(1/2, 1/2)$ distribution over the two items at a slightly lower price of $p' = \$5.057$. If the buyer buys this lottery, the seller tosses a coin and allocates the first item to her with probability $1/2$ and the second with probability $1/2$. A buyer that is nearly indifferent between the two items would prefer to buy the lottery because of its lower cost, than either one of the items. While the seller loses some revenue by selling the lower priced lottery with some probability, he gains by selling to a larger segment of the market (those that cannot afford either of the individual items but can afford the lower priced lottery). In this example the gain is more than the loss, so that introducing the lottery improves the seller's revenue. As this example indicates, lotteries help in optimal mechanism design by giving the seller more latitude to price discriminate among buyers with different preferences.

In general, a randomized mechanism can offer to the buyer a menu of prices for arbitrarily many lotteries. We call such a menu a *lottery pricing*, and likewise a deterministic pricing an *item pricing*. While in multiple




Figure 2: An example of a product distribution for valuations contrasted against an example of an additive distribution.

agent settings randomized mechanisms can be more complicated, we show that any such mechanism can be interpreted as offering to each agent simultaneously a lottery pricing that is a function of values of other agents.

The question of whether and to what extent randomization helps in Bayesian optimal mechanism design is not merely a pedantic one. Mechanisms similar to lottery pricings are seen in practice. For example, the website priceline.com routinely sells airline tickets to customers without disclosing at the time of sale crucial details such as the time of travel, carrier, etc. While customers are unaware of the distribution from which the final service is picked, the tradeoffs for customers are similar—the uncertainty in the quality of the final item against the cheaper price. Travel agencies offering vacation packages use similar devices.

Until recently, the largest gap known between item pricings and lottery pricings for a single agent was a gap of $3/2$ due to Pavlov [15]; For the special case where values for different items are independent, Thanassoulis gave the best gap example with a gap of 1.1. Recently Briest et al. [7] showed that in single-agent settings in fact the gap between lottery pricings and item pricings can be unbounded even with only 4 items. However the value distributions for which such gaps are achieved are quite unnatural with the values of different items being highly correlated. In this paper we show that the gap between lottery pricings and item pricings is small for distributions involving limited correlation between items.

We further extend these results to the multiple-agent setting with the seller facing a general feasibility constraint, obtaining the first results of this kind. Mechanism design in the multiple-agent multi-parameter setting is poorly understood [19]. Until recently there were no general characterizations for optimal or approximately optimal mechanisms similar to Myerson’s for the single-parameter case. Chawla et al. [9] recently developed constant-factor approximations to optimal *deterministic* mechanisms in this setting for a certain class of feasibility constraints (namely matroids and related set systems). We extend their results to show that their (deterministic) mechanisms achieve a constant factor approximation with respect to the optimal randomized mechanism as well, again implying a small gap between randomized and deterministic mechanisms.

Our results and techniques

We follow a technique introduced in [8] for relating multi-parameter mechanisms to mechanisms for a related single-parameter problem. Chawla, Hartline and Kleinberg [8] relate a single unit-demand agent m -item mechanism design problem to an m -agent single-item auction setting, by “splitting” the unit-demand agent into m independent “copies”. They argue that the increased competition among copies benefits the seller and leads to higher revenue. Formally, given an item pricing \mathbf{p} they construct a truthful mechanism

A^P that allocates the item to agent i whenever \mathbf{p} allocates item i to the multi-parameter agent (that is, A^P has the “same” allocation rule as \mathbf{p}). They then argue that the price that A^P charges is no less than the price that \mathbf{p} charges for any instantiation of values. Therefore, *the expected revenue of the optimal multi-parameter mechanism is bounded above by the expected revenue of Myerson’s mechanism for the related single-parameter problem with copies*. Chawla et al. use this upper bound to design an item pricing for the multi-parameter problem with revenue within a factor of 3 of the expected revenue of Myerson’s mechanism for the instance with copies, thereby obtaining a 3-approximation to the optimal deterministic mechanism for the single-agent problem.

Unfortunately the upper bound of the expected revenue of Myerson’s mechanism does not hold for randomized mechanisms. The appendix gives an example where the revenue of Myerson’s mechanism for the instance with copies is a factor of 1.13 smaller than that of the optimal lottery pricing for the multi-parameter problem. In fact, the mechanism A^L with the “same” allocation rule as a lottery pricing \mathcal{L} may obtain zero revenue even when the lottery pricing obtains non-zero revenue. Our main result is that this gap between Myerson’s mechanism and the optimal lottery pricing is no larger than a factor of 2. Specifically, given a lottery pricing, we can construct two mechanisms, one being A^L and the other a Vickrey auction, such that the sum of the revenues of the two mechanisms is an upper bound on the revenue of the lottery pricing. Combining this with the result of Chawla, Hartline and Kleinberg (and an improvement over it in [9]), we get that for a single unit-demand agent multi-parameter problem, the gap between lottery pricings and item pricings is at most 4.

Chawla et al.’s result as well as our factor-of-4 gap holds for instances where the values of the agent for different items are independent. For a unit-demand agent, this independence assumption is unrealistic. However, on the other end of the spectrum, Briest et al. show that with arbitrary correlations between item values, the gap can be unbounded. We therefore examine the following natural model for values involving limited correlation. The type of the unit-demand agent is $m + 1$ dimensional — (t_0, t_1, \dots, t_m) ; the agent’s value for item i is $v_i = t_0 + t_i$. Here t_0 can be thought of as the buyer’s “base” value for obtaining any of the items, and the t_i ’s represent the buyer’s perceived quality of the different items. This additive value distribution introduces a positive correlation between values of different items¹. Figure 2 shows an example of one such discrete distribution contrasted against a product distribution.

In this additive distribution setting we show that the gap between randomized and deterministic mechanisms is at most a factor of 8. Once again our approach is to start with an optimal lottery pricing for the multi-parameter instance, construct an ensemble of mechanisms based on it for the related single-parameter instance, and then construct a pricing for the multi-parameter instance based on the mechanisms.

Our results extend to multi-agent settings as well. The simplest multi-agent setting we consider involves n agents and m items (with copies), where the seller faces a supply constraint for each of the items. A feasible allocation is a matching between agents and items that respects multiplicities of items. More generally, we consider settings where the seller faces a matroid feasibility constraint—any feasible allocation must be an independent set in a given matroid in addition to allocating at most one item per agent (see Section 5.1 for the definition of a matroid). In both these cases we show that the gap between the expected revenue of the optimal randomized and the optimal deterministic mechanisms is a small constant factor. Once again we rely on the approach of relating the multi-parameter instance to a single-parameter instance where each unit-demand agent is split into multiple selfish “pseudo-agents”. This approach was first developed in [9]. In particular we showed in [9] that for the settings described above, there exist deterministic mechanisms that obtain revenue within a constant factor of the revenue of Myerson’s mechanism for the related single-parameter instance. In Section 5 we show that the revenue of any randomized mechanism for

¹This model is similar to “multiplicative” value distributions that have been studied previously in the context of bundle pricing problems (see, e.g., [2]).

these settings can be bounded from above by 5 times the revenue of Myerson’s mechanism for the single-parameter instance. The challenge in these settings is to ensure that the mechanisms that we construct satisfy the non-trivial feasibility constraint that the seller faces.

Related work

As mentioned earlier, randomness is used extensively in prior-free mechanism design (see, e.g., [11] and references therein). While symmetric deterministic mechanisms provably cannot obtain any guarantees on revenue in that setting, Aggarwal et al. [1] show that by exploiting asymmetry prior-free mechanisms can be derandomized at a constant factor loss in revenue.

Our mechanism design setting with unit-demand agents is closely related to the standard setting for envy-free pricing problems considered in literature [10, 5, 4, 6, 8]; Those works study the single-agent problem with a correlated value distribution and aim to approximate the optimal deterministic mechanism (item pricing). Our single-agent setting is most closely related to the work of Chawla, Hartline and Kleinberg [8] who gave a 3 approximation to the optimal deterministic mechanism for single-agent product-distribution instances, and builds upon techniques developed in that work.

