

Chawla, Shuchi; Hartline, Jason; Malec, David; Sivan, Balasubramanian

**Working Paper**

## Multi-parameter mechanism design and sequential posted pricing

Discussion Paper, No. 1486

**Provided in Cooperation with:**

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

*Suggested Citation:* Chawla, Shuchi; Hartline, Jason; Malec, David; Sivan, Balasubramanian (2010) : Multi-parameter mechanism design and sequential posted pricing, Discussion Paper, No. 1486, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/59680>

**Standard-Nutzungsbedingungen:**

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

**Terms of use:**

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*

# Multi-parameter Mechanism Design and Sequential Posted Pricing

Shuchi Chawla\*    Jason D. Hartline†    David Malec‡    Balasubramanian Sivan§

## Abstract

We consider the classical mathematical economics problem of *Bayesian optimal mechanism design* where a principal aims to optimize expected revenue when allocating resources to self-interested agents with preferences drawn from a known distribution. In single-parameter settings (i.e., where each agent’s preference is given by a single private value for being served and zero for not being served) this problem is solved [19]. Unfortunately, these single parameter optimal mechanisms are impractical and rarely employed [1], and furthermore the underlying economic theory fails to generalize to the important, relevant, and unsolved multi-dimensional setting (i.e., where each agent’s preference is given by multiple values for each of the multiple services available) [24].

In contrast to the theory of optimal mechanisms we develop a theory of sequential posted price mechanisms, where agents in sequence are offered take-it-or-leave-it prices. We prove that these mechanisms are approximately optimal in single-dimensional settings. These posted-price mechanisms avoid many of the properties of optimal mechanisms that make the latter impractical. Furthermore, these mechanisms generalize naturally to multi-dimensional settings where they give the first known approximations to the elusive optimal multi-dimensional mechanism design problem. In particular, we solve multi-dimensional multi-unit auction problems and generalizations to matroid feasibility constraints. The constant approximations we obtain range from 1.5 to 8. For all but one case, our posted price sequences can be computed in polynomial time.

This work can be viewed as an extension and improvement of the single-agent algorithmic pricing work of [10] to the setting of multiple agents where the designer has combinatorial feasibility constraints on which agents can simultaneously obtain each service.

---

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

†EECS, Northwestern University. hartline@eecs.northwestern.edu.

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

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

# 1 Introduction

Suppose the local organizers for a prominent symposium on computer science need to arrange for suitable hotel accommodations in the Boston area for the attendees of the conference. There are a number of hotel rooms available with different features and attendees have preferences over the rooms. The organizers need a mechanism for soliciting preferences, assigning rooms, and calculating payments. Fortunately, they have distributional knowledge over the participants' preferences (e.g., from similar conferences). This is a stereotypical multi-dimensional setting for mechanism design that, for instance, also arises in most resource allocation problems in the Internet. What mechanism should the organizers employ to maximize their objective (e.g., revenue)?

The economic theory of optimal mechanism design is elegant and predictive in single-dimensional settings. Here Myerson's theory of virtual valuations and characterizations of incentive constraints via monotonicity guide the design of optimal truthful mechanisms [19] with practical (often non-truthful) implementations [1]. The challenge of multi-dimensional settings (e.g., in the likely case that conference attendees, i.e., agents, have different values for different hotel rooms) is two-fold. First, multi-dimensional settings are unlikely to permit succinct descriptions of optimal mechanisms [18, 21, 24]. Second, optimal mechanisms in multi-dimensional settings are unlikely to have practical implementations – even asking agents to report their true types across the many possible outcomes of the mechanism may be impractical. In summary, theory and practical considerations from optimal mechanism design in single-dimensional settings fail to generalize to multi-dimensional settings.

This paper approaches these issues through the lens of approximation. Our main results are simple, practical, approximately optimal mechanisms for a large class of multi-dimensional settings. We consider the multi-dimensional setting through a single dimensional analogy wherein each multi-dimensional agent is represented by many independent single-dimensional agents (e.g., one for each hotel room). The optimal revenue for this single-dimensional setting is well understood and, due to increased competition among agents, upper-bounds that of the original multi-dimensional setting. We describe a “sequential posted price” mechanism for the single-dimensional setting that is practical and approximately optimal and, in contrast to the optimal single-dimensional mechanism, achieves its approximation without inter-agent competition. This gives a robustness to deviations in modeling assumptions and, for instance, the same mechanism continues to be approximately optimal in the original multi-dimensional setting. Therefore, our theory for approximately optimal single-dimensional mechanisms generalizes to multi-dimensional settings.

In the context of computer science literature this work is an extension of *algorithmic pricing* (e.g., [12]) to settings with multiple agents; it is unrelated to the standard computational questions of *algorithmic mechanism design* (e.g., [17, 20]). The central problem in algorithmic pricing can be viewed (for the most part) as Bayesian revenue maximization in a single agent setting (e.g., [12]). Algorithmic pricing is hard to approximate when the agent's values for different outcomes are generally correlated [8]; however, when the values are independent there is a 3-approximation [10]. In this context, our results improve and extend the independent case to settings with multiple agents and combinatorial feasibility constraints. Notice that the challenge in these problems is one imposed by the multi-dimensional incentive constraints and not one from an inherent complexity of an underlying non-game-theoretic optimization problem. (E.g., in the hotel example the underlying optimization problem is simply maximum weighted matching.) In contrast, most work in algorithmic mechanism design addresses settings where economic incentives are well understood but the underlying optimization problem is computationally intractable (e.g., combinatorial auctions [17]).

While our exposition focuses on revenue maximization, all of our techniques and results apply equally well to *social welfare*. Social welfare is unique among objectives in that designing optimal mechanisms in multi-dimensional settings is solved (by the VCG mechanism). Therefore, the interesting implication of our work on social welfare maximization is that sequential posted pricing approximates the welfare of the VCG mechanism and may be more practical.

**Sequential Posted Pricing.** Consider a single-parameter setting where each agent has a private value for service and there is a combinatorial feasibility constraint on the set of agents that can be simultaneously served. For this setting a *sequential posted pricing* (SPM) is a mechanism defined by a price for each agent, a sequence on agents, and the semantics that each agent is offered their corresponding price in sequence as a take-it-or-leave-it while-supplies-last offer. Meaning: if it is possible to serve the agent given the set of agents already being served then the agent is offered the price. A rational agent will accept if and only if the price is no more than their private value for service. That prices are associated with the agents and not the sequence reflects the possibility that agents may play asymmetric roles for a given feasibility constraint or value distribution.

Consider the following hotel rooms example with one room, two attendees, and attendee values independently and identically distributed uniformly between \$100 and \$200. The optimal mechanism is the Vickrey auction and its expected revenue is \$133. The optimal sequential posted pricing is for the organizers to offer the room to attendee 1 at a price of \$150. If the attendee accepts, then the room is sold, otherwise it is offered to attendee 2 for \$100. The expected revenue of this SPM is \$125.

We are interested in comparing the optimal mechanism to the optimal posted pricing in general settings. A special class of SPMs is one where mechanisms have provable performance guarantees for any sequence of the agents. These *order-oblivious posted pricings* (OPM) are mechanisms defined by a price for each agent and the semantics that each agent is offered their corresponding price in some arbitrary sequence as a take-it-or-leave-it while-supplies-last offer.

In single-dimensional settings, the advantages of sequential posted pricings speak to the many reasons optimal auctions are rarely seen in practice [1], and explain why posted pricings are ubiquitous [14]. First, take-it-or-leave-it offers result in trivial game dynamics: truthful responding is a dominant strategy. Second, SPMs satisfy strong notions of collusion resistance, e.g., *group strategyproofness* (see [11]): the only way in which an agent can “help” another agent is to decline an offer that he could have accepted, thereby hurting his own utility. Third, agents do not need to precisely know or report their value, they must only be able to evaluate their offer; therefore, they risk minimal exposure of their private information. Fourth, agents learn immediately whether they will be served or not. In conclusion, the robustness of SPMs in single-dimensional settings makes their approximation of optimal mechanisms independently worthy of study.

The final robustness property of SPMs, which is of paramount importance to our study of the multi-dimensional setting, is that they minimize the role of agent competition which implies that single-dimensional SPMs can be used “as-is” in multi-dimensional settings with only a constant factor loss in performance. In our translation from the multi-dimensional setting to the single-dimensional setting, each multi-dimensional agent has many single-dimensional representatives. A good OPM for the single-dimensional setting can be viewed as an OPM for the multi-dimensional setting by grouping all representatives of an agent together and making their offers simultaneously to the agent. The agent will of course accept the offer that maximizes their utility. The resulting mechanism is truthful and achieves the same performance guarantee as the single-parameter OPM. For SPMs where we are not free to group each multi-dimensional agent’s single-dimensional representative together, an agent possibly faces a strategic dilemma of whether to accept an offer (e.g., for one hotel room) early on or wait for a later offer (e.g., another hotel room) which may or may not still be available. Our guarantee is that regardless of the actions of any agent with such a strategic option (i.e., *implementation in undominated strategies*, see, e.g., [4]) our performance is a constant fraction of the original SPM’s performance. Given the advantages of SPMs over truthful mechanisms, such a non-truthful SPM may be more practically relevant than a truthful implementation.

Finally, we note that most of our results for posted pricings are constructive and give efficient algorithms for them. A posted price mechanism has two components where computation is necessary: an offline computation of the prices to post (and for SPMs, the sequence of agents) and an online while-supplies-last offering of said prices.<sup>1</sup> The agents are only present for the online part where the mechanism is trivial. All of the computational burden for an SPM is in the offline part. The offline computation of our posted price mechanisms is based on a subroutine that repeatedly samples the distribution of agent values and simulates Myerson’s mechanism on the sample. This clearly requires more computation than just running Myerson’s mechanism on the real agents in the first place; however, we benefit from the robustness that comes from the trivial online implementation of posted pricings.

**Related work.** See [24] and references therein for work in economics on optimal multi-dimensional mechanism design. See [10] and references therein for work in computer science on multi-dimensional pricing for a single agent. We extend the setting from [10] to multiple agents and improve their approximation for a single agent from 3 to 2.

Sequential posted price mechanisms have been considered previously in single-dimensional settings. Sandholm and Gilpin [23] show experimentally that these mechanisms compare favorably to Myerson’s optimal mechanisms. Blumrosen and Holenstein [7] show how to compute the optimal posted prices in the special case where agents’ values are distributed identically, and also show that in this case the revenue of these mechanisms approaches the optimal revenue asymptotically. Several papers study revenue maximization through online posted pricings in the context of adversarial values, albeit in the simpler context of multi-unit auctions [6, 15, 5].

---

<sup>1</sup>This is similar, for example, to nearest neighbor algorithms, where one distinguishes the time taken to construct a database, and the time taken to compute nearest neighbors over that database given a query.

Feasibility constraint	Type of mechanism	Gap from optimal		Reference
		upper bound	lower bound	
General matroid	SPM	2	$\sqrt{\pi}/2 \approx 1.25$	(§ 4.1, § C.1, [7])
	OPM	$O(\log k)$	2	(§ 5.1, § D.2)
	VCG	2	-	(§ I)
$k$ -uniform matroid, partition matroid	SPM	$e/(e-1) \approx 1.58$	1.25	(§ 4.2, § C.2)
	OPM	2	2	(§ 5.2, § D.2)
Graphical matroid	OPM	3	2	(§ 5.2, § D.3)
Intersection of two matroids	SPM	3	1.25	(§ 4.2, § C.3)
Intersection of two partition matroids	OPM	6.75	2	(§ 5.3, § D.4)
Non-matroid downward closed	SPM, OPM	-	$\Omega(\log n / \log \log n)$	(§ C.5)

Table 1: A summary of our results for single-dimensional preferences. Here  $n$  is the number of agents, and  $k$  is the size of the largest feasible set.

Feasibility constraint	Solution concept	Mechanism	Gap from optimal
multi-unit multi-item with unit-demand	dominant strategy truthful	OPM	6.75
Graphical matroid with unit-demand	dominant strategy truthful	OPM	32/3
General matroid intersection	alg. imp. in undominated strategies	SPM	8
Combinatorial auction with small bundles	alg. imp. in undominated strategies	SPM	8

Table 2: A summary of our results for multi-dimensional preferences (§ 6).

The question of whether simple mechanisms can achieve near-optimal revenue was considered recently by Hartline and Roughgarden [13]. Except for their result on single-item auctions with anonymous reserve prices, their VCG based mechanisms are likely to suffer the same impracticality criticisms as the optimal mechanism. The essay “The Lovely but Lonely Vickrey Auction” by Ausubel and Milgrom [1] discusses why this is the case. As a consequence of the near-optimality of sequential posted prices, we answer one of their open questions in the positive, namely, that the gap between the revenue optimal mechanism and a VCG mechanism with appropriate reserve prices is a constant (i.e., 2) in matroid settings but with arbitrary valuation distributions. This bound matches their result for regular distributions.

Our setting of sequential posted pricing with a matroid constraint is very closely related to the so-called matroid secretary problem [2, 3, 16], but there are two important differences: (a) they assume that agents’ values are adversarial, whereas in our setting they are drawn from known distributions, and (b) in their setting agents arrive in random order, whereas we consider optimized and adversarial orderings. Some of our results are reminiscent of that work, but our techniques are necessarily different.

