

Multi-parameter Mechanism Design and Sequential Posted Pricing

Shuchi Chawla*
Computer Sciences Dept.
Univ. of Wisconsin - Madison
shuchi@cs.wisc.edu

David L. Malec*
Computer Sciences Dept.
Univ. of Wisconsin - Madison
dmalec@cs.wisc.edu

Jason D. Hartline†
EECS
Northwestern University
hartline@eecs.northwestern.edu

Balasubramanian Sivan*
Computer Sciences Dept.
Univ. of Wisconsin - Madison
balu2901@cs.wisc.edu

ABSTRACT

We study the classic 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 [20]. 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) [25].

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 [9] to the setting of multiple agents where the designer has combinatorial feasibility constraints on which agents can simultaneously obtain each service.

*This author was supported in part by NSF award number CCF-0830494 and in part by a Sloan Foundation fellowship.

†This author was supported in part by NSF award number CCF-0830773.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

STOC'10, June 5–8, 2010, Cambridge, Massachusetts, USA.
Copyright 2010 ACM 978-1-4503-0050-6/10/06 ...\$10.00.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics; F.2.0 [Analysis of Algorithms and Problem Complexity]: General

General Terms

Algorithms, Economics, Theory

Keywords

Bayesian mechanism design

1. INTRODUCTION

Suppose that 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 incentive-compatible mechanisms [20] with practical (often non-incentive-compatible) implementations [1]. The challenge posed by 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 [19, 22, 25]. Second, in multi-dimensional settings optimal mechanisms 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., [13]) to settings with multiple agents; it is unrelated to the standard computational questions of *algorithmic mechanism design* (e.g., [18, 21]). The central problem in algorithmic pricing can be viewed (for the most part) as Bayesian revenue maximization in a single agent setting (e.g., [13]). Algorithmic pricing is hard to approximate when the agent’s values for different outcomes are generally correlated [7]; however, when the values are independent there is a 3-approximation [9]. 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 [18]).

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 [15]. 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 [12]): 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, implying 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 incentive compatible 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 representatives 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 robust to the actions of any agent with such a strategic option; if all agents with dominant strategies follow said strategies then our performance is a constant fraction of the original SPM’s performance. (This is a non-standard notion of *dominant strategy implementation*.) Given the advantages of SPMs over standard dominant strategy mechanisms, these partially dominant strategy mechanisms may be more practically relevant.

Finally, we note that most of our results for posted pricings are constructive and there are 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.¹ 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

¹This 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.

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 [25] and references therein for work in economics on optimal multi-dimensional mechanism design. See [9] and references therein for work in computer science on multi-dimensional pricing for a single agent. We extend the setting from [9] 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 [24] show experimentally that these mechanisms compare favorably to Myerson’s optimal mechanisms. Blumrosen and Holenstein [6] 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 [5, 16, 4].

The question of whether simple mechanisms can achieve near-optimal revenue was considered recently by Hartline and Roughgarden [14]. 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, 17], 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.

2. PRELIMINARIES

2.1 Bayesian settings for mechanism design

The Bayesian multi-parameter unit-demand mechanism design problem (BMUMD for short) is an abstraction of the setting where a seller can provide a number of services to agents where each agent desires at most one service. An agent’s value for each service may be different, but is drawn independently from known distributions. Formally we denote an instance of this problem as $(J, \mathcal{S}, \Pi, \mathbf{F})$ given as follows:

- $J = [m]$ is a set of m services.

- $\Pi = (J_1, \dots, J_n)$ is a partitioning of the services among the n agents.² The services in J_i are the ones being targeted at agent i . These agents are *unit demand* in that they each desire at most one service from their partition.
- $\mathcal{S} \subset 2^J$ is a feasibility constraint. It specifies the sets of services the seller can simultaneously provide. We assume \mathcal{S} is downward closed, i.e., for $S \in \mathcal{S}$ all subsets $S' \subset S$ are in \mathcal{S} . We assume \mathcal{S} respects the partitioning Π and the unit-demand constraint, i.e., $S \in \mathcal{S}$ and $i \in [n]$ satisfies $|S \cap J_i| \leq 1$.
- $\mathbf{F} = F_1 \times \dots \times F_m$ is the joint product distribution on agent values over the m services. I.e., v_j is drawn from distribution F_j with density function f_j .

The Bayesian single-parameter setting (BSMD for short) is an abstraction of the setting where each agent has a single private value for any “good outcome” of the mechanism. In this setting the mechanism can produce a good outcome for agent i in which case their valuation is v_i or a bad outcome in which case their valuation is zero. This can be represented as the special case of the multi-dimensional setting where there is exactly one service available to each agent, i.e., $n = m$ and $J_i = \{i\}$. We denote an instance of BSMD by the tuple $(J, \mathcal{S}, \mathbf{F})$.