In economics literature, the study of Bayesian optimal mechanisms has focused on deterministic mechanisms. It is well-known that for single-parameter instances the optimal mechanism is deterministic [14, 16]. Following Myerson’s result [14] for single-parameter mechanisms, there were a number of attempts to obtain simple characterizations of optimal mechanisms in the multi-parameter setting [13, 17, 19], however no general-purpose characterization of such mechanisms is known [19]. Recently Chawla et al. [9] gave the first approximations to optimal deterministic mechanisms for a large class of multi-parameter problems. This paper extends techniques developed in that work and one of the implications of our work is that the mechanisms developed in [9] are approximately-optimal with respect to the optimal randomized mechanisms as well.

The study of the benefit of randomness in multi-parameter mechanism design was initiated by Thanassoulis [18] who presented single-parameter instances with valuations drawn from product distributions where randomness helps increase the revenue by about 8-10%. Manelli and Vincent [12] and Pavlov [15] presented other examples with small gaps. Briest et al. [7] were the first to uncover the extent of the benefit of randomization. They showed that lottery pricings can be arbitrarily better than item pricings in terms of revenue even for the case of 4 items offered to a single agent.

2 Definitions and problem set-up

2.1 Bayesian optimal mechanism design

We study the following mechanism design problem. There is one seller and n buyers (agents) indexed by the set I . The seller offers m different services indexed by the set J . Agents are risk-neutral and are each interested in buying any one of the m services. Agent i has value v_{ij} for service j which is a random variable. We use \mathbf{v}_{-i} to denote the vector of values of all agents except agent i . The seller faces no costs for providing service, but must satisfy certain feasibility constraints (e.g. supply constraints in a limited supply setting). We represent these feasibility constraints as a set system \mathcal{J} over pairs (i, j) , that is, $\mathcal{J} \subseteq 2^{I \times J}$. Each subset of $I \times J$ in \mathcal{J} is a feasible allocation of services to agents.

The seller’s goal is to maximize her revenue in expectation over the buyers’ valuations. We call this problem the *Bayesian multi-parameter unit-demand (optimal) mechanism design problem* (BMUMD). A deterministic mechanism for this problem maps any set of bids \mathbf{b} to an allocation $M(\mathbf{b}) \in \mathcal{J}$ and a pricing

$\pi(\mathbf{b})$ with a price π_i to be paid by agent i . A randomized mechanism maps a set of bids to a distribution over \mathcal{J} ; we use $M(\mathbf{b})$ to denote this distribution over $I \times J$.

We focus on the class of incentive compatible mechanisms and will hereafter assume that $\mathbf{b} = \mathbf{v}$. We use $\mathcal{R}^M(\mathbf{v})$ to denote the revenue of a mechanism M at valuation vector \mathbf{v} : $\mathcal{R}^M(\mathbf{v}) = \sum_{i \in I} \pi_i(\mathbf{v})$ where π is the pricing rule for M . To aid disambiguation, we sometimes use $\mathcal{R}_i^M(\mathbf{v})$ to denote $\pi_i(\mathbf{v})$ for M . The expected revenue of a mechanism is $\mathcal{R}^M = \mathbb{E}_{\mathbf{v}}[\mathcal{R}^M(\mathbf{v})]$.

We consider the following special cases of the BMUMD:

Setting 1: Single agent with independent values. The agent values item j at v_j , which is an independent random variable with distribution F_j and density f_j .

Setting 2: Single agent with additive values. There are m items, and the agent’s type, $\{t_0, \dots, t_m\}$, is $m + 1$ dimensional. t_j is distributed independently according to F_j . The agent’s value for item j is $v_j = t_0 + t_j$.

Setting 3: Multiple agents and multiple items with independent values. There are n agents and m items. Agent i ’s value for item j , v_{ij} , is distributed independently according to F_{ij} . Any matching between items and agents is feasible.

Setting 4: Multiple unit-demand agents with matroid feasibility constraint. There are n agents and m services. Agent i ’s value for item j , v_{ij} , is distributed independently according to F_{ij} . The set system \mathcal{J} is an intersection of a matroid with the unit-demand constraints for the agents and is thus a generalization of the previous matching setting. (See Section 5.1 for the definition of a matroid.)

Single-parameter mechanism design

The single-parameter version of the Bayesian optimal mechanism design problem (abbreviated BSMD) is stated as follows. There are n single-parameter agents and a single seller providing a certain service. Agent i ’s value v_i for getting served is a random variable. We use \mathbf{v}_{-i} to denote the vector of values of all agents except agent i . The seller faces a feasibility constraint specified by a set system $\mathcal{J} \subseteq 2^{[n]}$, and is allowed to serve any set of agents in \mathcal{J} . As in the multi parameter case, a mechanism M for this problem is a function that maps a vector of values \mathbf{v} to an *allocation* $M(\mathbf{v}) \in \mathcal{J}$ and a *pricing* $\pi(\mathbf{v})$. Myerson’s seminal work describes the revenue maximizing mechanism for BSMD; this optimal mechanism is deterministic.

2.2 Relating multi-parameter MD to single-parameter MD

In previous work [9] we presented a general reduction from the multi-parameter optimal mechanism design problem to the single-parameter setting. This approach begins with defining an instance $\mathcal{I}^{\text{copies}}$ of the BSMD given an instance \mathcal{I} of the BMUMD. Our previous work then shows that for several kinds of feasibility constraints there exists a deterministic mechanism for \mathcal{I} with revenue at least a constant fraction of that of the optimal mechanism for $\mathcal{I}^{\text{copies}}$. We state these results below without proof.

We begin by describing the instance $\mathcal{I}^{\text{copies}}$. Let \mathcal{I} be an instance of the BMUMD with n agents and a single seller providing m different services, and with feasibility constraint \mathcal{J} . We define a new instance of the BSMD in the following manner. We split each agent in \mathcal{I} into m distinct agents (hereafter called “copies” or “pseudo-agents”). Each pseudo-agent is interested in a single item $j \in [m]$ and behaves independently of (and potentially to the detriment of) other pseudo-agents. Formally, the instance has mn distinct pseudo-agents each interested in a single service; pseudo-agent (i, j) ’s value for getting served, v_{ij} , is distributed according to F_{ij} . The mechanism again faces a feasibility constraint given by the set system \mathcal{J} .

$\mathcal{I}^{\text{copies}}$ is similar to \mathcal{I} except that it involves more competition (among different pseudo-agents corresponding to the same multi-parameter agent). Therefore it is natural to expect that a seller can obtain more revenue in the instance $\mathcal{I}^{\text{copies}}$ than in \mathcal{I} . The following results show that in Settings 1 and 3 it cannot obtain too much more.

Theorem 1 (Theorem 4 and 10 in [9]) *Given an instance \mathcal{I} of the single agent BMUMD (Setting 1), there exists a truthful deterministic mechanism for \mathcal{I} , whose revenue is at least 1/2 of the revenue of any truthful mechanism for the instance $\mathcal{I}^{\text{copies}}$.*

Theorem 2 (Theorem 14 in [9]) *Given an instance \mathcal{I} of the BMUMD with multiple agents and multiple items (Setting 3), there exists a truthful deterministic mechanism for \mathcal{I} , whose revenue is at least 4/27th of the revenue of any truthful mechanism for the instance $\mathcal{I}^{\text{copies}}$.*

In Setting 4, [9] obtain a somewhat weaker result comparing the revenue of an incentive-compatible mechanism for $\mathcal{I}^{\text{copies}}$ to that of a deterministic mechanism for \mathcal{I} that is not truthful but is an *implementation in undominated strategies* [3]. Formally, for an agent i , a strategy s_i is said to be dominated by a strategy s'_i if for all strategies s_{-i} of other agents, the utility that i obtains from using s_i is no better than that from using s'_i , and for some strategy s_{-i} , it is strictly worse. A mechanism is an algorithmic implementation of an α -approximation in undominated strategies if for every outcome of the mechanism where every agent plays an undominated strategy, the objective function value of the mechanism is within a factor of α of the optimal, and every agent can easily compute for any dominated strategy a strategy that dominates it.