Finally, our results for OPMs in the multi-unit auction setting are based on work on prophet inequalities from optimal stopping theory. While that work applies directly to the analysis of OPMs in the single-item auction setting, we show that it extends to  $k$ -unit auctions with no loss in approximation factor.

**Our results.** Our results are summarized in Tables 1 and 2. Our approximation factors in both the single-dimensional and multi-dimensional settings depend on the kind of feasibility constraint that the seller faces. In the single-dimensional setting, the feasibility constraint is a set system over agents specifying the sets of agents that can be simultaneously served. In the multi-dimensional setting, each agent is interested in buying one of multiple kinds of items or services and we assume that agents’ values for the different services are independent. The feasibility constraint is a set system over (agent, service) pairs. In both cases we assume that the set system is downward closed, i. e., any subset of a feasible set is also feasible. All of the mechanisms we develop can be computed efficiently, except for the  $O(\log k)$  approximate OPM in general matroid settings.

## 2 Problem set-up and preliminaries

### 2.1 Bayesian optimal mechanism design

In the single-parameter setting, the mechanism design problem we study (hereafter abbreviated BSMD for Bayesian single-parameter mechanism design) 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 distributed independently according to distribution function  $F_i$  with density  $f_i$ . 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  $\mathcal{J}$ . We assume that the set system  $\mathcal{J}$  is downward closed. That is, for any  $A \subset B \subseteq 2^{[n]}$ ,  $B \in \mathcal{J}$  implies  $A \in \mathcal{J}$ . 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*  $\pi(\mathbf{v})$  with a price  $\pi_i$  to be paid by agent  $i$ .

In the Bayesian multi-parameter unit-demand setting (BMUMD for short), we again have  $n$  buyers and one seller. The seller offers a number of different services indexed by set  $J$ . The set  $J$  is partitioned into groups  $J_i$ , with the services in  $J_i$  being targeted at agent  $i$ .<sup>2</sup> Each agent  $i$  is interested in getting any one of the services in  $J_i$  (that is, consumers are unit-demand agents). In the hotel rooms example, the set  $J_i$  would contain all the rooms that agent  $i$  may be interested in and the feasibility constraint ensures that each room is allocated to at most one agent. Another setting with a more general feasibility constraint arises in the context of airline ticket sales: we have a directed graph with capacities on edges owned by a seller, and a number of agents. Each agent is interested in a path of at most two hops from some source to some destination in the graph (agents want to buy airline tickets for an itinerary with at most two legs), and  $J_i$  contains all such paths. The feasibility constraint ensures that each leg is allocated upto its capacity and no more.

Agent  $i$  has value  $v_j$  for service  $j \in J_i$ .  $v_j$  is independent of all other values and is drawn from distribution  $F_j$ . Once again the seller faces a feasibility constraint specified by a set system  $\mathcal{J} \subseteq 2^J$ . Note that for every  $S \in \mathcal{J}$  and  $i \in [n]$ ,  $|S \cap J_i| \leq 1$ , that is each agent gets at most one service. As in the single-parameter case, a mechanism for this problem maps any set of bids  $\mathbf{v}$  to an allocation  $M(\mathbf{v}) \in \mathcal{J}$  and a pricing  $\pi(\mathbf{v})$ .

### 2.2 Posted-price mechanisms

A sequential posted-price mechanism (SPM),  $\mathcal{S}$ , is defined by an ordering  $\sigma$  over agents and a collection of prices  $p_i$  for  $i \in [n]$ . The mechanism is run as follows:

1. Initialize  $A \leftarrow \emptyset$ .
2. For  $i = 1$  through  $n$ , do:
  - (a) If  $A \cup \{\sigma(i)\} \in \mathcal{J}$ , offer to serve agent  $\sigma(i)$  at price  $p_i$ .
  - (b) If the agent accepts,  $A \leftarrow A \cup \{\sigma(i)\}$ .
3. Serve the agents in  $A$ .

Let  $c_i$  denote the probability taken over values of agents  $\sigma(1), \dots, \sigma(i-1)$  that the mechanism offers to serve agent  $i$ , and let  $q_i = 1 - F_i(p_i)$ . Then the expected revenue of the sequential mechanism,  $\mathcal{R}_{\mathbf{p}}^\sigma$ , is given by  $\sum_i c_i q_i p_i$ .

**Order-oblivious posted prices.** As mentioned earlier, we also study posted-price mechanisms where the order of offers is picked adversarially. We estimate (pessimistically) the expected revenue of this mechanism as follows:

$$\mathcal{R}_{\mathbf{p}}^{\text{obl}} = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \min_{S \in \mathcal{S}_{\mathbf{v}}} \sum_{i \in S} p_i$$

Here the minimization is over the class  $\mathcal{S}_{\mathbf{v}}$  of sets  $S$  that are **maximal feasible subsets** of agents that “desire” service given values  $\mathbf{v}$  and prices  $\mathbf{p}$ : (1)  $S \subseteq \mathcal{J}$ , (2)  $v_i \geq p_i$  for all  $i \in S$ ,  $S \in \mathcal{S}_{\mathbf{v}}$ , and, (3) for every feasible superset  $S'$  of a set  $S \in \mathcal{S}_{\mathbf{v}}$ ,  $S'$  contains some agent  $i$  with  $v_i < p_i$ .

For some instances, we allow a strengthening of OPMs to posted-price mechanisms where the seller is allowed to deny service to an agent even when the agent can be feasibly served alongside previously served agents. Formally,

<sup>2</sup>Since we allow for an arbitrary feasibility constraint over the set  $J$ , the assumption that the sets  $J_i$  are disjoint is without loss of generality.

the mechanism selects a pricing  $\mathbf{p}$ , and a set system  $\mathcal{J}' \subseteq \mathcal{J}$ , and runs the OPM using the prices  $\mathbf{p}$  but determining feasibility according to  $\mathcal{J}'$  instead of  $\mathcal{J}$ . Crucially, the system  $\mathcal{J}'$  is determined based only on the distributions of agents' values and not the values themselves. Therefore, this more general mechanism (that we call an OPM with a restricted feasibility constraint) retains all of the good properties of OPMs.

OPMs in multi-dimensional settings are similar: agents are approached in turn (in an arbitrary order); each agent  $i$  gets a price-menu over the subset of services in  $J_i$  that can be feasibly allocated to the agent. However, we define SPMs differently: agents are approached in turn (according to an optimal ordering) and offered individual items at a time. Offers to a single agent are not necessarily contiguous. These mechanisms are not truthful, but we show in Section 6.2 that they can nonetheless be useful in approximating the BMUMD.

### 2.3 Myerson's optimal mechanism

Myerson's seminal work describes the revenue maximizing mechanism for the Bayesian single-parameter mechanism design problem. When the value distributions  $F_i$  are regular, in Myerson's mechanism the seller first computes so-called virtual values for each agent, and then allocates to a feasible subset of agents that maximizes the "virtual surplus"—the sum of the virtual values of agents in the set minus the cost of serving that set of agents. We define these quantities formally in Appendix A. For our analyses, we mainly require the following two characterizations of the expected revenue of any truthful mechanism when the value distributions are regular. Similar characterizations hold in the non-regular case. These and their extensions to the non-regular case are proved in Appendix A.

**Proposition 1** *When all input distributions  $F_i$  are regular, the expected revenue of any truthful single-parameter mechanism  $M$  is equal to its expected virtual surplus.*

**Lemma 2** *If  $F_i$  is regular for each  $i$ , for any truthful mechanism  $M$  over the  $n$  agents, the revenue of  $M$  is bounded from above by  $\sum_i p_i^M q_i^M$  where  $q_i^M$  is the probability (over  $v_1, \dots, v_n$ ) with which  $M$  allocates to agent  $i$  and  $p_i^M = F_i^{-1}(1 - q_i^M)$ .*

*Furthermore for every  $i$  (with a regular or non-regular value distribution), there exist two prices  $\underline{p}_i$  and  $\overline{p}_i$  with corresponding probabilities  $\underline{q}_i$  and  $\overline{q}_i$ , and a number  $x_i \leq 1$ , such that  $x_i \underline{q}_i + (1 - x_i) \overline{q}_i = q_i^M$ , and the expected revenue of  $M$  is no more than  $\sum_i x_i \underline{p}_i \underline{q}_i + (1 - x_i) \overline{p}_i \overline{q}_i$ .*

## 3 A reduction from multi-parameter MD to single-parameter MD

We now present a general reduction from the multi-parameter optimal mechanism design problem to the single-parameter setting. Understanding the properties of optimal mechanisms in multi-parameter settings is tricky. Our approach begins with an upper bound on the optimal revenue in terms of the optimal revenue for a related single-parameter problem following an approach of [10]. We describe this first.

**An upper bound via copies.** Consider an instance  $\mathcal{I}$  of the BMUMD with  $n$  agents, a set  $J$  of available services (with group  $J_i$  of services targeted at agent  $i$ ), and a feasibility constraint  $\mathcal{J}$ . We will define a new instance of the BSMD in the following manner. We split each agent in  $\mathcal{I}$  into  $|J_i|$  distinct agents (called "copies"). Each copy is interested in a single item  $j \in J_i$  and behaves independently of (and potentially to the detriment of) other copies. We call this instance  $\mathcal{I}^{\text{copies}}$ . Formally, the instance has  $|J|$  distinct agents interested in a single service; agent  $j$ 's value for getting served,  $v_j$ , is distributed independently according to  $F_j$ . 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 copies of 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 lemma formalizes this (see Appendix B for a proof).

**Lemma 3** *Let  $\mathcal{A}$  be any individually rational and truthful deterministic mechanism for instance  $\mathcal{I}$  of the BMUMD. Then the expected revenue of  $\mathcal{A}$ ,  $\mathcal{R}^{\mathcal{A}}$ , is no more than the expected revenue of Myerson's mechanism for the single-parameter instance with copies,  $\mathcal{I}^{\text{copies}}$ .*

**A reduction to single-dimensional OPMs.** Next we show that if we can construct a good OPM for the setting with copies, we can construct a good OPM for the multi-dimensional setting as well. (Again, see Appendix B for a proof).

**Theorem 4** *Given an instance  $\mathcal{I}$  of the BMUMD specified by the set system  $(J, \mathcal{J})$ , there exists a truthful posted price mechanism for  $\mathcal{I}$  which achieves an  $\alpha$ -approximation to the revenue achievable by an optimal deterministic truthful mechanism, whenever there exists an OPM for the corresponding BSMD instance  $\mathcal{I}^{\text{copies}}$  that achieves an  $\alpha$ -approximation to the optimal revenue for  $\mathcal{I}^{\text{copies}}$ .*

## 4 Sequential posted-price mechanisms

In this section we focus on the BSMD and present approximations to optimal revenue via sequential posted price mechanisms for several kinds of feasibility constraints, most notably matroids and matroid intersections. Our exposition focuses on describing and analysing the approximately-optimal SPMs, and we defer a discussion of efficiently computing the SPMs to Appendix F. While our focus is on revenue, our techniques extend to a large class of objective functions, namely those that are linear in the valuations of the served agents and the payment received by the mechanism (see Appendix H).

While our analysis of the approximation factor depends closely on the feasibility constraint, we use the same approach for constructing the SPM in each case. We describe this next.

Suppose first that all the distributions  $F_i$  are regular and do not contain any point masses. Let  $q_i = q_i^M$  denote the probability that Myerson's mechanism serves agent  $i$ , and let  $p_i = F_i^{-1}(1 - q_i)$  for all  $i$ . The SPM sets a price of  $p_i$  for agent  $i$  and offers to serve the agents in decreasing order of their prices. If offered service, agent  $i$  accepts with probability exactly  $q_i$ . If the distribution  $F_i$  contains point masses, we modify the mechanism so that agent  $i$  is offered the price  $p_i$  with probability  $q_i/(1 - F_i(p_i))$ , and again has a probability exactly  $q_i$  of accepting the offer. We denote this mechanism by  $\mathcal{S}$ . We note that by Lemma 2, the revenue of Myerson's mechanism is at most  $\sum_i p_i q_i$ , and we will compare the revenue of  $\mathcal{S}$  to this upper bound. Finally, let the rank of a subset  $S$  of agents,  $\text{rank}(S)$ , denote the size of the largest feasible subset in  $S$ , that is,  $\text{rank}(S) = \max_{S' \subseteq S, S' \in \mathcal{J}} |S'|$ . Then, by definition,  $\sum_{i \in S} q_i \leq \sum_{i \in S} q_i^M \leq \text{rank}(S)$ .

When the distributions are not regular, we pick prices  $\mathbf{p}$  randomly as suggested by Lemma 2. In Appendix E we sketch the modifications required to the analysis to obtain the same approximation factors for this case as in the regular case. We now present analyses for the expected revenue of  $\mathcal{S}$  when all the distributions are regular.

### 4.1 A 2 approximation for matroids

We first consider the setting where the set system  $([n], \mathcal{J})$  is a matroid. Precisely, it satisfies the following conditions:

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

Sets in  $\mathcal{J}$  are called independent, and maximal independent sets are called *bases*. A simple consequence of the above properties is that all bases are equal in size. Therefore, the rank of a set  $S \subseteq [n]$ , is equal to the size of any maximal independent subset of  $S$ . The *span* of a set  $S \subseteq [n]$ ,  $\text{span}(S)$ , is the maximal set  $T \supseteq S$  with  $\text{rank}(T) = \text{rank}(S)$ .