Theorem 3 (Theorem 17 in [9]) *Given an instance \mathcal{I} of the BMUMD with unit-demand agents and a general matroid constraint (Setting 4), there exists a deterministic mechanism for \mathcal{I} implemented in undominated strategies, whose revenue is at least 1/8th of the revenue of any truthful mechanism for the instance $\mathcal{I}^{\text{copies}}$.*

3 Lotteries and randomized mechanisms

We now define a class of mechanisms for the BMUMD that will be useful in our analysis. The following subsection shows that this class encompasses arbitrary randomized mechanisms.

3.1 Lotteries or random allocations

An m -dimensional *lottery* is a vector $\ell = (q_1, \dots, q_m, p)$ where p is the price of the lottery and (q_1, \dots, q_m) is a probability distribution over m items, $\sum_{j \in [m]} q_j \leq 1$. A *lottery pricing* $\mathcal{L} = \{\ell_1, \ell_2, \dots\}$ is a randomized selling mechanism for m items targeted towards a single unit-demand buyer where the buyer is offered a collection of (an arbitrary number of) lotteries. The buyer can select any one or no lottery from the collection, and is then allocated an item drawn from the probability distribution defined by the lottery and charged the price of the lottery. A rational risk-neutral buyer selects the lottery that maximizes her utility: $\sum_{j \in [m]} q_j v_j - p$.

A *lottery-based mechanism* $M^{\mathcal{L}}$ for m services targeted towards n agents is a randomized selling mechanism defined through an ensemble of lottery pricings \mathcal{L} . $M^{\mathcal{L}}$ and \mathcal{L} satisfy the following properties:

1. For every instantiation of values of the agents \mathbf{v} , \mathcal{L} contains n lottery pricings, $\mathcal{L}_1(\mathbf{v}), \dots, \mathcal{L}_n(\mathbf{v})$, where $\mathcal{L}_i(\mathbf{v})$ is an m -dimensional lottery pricing targeted toward agent i .
2. $\mathcal{L}_i(\mathbf{v})$ is a function of \mathbf{v}_{-i} , the values of all agents other than agent i .

3. The mechanism $M^{\mathcal{L}}$ is implemented as follows. It first elicits bids \mathbf{b} from agents, and then offers to agent i (simultaneously with other agents) the lottery pricing $\mathcal{L}_i(\mathbf{b})$. Let $\ell_i(\mathbf{b})$ denote the lottery picked by agent i and let $q_{ij}(\mathbf{b})$ denote the probability with which lottery $\ell_i(\mathbf{b})$ offers service j to agent i . Agent i is allocated item j with probability $q_{ij}(\mathbf{b})$.²
4. The probabilities $q_{ij}(\mathbf{v})$ satisfy the following feasibility constraint:

$$\sum_{(i,j) \in S} q_{ij}(\mathbf{v}) \leq r(S), \forall S \subseteq I \times J, \forall \mathbf{v}$$

where $r(S)$ is the cardinality of some maximum sized feasible subset of S .³

3.2 Randomized mechanisms as lotteries

We now show that every truthful randomized mechanism for the BMUMD can be interpreted as a truthful lottery-based mechanism.

Lemma 4 *Every incentive-compatible randomized mechanism for a multi-agent BMUMD problem is equivalent to a lottery-based mechanism.*

Proof: Given a mechanism M with randomized allocation rule $M(\mathbf{v})$ and pricing rule $\pi(\mathbf{v})$ we define a lottery-based mechanism as follows. Consider an agent i and a fixed instantiation of \mathbf{v}_{-i} . Then for every instantiation of \mathbf{v}_i , consider the probabilities with which M allocates service j to agent i , as well as the prices that M charges. Each such probability vector along with the corresponding price forms a lottery in $\mathcal{L}_i(\mathbf{v}_{-i})$ in the new mechanism. Formally, $\mathcal{L}_i(\mathbf{v}_{-i}) = \{(\mathbf{q}_i, p_i) \mid \exists \mathbf{v}_i \text{ with } \mathbf{q}_i = M_i(\mathbf{v}_{-i}, \mathbf{v}_i) \text{ and } p_i = \pi_i(\mathbf{v}_{-i}, \mathbf{v}_i)\}$.

We now claim that the allocation rule and pricing rule of the new mechanism is precisely the same as the old mechanism. Suppose not. Then at some valuation vector \mathbf{v} and for some agent i , $(\mathbf{q}_i(\mathbf{v}_{-i}, \mathbf{v}_i), p_i(\mathbf{v}_{-i}, \mathbf{v}_i)) \neq (M_i(\mathbf{v}_{-i}, \mathbf{v}_i), \pi_i(\mathbf{v}_{-i}, \mathbf{v}_i))$, where the former is the allocation and price rule for the lottery-based mechanism and the latter the allocation and price rule for the original mechanism M . But, given our construction, $(\mathbf{q}_i(\mathbf{v}_{-i}, \mathbf{v}_i), p_i(\mathbf{v}_{-i}, \mathbf{v}_i)) = (M(\mathbf{v}_{-i}, \mathbf{v}'_i), \pi_i(\mathbf{v}_{-i}, \mathbf{v}'_i))$ for some other value vector \mathbf{v}'_i for agent i . But this implies that in M agent i can benefit from lying and reporting \mathbf{v}'_i when the true value vector is \mathbf{v} . This contradicts the incentive compatibility of M . ■

3.3 A mechanism for $\mathcal{I}^{\text{copies}}$ based on lotteries

As noted earlier, our main technique is to relate the revenue of lottery-based mechanisms for an instance \mathcal{I} of the BMUMD to the optimal mechanism for a related instance $\mathcal{I}^{\text{copies}}$ of the BSMD. We now describe a mechanism for $\mathcal{I}^{\text{copies}}$ based on a given lottery-based mechanism for \mathcal{I} .

Consider an instance \mathcal{I} of the BMUMD. Given a lottery-based mechanism $M^{\mathcal{L}}$ for \mathcal{I} that uses the ensemble of lottery pricings \mathcal{L} , we define a mechanism $\mathcal{A}^{\mathcal{L}}$ for the instance $\mathcal{I}^{\text{copies}}$.

Based on \mathcal{L} , the mechanism $\mathcal{A}^{\mathcal{L}}$ forms a one dimensional lottery pricing for each of the mn pseudo-agents. The lottery pricing offered to pseudo-agent (i, j) , which we denote \mathcal{L}_{ij} , is a function of \mathbf{v}_{-ij}

²Note that these allocations to agents are not necessarily done independently; The feasibility constraint may require correlations between items allocated to different agents. However these details do not affect our analysis, so we ignore them.

³This condition is weaker than may be necessary for certain kinds of set systems, but suffices for our purpose.

and is derived from the lottery pricing $\mathcal{L}_i \in \mathcal{L}$ as follows. Given a valuation vector \mathbf{v}_{-ij} , for each $\ell = (q_{i1}, q_{i2}, \dots, q_{im}, p) \in \mathcal{L}_i(\mathbf{v}_{-i})$, $\mathcal{A}^{\mathcal{L}}$ adds a lottery $\ell_j = (q', p')$ to \mathcal{L}_{ij} defined by

$$q' = q_{ij}; \text{ and}$$

$$p' = p - \sum_{k \neq j} q_{ik} v_{ik} + u_{ij}(\mathbf{v}_{-ij}),$$

where the term $u_{ij}(\mathbf{v}_{-ij}) \geq 0$ is chosen to be the least value ensuring that the lottery preferred by pseudo-agent (i, j) when $v_{ij} = 0$ (if any) has a non-negative price.