We now show that for matroid set systems the SPM described above approximates the expected revenue of the optimal mechanism within a factor of 2.

**Theorem 5** *Let  $\mathcal{I}$  be an instance of the BSMD with a matroid feasibility constraint. Then, the mechanism  $\mathcal{S}$  described above 2-approximates the revenue of Myerson's mechanism for  $\mathcal{I}$ .*

*Proof:* We show that the mechanism  $\mathcal{S}$  obtains an expected revenue of at least  $\frac{1}{2} \sum_i p_i q_i$ . Note that if the mechanism ignored the feasibility constraint, and offered the prices  $\mathbf{p}$  to all agents, serving any agent that accepted its offered price, then its expected revenue would be exactly  $\sum_i p_i q_i$ . So our proof accounts for the total revenue lost due to agents "blocked" from getting an offer by previously served agents.

Formally, let  $S = \{i_1 < i_2 < \dots < i_\ell\}$  be the set of agents served, and let  $S_j$  denote the first  $j$  elements of  $S$ . Define the sets  $B_j = \text{span}(S_j) \setminus \text{span}(S_{j-1})$ . Note that the sets  $B_j$  partition the set of blocked agents. Moreover,  $B_j \subseteq \{i : i \geq i_j\}$ , since we condition on serving  $S$ , and so,  $p_i \leq p_{i_j}$  for all  $i \in B_j$ .



Denote the price offered to agent  $i_j$  by  $p^j$ . Then, the expected revenue lost given that  $S$  is served is

$$\begin{aligned}
\sum_{1 \leq j \leq \ell} \sum_{i \in B_j} p_i q_i &\leq p^1 \left( \sum_{i \in \text{span}(S_1)} q_i \right) + \sum_{1 < j \leq \ell} p^j \left( \sum_{i \in \text{span}(S_j)} q_i - \sum_{i \in \text{span}(S_{j-1})} q_i \right) \\
&= \sum_{1 \leq j < \ell} \left( (p^j - p^{j+1}) \sum_{i \in \text{span}(S_j)} q_i \right) + p^\ell \left( \sum_{i \in \text{span}(S_\ell)} q_i \right) \\
&\leq \sum_{1 \leq j < \ell} (p^j - p^{j+1}) \cdot j + p^\ell \cdot \ell = \sum_{1 \leq j < \ell} p^j,
\end{aligned}$$

which is the revenue obtained by serving  $S$ . Here we used  $\sum_{i \in \text{span}(S_j)} q_i \leq \text{rank}(S_j) \leq |S_j| = j$ . Therefore,

$$\mathbb{E}[\text{revenue lost}] \leq \sum_S \sum_{j \in S} p^j \cdot \Pr[S \text{ is served}] = \mathcal{R}_{\mathbf{p}}^\sigma,$$

and so it follows that  $\sum_i p_i q_i \leq 2\mathcal{R}_{\mathbf{p}}^\sigma$ . ■

Blumrosen et al. [7] show that the gap between the optimal SPM and Myerson's mechanism can be as large as  $\sqrt{\pi/2} \approx 1.253$  even in the single item auction case with i. i. d. agents. We describe this gap example in Appendix C.1.

## 4.2 Constant factor approximations for other feasibility constraints

We now present improved approximations for special matroids, as well as constant factor approximations for special non-matroid feasibility constraints. The theorems below are proved in Appendix C.

**Uniform matroids and partition matroids.** A matroid is  $k$ -uniform if all subsets of size at most  $k$  are feasible. An example of a  $k$ -uniform matroid constraint is a multi-unit auction where the seller has  $k$  units of an item on sale. We show that we can obtain an improved  $e/(e-1) \approx 1.58$  approximation in this case. We show in Appendix C.2 that this analysis is tight. This result extends also to partition matroids, i.e. disjoint unions of uniform matroids.

**Theorem 6** *Let  $\mathcal{I}$  be an instance of the BSMD with a partition-matroid feasibility constraint. Then, the mechanism  $\mathcal{S}$  described above  $e/(e-1)$ -approximates the revenue of Myerson's mechanism for  $\mathcal{I}$ .*

**Matroid intersections.** An intersection of  $m$  matroids,  $\mathcal{M}_1, \dots, \mathcal{M}_m$ , is a set system where a set is feasible if and only if it is feasible in each of the  $m$  matroids. An example of an intersection of two matroids is a matching. We show that the mechanism described above is an  $m+1$  approximation for intersections of  $m$  matroids.

**Theorem 7** *Let  $\mathcal{I}$  be an instance of the BSMD with a feasibility constraints that is an intersection of  $m$  matroids. Then, the mechanism  $\mathcal{S}$  described above  $(m+1)$ -approximates the revenue of Myerson's mechanism for  $\mathcal{I}$ .*

**Combinatorial auctions with small bundles.** Consider a situation where the seller has multiple copies of a number of items on sale, and each agent is interested in some (commonly known) bundles over items (and has a common value for all of these bundles). When each desired bundle is of size at most  $m$ , we call this setting a single-parameter combinatorial auction with known bundles of size  $m$ . In this case the SPM described above achieves an  $m+1$  approximation.

**Theorem 8** *Let  $\mathcal{I}$  be an instance of a single-parameter combinatorial auction with known bundles of size  $m$ . Then, the mechanism  $\mathcal{S}$  described above  $(m+1)$ -approximates the revenue of Myerson's mechanism for  $\mathcal{I}$ .*

**The general non-matroid case.** We show in Appendix C.5 that the approximations described above cannot extend to general non-matroid set systems. In particular, the example we construct describes a family of instances with i. i. d. agents and a symmetric non-matroid constraint for which the ratio between the expected revenue of Myerson’s mechanism and that of the optimal SPM is  $\Omega(\log n / \log \log n)$  where  $n$  is the number of agents. The same example also shows that while in many single-parameter pricing problems when the values are distributed in the range  $[1, h]$  it is possible to obtain an  $O(\log h)$  approximation to social welfare, the same does not hold in our general setting, and the gap can be  $\Omega(h)$ . On the other hand, the gap is always bounded by  $O(h)$  and is achieved by an SPM that charges each agent a uniform price of 1.

## 5 Order-oblivious posted-prices

The approximations designed in Section 4 rely heavily on a specific ordering of the agents. A natural question is whether the seller can obtain good revenue when he has no control over the ordering. In such a case the seller picks a set of prices in advance, and then offers them to the agents on a first-come first-served basis. We show that in many setting it is possible to determine a set of prices for which such “order-oblivious” mechanisms (OPMs) perform well.

As described in Section 2, an OPM specifies the prices to charge every agent, as well as a feasibility constraint (potentially different from  $\mathcal{J}$ ) to determine whether or not to make an offer to an agent. To pick the prices, we follow the approach taken in Section 4. The prices in the OPM are either set to be the same as for the corresponding approximately-optimal SPM, or set to infinity (effectively dropping the respective agent from consideration). We now present the details for different kinds of feasibility constraints.

### 5.1 An $O(\log k)$ approximation for general matroids

For general matroids we give an  $O(\log k)$  bound on the gap below, where  $k$  is the rank of the matroid. We remark that a similar result was obtained by Babaioff et al. [2] for the related matroid secretary problem. However, we show in Appendix D.1 that their approach cannot give a non-trivial approximation in our setting.

**Theorem 9** *Let  $\mathcal{I}$  be an instance of the BSMD with a matroid feasibility constraint. Then, there exists a set of prices  $\mathbf{p}$  such that  $\mathcal{R}_{\mathbf{p}}^{obl}$   $O(\log k)$ -approximates  $\mathcal{R}^M$  for  $\mathcal{I}$ .*

*Proof:* We present the proof for regular distributions. Appendix E presents the extension to the non-regular case. Note that since the feasibility constraint is a matroid, for any instantiation of values, the worst (least revenue) allocation is achieved when agents arrive in the order of increasing prices. Hereafter we assume that agents always arrive in that order. Let  $c_i$  be as defined in Section 2.2; recall that the expected revenue may be expressed as  $\sum_i c_i p_i q_i$ .

Now consider a hypothetical situation where the prices are all equal to 1 but the probabilities with which the agents accept the offered prices are still  $q_i$ . Then, the expected revenue of this hypothetical mechanism would be given by  $\sum_i c_i q_i$  which is at least  $1/2 \sum_i q_i$  by the argument in Theorem 5. In other words, the weighted average of the  $c_i$ s is at least  $1/2$ , weighted by the  $q_i$ s. We get the following sequence of implications.

$$(1/2) \sum_i q_i \leq \sum_i c_i q_i \leq \sum_{i: c_i < 1/4} q_i / 4 + \sum_{i: c_i \geq 1/4} q_i = (1/4) \sum_i q_i + (3/4) \sum_{i: c_i \geq 1/4} q_i \Rightarrow \sum_{i: c_i \geq 1/4} q_i \geq (1/3) \sum_i q_i$$

This means that the probability mass of elements having  $c_i \geq 1/4$  is at least a third of the total. Let  $G = \{i | c_i \geq 1/4\}$ ; the revenue obtained from serving only the agents in  $G$  is

$$\sum_{i \in G} c_i p_i q_i \geq 1/4 \sum_{i \in G} p_i^M q_i^M. \quad (1)$$

Consider recursively applying the above argument to the elements outside  $G$ . At step  $j$ , let  $G_j$  be the newly found  $G$ , and let  $E_j$  be the set of agents still under consideration, defined as  $E_1 = [n]$  and  $E_j = E_{j-1} - G_{j-1}$  for  $j > 1$ . Now, at each stage,  $G_j$  contains at least one third of the total probability mass of the remaining elements; thus, at stage  $\ell = \lceil 1 + \log_{3/2} k \rceil$ , we would have reduced the total probability mass to less than  $3/4$ ; by noting that any singleton set is independent in a matroid and applying Markov’s inequality we may see that  $G_\ell = E_\ell$ . Since the collection of  $G_j$ ’s

form a size  $O(\log k)$  partition of  $[n]$ , and summing (1) over the collection gives a total expected revenue of  $\mathcal{R}^{\mathcal{M}}/4$ , we may conclude that there is some  $G_j$  which gives a  $\Omega(1/\log k)$ -fraction of  $\mathcal{R}^{\mathcal{M}}$  regardless of ordering. ■

We remark that while the 2-approximate SPM in Section 4 can be computed efficiently, we do not know of an efficient algorithm for computing an  $O(\log k)$ -approximate order-oblivious pricing.

## 5.2 Improved approximations for special matroids

We first note that for the case of uniform matroids (where every set of size at most  $k$  is independent), an approximation of 3 can be obtained by employing techniques developed by Chawla, Hartline and Kleinberg [10] for pricing problems in multi-parameter settings. We can further improve this approximation factor to 2 via techniques developed by Samuel-Cahn [22] in the context of prophet inequalities in optimal stopping theory. We describe this approach in Appendix D.2 and show that this approximation factor is tight.

**Theorem 10** *Let  $\mathcal{I}$  be an instance of the BSMD with a uniform matroid feasibility constraint. Then, there exists a set of prices  $\mathbf{p}$  such that  $\mathcal{R}_{\mathbf{p}}^{\text{obl}}$  2-approximates  $\mathcal{R}^{\mathcal{M}}$  for  $\mathcal{I}$ .*

**Corollary 11** *Let  $\mathcal{I}$  be an instance of the BSMD with a partition matroid feasibility constraint. Then, there exists a set of prices  $\mathbf{p}$  such that  $\mathcal{R}_{\mathbf{p}}^{\text{obl}}$  2-approximates  $\mathcal{R}^{\mathcal{M}}$  for  $\mathcal{I}$ .*

For graphical matroids, Babaioff et al. [3] and Korula and Pál [16] develop approaches for reducing this case to a partition matroid that in our setting yield a 4-approximation to the optimal revenue; in Appendix D.3 we use a similar approach but exploit the connection between prophet inequalities and partition matroids to obtain a 3-approximation.

**Theorem 12** *Let  $\mathcal{I}$  be an instance of the BSMD with a graphical matroid feasibility constraint. Then, there exists an OPM with a restricted feasibility constraint that is a partition matroid, that 3-approximates  $\mathcal{R}^{\mathcal{M}}$  for  $\mathcal{I}$ .*

## 5.3 OPMs for matroid intersections

As for SPMs, our techniques for approximately-optimal OPMs in matroid settings extend to intersections of few matroids. For intersections of two partition matroids we get an 6.75-approximation (see theorem below, and proof in Appendix D.4). For intersections of  $m$  arbitrary matroids, our techniques imply an  $O(m \log k)$  approximation where  $k$  is the maximum size of a feasible set (and so is bounded by the least matroid rank); we omit the proof for brevity.

**Theorem 13** *Let  $\mathcal{I}$  be an instance of the BSMD with a feasibility constraint given by the intersection of two partition matroids. Then, there exists a set of prices  $\mathbf{p}$  such that  $\mathcal{R}_{\mathbf{p}}^{\text{obl}}$  6.75-approximates  $\mathcal{R}^{\mathcal{M}}$  for  $\mathcal{I}$ .*

**The non-matroid case.** Example 1 in Appendix C.5 already implies that order-oblivious pricings cannot obtain more than an  $O(\log n / \log \log n)$  fraction of the revenue of Myerson’s mechanism in general in non-matroid settings. How do they compare to the optimal SPM? We show in Appendix D.5 that the gap between the optimal order-oblivious pricing and the optimal SPM can be large —  $\Omega(\log n / \log \log n)$  — in the non-matroid setting.

# 6 Approximations for the multi-parameter setting

We now present approximations for various versions of the BMUMD.

## 6.1 Approximation through truthful mechanisms

We first note that for the hotel rooms example discussed in the introduction, and indeed for any setting with unit-demand agents and multiple units of multiple items on sale, a 6.75-approximation follows from Theorems 4 and 13.

**Theorem 14** *Consider an instance of the BMUMD where the seller has multiple copies of  $n$  items on sale, and agents are unit-demand and have independently distributed values for each item. Then there exists an 6.75-approximate OPM for this instance. The prices for this mechanism can be computed in polynomial time.*

A similar result for graphical matroids follows from Theorems 4, 12 and 13 (see Appendix G for a proof).

**Theorem 15** *Consider an instance of the BMUMD based on a graph  $G = (V, E)$  where the agents have independent values for different edges and are interested in buying one edge each. The seller can allocate any forest in the graph. Then there exists a 10.67 approximate OPM for this instance. The prices for this mechanism can be computed in polynomial time.*

## 6.2 Approximation through implementation in undominated strategies

Both of the results above involve feasibility constraints that admit good OPMs in single-dimensional settings. Can we design good multi-dimensional mechanisms for set systems that admit good SPMs in the single-dimensional setting, but for which we do not know constant approximate OPMs? Two examples are general matroid intersections and combinatorial auctions with small bundles (e.g. the airline tickets setting described in Section 2).

We now show that this can be done if we relax truthfulness to *implementation in undominated strategies*. (See formal definition in Appendix G.) Our mechanism for both of the cases above is an SPM specified by a set of prices, one for each service in  $J$ , and an ordering over services. It begins by announcing the prices to the agents. Then, as for single-dimensional instances, it considers offering the services to the agents in turn: at every step, depending on the services allocated so far, it determines whether or not it is feasible to allocate the next service  $j \in J_i$  to the corresponding agent  $i$ , and if so, offers a price  $p_j$  to  $i$ . This mechanism is not truthful. For example, an agent may reject an offer for a service  $j$  even if his value for  $j$  exceeds its price, if he anticipates obtaining a more profitable offer in the future. Nevertheless we can infer some properties about rational agent behavior in such a mechanism.

**Lemma 16** *Consider an instance  $I$  of the BMUMD and an SPM as defined above with prices  $\mathbf{p}$  and ordering  $\sigma$ . Then, the following holds for any undominated strategy of any agent: if an agent  $i$  desires only one service at prices  $\mathbf{p}$ , that is,  $v_j \geq p_j$  for only one  $j \in J_i$ , then the agent must accept  $j$  if offered the service.*

This lemma and the following theorem are proved in Appendix G.

**Theorem 17** *Given an instance of the BMUMD with a general matroid intersection constraint, there exists an SPM that implements a 8-approximation in undominated strategies. Given an instance of a combinatorial auction with known bundles of size 2, there exists an SPM that implements a 8-approximation for the instance in undominated strategies.*

## 7 Discussion

We presented constant factor approximations to revenue for several classes of multi-dimensional mechanism design problems by designing approximately-optimal posted price mechanisms for single-dimensional settings. This approach does not extend beyond matroid and matroid-like settings. However, it is possible that there is some other class of simple near-optimal mechanisms for non-matroid single-dimensional settings that do not exploit competition among agents. Such mechanisms may lead to approximately-optimal multi-dimensional mechanisms for a broader class of feasibility constraints.

More generally, two important assumptions underlie our work: (1) agents are unit-demand, and (2) their values for different services are distributed independently. In the absence of either of these assumptions the upper bound on the optimal revenue based on the setting with copies does not remain valid. An important open question is to design a reasonably tight upper bound in those cases, and use it to approximate the optimal mechanism.

## References

- [1] Lawrence M. Ausubel and Paul Milgrom. The lovely but lonely vickrey auction. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, chapter 1, pages 17–40. MIT Press, Cambridge, MA, 2006.
- [2] M. Babaioff, N. Immorlica, and R. Kleinberg. Matroids, secretary problems, and online mechanisms. In *Proc. 19th ACM Symp. on Discrete Algorithms*, 2007.

- [3] Moshe Babaioff, Michael Dinitz, Anupam Gupta, Nicole Immorlica, and Kunal Talwar. Secretary problems: weights and discounts. In *SODA '09: Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1245–1254, Philadelphia, PA, USA, 2009. Society for Industrial and Applied Mathematics.
- [4] Moshe Babaioff, Ron Lavi, and Elan Pavlov. Single-value combinatorial auctions and algorithmic implementation in undominated strategies. *J. ACM*, 56(1):1–32, 2009.
- [5] A. Blum and J. Hartline. Near-optimal online auctions. In *Proc. 16th ACM Symp. on Discrete Algorithms*. ACM/SIAM, 2005.
- [6] A. Blum, V. Kumar, A. Rudra, and F. Wu. Online learning in online auctions. *Theoretical Computer Science*, 324:137–146, 2004.
- [7] L. Blumrosen and T. Holenstein. Posted prices vs. negotiations: an asymptotic analysis. In *Proc. 10th ACM Conf. on Electronic Commerce*, 2008.
- [8] P. Briest. Towards hardness of envy-free pricing. Technical Report TR06-150, ECCC, 2006.
- [9] J. Bulow and J. Roberts. The simple economics of optimal auctions. *The Journal of Political Economy*, 97:1060–90, 1989.
- [10] S. Chawla, J. Hartline, and R. Kleinberg. Algorithmic pricing via virtual valuations. In *Proc. 9th ACM Conf. on Electronic Commerce*, pages 243–251, 2007.
- [11] A. Goldberg and J. Hartline. Collusion-resistant mechanisms for single-parameter agents. In *Proc. 16th ACM Symp. on Discrete Algorithms*, 2005.
- [12] 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.
- [13] J. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *Proc. 11th ACM Conf. on Electronic Commerce*, 2009.
- [14] Catherine Holahan. Auctions on eBay: A Dying Breed.  
[http://www.businessweek.com/technology/content/jun2008/tc2008062\\_112762.htm](http://www.businessweek.com/technology/content/jun2008/tc2008062_112762.htm), June 3 2008.
- [15] R. Kleinberg and T. Leighton. The value of knowing a demand curve: Bounds on regret for on-line posted-price auctions. In *Proc. 44th IEEE Symp. on Foundations of Computer Science*, 2003.
- [16] Nitish Korula and Martin Pál. Algorithms for secretary problems on graphs and hypergraphs. In *ICALP '09: Proceedings of the 36th International Colloquium on Automata, Languages and Programming*, pages 508–520, Berlin, Heidelberg, 2009. Springer-Verlag.
- [17] D. Lehmann, L. I. O’Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. In *Proc. 1st ACM Conf. on Electronic Commerce*, pages 96–102. ACM Press, 1999.
- [18] R.P. McAfee and J. McMillan. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory*, 46(2):335–354, 1988.
- [19] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [20] N. Nisan and A. Ronen. Algorithmic mechanism design. In *Proc. 31st ACM Symp. on Theory of Computing*, pages 129–140. ACM Press, 1999.
- [21] Jean-Charles Rochet and Philippe Chone. Ironing, Sweeping, and Multidimensional Screening. *Econometrica*, 66(4):783–826, 1998.
- [22] Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *The Annals of Probability*, 12(4):1213–1216, 1984.

- [23] Tuomas Sandholm and Andrew Gilpin. Sequences of take-it-or-leave-it offers: near-optimal auctions without full valuation revelation. In *AAMAS '06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 1127–1134, New York, NY, USA, 2006. ACM.
- [24] D. Vincent and A. Manelli. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory*, 137(1):153–185, 2007.
- [25] Robert B. Wilson. *Nonlinear Pricing*. Oxford University Press, 1997.

## A Myerson’s mechanism and revenue bounds for truthful mechanisms

Myerson’s seminal work describes the revenue maximizing mechanism for the Bayesian single-parameter mechanism design problem, BSMD, described in Section 2.1. In Myerson’s mechanism the seller first computes so-called virtual values for each agent, and then allocates to a feasible subset of agents that maximizes the “virtual surplus”—the sum of the virtual values of agents in the set minus the cost of serving that set of agents. These quantities are formally defined as follows.

**Definition 1** For a valuation  $v_i$  drawn from  $F_i$ , the virtual valuation of agent  $i$  is given by

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$$

The virtual surplus of a set  $S$  of agents is defined as  $\Phi(S, \mathbf{v}) = \sum_{i \in S} \phi_i(v_i)$ .

Myerson’s optimal mechanism is based on the following observation.

**Proposition 1** The expected revenue of any truthful single-parameter mechanism  $M$  is equal to its expected virtual surplus.

A direct consequence of Proposition 1 is that the expected revenue maximizing mechanism would be one that maximizes expected virtual surplus. Given a vector  $\mathbf{v}$  of values, Myerson’s mechanism serves the set  $\operatorname{argmax}_S \Phi(S, \mathbf{v})$ . This mechanism is truthful when the virtual valuation function is monotone non-decreasing for every  $i$ , or in other words, the distribution  $F_i$  is *regular*. Note that we do not explicitly specify the prices charged by the mechanism. These are uniquely determined by the allocation rule assuming that agents that are not served pay nothing.

**Definition 2** A one dimensional distribution  $F$  is *regular*, if  $\phi(v)$  is monotone non-decreasing in  $v$ .

**Irregular distributions and ironed virtual values.** When the distributions  $F_i$  are irregular, that is, Definition 2 does not hold, Myerson’s mechanism as described above will no longer be truthful. Myerson fixed this case by “ironing” the virtual valuation function and converting it into a monotone non-decreasing function. We skip the description of this procedure; the reader is referred to [9, 10] for details. Let us denote the ironed virtual value of an agent with value  $v_i$  by  $\bar{\phi}_i(v_i)$ . We then note the following.

**Proposition 18** The expected revenue of any truthful single-parameter mechanism  $M$  is no more than its expected ironed virtual surplus. If the probability with which the mechanism serves agent  $i$ , as a function of  $v_i$ , is constant over any valuation range in which the ironed virtual value of  $i$  is constant, the expected revenue is equal to expected ironed virtual surplus.

Myerson’s mechanism serves a subset of agents that maximizes the ironed virtual surplus, breaking ties in an arbitrary but consistent manner. Denoting the revenue of a mechanism  $\mathcal{A}$  by  $\mathcal{R}^{\mathcal{A}}$  and the revenue of Myerson’s mechanism  $\mathcal{M}$  by  $\mathcal{R}^{\mathcal{M}}$  we get the following:

**Theorem 19**  $\mathcal{R}^{\mathcal{M}} \geq \mathcal{R}^{\mathcal{A}}$  for every truthful mechanism  $\mathcal{A}$ .

**Bounding the revenue of the Bayesian optimal mechanism.** Propositions 1 and 18 give one approach of bounding the expected revenue of Myerson's mechanism. We now develop a different bound that is useful in proving performance guarantees for posted-price mechanisms.

**Lemma 2** *If  $F_i$  is regular for each  $i$ , for any truthful mechanism  $M$  over the  $n$  agents, the revenue of  $M$  is bounded from above by  $\sum_i p_i^M q_i^M$  where  $q_i^M$  is the probability (over  $v_1, \dots, v_n$ ) with which  $M$  allocates to agent  $i$  and  $p_i^M = F_i^{-1}(1 - q_i^M)$ .*

*Furthermore for every  $i$  (with a regular or non-regular value distribution), there exist two prices  $\underline{p}_i$  and  $\overline{p}_i$  with corresponding probabilities  $\underline{q}_i$  and  $\overline{q}_i$ , and a number  $x_i \leq 1$ , such that  $x_i \underline{q}_i + (1 - x_i) \overline{q}_i = q_i^M$ , and the expected revenue of  $M$  is no more than  $\sum_i x_i \underline{p}_i \underline{q}_i + (1 - x_i) \overline{p}_i \overline{q}_i$ .*

*Proof:* We prove the regular case first. Consider the revenue that  $M$  draws from serving agent  $i$ . This is clearly bounded above by the optimal mechanism that sells to only  $i$ , but with probability at most  $q_i^M$ . By Proposition 1, such a mechanism should sell to agent  $i$  with probability 1 whenever the value of the agent is above  $F_i^{-1}(1 - q_i^M)$  and with probability 0 otherwise. The revenue of the optimal such mechanism is therefore  $p_i^M q_i^M$ .