We note the following properties of $\mathcal{A}^{\mathcal{L}}$:

1. (truthfulness) That $\mathcal{A}^{\mathcal{L}}$ is truthful follows immediately from the fact that the one dimensional lottery pricing \mathcal{L}_{ij} offered to pseudo-agent (i, j) does not depend on v_{ij} , and the pseudo-agent may choose any lottery from \mathcal{L}_{ij} .
2. (allocation rule) Suppose first that for (i, j) and some \mathbf{v}_{-ij} , $u_{ij}(\mathbf{v}_{-ij}) = 0$. Then for any v_{ij} , the utility of pseudo-agent (i, j) from lottery $\ell_j \in \mathcal{L}_{ij}$ is the same as utility of agent i from lottery $\ell \in \mathcal{L}_i$. Therefore with $u_{ij}(\mathbf{v}_{-ij}) = 0$, in $M^{\mathcal{L}}$ agent i purchases lottery $\ell \in \mathcal{L}_i$ if and only if, in $\mathcal{A}^{\mathcal{L}}$ the pseudo-agent (i, j) purchases lottery $\ell_j \in \mathcal{L}_{ij}$. Moreover, since the price shifts $u_{ij}(\mathbf{v}_{-ij})$ we apply are the same for every lottery offered to (i, j) , the only manner in which preferences can change is if the pseudo-agent obtains negative utility from his preferred lottery, in which case he chooses to buy no lottery at all. However, our choice of $u_{ij}(\mathbf{v}_{-ij})$ ensures that the agent obtains non-negative utility at $v_{ij} = 0$ and thus also at arbitrary v_{ij} , and so the allocation rule of $\mathcal{A}^{\mathcal{L}}$ is identical to that of $M^{\mathcal{L}}$.
3. (feasibility) Feasibility follows immediately from the fact that $M^{\mathcal{L}}$ satisfies feasibility and the allocation rules of the two mechanisms are identical.
4. (nonnegative revenue) Our choice of $u_{ij}(\mathbf{v}_{-ij})$ ensures that the revenue $\mathcal{A}^{\mathcal{L}}$ receives from each agent is always nonnegative; this is critical in later arguments, since it allows us to claim that the revenue that $\mathcal{A}^{\mathcal{L}}$ obtains from any subset of the pseudo-agents is bounded from above by the total expected revenue of $\mathcal{A}^{\mathcal{L}}$.

We now relate the revenues of $M^{\mathcal{L}}$ and $\mathcal{A}^{\mathcal{L}}$. Let a be any function carrying valuation vectors to sets of pseudo-agents which respects the unit-demand constraint, i.e. for any valuation vector \mathbf{v} , for each $i \in I$ there exists at most one $j \in J$ such that $(i, j) \in a(\mathbf{v})$. We call such a function a unit-demand allocation function. Then we get the following lemma.

Lemma 5 *For any valuation vector \mathbf{v} and any unit-demand allocation function $a(\mathbf{v})$, we have*

$$\begin{aligned} \mathcal{R}^{M^{\mathcal{L}}}(\mathbf{v}) &\leq \sum_{(i,j) \in a(\mathbf{v})} \mathcal{R}_{ij}^{\mathcal{A}^{\mathcal{L}}}(\mathbf{v}) + \sum_{(i,j) \notin a(\mathbf{v})} q_{ij}(\mathbf{v}) v_{ij} \\ &\leq \mathcal{R}^{\mathcal{A}^{\mathcal{L}}}(\mathbf{v}) + \sum_{(i,j) \notin a(\mathbf{v})} q_{ij}(\mathbf{v}) v_{ij}, \end{aligned}$$

where $\ell_i(\mathbf{v}) = (q_{i1}(\mathbf{v}), \dots, q_{im}(\mathbf{v}), p_i(\mathbf{v}))$ is the lottery purchased by agent i at valuation \mathbf{v} in the mechanism $M^{\mathcal{L}}$.

Proof: The revenue $\mathcal{R}^{M^\mathcal{L}}(\mathbf{v})$ of the lottery-based mechanism $M^\mathcal{L}$ at \mathbf{v} can be written as the sum of the revenues from the constituent lottery pricings:

$$\mathcal{R}^{M^\mathcal{L}}(\mathbf{v}) = \sum_{i=1}^n \mathcal{R}_i^{M^\mathcal{L}}(\mathbf{v}).$$

If we define $\ell_i(\mathbf{v}) = (q_{i1}(\mathbf{v}), \dots, q_{im}(\mathbf{v}), p_i(\mathbf{v})) \in \mathcal{L}_i$ to be the lottery chosen by agent i at \mathbf{v} , then $\mathcal{R}_i^{M^\mathcal{L}}(\mathbf{v})$, which is just the price $p_i(\mathbf{v})$, can be written as

$$\begin{aligned} \mathcal{R}_i^{M^\mathcal{L}}(\mathbf{v}) &= \left(p_i(\mathbf{v}) - \sum_{k \neq j} q_{ik}(\mathbf{v})v_{ik} \right) + \sum_{k \neq j} q_{ik}(\mathbf{v})v_{ik} \\ &\leq \mathcal{R}_{ij}^{\mathcal{A}^\mathcal{L}}(\mathbf{v}) + \sum_{k \neq j} q_{ik}(\mathbf{v})v_{ik}, \end{aligned} \tag{1}$$

for any j , where $\mathcal{R}_{ij}^{\mathcal{A}^\mathcal{L}}(\mathbf{v})$ is the revenue of mechanism $\mathcal{A}^\mathcal{L}$ from the pseudo-agent (i, j) . Furthermore, since agent i would never elect to purchase a lottery yielding negative utility, we also have that

$$\mathcal{R}_i^{M^\mathcal{L}}(\mathbf{v}) \leq \sum_k q_{ik}(\mathbf{v})v_{ik}. \tag{2}$$

Note that we designed $\mathcal{A}^\mathcal{L}$ such it receives nonnegative revenue from every pseudo-agent, and a contains at most one pseudo-agent (i, j) for any i ; so by applying one of (1) or (2) for each i according to which pseudo-agents $a(\mathbf{v})$ contains, we get that

$$\begin{aligned} \mathcal{R}^{M^\mathcal{L}}(\mathbf{v}) &\leq \sum_{(i,j) \in a(\mathbf{v})} \mathcal{R}_{ij}^{\mathcal{A}^\mathcal{L}}(\mathbf{v}) + \sum_{(i,j) \notin a(\mathbf{v})} q_{ij}(\mathbf{v})v_{ij} \\ &\leq \mathcal{R}^{\mathcal{A}^\mathcal{L}}(\mathbf{v}) + \sum_{(i,j) \notin a(\mathbf{v})} q_{ij}(\mathbf{v})v_{ij}, \end{aligned}$$

the claimed bound. ■

4 Single-agent setting

In this section we focus on instances of the BMUMD involving a single agent and m items. In the single agent setting, randomized and deterministic mechanisms become simply lotteries and pricings, respectively. Briest et al. [7] demonstrated that when values for different items are arbitrarily correlated, it is possible to construct examples where the ratio between the optimal expected revenues from lotteries and pricings is unbounded. We show that in the absence of such correlation this ratio is small. Specifically, when values are distributed independently, the ratio is no more than 4 (Section 4.1). Moreover, when values have a certain kind of positive correlation (additive values; Setting 2 described in Section 2.1), the ratio is at most 8 (Section 4.2).

4.1 Independent values (Setting 1)

We first analyze Setting 1, that is where the value of the agent for item i , v_i , is independently distributed according to c.d.f. F_i . Given an instance \mathcal{I} of the single agent BMUMD, consider the form of the associated

instance $\mathcal{I}^{\text{copies}}$. Note that while each pseudo-agent desires a different item, the fact that only one item may be sold means they are effectively competing for the same thing, the privilege of being served. Thus, $\mathcal{I}^{\text{copies}}$ can be thought of as being in a single-item auction setting. This observation leads to the following lemma.