In the non-regular case, note that the value  $p_i^M$  may fall in a valuation range that has constant ironed virtual value. Let  $\underline{p}_i$  denote the infimum  $\inf\{v : \bar{\phi}_i(v) = \bar{\phi}_i(p_i^M)\}$  of this range and  $\overline{p}_i$  denote the supremum  $\sup\{v : \bar{\phi}_i(v) = \bar{\phi}_i(p_i^M)\}$ . Let  $\underline{q}_i = 1 - F_i(\underline{p}_i)$  and  $\overline{q}_i = 1 - F_i(\overline{p}_i)$ . Then,  $\overline{q}_i \leq q_i^M \leq \underline{q}_i$ , and there exists an  $x_i$  such that  $x_i \underline{q}_i + (1 - x_i) \overline{q}_i = q_i^M$ . Now an easy consequence of Proposition 18 is that the optimal mechanism with selling probability  $q_i^M$  sells to the agent with probability  $x_i$  if the agent's value is between  $\underline{p}_i$  and  $\overline{p}_i$ , and with probability 1 if the value is above  $\overline{p}_i$ . The revenue of this mechanism is exactly  $x_i \underline{p}_i \underline{q}_i + (1 - x_i) \overline{p}_i \overline{q}_i$ . ■

## B Reducing BMUMD to BSMD

In this section we prove Lemma 3 and Theorem 4.

*Proof of Lemma 3.* We first note that a mechanism is individually rational if we have  $\pi_i \leq v_j$  for  $j \in S \cap J_i$ , and  $\pi_i = 0$  if  $S \cap J_i = \emptyset$ . Truthful mechanisms in multi-parameter settings satisfy the weak monotonicity condition defined below.

**Definition 3** *A mechanism  $(M, \pi)$  satisfies weak monotonicity if for any agent  $i$  and any two types (value vectors)  $\mathbf{v}^1$  and  $\mathbf{v}^2$  with  $v_j^1 = v_j^2$  for all  $j \in J \setminus J_i$ , the following holds:*

$$v_{M_i(\mathbf{v}^1)}^1 + v_{M_i(\mathbf{v}^2)}^2 \geq v_{M_i(\mathbf{v}^2)}^1 + v_{M_i(\mathbf{v}^1)}^2$$

Here  $M_i(\mathbf{v})$  denotes the unique index in  $M(\mathbf{v}) \cap J_i$ .

We show that we can construct a truthful mechanism  $\mathcal{A}^{\text{copies}}$  for the  $\mathcal{I}^{\text{copies}}$  with revenue no less than that of  $\mathcal{A}$ . The lemma then follows from the optimality of Myerson's mechanism. Given a vector of values  $\mathbf{v}$ , the mechanism  $\mathcal{A}^{\text{copies}}$  allocates to the set that  $\mathcal{A}$  allocates to in  $\mathcal{I}$  given the same value vector. We first claim that the allocation rule of  $\mathcal{A}^{\text{copies}}$  is monotone non-decreasing in any  $v_j$ , implying that there exists a payment rule that makes the mechanism truthful. To prove the claim, fix any agent  $i$  and  $j \in J_i$ , and consider two value vectors  $\mathbf{v}^1$  and  $\mathbf{v}^2$  with  $v_j^1 = x$ ,  $v_j^2 = y$ , and  $v_{j'}^1 = v_{j'}^2$  for  $j' \neq j$ . Let  $\alpha_x$  and  $\alpha_y$  denote the probabilities of serving agent  $i$  with service  $j$  under the two value vectors respectively, and let  $\beta_x$  and  $\beta_y$  denote the total value that agent  $i$  obtains from other services  $j' \in J_i, j' \neq j$ , in the two cases respectively. Then the weak-monotonicity (Definition 3) of  $\mathcal{A}$  implies that

$$(x\alpha_x + \beta_x) + (y\alpha_y + \beta_y) \geq (x\alpha_y + \beta_y) + (y\alpha_x + \beta_x)$$

or,

$$(x - y)(\alpha_x - \alpha_y) \geq 0$$

Therefore the claim holds.

It remains to prove that the expected revenue of  $\mathcal{A}^{\text{copies}}$  given  $\mathcal{I}^{\text{copies}}$  is no less than the expected revenue of  $\mathcal{A}$  given  $\mathcal{I}$ . Note that any deterministic multi-parameter mechanism can be interpreted as offering a price menu with one price for each item or service to each agent as a function of other agents' bids [25]. The agent then chooses the item or service that brings her the most utility. Given this characterization, suppose that for a fixed set  $\mathbf{v}$  of values, mechanism  $\mathcal{A}$  offers a price menu with prices  $\mathbf{p}$  to agent  $i$ . Then, it draws a revenue of  $p_j$  from  $i$  whenever service  $j$  is offered. On the other hand, mechanism  $\mathcal{A}^{\text{copies}}$  charges the agent  $j$  the minimum amount it needs to bid to be served, which is no less than  $p_j$ , as  $\mathcal{A}$  is individually rational.

*Proof of Theorem 4.* Consider an  $\alpha$ -approximate OPM for  $\mathcal{I}^{\text{copies}}$  with prices  $\mathbf{p}$ . The  $\alpha$ -approximate mechanism for  $\mathcal{I}$  is described as follows. It serves the agents in the order in which they arrive. When agent  $i$  arrives, depending on the set of services already allocated, the mechanism determines the subset  $J'_i$  of services in  $J_i$  that can be feasibly allocated to  $i$ , and offers a price menu of  $\{p_j\}_{j \in J'_i}$  to  $i$ . Agent  $i$  then chooses a service from the menu and this service is allocated to it. Truthfulness follows from the definition. In order to argue that the mechanism is  $\alpha$ -approximate, we will show that its revenue is no less than the revenue of the OPM for  $\mathcal{I} - \mathcal{R}_{\mathbf{p}}^{\text{obl}}$ . Then the result follows from Lemma 3. To see that the expected revenue of the mechanism is at least  $\mathcal{R}_{\mathbf{p}}^{\text{obl}}$ , we claim that the mechanism allocates a maximal feasible set of services. If not, then there exists an agent  $i$  and a service  $j$  such that it is feasible to allocate  $j$  to  $i$  (that is,  $j \in J'_i$ , and  $i$  has not been allocated any service), and the value of  $i$  for  $j$  exceeds its price. Then, at the time that  $i$  is offered a price menu, it must have been the case that  $i$  chose  $j$  or some other service in  $J'_i$  and got allocated that service, and we get a contradiction. This concludes the proof.

## C Approximations via SPMs

In this section we present missing proofs from Section 4. In particular, we prove that the SPMs described in Section 4 are  $e/(e-1)$  approximate for partition matroids,  $m$  approximate for intersections of  $m$  matroids, and  $m$  approximate for combinatorial auctions with known bundles of size  $m$ .

### C.1 A lower bound example for 1-uniform matroids

Blumrosen and Holenstein [7] give an example where the gap between the revenue of the optimal SPM and that of Myerson's mechanism is a factor of  $\sqrt{\pi/2} \approx 1.253$ . We reproduce the example here for completeness. There are  $n$  agents, each with a value distributed independently according to function  $F(v) = 1 - 1/v^2$ . The seller has one item to sell. Then, the expected revenue of Myerson's mechanism is  $\Gamma(1/2)\sqrt{n}/2$ , where  $\Gamma(\cdot)$  is the Gamma function. On the other hand, the expected revenue of the optimal SPM can be computed to be  $\sqrt{n}/2$ . Therefore, we get a gap of  $\Gamma(1/2)/\sqrt{2} = \sqrt{\pi/2} \approx 1.253$ .

### C.2 Proof of Theorem 6: an $\frac{e}{e-1}$ approximation for uniform and partition matroids

We first prove Theorem 6 for 1-uniform matroids. The Revenue  $\mathcal{R}^S$  of the SPM described in Section 4 can be written as

$$\mathcal{R}^S = \sum_{i=1}^n c_i p_i^M q_i^M = \sum_{i=1}^n \prod_{j=1}^{i-1} (1 - q_j) p_i^M q_i^M,$$

where  $c_i = \prod_{j=1}^{i-1} (1 - q_j)$  is the probability that agent  $i$  is offered service. Note that  $c_i \geq c_j$  for  $i \leq j$ . Let  $p$  be the price satisfying the equation

$$\sum_i p_i q_i = p \sum_i q_i. \quad (2)$$

We now prove that among the set of all product distributions  $G = (G_1 \times G_2 \times \dots \times G_n)$  which satisfy

- $\Pr[\text{Myerson's mechanism serves agent } i] = q_i$ ; and
- $\sum_i G_i^{-1} (1 - q_i) q_i = \sum_i p_i q_i$ ,



the revenue obtained is lowest when  $G$  is the distribution for which all the prices are equal, i.e.  $G_i^{-1}(1 - q_i) = p$  for all agents  $i$ . Let  $\mathcal{R}_{eq}^S$  denote the revenue of the SPM whose prices are equal to  $p$  for all agents.

**Lemma 20** *It is always the case that  $\mathcal{R}^S \geq \mathcal{R}_{eq}^S$ .*

*Proof:* Let  $\delta_i = q_i(p_i - p)$ . So we have

$$\mathcal{R}^S = \sum_{i=1}^n c_i p_i q_i = \sum_{i=1}^n c_i (p q_i + \delta_i) = \mathcal{R}_{eq}^S + \sum_{i=1}^n c_i \delta_i \geq \mathcal{R}_{eq}^S,$$

where the inequality follows from observing that:

- $c_i$ 's are in descending order;
- $\exists j$  such that  $\delta_i$  is non-negative for all  $i \leq j$  and negative otherwise; and
- $\sum_i \delta_i = 0$  (By (2)).

■

**Theorem 21** *The revenue  $\mathcal{R}^S$  of the SPM  $\mathcal{S}$  is a  $\frac{e}{e-1}$  approximation to the expected revenue of Myerson's mechanism in the case of a 1-uniform matroid.*

*Proof:* Let  $\sum_i q_i = s (\leq 1)$ . Lemma 20 implies the theorem when  $\mathcal{R}_{eq}^S \geq (1 - \frac{1}{e})ps$ . We can see this holds, since

$$\begin{aligned} \mathcal{R}_{eq}^S &= p(\Pr[\text{Some agent is served}]) = p(1 - \Pr[\text{No agent is served}]) \\ &= p \left( 1 - \prod_{i=1}^n (1 - q_i) \right) \\ &\geq p \left( 1 - \prod_{i=1}^n (1 - s/n) \right) \\ &\geq (1 - 1/e)ps, \end{aligned} \tag{3}$$

where (3) follows since the product is maximized when the  $q_i$ 's are all equal. ■

Next we consider the  $k$ -uniform case.

**Theorem 22** *The revenue  $\mathcal{R}^S$  of the SPM  $\mathcal{S}$  is a  $\frac{e}{e-1}$  approximation to the expected revenue of Myerson's mechanism in the case of a  $k$ -uniform matroid.*

*Proof:* Our proof technique is closely related to the proof for 1-uniform matroids. If we define  $\mathcal{R}_{eq}^S$  as defined above for the 1-uniform case, then the proof of Lemma 20 extends to  $k$ -uniform matroids also. Thus it would be enough to argue that  $\mathcal{R}_{eq}^S$  achieves a  $e/(e-1)$  approximation. Let  $p$  be the common price for all agents which satisfies (2).

For any set of probabilities  $\{q_i\}$  in the  $k$ -item case, let us define  $q_i' = q_i/k$ . Note that the probabilities  $\{q_i'\}$  form a valid set of probabilities for a 1-item case because

$$\sum_i q_i' = \sum_i q_i/k \leq 1$$

Let  $c_i'$  denote the probability that agent  $i$  is considered for service in the 1-item case. We can come up with distributions  $F_i'$  for the 1-item case such that the price  $F_i'^{-1}(1 - q_i')$  is the same for all agents and is equal to  $p$ . By Theorem 21, we know that the revenue in this 1-item case is at least  $(1 - 1/e) \sum_i p q_i'$ . We prove that we get the same approximation factor of  $e/(e-1)$  for the  $k$ -item case by the following induction. We assume that for  $j-1 \leq n$ , the revenue  $R_{j-1}$  from the first  $j-1$  items is at least  $k$  times the revenue  $R'_{j-1}$  from the first  $j-1$  items in the corresponding 1-item case i.e.  $\sum_{i=1}^{j-1} c_i p q_i \geq k \cdot \sum_{i=1}^{j-1} c_i' p q_i'$ . We prove the same for  $j$  through two cases.

1. If  $c_j \geq c'_j$ , then we are done, because we know that revenue  $R_j$  from the first  $j$  items can be written as

$$R_j = R_{j-1} + c_j p q_j \geq k(R'_{j-1}) + k c'_j p q'_j = k R'_j.$$

The inequality uses the induction hypothesis.

2. If  $c_j < c'_j$ , we show that the revenue obtained is better than when  $c_j = c'_j$  and then we will be done. To see this observe that the revenue  $R_j$  can be written as being conditioned on whether or not  $k$  items were sold in the first  $j - 1$  items. So we have

$$R_j = (1 - c_j)kp + c_j(pq_j + \mathbb{E}[\text{Revenue from first } j - 1 \mid \text{at most } k - 1 \text{ of first } j - 1 \text{ are served}]);$$

since

$$kp \geq (pq_j + \mathbb{E}[\text{Revenue from first } j - 1 \mid \text{at most } k - 1 \text{ of first } j - 1 \text{ are served}]),$$

the revenue only decreases by increasing  $c_j$  to  $c'_j$ .

Thus in either case, the  $k$ -item case has a better revenue, guaranteeing an approximation factor of  $\frac{e}{e-1}$ .  $\blacksquare$

**Corollary 23** *The revenue  $\mathcal{R}^S$  of the SPM  $S$  is a  $\frac{e}{e-1}$  approximation to the expected revenue of Myerson's mechanism in the case of a partition-matroid.*

We note that this analysis is tight. In particular, consider an example with  $n$  i. i. d. agents and a seller with one item. Suppose that the value of each agent is independently 1 with probability  $1 - \epsilon$  and 0 otherwise for some small  $\epsilon > 0$ . Then, the expected revenue of Myerson's mechanism is equal to the probability that at least one agent has value 1, which is  $1 - o(1)$ . The probability with which Myerson's mechanism serves a particular agent is  $(1 - o(1))/n$ . Therefore, our mechanism sets a price of 1 for each agent, and offers the price to each agent with probability roughly  $1/n$  until some agent accepts. So the revenue of our mechanism is roughly  $1 - (1 - 1/n)^n = 1 - 1/e$ . A similar gap can be obtained even if the SPM decides to offer a price to the last agent with certainty when no other agent accepts the item. Note that there is a simple SPM that in this case obtains the optimal revenue—offer a price of 1 to each agent in turn with certainty until the item is sold.

### C.3 Proof of Theorem 7: an $m + 1$ approximation for intersections of $m$ matroids

Let the  $m$  matroids be denoted by  $\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_m$ . Let  $\text{rank}_a(S)$  and  $\text{span}_a(S)$  denote respectively the rank and span of set  $S$  in the matroid  $\mathcal{M}_a$ . Note that for any subset  $S$  and any  $a \in [m]$ , we have  $\sum_{i \in S} q_i^M \leq \text{rank}_a(S)$ .

Once again, let  $S = \{i_1 < i_2 < \dots < i_\ell\}$  denote the set of agents served. We prove the theorem by showing that the expected revenue of  $S$  is at least  $1/(m + 1) \sum_i p_i q_i$ , by arguing that the total price paid by agents in  $S$  is at least  $1/m$  times the expected revenue from agents that are “blocked” by  $S$ .

Let  $S_j$  denote the first  $j$  elements of  $S$ . For each  $1 \leq a \leq m$ , define sets  $B_j^a$  with respect to matroid  $\mathcal{M}_a$  as in the proof of Theorem 5. That is,  $B_j^a = \text{span}_a(S_j) \setminus \text{span}_a(S_{j-1})$ . Denote the price of item  $i_j$  by  $p^j$ . Then, if we let  $B_j = \cup_{a=1}^m B_j^a$ , we can upper bound the expected revenue lost when  $S$  is served by

$$\sum_{1 \leq j \leq \ell} \sum_{i \in B_j} p_i q_i \leq \sum_{a=1}^m \sum_{1 \leq j \leq \ell} \sum_{i \in B_j^a} p_i q_i \leq m \sum_{1 \leq j < \ell} p^j.$$

Here we used the same algebraic transformation as in the proof of Theorem 5 along with the fact that  $\sum_{i \in B_j^a} q_i \leq \sum_{i \in \text{span}_a(S_j)} q_i \leq j$ .

Therefore as before we get  $\sum_i p_i q_i \leq (m + 1) \mathcal{R}_p^\sigma$ .

## C.4 Proof of Theorem 8: an $m + 1$ approximation for combinatorial auctions with known bundles of size $m$

Let  $A$  denote the set of items available to the seller, each with some multiplicity. First suppose that each agent is single-minded, that is, each agent is interested in only one bundle of items, the bundle being of size at most  $m$ . Then, the feasibility constraint is an intersection over  $|A|$  uniform matroids, one corresponding to each item, with each agent participating in only  $m$  of the matroids. Now it is easy to adapt the proof of Theorem 7 to obtain an  $m + 1$  approximation.

More generally suppose that every agent is interested in a collection of bundles, each of size at most  $m$ , and modify the mechanism  $\mathcal{S}$  so that in addition to deciding whether or not to serve an agent, it also arbitrarily allocates any available desired bundle to every agent it serves. Then we can argue that for any set  $S$ , and set  $B$  blocked by the agents in  $S$ , the sum of the probabilities  $q_i$  over the set  $B$  is no more than  $m$  times the size of  $S$ . Therefore, once again following along the proof of Theorem 7, we get an  $m + 1$  approximation.

## C.5 Bad gap example for general non-matroids

We now show that the approximations described above cannot extend to general non-matroid set systems. In particular, the example below describes a family of instances with i. i. d. agents and a symmetric non-matroid constraint for which the ratio between the expected revenue of Myerson's mechanism and that of the optimal SPM is  $\Omega(\log n / \log \log n)$  where  $n$  is the number of agents.

**Example 1** For a given  $m$  set  $n = m^{m+1}$ . Partition  $[n]$  into  $m^m$  groups  $G_1, \dots, G_{m^m}$  of size  $m$  each, with  $G_i \cap G_j = \emptyset$  for all  $i \neq j$ . The set system  $\mathcal{J}$  contains all subsets of groups  $G_i$ , that is,  $\mathcal{J} = \{A : \exists i \text{ with } A \subseteq G_i\}$ . Each agent has a value of 1 with probability  $1 - 1/m$  and  $m$  with probability  $1/m$ .

For any given valuation profile, let us call the agents with a value of  $m$  to be good agents and the rest to be bad agents. The probability that a group contains  $m$  good agents is  $m^{-m}$ . Therefore in expectation one group has  $m$  good agents and Myerson's mechanism can obtain revenue  $m^2$  from such a group:  $\mathcal{R}^{\mathcal{M}} = \Omega(m^2)$ .

Next consider any SPM. The mechanism can serve at most  $m$  agents. If all the served agents are bad, the mechanism obtains a revenue of at most  $m$ . On the other hand, once the mechanism commits to serving a good agent, it can only serve agents within the same group in the future. These have a total expected value less than  $2m$ . Therefore, the revenue of any SPM is at most  $3m$ , and we get a gap of  $\Omega(m) = \Omega(\log n / \log \log n)$ .

The above example also shows that while in many single-parameter pricing problems when the values are distributed in the range  $[1, h]$  it is possible to obtain a  $\log h$  approximation to social welfare, the same does not hold in our general setting. In the example we have  $h = m$  and the gap between the expected revenue of the optimal SPM and that of Myerson's mechanism is  $\Omega(h)$ . On the other hand, the gap is always bounded by  $O(h)$  and is achieved by an SPM that charges each agent a uniform price of 1.

# D Approximations via OPMs

## D.1 General matroids.

In Section 5.1 we design an  $O(\log k)$  approximate OPM for general matroids. We remark that a similar result was obtained by Babaioff et al. [2] for the related matroid secretary problem. In Babaioff et al.'s setting agents arrive in a random order but their values are adversarial. They present an  $O(\log k)$  approximation by picking a price uniformly at random in the set  $\{h/k, 2h/k, \dots, h\}$  and charging it to every agent; here  $h$  is the largest among all values. In our setting such an approach does not work: the example below shows that no uniform pricing can achieve an  $o(\log h)$  approximation even for  $k = 1$ .

**Example 2** Let  $k = 1$  and consider a group of  $h$  agents where agent  $i$  has a value of  $i$  with probability  $1/2i^2$  and zero otherwise. Then an SPM that sets a price of  $i$  for agent  $i$  obtains an expected revenue of  $\Omega(\log h)$ . On the other hand, an SPM that uses a uniform price of  $c$  only obtains expected revenue  $\sum_{i \in [c, h]} c/2i^2 < c/2c = 1/2$ .

## D.2 Uniform and partition matroids

Consider the following setting from [22]: a gambler is presented samples from  $n$  distributions in order,  $X_1, \dots, X_n$ . For each sample, the gambler must decide whether to pick this sample (and end the game) or skip it (to never return to it). The gambler can choose at most one of the samples, and obtains a reward equal to the value of the sample. Can the gambler do nearly as well in expectation as a prophet that knows the maximum value in the sample? Samuel-Cahn [22] shows that there is a simple threshold rule for picking samples that uses a common threshold for each random variable, such that the expected value of the gambler is within a factor of 2 of the expected value obtained by the prophet. We first extend this result of [22] to the case where both gambler and prophet can pick  $k$  values, and then describe how it applies to our setting of maximizing revenue.

We begin with some definitions. Given a collection of  $n$  independent, nonnegative random variables  $X_1, \dots, X_n$ , we consider extending the prophet inequalities to the case where the gambler and the prophet are each allowed  $k$  choices. Let  $X_{(1)} \geq \dots \geq X_{(n)}$  be the order statistics for  $X_1, \dots, X_n$ . For a value  $x$ , let  $(x)^+$  denote the positive portion of  $x$ , i.e.  $(x)^+ = \max(0, x)$ .

For a constant  $c$ , let  $t_1(c), \dots, t_k(c)$  denote the  $k$  indices selected by a threshold stopping rule using  $c$ , i.e.  $t_i(c)$  is the lesser  $(n - k + i)$  and the  $i^{\text{th}}$  smallest index  $j$  such that  $X_j \geq c$  (or simply the former when the latter does not exist).

Let  $a^*$  and  $b^*$  be the unique solutions to the equations

$$a = \sum_{i=1}^k \mathbb{E} (X_{(i)} - a/k)^+, \quad \text{and} \quad b = \sum_{i=1}^n \mathbb{E} (X_i - b/k)^+,$$

respectively. Then it must be the case that  $a^* \leq b^*$ , and we get the following theorem.

**Theorem 24** *For  $a^* \leq k \cdot c \leq b^*$ , we have that  $\sum_{i=1}^k \mathbb{E} [X_{(i)}] \leq 2 \sum_{i=1}^k \mathbb{E} [X_{t_i(c)}]$ .*

*Proof:* First, we note that for any threshold  $t$

$$\sum_{i=1}^k \mathbb{E} [X_{(i)}] \leq \sum_{i=1}^k \mathbb{E} [t + (X_{(i)} - t)^+] = k \cdot t + \sum_{i=1}^k \mathbb{E} (X_{(i)} - t)^+,$$

which implies  $\sum_{i=1}^k \mathbb{E} [X_{(i)}] \leq 2a^*$  with the substitution  $t = a^*/k$ . Now, any time  $t_k(c) < n$ , we know there are at least  $k$   $X_i$  at or above our threshold  $c$ , and so

$$\sum_{i=1}^k \mathbb{E} [X_{t_i(c)}] \geq kc \cdot \Pr[t_k < n] + \sum_{i=1}^k \mathbb{E} (X_{t_i(c)} - c)^+.$$

Let  $I(\mathcal{E})$  denote the indicator random variable for event  $\mathcal{E}$ . Considering the second term above, we see that

$$\begin{aligned} \sum_{i=1}^k \mathbb{E} (X_{t_i(c)} - c)^+ &= \sum_{i=1}^k \sum_{j=1}^n \mathbb{E} [(X_j - c)^+ I(t_i(c) = j)] \\ &= \sum_{j=1}^n \mathbb{E} [(X_j - c)^+ I(t_k(c) > j - 1)] \\ &\geq \sum_{j=1}^n \mathbb{E} (X_j - b^*/k)^+ \cdot \Pr[t_k(c) > n - 1]. \end{aligned}$$

Since the sum in the last line is precisely  $b^*$ , we see our choice of  $c$  gives  $\sum_{i=1}^k \mathbb{E} [X_{t_i(c)}] \geq kc \geq a^*$  as claimed. ■

We now have the necessary results in place to proceed with the proof of Theorem 10.

*Proof of Theorem 10.* We prove our revenue bound via virtual values. We assume that all the  $F_i$  are regular. In fact, we see that we choose prices which do not distinguish (for a given agent) between differing valuations that yield the

same virtual value, and so by Proposition 18 may use ironed virtual values in the irregular case to achieve the same result.

We cannot immediately apply Theorem 24, since virtual values can, in general, be negative; so we consider  $(\phi_i)^+$  in place of  $\phi_i$ . (Note that Myerson's mechanism never selects an agent with a negative virtual value, and neither will our mechanism.)

Let  $c$  be the threshold from applying Theorem 24 to the random variables  $(\phi_i)^+$ , and  $p_i = \inf\{v : \phi_i(v) \geq c\}$ . Then, we can see that the expected revenue of an OPM using these prices is

$$\mathcal{R}_{\mathbf{p}}^{\text{obl}} = \mathbb{E} \left[ \sum_{i=1}^k (\phi_{t_i(c)})^+ \right] \geq \frac{1}{2} \mathbb{E} \left[ \sum_{i=1}^k (\phi_{(i)})^+ \right] = \mathcal{R}^{\mathcal{M}}/2.$$

**An example with a gap of 2.** We now show that OPMs cannot approximate the optimal revenue to within a factor better than 2 even in the single-item setting. Consider a seller with one item and two agents. The first agent has a fixed value of 1. The second has a value of  $1/\epsilon$  with probability  $\epsilon$  and 0 otherwise, for some small constant  $\epsilon > 0$ . Then, the optimal mechanism can obtain a revenue of  $2 - \epsilon$  by first offering a price of  $1/\epsilon$  to the second agent, and then a price of 1 to the first if the second declines the item. On the other hand, if the mechanism is forced to offer the item to the first agent first, then it has two choices: (1) offer the item at price 1 to agent 1; the agent always accepts, and (2) skip agent 1 and offer the item at price  $1/\epsilon$  to agent 2; the agent accepts with probability  $\epsilon$ . In either case, the expected revenue of the mechanism is 1.