Lemma 6 *For any instance \mathcal{I} of the BMUMD in Setting 1, the revenue of the optimal deterministic mechanism is at least one-fourth the revenue of the optimal randomized mechanism.*

Proof: As previously observed, any randomized mechanism in the single-agent setting is precisely a lottery pricing \mathcal{L} . Let the mechanism $\mathcal{A}^{\mathcal{L}}$ be as described in Section 3.3. Applying Lemma 5 with $a(\mathbf{v}) = i^* = \operatorname{argmax}_i(v_i)$ yields

$$\begin{aligned} \mathcal{R}^{\mathcal{L}}(\mathbf{v}) &\leq \mathcal{R}^{\mathcal{A}^{\mathcal{L}}}(\mathbf{v}) + \sum_{i \neq i^*} q_i(\mathbf{v})v_i \\ &\leq \mathcal{R}^{\mathcal{A}^{\mathcal{L}}}(\mathbf{v}) + \max_{i \neq i^*} v_i, \end{aligned}$$

since the $q_i(\mathbf{v})$'s sum to at most one. The key observation is that the second term is precisely the revenue that the Vickrey auction \mathcal{V} would achieve in the instance $\mathcal{I}^{\text{copies}}$ given bids \mathbf{v} ; so we get that in expectation

$$\mathcal{R}^{\mathcal{L}} \leq \mathcal{R}^{\mathcal{A}^{\mathcal{L}}} + \mathcal{R}^{\mathcal{V}},$$

and need only apply Theorem 1 to prove the lemma. ■

4.2 Additive values (Setting 2)

We demonstrate that a result similar to that of the previous section holds even in the presence of certain types of correlation. Consider again the single agent setting; since the agent is unit demand, it makes sense to think of the services being offered as perfect substitutes. A natural form of correlation, then, would be for the agent to have some “base” value for being served (regardless of which service is received), plus an additive value specific to the particular service received.

The setting we consider modifies the single-agent setting by making agent types consist of $(m + 1)$ independently distributed values $\{t_0, t_1, \dots, t_m\}$; now, the agent’s value for item i becomes $v_i = t_i + t_0$.

Let \mathcal{L} be a lottery system over m items in the additive setting described. We have the following lemma.

Lemma 7 *Given an instance \mathcal{I} of the BMUMD in Setting 2, the revenue of any lottery system \mathcal{L} for \mathcal{I} satisfies $\mathcal{R}^{\mathcal{L}} \leq 8\mathcal{R}^{\mathcal{P}}$, for some pricing \mathbf{p} for \mathcal{I} .*

Proof: We begin by demonstrating a bound with a weaker multiplicative factor of 9 and then show how to improve it to a factor of 8. Our main technique is to consider an uncorrelated setting \mathcal{I}' derived from \mathcal{I} . We define \mathcal{I}' to be a single agent setting with $(m + 1)$ items, and interpret the values $\{t_0, \dots, t_m\}$ making up an agent’s type in \mathcal{I} as being the values of the agent in setting \mathcal{I}' for the $(m + 1)$ items. In keeping with \mathcal{I} , the feasibility constraint we associate with \mathcal{I}' is that we may sell item 0, and at most one additional item from among items $1, \dots, m$. Note that the agent in \mathcal{I}' is not a unit-demand agent.

We now construct a lottery system \mathcal{L}' for instance \mathcal{I}' from \mathcal{L} . Let $\ell = (q_1, \dots, q_m, p)$ be a lottery in \mathcal{L} . Define $q_0 = \sum_{i=1}^m q_i$, and construct $\ell' = (q_0, \dots, q_m, p)$. Note that ℓ' does not necessarily satisfy the requirement that the q_i 's sum to at most one; it does, however, satisfy the feasible constraint indicated for \mathcal{I}' . We may thus still apply the same technique as in the proof of Lemma 6, albeit with a worsened constant.

Let \mathcal{L}' be the system over $m + 1$ services consisting of all of the ℓ' defined as above based on $\ell \in \mathcal{L}$. Now, for any setting of t_0, \dots, t_m , note that the utility an agent in \mathcal{I} receives from a particular lottery $\ell \in \mathcal{L}$ is

$$\sum_{i=1}^m q_i v_i - p = \sum_{i=1}^m q_i (t_i + t_0) - p = \sum_{i=0}^m q_i t_i - p,$$

precisely the utility a corresponding agent in \mathcal{I}' would receive from the corresponding $\ell' \in \mathcal{L}'$. We thus have $\mathcal{R}^{\mathcal{L}} = \mathcal{R}^{\mathcal{L}'}$.

Consider applying the proof of Lemma 6 to \mathcal{L}' . Due to the less restrictive feasibility constraint ($\sum_{i=0}^m q_i \leq 2$) we get

$$\begin{aligned} \mathcal{R}^{M^{\mathcal{L}'}} &\leq \mathcal{R}^{\mathcal{A}^{\mathcal{L}'}} + 2\mathcal{R}^{\mathcal{V}'} \\ &\leq 3\mathcal{R}^{\mathcal{M}'}, \end{aligned}$$

where the mechanisms $\mathcal{A}^{\mathcal{L}'}$ and \mathcal{V}' are interpreted as being in the copies setting $\mathcal{I}^{\text{copies}'}$ associated with \mathcal{I}' , and \mathcal{M}' is the optimal mechanism in this setting. In order to prove a bound of the form desired, however, we need to relate a mechanism in the setting $\mathcal{I}^{\text{copies}'}$ to a deterministic one (a pricing) in \mathcal{I} .

The key observation is that our feasibility constraint in $\mathcal{I}^{\text{copies}'}$ (carried over from \mathcal{I}') means that \mathcal{M}' may make decisions about allocations and prices for pseudo-agent 0 separately from those for pseudo-agents $1, \dots, m$; as such, \mathcal{M}' effectively consists of two mechanisms, one serving pseudo-agent 0 and another serving pseudo-agents $1, \dots, m$, both under a unit-demand constraint. Now, the optimal mechanism for serving the lone single-parameter pseudo-agent is a pricing, and Theorem 1 gives us that a mechanism serving pseudo-agents $1, \dots, m$ is within a factor of 2 of a pricing on m items; so recalling that an agent in setting \mathcal{I} has a value of $v_i = t_i + t_0$ for item i , we can see that

$$\mathcal{R}^{\mathcal{L}} \leq 3\mathcal{R}^{\mathcal{M}} \leq 9\mathcal{R}^{\mathbf{p}},$$

where \mathbf{p} is the optimal pricing for the setting \mathcal{I} .

In order to improve the factor from 9 to 8, we need to consider the revenue a mechanism M in the setting $\mathcal{I}^{\text{copies}'}$ obtains from pseudo-agent 0 and from pseudo-agents $1, \dots, m$; at a particular valuation vector \mathbf{t} denote these quantities as $\mathcal{R}_0^M(\mathbf{t})$ and $\mathcal{R}_{-0}^M(\mathbf{t})$, respectively. Now, as previously noted, the optimal mechanism \mathcal{M} in $\mathcal{I}^{\text{copies}'}$ may treat pseudo-agent 0 independently from pseudo-agents $1, \dots, m$; thus, we have that any mechanism M in this setting must satisfy both $\mathcal{R}_0^M(\mathbf{t}) \leq \mathcal{R}_0^{\mathcal{M}}(\mathbf{t})$ and $\mathcal{R}_{-0}^M(\mathbf{t}) \leq \mathcal{R}_{-0}^{\mathcal{M}}(\mathbf{t})$.

Since we know that $\sum_{i=1}^m q_i \leq 1$, when t_0 is the maximum among all the t_i , Lemma 5 implies

$$\mathcal{R}^{\mathcal{L}}(\mathbf{t}) \leq \mathcal{R}_0^{\mathcal{A}^{\mathcal{L}}}(\mathbf{t}) + \mathcal{R}_0^{\mathcal{V}}(\mathbf{t});$$

On the other hand, when one of t_1, \dots, t_m takes on the maximum value, we end up with, for some i ,

$$\mathcal{R}^{\mathcal{L}}(\mathbf{t}) \leq \mathcal{R}_i^{\mathcal{A}^{\mathcal{L}}}(\mathbf{t}) + 2\mathcal{R}_i^{\mathcal{V}}(\mathbf{t}),$$

Combining these two gives us a pointwise guarantee of

$$\begin{aligned} \mathcal{R}^{\mathcal{L}}(\mathbf{t}) &\leq \mathcal{R}_0^{\mathcal{A}^{\mathcal{L}}}(\mathbf{t}) + \mathcal{R}_0^{\mathcal{V}}(\mathbf{t}) + \mathcal{R}_{-0}^{\mathcal{A}^{\mathcal{L}}}(\mathbf{t}) + 2\mathcal{R}_{-0}^{\mathcal{V}}(\mathbf{t}) \\ &\leq 2\mathcal{R}_0^{\mathcal{M}}(\mathbf{t}) + 3\mathcal{R}_{-0}^{\mathcal{M}}(\mathbf{t}). \end{aligned}$$

Therefore,

$$\mathcal{R}^{\mathcal{L}} \leq 2\mathcal{R}_0^{\mathcal{M}} + 3\mathcal{R}_{-0}^{\mathcal{M}} \leq 2\mathcal{R}^{\mathbf{p}} + 6\mathcal{R}^{\mathbf{p}}$$

implying the claimed bound of 8. ■

5 Multi-agent setting

In this section we study multi-agent versions of the BMUMD and once again bound the gap between deterministic and randomized mechanisms with respect to expected revenue for this setting. The starting point for our bounds is the observation in Section 3.2 that randomized mechanisms for this problem can be interpreted as lottery-based mechanisms. We first discuss Setting 3, namely instances with multiple agents and multiple items and a “matching” feasibility constraint. The following subsection contains a more general version with a matroid intersection feasibility constraint (Setting 4).

5.1 The multi-item auction setting (Setting 3)

We consider instances of the BMUMD where the seller has m different items, with k_j copies of item j , and each of the n unit-demand buyers have independently distributed values for each item. The seller’s constraint is to allocate item j to no more than k_j agents, and to allocate at most one item to each agent.

We note that the set system defined by this feasibility constraint is a matroid intersection. A set system (E, \mathcal{F}) where E is the ground set of elements ($E = I \times J$ in our setting) is a matroid if it satisfies the following properties.

1. **(heredity)** For every $A \in \mathcal{F}$, $B \subset A$ implies $B \in \mathcal{F}$.
2. **(augmentation)** For every $A, B \in \mathcal{F}$ with $|A| > |B|$, there exists an $e \in A \setminus B$ such that $B \cup \{e\} \in \mathcal{F}$.

The sets in a matroid set system are called independent sets.

A matroid intersection set system \mathcal{F} is an intersection of two matroids: $\mathcal{F} = \mathcal{F}_1 \cap \mathcal{F}_2$ where \mathcal{F}_1 and \mathcal{F}_2 are matroids. The unit-demand constraint and the supply constraints for each item are each instances of a partition matroid. Thus the system \mathcal{J} in this setting can be seen to be an intersection of two partition matroids. We use \mathcal{J}_1 and \mathcal{J}_2 to denote the two constituent matroids, and the term matching to refer to any allocation or set in \mathcal{J} .

We will need the following facts about matroids.

Proposition 8 *Let B_1 and B_2 be any two independent sets of equal size in some matroid set system \mathcal{E} . Then there is a bijective function $g : B_1 \setminus B_2 \rightarrow B_2 \setminus B_1$ such that for all $e \in B_1 \setminus B_2$, $B_1 \setminus \{e\} \cup \{g(e)\}$ is independent in \mathcal{E} .*

Corollary 9 *Let B_1 and B_2 be arbitrary independent sets in some matroid set system \mathcal{E} . Then there exists a set $B_2' \subseteq B_2$ and a one to one function $g : B_2' \rightarrow B_1$ such that for all $e \in B_2'$, $B_1 \setminus \{g(e)\} \cup \{e\}$ is independent in \mathcal{E} , and for all $e \in B_2 \setminus B_2'$, $B_1 \cup \{e\}$ is independent in \mathcal{E} .*

Proof: In order to apply Proposition 8 we need independent sets of equal size. So we begin by repeatedly applying the augmentation property to whichever of B_1 and B_2 is smaller in order to end up with two sets $\bar{B}_1 \supset B_1$ and $\bar{B}_2 \supset B_2$ such that $|\bar{B}_1| = |\bar{B}_2|$. Now, Proposition 8 guarantees us a bijection $g : \bar{B}_2 \setminus \bar{B}_1 \rightarrow \bar{B}_1 \setminus \bar{B}_2$ such that $\forall e \in \bar{B}_2 \setminus \bar{B}_1$, $\bar{B}_1 \setminus \{g(e)\} \cup \{e\}$ is independent.

Set $B_2' = B_2 \setminus \bar{B}_1 \subset \bar{B}_2 \setminus \bar{B}_1$; note that since $\bar{B}_1 \subset B_1 \cup B_2$, we have $\bar{B}_1 \setminus \bar{B}_2 \subset B_1$. Thus, we may view g as a one to one function $g : B_2' \rightarrow B_1$. It retains the first specified property, since for any $e \in B_2'$, $B_1 \setminus \{g(e)\} \cup \{e\} \subset \bar{B}_1 \setminus \{g(e)\} \cup \{e\}$ is independent. Furthermore, $e \in B_2 \setminus B_2' \subset \bar{B}_1$ implies $B_1 \cup \{e\} \subset \bar{B}_1$ is independent, and so the second specified property holds as well. ■

Our proof consists of three steps:

1. From Lemma 4, we note that any randomized mechanism for this problem can be seen as a lottery-based mechanism.
2. We bound the revenue of any lottery-based mechanism for an instance \mathcal{I} of the BMUMD by those of a collection of three truthful deterministic mechanisms for the corresponding single-parameter instance with copies, $\mathcal{I}^{\text{copies}}$.
3. We use the result in [9] (Theorem 2) to construct a truthful deterministic mechanism for \mathcal{I} whose revenue is within a factor of $4/27$ of the optimal revenue for $\mathcal{I}^{\text{copies}}$.

Lemma 10 *Consider an instance \mathcal{I} of the BMUMD in Setting 3. The revenue from any lottery-based mechanism $M^{\mathcal{L}}$ for \mathcal{I} is at most five times the expected revenue of Myerson's mechanism for the instance $\mathcal{I}^{\text{copies}}$.*

Proof: We define three truthful deterministic mechanisms M_1, M_2, M_3 for $\mathcal{I}^{\text{copies}}$, all facing the same feasibility constraint \mathcal{J} as the set of lottery pricings \mathcal{L} , such that

$$\begin{aligned} \mathcal{R}^{M^{\mathcal{L}}}(\mathbf{v}) &\leq \mathcal{R}^{M_1}(\mathbf{v}) + 2(\mathcal{R}^{M_2}(\mathbf{v}) + \mathcal{R}^{M_3}(\mathbf{v})) \\ &\leq 5\mathcal{R}^{\mathcal{M}}(\mathbf{v}), \end{aligned} \tag{3}$$

The second inequality follows from the optimality of Myerson's mechanism for single parameter settings (Myerson's mechanism also faces the feasibility constraint \mathcal{J}).

Consider the $\mathcal{I}^{\text{copies}}$ setting and fix an instantiation of values \mathbf{v} . Let $A_1(\mathbf{v})$ denote the set of pseudo-agents that belong to the maximum-valued matching (we drop the argument wherever it is obvious). Among the remaining $(I \times J) \setminus A_1$ pseudo-agents, again let A_2 denote the set of pseudo-agents that belong to the maximum-valued matching i.e.

$$A_2(\mathbf{v}) = \underset{\substack{S \subseteq [mn], S \cap A_1(\mathbf{v}) = \emptyset \\ S \text{ is a matching}}}{\text{argmax}} v(S).$$

We may assume without loss of generality that A_1 and A_2 are defined uniquely.