### D.3 Graphical matroids

*Proof of Theorem 12.* Our technique here is to partition the elements of the matroid such that we may treat each part as a 1-uniform matroid yet still respect the original feasibility constraint, and achieve good revenue while doing so.

Let  $G = (V, [n])$  be the graph defining our matroid constraint. Let  $\delta(v)$  denote the set of edges incident on a vertex  $v$ , and for each  $v \in V$  define  $q_v = \sum_{i \in \delta(v)} q_i$ . Now, we can see that

$$\sum_{v \in V} q_v = \sum_{i \in E} 2q_i \leq 2(|V| - 1),$$

which can only hold if there exists  $v$  such that  $q_v \leq 2$ ; let  $\delta(v)$  be one of our partitions. Furthermore, the edge set  $\delta(v)$  forms a cut in  $G$ , and so given an independent set of edges from  $E - \delta(v)$  we may add any single edge from  $\delta(v)$  while retaining independence. We apply this argument recursively to  $(V - v, E - \delta(v))$  to form the rest of our partition. At the end, we have a partition of  $E$  such that each part has total mass no more than 2, and any collection of edges using no more than one edge from each part is independent.

We now show that within each part  $P$ , we can achieve expected revenue at least a third of what Myerson's mechanism received from that part, via an application of Theorem 24. Note that the revenue achieved by offering an agent  $i$  a price of  $p_i$  is a random variable, and these random variables are nonnegative and independent. Furthermore, we can successfully apply a threshold rule to them – we only offer to an agent if  $p_i$  is above our threshold, and they accept if and only if our stopping rule would have chosen this agent.

Let  $p = \sum_{i \in P} q_i (p_i - p)^+$ . Then  $p$  is precisely the upper bound on acceptable thresholds for Theorem 24 applied to our specified random variables, allowing one choice. From the proof of that theorem, we can see that applying the threshold  $p$  results in an expected revenue of at least  $p$ ; on the other hand,

$$\sum_{i \in P} p_i q_i \leq \sum_{i \in P} q_i (p + (p_i - p)^+) = p(1 + \sum_{i \in P} q_i) \leq 3p.$$

### D.4 Matroid intersections

*Proof of Theorem 13.* We describe the mechanism which achieves a 6.75-approximation when the distributions are regular. Appendix E sketches the extension to the non-regular case.

Let  $q_i = q_i^M/3$  and  $p_i = F_i^{-1}(1 - q_i)$ . Note that  $p_i \geq p_i^M$ . The mechanism serves agents in any arbitrary order, but offers a price of  $p_i$  for agent  $i$ .

Let  $c_i$  denote the probability that agent  $i$  is considered for service. We prove that  $c_i \geq 4/9$  for all  $i$ . This would prove the theorem, as the expected revenue is

$$\mathcal{R}_{\mathbf{p}}^{\text{obl}} = \sum_i c_i p_i q_i \geq \sum_i (4/9) p_i^M (q_i^M/3) \geq \sum_i (1/6.75) p_i^M q_i^M.$$

Let  $\mathcal{M}_1, \mathcal{M}_2$  be the two partition matroids. Let agent  $i$  be in partition  $P_1$  of  $\mathcal{M}_1$  and in partition  $P_2$  of  $\mathcal{M}_2$ . Let  $k_1$  be the maximum number of elements in  $P_1$  that can be present in an independent set of  $\mathcal{M}_1$  and let  $k_2$  be the maximum number of elements in  $P_2$  that can be present in an independent set of  $\mathcal{M}_2$ . We then have that for  $j = 1, 2$  that the expected number of agents in  $P_j$  desiring service is

$$\sum_{i \in P_j} q_i \leq k_j/3.$$

Define  $\mathcal{E}_j$  to be the event that at most  $k_j - 1$  agents from  $P_j$  desire service for  $j = 1, 2$ ; then agent  $i$  is always considered for service when events  $\mathcal{E}_1$  and  $\mathcal{E}_2$  both happen. By Markov's inequality, it is clear that  $\Pr[\bar{\mathcal{E}}_1] \leq 1/3$  and  $\Pr[\bar{\mathcal{E}}_2] \leq 1/3$ . So we may conclude that

$$c_i \geq \Pr[\mathcal{E}_1 \cap \mathcal{E}_2] = \Pr[\mathcal{E}_1] \cdot \Pr[\mathcal{E}_2 | \mathcal{E}_1] \geq (2/3) \cdot (2/3),$$

and the claim follows.

## D.5 Order-oblivious pricings in the non-matroid setting

In this section we present an example with a non-matroid constraint for which the revenue obtained by ordering the agents in the optimal way is a factor of  $\Omega(\log n / \log \log n)$  larger than that obtained by ordering the agents in the least optimal way.

**Lemma 25** *There exists an instance of the single-parameter mechanism design problem with a non-matroid feasibility constraint, along with two orderings  $\sigma_1$  and  $\sigma_2$  such that the revenue of the optimal SPM using ordering  $\sigma_1$  is a factor of  $\Omega(\log n / \log \log n)$  larger than that of the optimal SPM using ordering  $\sigma_2$ .*

*Proof:* Consider the following example. Construct a complete  $m$ -ary tree of height  $m + 1$ , and place a single agent at each node other than the root. The agents' valuations are i. i. d. , where any agent has a valuation of  $m$  with probability  $1/m$ , and a valuation of 0 otherwise. Our constraint on serving the agents is that we may serve any set of agents that lie along a single path from the root of the tree to some leaf – it is easy to verify that this is downward-closed.

Consider what happens when we may serve the agents in order by level from the root of the tree to the leaves. At each level of the tree, we may offer to serve at least  $m$  different agents, regardless of the outcome on previous levels. Since we may never sell to more than one agent per level, our revenue is either 0 or  $m$  on each level. We get a revenue of 0 if and only if every agent has a valuation of 0; this occurs with probability at most

$$(1 - 1/m)^m \leq 1/e,$$

and thus our expected revenue overall is at least

$$m^2 \cdot (1 - 1/e) = \Omega(m^2).$$

On the other hand, if we must serve the agents in order by level from the leaves of the tree to the root, then the first agent we serve commits us to a specific path. So we cannot hope to achieve revenue better than  $m$  for this specific node, plus the revenue expected revenue for an arbitrarily chosen path. Since each agent has an expected valuation of 1, this is bounded by

$$m + (m - 1) \cdot 1 = O(m).$$

Thus, the difference in revenue between the described orderings is  $\Omega(m)$ ; since the total number of agents is  $n = O(m^m)$ , in terms of  $n$  this gap is  $\Omega(\log n / \log \log n)$ . ■

## E Approximation in the non-regular case

We now sketch changes required to the theorems proved in Sections 4 and 5 to obtain the same approximations in the non-regular case.

From Lemma 2 we know that for every  $i$ , there exist prices  $\underline{p}_i$  and  $\overline{p}_i$  with corresponding probabilities  $\underline{q}_i = 1 - F_i(\underline{p}_i)$  and  $\overline{q}_i = 1 - F_i(\overline{p}_i)$ , as well as a number  $x_i$  such that  $x_i \underline{q}_i + (1 - x_i) \overline{q}_i = q_i^M$ , and Myerson's expected revenue is bounded by  $\sum_i (x_i \underline{p}_i \underline{q}_i + (1 - x_i) \overline{p}_i \overline{q}_i)$ . In fact this holds more generally. Let  $q_i$  be any probability less than 1. Then there exist probabilities  $\underline{q}_i$  and  $\overline{q}_i$ , and a number  $x_i \in [0, 1]$  with  $q_i = x_i \underline{q}_i + (1 - x_i) \overline{q}_i$ , such that for  $\underline{p}_i = F_i^{-1}(1 - \underline{q}_i)$ ,  $\overline{p}_i = F_i^{-1}(1 - \overline{q}_i)$ , and  $p_i$  defined as

$$\frac{x_i \underline{p}_i \underline{q}_i + (1 - x_i) \overline{p}_i \overline{q}_i}{q_i},$$

the optimal revenue achievable by selling an item with probability  $q_i$  to agent  $i$  is no more than  $p_i q_i$ .

Now consider a hypothetical situation in which the probability that agent  $i$  accepts a price of  $p_i$  is exactly  $q_i$ , and consider running an SPM/OPM with prices  $p_i$  that is  $\alpha$ -approximate with respect to  $\sum_i p_i q_i$ . We claim that we obtain an  $\alpha$ -approximation by using the same SPM/OPM but instead picking a price of  $\underline{p}_i$  with probability  $x_i$  and  $\overline{p}_i$  with probability  $1 - x_i$ .

To prove the claim we first note that we can defer the process of picking a price for agent  $i$  until the mechanism decides to offer the agent some price. In this case, the probability that the agent accepts the offered price is exactly  $q_i$ , and the revenue obtained from the agent conditioned on serving him is exactly  $p_i$ . Therefore, the probability that the mechanism makes an offer to an agent is also identical to the corresponding probability in the hypothetical deterministic mechanism, and the expected revenue of the mechanism is exactly the same as that of the hypothetical mechanism.

## F Computing the near-optimal posted-price mechanisms

We now describe how to compute the approximately optimal OPMs and SPMs designed in Sections 4 and 5. We assume that we are given access to the following oracles and algorithms:

- An algorithm to compute the optimal price to charge to a single-parameter agent given the agent's value distribution. Note that given such an algorithm and some value  $x$ , we can modify it to return the optimal price in the range  $[x, \infty)$  to charge the agent.
- An oracle that given a value  $v$  and index  $i$  returns  $F_i(v)$  and  $f_i(v)$ , as well as, given a probability  $\alpha$  returns  $F_i^{-1}(\alpha)$ . Note that the oracle can be used to compute the virtual value  $\phi_i(v)$ .
- An oracle for computing ironed virtual values in order to compute the approximately optimal SPM for non-regular distributions.
- An algorithm to maximize social welfare over the given feasibility constraint in order to be able to compute the outcome of Myerson's mechanism.

All of the mechanisms designed by us require computing the probabilities  $q_i^M$ . We first show how to estimate these probabilities within small constant factors:

1. Let  $\epsilon = 1/3n$ . Sample  $N = 4n^4 \log n / \epsilon^2$  value profiles from  $F_1 \times F_2 \times \dots \times F_n$ . For each sample, compute the (ironed) virtual value for each agent, and use these to compute the outcome of Myerson's mechanism for that value profile.
2. Estimate the probabilities  $q_i^M$  using the samples. Call the estimates  $\widehat{q}_i^M$ .
3. If  $\widehat{q}_i^M < 1/n^2$ , set  $\widehat{q}_i = 1/n^2$ , else set  $\widehat{q}_i = \widehat{q}_i^M / (1 - \epsilon)$ . Compute for each  $i$  the value  $\widehat{p}_i = F_i^{-1}(1 - \widehat{q}_i)$ .
4. Find the optimal price in the range  $[\widehat{p}_i, \infty)$  to charge to agent  $i$ . Call it  $p_i$ . Let  $q_i = 1 - F_i(p_i)$ .

5. Output the prices computed in the last step and order the agents in order of decreasing prices.

In order to analyse the performance of this approach, we compare it to a mechanism that charges agent  $i$  the price  $p_i^M = F_i^{-1}(1 - q_i^M)$  but uses the same ordering as the mechanism above. We first show that the probabilities  $q_i$  closely estimate the probabilities  $q_i^M$ .

**Lemma 26** *With probability at least  $1 - 2/n$ , we have  $\widehat{q}_i \in [q_i^M, (1 + 3\epsilon)q_i^M + 2/n^2]$ .*

*Proof:* First, for any  $i$  with  $q_i^M \geq 1/n^4$ , using Chernoff bounds we get that

$$\Pr[|\widehat{q}_i^M - q_i^M| \geq \epsilon q_i^M] \leq 2e^{-\epsilon^2 q_i^M N/2} \leq 2/n^2$$

$\widehat{q}_i^M \in (1 \pm \epsilon)q_i^M$  in turn implies by definition that  $q_i^M \leq \widehat{q}_i \leq (1 + \epsilon)/(1 - \epsilon)q_i^M \leq (1 + 3\epsilon)q_i^M$ . Therefore we have  $\widehat{q}_i \in [q_i^M, (1 + 3\epsilon)q_i^M]$ . On the other hand, for  $q_i^M < 1/n^4$ , by Markov's inequality, with probability  $1 - 1/n^2$ ,  $\widehat{q}_i^M < 1/n^2$ , and so  $\widehat{q}_i \in [q_i^M, 1/n^2]$ . The lemma now follows by employing the union bound. ■

Furthermore, conditioned on the event defined in the statement of the above lemma (call it  $\mathcal{E}$ ), since  $p_i^M$  lies in the range  $[\widehat{p}_i, \infty)$ , we have that  $q_i^M p_i^M \leq q_i p_i$ . This implies that the prices  $p_i$  give a good estimate on the revenue of Myerson's mechanism.