Note that $A_1(\mathbf{v})$ is a unit-demand allocation function. Therefore, Lemma 5 implies that

$$\mathcal{R}^{M^{\mathcal{L}}}(\mathbf{v}) \leq \underbrace{\mathcal{R}^{\mathcal{A}^{\mathcal{L}}}(\mathbf{v})}_{\text{Term}_1} + \underbrace{\sum_{(i,j) \notin A_1(\mathbf{v})} q_{ij}(\mathbf{v})v_{ij}}_{\text{Term}_2}. \tag{4}$$

We now define the three mechanisms M_1, M_2 and M_3 for $\mathcal{I}^{\text{copies}}$. Mechanism M_1 is $\mathcal{A}^{\mathcal{L}}$ and so \mathcal{R}^{M_1} is exactly Term₁. Mechanisms M_2 and M_3 are defined in such a way that $2(\mathcal{R}^{M_2} + \mathcal{R}^{M_3})$ is at least Term₂. This would prove (3).

Now, Corollary 9 implies the existence of two one to one partial functions with the following properties.

$$\begin{array}{ll} g_1 : A_2 \rightarrow A_1 & \text{s.t. } \forall e \in A_2 : \\ & g_1(e) \text{ is undefined and } A_1 \cup \{e\} \in \mathcal{J}_1, \text{ or} \\ & g_1(e) \text{ is defined and } A_1 \setminus \{g_1(e)\} \cup \{e\} \in \mathcal{J}_1 \\ g_2 : A_2 \rightarrow A_1 & \text{s.t. } \forall e \in A_2 : \\ & g_2(e) \text{ is undefined and } A_1 \cup \{e\} \in \mathcal{J}_2, \text{ or} \\ & g_2(e) \text{ is defined and } A_1 \setminus \{g_2(e)\} \cup \{e\} \in \mathcal{J}_2 \end{array}$$

Note that the maximality of A_1 implies that every element of A_2 has an image under either g_1 or g_2 or both. We define the mechanisms M_2 and M_3 by specifying their allocation rules. Given a valuation vector \mathbf{v} , the mechanism M_2 serves only those pseudo-agents (i, j) that belong to A_1 and for which $v_{ij} \geq v_{g_1^{-1}(i,j)}/2$ (if g_1^{-1} is defined at that point). Likewise, mechanism M_3 serves only those pseudo-agents $(i, j) \in A_1$ that have $v_{ij} \geq v_{g_2^{-1}(i,j)}/2$ (if defined). We note that M_2 and M_3 have monotone allocation rules, and are therefore truthful. Truthful payments can be defined appropriately. They also satisfy the feasibility constraint \mathcal{J} .

We now prove the revenue guarantee for M_2 and M_3 through the following two claims.

Claim 1 *Twice the combined revenue of mechanisms M_2 and M_3 is no less than the sum of values of all pseudo-agents in A_2 , i.e.,*

$$2(\mathcal{R}^{M_2}(\mathbf{v}) + \mathcal{R}^{M_3}(\mathbf{v})) \geq \sum_{(i,j) \in A_2} v_{ij}.$$

Proof: Consider any pseudo-agent $(i, j) \in A_2$, and the pseudo-agents $g_1(i, j)$ and $g_2(i, j) \in A_1$ if defined. Note that $A'_1 = A_1 \cup (i, j) \setminus \{g_1(i, j), g_2(i, j)\}$ is feasible. Suppose both $v_{g_1(i,j)}$ and $v_{g_2(i,j)}$ are less than $v_{ij}/2$; then the matching A'_1 is a valid matching and $v(A'_1) > v(A_1)$ which is a contradiction to the optimality of A_1 . Thus one of $v_{g_1(i,j)}$ or $v_{g_2(i,j)}$ must be at least $v_{ij}/2$ and we get this amount in M_2 or M_3 respectively. ■

Claim 2 *The sum of values of all pseudo-agents in A_2 is no less than Term₂:*

$$\sum_{(i,j) \in A_2(\mathbf{v})} v_{ij} \geq \sum_{(i,j) \notin A_1(\mathbf{v})} q_{ij}(\mathbf{v})v_{ij}.$$

Proof: Consider the $n \times m$ matrix of all probabilities $q_{ij}(\mathbf{v})$. This matrix arose from a feasible randomized mechanism; it therefore represents a probability distribution over matchings and can be represented as a convex combination of matchings. In this probability matrix, replace with zeros all the entries $(i, j) \in A_1$. The newly obtained matrix can be represented as a convex combination of matchings all of which have a zero entry for every $(i, j) \in A_1$. Then the claim follows by the definition of A_2 . ■

Claims 1 and 2 together with Equations (3) and (4) complete the proof. ■

Theorem 11 *The revenue of any randomized mechanism for an instance of the BMUMD in Setting 3 is at most 33.75 times the revenue of the optimal truthful deterministic mechanism for the instance.*

Proof: The proof follows from Lemmas 4 and 10, and Theorem 2. ■

5.2 The general matroid setting (Setting 4)

We now show that Theorem 11 extends to the general matroid intersection version of the BMUMD as well. While Lemma 10 extends to this more general setting almost exactly, the counterpart of Theorem 2 for this setting is somewhat weaker. So we can only bound the gap between the revenue of an optimal randomized incentive-compatible mechanism and that of an optimal deterministic implementation in undominated strategies (see Theorem 3) for this setting.

As defined earlier, in Setting 4, the seller faces a feasibility constraint specified by the set system $\mathcal{J} \subseteq 2^{I \times J}$, where I is the set of agents and J is the set of services, \mathcal{J} is the intersection of a general matroid

constraint (given by \mathcal{J}_1) and the unit demand constraint (that we denote using \mathcal{J}_2); $\mathcal{J} = \mathcal{J}_1 \cap \mathcal{J}_2$. Note that \mathcal{J}_2 is also a matroid.

We use the same three step approach as for the matching version to bound the revenue of the randomized mechanism

Lemma 12 *Consider an instance \mathcal{I} of the BMUMD in Setting 4. The revenue from any lottery-based mechanism $M^{\mathcal{L}}$ for instance \mathcal{I} is at most five times the expected revenue of Myerson's mechanism for the single parameter instance with copies $\mathcal{I}^{\text{copies}}$.*

Proof: We will prove this Lemma along the lines of our proof for Lemma 10. We define three truthful deterministic mechanisms M_1, M_2, M_3 for $\mathcal{I}^{\text{copies}}$ so that

$$\begin{aligned} \mathcal{R}^{M^{\mathcal{L}}}(\mathbf{v}) &\leq \mathcal{R}^{M_1}(\mathbf{v}) + 2(\mathcal{R}^{M_2}(\mathbf{v}) + \mathcal{R}^{M_3}(\mathbf{v})) \\ &\leq 5\mathcal{R}^{\mathcal{M}}(\mathbf{v}). \end{aligned} \tag{5}$$

As before, given an instantiation of values \mathbf{v} , let $A_1(\mathbf{v})$ denote the set of pseudo-agents that belong to the maximum valued feasible set. Among the remaining pseudo-agents, let $A_2(\mathbf{v})$ denote the set of pseudo-agents that belong to the maximum valued feasible set i.e.

$$A_2(\mathbf{v}) = \operatorname{argmax}_{S \in \mathcal{J} - A_1(\mathbf{v})} v(S)$$

Lemma 5 implies

$$\mathcal{R}^{M^{\mathcal{L}}}(\mathbf{v}) \leq \underbrace{\mathcal{R}^{\mathcal{A}^{\mathcal{L}}}(\mathbf{v})}_{\text{Term}_1} + \underbrace{\sum_{(i,j) \notin A_1(\mathbf{v})} q_{ij}(\mathbf{v})v_{ij}}_{\text{Term}_2}.$$

Therefore, once again we define M_1 to be $\mathcal{A}^{\mathcal{L}}$ and define M_2 and M_3 in such a way that twice their revenue combined is no less than Term₂.

As before we can define partial one to one functions from A_2 to A_1 satisfying

$$\begin{aligned} g_1 : A_2 &\rightarrow A_1 && \text{s.t. } \forall e \in A_2 : \\ &&& g_1(e) \text{ is undefined and } A_1 \cup \{e\} \in \mathcal{J}_1, \text{ or} \\ &&& g_1(e) \text{ is defined and } A_1 \setminus \{g_1(e)\} \cup \{e\} \in \mathcal{J}_1 \\ g_2 : A_2 &\rightarrow A_1 && \text{s.t. } \forall e \in A_2 : \\ &&& g_2(e) \text{ is undefined and } A_1 \cup \{e\} \in \mathcal{J}_2, \text{ or} \\ &&& g_2(e) \text{ is defined and } A_1 \setminus \{g_2(e)\} \cup \{e\} \in \mathcal{J}_2 \end{aligned}$$

The mechanisms M_2 and M_3 are also defined as before: M_2 serves only those pseudo-agents (i, j) in A_1 for which $v_{ij} \geq v_{g_1^{-1}(i,j)}/2$ (if defined), and M_3 serves only those pseudo-agents $(i, j) \in A_1$ that have $v_{ij} \geq v_{g_2^{-1}(i,j)}/2$ (if defined). We note that every element in A_2 gets mapped to at least one and at most two elements under the partial functions defined above. Therefore, we can extract a revenue of at least $1/2 \sum_{(i,j) \in A_2} v_{ij}$ from M_2 and M_3 together. Claim 2 now implies the result. \blacksquare

Theorem 13 *The revenue of any incentive compatible randomized mechanism for an instance \mathcal{I} of the BMUMD in Setting 4 is at most 40 times the revenue of the optimal deterministic mechanism for \mathcal{I} implemented in undominated strategies.*

Proof: The proof follows from Lemmas 4 and 12, and Theorem 3. \blacksquare

6 Discussion and open problems

We show that in multi-parameter Bayesian optimal mechanism design the benefit of randomness is only a small constant factor when agents are unit-demand and their values for different items have little or no correlation. We believe that this result should extend to instances involving arbitrary positive correlation between values of a single agent for items that are substitutes (the unit-demand constraint). For example, it would be interesting to extend our result to the multiplicative values model of Armstrong [2]. Another open problem is to extend our techniques beyond the unit-demand setting. This may lead to a better understanding of and approximations to optimal mechanism design in those settings, for which nothing is known as yet.

Acknowledgements

We thank Jason Hartline and Robert Kleinberg for many helpful discussions.

References

- [1] G. Aggarwal, A. Fiat, A. Goldberg, J. Hartline, N. Immorlica, and M. Sudan. Derandomization of Auctions. In *Proc. 37th ACM Symposium on the Theory of Computing*. ACM Press, New York, 2005.
- [2] Mark Armstrong. Price discrimination by a many-product firm. *Review of Economic Studies*, 66(1):151–68, January 1999.
- [3] Moshe Babaioff, Ron Lavi, and Elan Pavlov. Single-value combinatorial auctions and algorithmic implementation in undominated strategies. *J. ACM*, 56(1):1–32, 2009.
- [4] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Item pricing for revenue maximization. In *EC '08: Proceedings of the 9th ACM conference on Electronic commerce*, pages 50–59, New York, NY, USA, 2008. ACM.
- [5] N. Balcan and A. Blum. Approximation algorithms and online mechanisms for item pricing. In *Proc. 8th ACM Conf. on Electronic Commerce*, 2006.
- [6] P. Briest. Towards hardness of envy-free pricing. Technical Report TR06-150, ECCC, 2006.
- [7] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. Pricing randomized allocations. In *Proc. 21st ACM Symp. on Discrete Algorithms*, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.
- [8] S. Chawla, J. Hartline, and R. Kleinberg. Algorithmic pricing via virtual valuations. In *Proc. 9th ACM Conf. on Electronic Commerce*, pages 243–251, 2007.
- [9] Shuchi Chawla, Jason D. Hartline, David Malec, and Balasubramanian Sivan. Sequential posted pricing and multi-parameter mechanism design. *CoRR*, abs/0907.2435, 2009.
- [10] V. Guruswami, J. Hartline, A. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *Proc. 16th ACM Symp. on Discrete Algorithms*, 2005.
- [11] Jason Hartline and Anna Karlin. Profit maximization in mechanism design. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, chapter 13, pages 331–362. Cambridge Press, 2007.

- [12] Alejandro M. Manelli and Daniel R. Vincent. Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory*, 127(1):1–35, March 2006.
- [13] R.P. McAfee and J. McMillan. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory*, 46(2):335–354, 1988.
- [14] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [15] Gregory Pavlov. Optimal mechanism for selling substitutes. Boston University - Department of Economics - Working Papers Series WP2006-014, Boston University - Department of Economics, February 2006.
- [16] John Riley and Richard Zeckhauser. Optimal selling strategies: When to haggle, when to hold firm. *The Quarterly Journal of Economics*, 98(2):267–289, 1983.
- [17] Jean-Charles Rochet and Philippe Chone. Ironing, Sweeping, and Multidimensional Screening. *Econometrica*, 66(4):783–826, 1998.
- [18] John E. Thanassoulis. Haggling Over Substitutes. *Journal of Economic Theory*, 117(2):217–245, 2004.
- [19] D. Vincent and A. Manelli. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory*, 137(1):153–185, 2007.

Gap between lottery pricings and Myerson’s mechanism

We give an example where the revenue of a lottery pricing for a single agent BMUMD instance \mathcal{I} is 1.13 times the revenue of Myerson’s mechanism for the instance $\mathcal{I}^{\text{copies}}$. The instance \mathcal{I} is defined as follows. There is a single agent with i. i. d. valuations for two items, distributed according to the *equal-revenue* distribution, bounded at n . Formally, the valuations v_1 and v_2 for items 1 and 2 have cdfs F_1 and F_2 such that

$$F_1(x) = F_2(x) = \begin{cases} 1 - 1/x & 1 \leq x < n \\ 1 & x = n \end{cases}.$$

For the single parameter setting $\mathcal{I}^{\text{copies}}$, an upper bound on the expected revenue of any mechanism can be obtained by removing the feasibility constraint of allocating to a single agent at a time. Then, the optimal revenue with the feasibility constraint is no more than twice the optimal revenue that can be obtained by a single agent alone. The latter, for the equal revenue distribution, is 1 regardless of the price charged to the agent. Therefore, the optimal revenue for $\mathcal{I}^{\text{copies}}$ is bounded above by 2. The same bound also applies to the revenue of any item pricing for \mathcal{I} .

Now let us consider the following lottery pricing \mathcal{L} for \mathcal{I} .

$$\mathcal{L} = \left\{ (0.5, 0.5, 2.5), \left(1, 0, 2 + \frac{3n}{8}\right), \left(0, 1, 2 + \frac{3n}{8}\right) \right\}$$

The first two coordinates in every lottery denote the probabilities with which items 1 and 2 are offered by that lottery and the third coordinate is the price.

Figure 3 shows the allocation function of this lottery pricing. In particular, R_i for $i \in [3]$ is the set of valuations where lottery i is bought. The probability mass of regions R_2 and R_3 together can be computed to be $2(4/3n + O(\log n/n^2))$. The probability mass of region R_1 is $0.4 + 0.08 \ln 4 - o(1) \approx 0.51$. Therefore,

the revenue of \mathcal{L} can be computed to be $5/2 \cdot 0.51 + 3n/8 \cdot 8/3n + o(1) = 2.275 + o(1)$. This is a factor of 1.13 higher than the optimal revenue for $\mathcal{I}^{\text{copies}}$, or the revenue of any item pricing for \mathcal{I} .




Figure 3: The allocation function for the lottery pricing \mathcal{L} .