Next, we compare the real mechanism  $\mathcal{S}$  with prices  $p_i$  to the theoretically good mechanism  $\mathcal{S}'$  that charges prices  $p_i^M$ . Let  $S$  be the set of agents for which  $\widehat{q}_i^M < 1/n^2$ . The probability that any of these agents is offered service in  $\mathcal{S}$  is at most  $1/n$ . Conditioned on this event not happening, the probability that an agent is made an offer in  $\mathcal{S}$  is no smaller than its counterpart in  $\mathcal{S}'$ . Moreover, conditioned on being made an offer, the revenue from an agent  $i$  is  $q_i p_i \geq q_i^M p_i^M$ .

Therefore, conditioned on the event  $\mathcal{E}$ , the expected revenue of  $\mathcal{S}$  is at least a  $(1 - 1/n)$  fraction of the expected revenue of  $\mathcal{S}'$ . But the event  $\mathcal{E}$  happens with probability  $1 - 2/n$ , therefore, we get a  $(1 - o(1))$  approximation to the expected revenue of  $\mathcal{S}'$ .

## G Approximations for the BMUMD

We first prove that there exists a good OPM for instances of the BMUMD involving a feasibility constraint that is the intersection of a graphical matroid and the agents' unit-demand constraints.

*Proof of Theorem 15.* Note that though the feasibility constraint we are facing is the intersection of a graphical matroid and partition matroid (from the unit demand constraint), we can view the situation as if we were in the intersection of two partition matroids. This follows from the proof of Theorem 12, where we see that a graphical matroid can be seen as a union of 1-uniform matroids, which is a partition matroid. The total probability mass of the elements of each 1-uniform matroid is at most 2. Thus, if we sell at prices for which the probability of an agent  $i$  desiring the item is  $q_i^M/4$ , then with a probability of at least  $1/2$  no more than 1 agent will desire service in the 1-uniform matroid which contains  $i$  and with a probability of at least  $3/4$  no more than 1 item is desired by the agent  $i$ . Thus the revenue obtained gives an approximation factor of  $4 \cdot 4/3 \cdot 2 = 32/3 \approx 10.67$ .

Next we prove that for the two settings discussed in Section 6.2, we can design an SPM that achieves a good approximation via implementation in undominated strategies.

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  $\alpha$ -approximation in undominated strategies [4] 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  $\alpha$  of the optimal, and every agent can easily compute for any dominated strategy a strategy that dominates it.



*Proof of Lemma 16.* Note that if agent  $i$  desires only one service  $j \in J_i$ , and refuses the service when offered, the agent obtains a utility of 0 regardless of others' strategies. On the other hand, the strategy of accepting the service when offered has strictly positive utility for some strategy profiles of others, therefore it dominates the previous strategy.

*Proof of Theorem 17.* We consider the matroid intersection setting first and assume that the valuation distributions are regular. The non-regular case is similar. Our SPM in this setting considers the hypothetical single-dimensional instance  $\mathcal{I}^{\text{copies}}$  defined in Section 3 and computes the probabilities  $q_j^M$  with which Myerson's mechanism allocates the service  $j$ . We then set  $q_j = q_j^M/2$  and  $p_i = F_j^{-1}(1 - q_j)$ . Note that for any  $i$ ,  $\sum_{j \in J_i} q_j \leq 1/2$ . Therefore, with probability at least  $1/2$ ,  $i$  desires no service other than  $j$  (we say that  $j$  is uniquely desired by  $i$ ). Lemma 16 shows that in this case, in any undominated strategy implementation, if  $i$  is offered  $j$  and desires it, then  $i$  accepts  $j$ .

For any particular run of the mechanism, divide the set of all services into three groups— $S$ , the set of *sold* services,  $B$  the set of services that are desired by their corresponding agents but “blocked” by services in  $S$ , and  $U$  the set of services that are desired by their corresponding agents and not in sets  $S$  or  $B$ . Then Lemma 16 implies that services in  $U$  are not uniquely desired. Now, the expected total price in the union of the sets  $S$ ,  $B$  and  $U$  is exactly  $\sum_j p_j q_j$ . Moreover, the expected total price in  $U$  is at most  $1/2 \sum_j p_j q_j$ . Finally, following the proof of Theorem 5, the expected total price in  $B$  conditioned on  $S$  is at most the total price contained in  $S$ . Therefore, putting everything together we get that the expected total price obtained from  $S$  is at least  $1/4 \sum_j p_j q_j$ . By our choice of  $\mathbf{p}$  and  $\mathbf{q}$ , this is an 8-approximation.

The argument for the combinatorial auction setting is identical and based on Theorem 8. We omit it for brevity.

## H Approximating social welfare and other objectives via posted-price mechanisms

We now show that our approach from Sections 4 and 5 in fact extends to the problem of maximizing any objective that is linear in social value and revenue via SPMs.

We start with some definitions. For all  $i \in [n]$  let  $g^i(v, p) = \alpha_i v + \beta_i p$  denote an arbitrary linear function of  $v$  and  $p$ . For a mechanism  $\mathcal{A}$  with payment rule  $\mathbf{p}$ , let  $G(\mathcal{A}, \mathbf{p})$  be the expected value of  $g$  over the outcome of the mechanism, that is,  $G(\mathcal{A}, \mathbf{p}) = \mathbb{E}_{\mathbf{v}}[\sum_{i \in \mathcal{A}(\mathbf{v})} g^i(v_i, p_i)]$ . Define the virtual value of  $i$  with respect to  $g^i$  to be

$$\phi_i^G(v) = (\alpha_i + \beta_i)v - \beta_i \frac{1 - F_i(v)}{f_i(v)}$$

and the virtual surplus with respect to  $G$  of a set  $S$  of agents to be  $\Phi^G(S) = \sum_{i \in S} \phi_i^G(v_i)$ . Then, the lemma below follows from standard techniques, and allows us to ignore the payment function in trying to maximize  $G$ .

**Lemma 27** *For any truthful mechanism  $\mathcal{A}$  with payment rule  $\mathbf{p}$ , the expected virtual surplus with respect to  $G$  of  $\mathcal{A}$  is equal to the expected value of  $G$  for  $\mathcal{A}$ 's outcome. That is,*

$$G(\mathcal{A}, \mathbf{p}) = \mathbb{E}_{\mathbf{v}}[\Phi^G(\mathcal{A}(\mathbf{v}))]$$

The lemma suggests that a mechanism  $\mathcal{M}^G$  with allocation rule  $\mathcal{M}^G(\mathbf{v}) = \arg\max_S \Phi^G(S)$  maximizes  $G$  over the class of all truthful mechanisms. However, as for revenue-maximizing mechanisms, in order for this mechanism to be truthful, the distributions  $F_i$  must satisfy a certain regularity condition.

**Definition 4** *A one dimensional distribution  $F$  is regular with respect to function  $G$ , if  $\phi^G(v)$  is monotone non-decreasing in  $v$ .*

The following theorem is straightforward:

**Theorem 28** *If for all  $i$ ,  $F_i$  is regular with respect to  $G$ , the mechanism  $\mathcal{M}^G$  defined above is truthful and obtains the maximum value of  $G$  over the class of all truthful mechanisms.*

In order to optimize  $G$  over the class of SPMs in the matroid setting, we follow an approach similar to the one in Section 4. Other approximations are similar. We focus on the regular setting. Our approximately optimal mechanism is defined as follows. Let  $\mathcal{M}^G$  denote the optimal mechanism in Theorem 28 above. Let  $q_i^G$  denote the probability that  $\mathcal{M}^G$  serves agent  $i$ . Define for all  $i$

$$\begin{aligned} q_i &= q_i^G, \\ p_i &= F_i^{-1}(1 - q_i), \quad \text{and,} \\ \gamma_i &= \left( \int_{p_i}^{\infty} \phi_i^G(v_i) f_i(v_i) dv_i \right) / q_i \end{aligned}$$

The SPM sets a price of  $p_i$  for agent  $i$  and offers to serve the agents in decreasing order of their corresponding  $\gamma_i$ 's. The  $\gamma_i$  reflects the expected virtual value we get from agent  $i$  upon serving the agent. We denote this mechanism by  $\mathcal{S}^G$ .

We first note that the performance of  $\mathcal{S}^G$  can be bounded in terms of the  $\gamma_i$ 's. In particular, Lemma 27 and the definition of  $\gamma_i$  imply that

$$G(\mathcal{S}^G) = \mathbb{E}_{\mathbf{v}} \left[ \sum_{i \in S} \gamma_i \right]$$

where  $S$  is the set of agents that are allocated service. Following the argument for Theorem 5 we infer that since agents are ordered in decreasing order of  $\gamma_i$ ,  $\mathbb{E}_{\mathbf{v}} \left[ \sum_{i \in S} \gamma_i \right] \geq \frac{1}{2} \sum_i q_i \gamma_i$ . In order to complete our argument, we bound the performance of  $\mathcal{M}^G$  in terms of the  $\gamma_i$ 's.

**Lemma 29** *If for all  $i$ ,  $F_i$  is regular with respect to  $G$ , then  $G(\mathcal{M}^G) \leq \sum_i \gamma_i q_i$ .*

*Proof:* Let us consider the contribution of agent  $i$  to the objective function value for  $\mathcal{M}^G$ . This is no more than the objective function value achieved by an optimal mechanism that sells only to  $i$  and with probability at most  $q_i$ . By the definition of  $\Phi^G$  and using regularity, this is exactly  $\int_{p_i}^{\infty} \phi_i^G(v_i) f_i(v_i) dv_i$  where  $p_i = F_i^{-1}(1 - q_i)$ . Finally, the integral is exactly equal to  $\gamma_i q_i$  by the definition of  $\gamma_i$ . ■

We therefore have the following theorem:

**Theorem 30** *The mechanism  $\mathcal{S}^G$  defined above obtains a 2-approximation to the objective  $G$  in the matroid case when all the input distributions are regular with respect to  $G$ .*

Similar techniques prove analogues of other theorems in Sections 4 and 5 for arbitrary functions  $G$ . Finally, we note that if the distributions are not regular as defined in Definition 4, we can apply an ironing procedure to the virtual values in much the same way as in Myerson's approach. We leave the details to the reader.

## I Revenue maximization through VCG mechanisms

A consequence of our constant-factor approximation to revenue through SPMs is that in matroid settings VCG mechanisms with appropriate reserve prices are near-optimal in terms of revenue. This follows from noting, as we show below, that VCG mechanisms perform no worse in terms of expected revenue than SPMs with the same reserve prices. Although VCG mechanisms aim to maximize the social welfare of the outcome, setting high enough reserve prices allows them to also obtain good revenue.

Formally, a Vickrey-Clarke-Groves (VCG) mechanism  $\mathcal{V}^{\mathbf{p}}$  with reserve prices  $\mathbf{p}$  serves the set  $S$  of agents, with  $v_i \geq p_i$  for all  $i \in S$ , that maximizes  $\sum_{i \in S} v_i$ .

Hartline and Roughgarden [13] show that in several single-parameter settings the VCG mechanism with monopoly reserve prices gives a constant factor approximation to revenue. This result holds when all the value distributions satisfy the so-called monotone hazard rate condition, or with a matroid feasibility constraint when all the value distributions are regular. Their result does not extend to the case of matroids with general (non-regular) value distributions. One of the main questions left open by their work is whether there is some set of reserve prices (not necessarily equal

to the monopoly reserve prices) for which the VCG mechanism gives a constant factor approximation to revenue in the matroid setting with general value distributions. We answer this question in the positive. We use the following fact about matroids.

**Proposition 31** *Let  $B_1$  and  $B_2$  be any two independent sets of equal size in a matroid set system  $\mathcal{J}$ . 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  $\mathcal{J}$ .*

**Theorem 32** *For any instance of the single-parameter Bayesian mechanism design problem with a matroid feasibility constraint, there exists a set of reserve prices such that the expected revenue of the VCG mechanism with those reserve prices is at least half of the expected revenue of Myerson's mechanism.*

*Proof:* We prove that when the set system  $\mathcal{J}$  is a matroid, for any collection of prices  $\mathbf{p}$ , the revenue of the SPM  $\mathcal{S}^{\mathbf{p}}$  is no more than the revenue of the VCG mechanism  $\mathcal{V}^{\mathbf{p}}$ . The result then follows from Theorem 5.

Fix a value vector  $\mathbf{v}$  and let  $A$  denote the set served by  $\mathcal{S}^{\mathbf{p}}$  and  $B$  denote the set served by  $\mathcal{V}^{\mathbf{p}}$ . Then, since both mechanisms serve a maximal independent set among the set of agents with  $v_i \geq p_i$ , we have  $|A| = |B|$ . Proposition 31 then implies the existence of a bijection  $g$  such that for all  $e \in B \setminus A$ ,  $B \setminus \{e\} \cup \{g(e)\}$  is independent. This implies that  $\mathcal{V}^{\mathbf{p}}$  charges  $e$  a price of at least the value of  $g(e)$ , which is at least the reserve price  $p_{g(e)}$ . On the other hand, by definition, the price charged to any  $e \in B \cap A$  is at least  $p_e$ . Therefore, the revenue of  $\mathcal{V}^{\mathbf{p}}$  in this case is at least  $\sum_{e \in B \cap A} p_e + \sum_{e \in B \setminus A} p_{g(e)} = \sum_{e \in A} p_e$  which is equal to the revenue of  $\mathcal{S}^{\mathbf{p}}$ . ■