Examples. In the multi-parameter hotel rooms example the set J of services is the set of edges in a bipartite graph between attendees and hotel rooms, J_i are the edges incident on attendee i , and \mathcal{S} is the set of all matchings in the graph. In the single-parameter hotel rooms example, where each attendee has a set of desired rooms and the same valuation for each, the set J corresponds to attendees and the feasible sets \mathcal{S} are sets of attendees that can be simultaneously matched, i.e., \mathcal{S} is the set of all independent sets of the transversal matroid corresponding to matchings in the bipartite graph between attendees and their desired hotel rooms.

2.2 Posted-price mechanisms

We will consider sequential posted-price mechanisms based on the following high-level protocol that is parameterized by \mathbf{p} , a vector of prices, one for each service, and σ , an ordering over the services.

The generic sequential posted pricing protocol for (\mathbf{p}, σ) is as follows:

1. Initialize $A \leftarrow \emptyset$.
2. For $j = 1$ through m , do:
 - (a) If $A \cup \{\sigma(j)\} \in \mathcal{S}$, offer service $\sigma(j)$ at price p_j .
 - (b) If the agent accepts, $A \leftarrow A \cup \{\sigma(j)\}$.
3. Provide the services in A to the corresponding agents.

We denote the revenue of this mechanism on valuation profile \mathbf{v} by $\mathcal{R}_{(J, \mathcal{S}, \Pi, \mathbf{F})}^{(\mathbf{p}, \sigma)}(\mathbf{v})$. In the strategically-simple single-parameter setting this revenue can be calculated with the assumption that service j is accepted by the agent when offered if $v_j \geq p_j$. We defer the discussion of incentives in the more complicated multi-parameter setting to Section 2.3.

It is clear that it is always better from the designer’s point of view to be able to choose the ordering σ . Unfortunately, this may not always be possible. We therefore distinguish between the following two kinds of posted-price mechanisms.

²This partitioning is for notational convenience only. 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.

Sequenced Posted-Price Mechanisms. An SPM is given by (\mathbf{p}, σ) . Its expected revenue is:

$$\mathcal{R}_{(J,S,F)}^{(\mathbf{p},\sigma)} = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\mathcal{R}_{(J,S,F)}^{(\mathbf{p},\sigma)}(\mathbf{v})]$$

Order-oblivious Posted-pricing Mechanisms. An OPM is given simply by the pricing \mathbf{p} where we allow the order over the services to be picked adversarially after the valuations of the agents are drawn. This pessimistically bounds the worst possible revenue for a given pricing. Formally:

$$\mathcal{R}_{(J,S,F)}^{\mathbf{p}} = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\min_{\sigma} \mathcal{R}_{(J,S,F)}^{(\mathbf{p},\sigma)}(\mathbf{v})]$$

When it is clear from the context we will omit the subscript (J, S, Π, \mathbf{F}) or (J, S, \mathbf{F}) .

In some settings we consider randomized versions of SPMs and OPMs where the pricing \mathbf{p} is picked randomly. In this case, we assume that the prices are drawn first and then the order σ is determined based on the prices (adversarially or by the designer).

2.3 Incentives

Most of the literature on mechanism design (especially in computer science) focuses on sealed-bid single-round direct-revelation mechanisms. These are mechanism that consist of two steps: first agents report bids, to be interpreted as their preferences over possible outcomes of the mechanisms, and second the mechanism selects an outcome and agent payments. In this context a mechanism is *incentive compatible* if each agent has a (weakly) dominant strategy of truthful reporting. It is assumed that agents report their true preferences in an incentive compatible mechanism.

Our posted price mechanisms do not take this general single-round form. Instead our mechanism will offer each agent a sequence of prices (and these offers may be arbitrarily interleaved among agents). Strategically, a bidder i when offered price p_j for $j \in J_i$ has two options. They can accept or reject the offer. An agent with value v_j for service j is *sincere* if they accept offers $p_j \leq v_j$ and reject offers $p_j > v_j$.

Sincere bidding is a dominant strategy for an agent only when the ordering σ respects the agent’s incentives. Formally, we say that an ordering σ is J_i -respecting if for all $j_1, j_2 \in J_i$, $v_{j_1} - p_{j_1} > v_{j_2} - p_{j_2} \geq 0$ implies $\sigma^{-1}(j_1) < \sigma^{-1}(j_2)$. That is, the offers made to agent i are ordered by decreasing utility for the agent (although they may be interleaved arbitrarily with offers for other agents). An ordering is Π -respecting if it is J_i -respecting for all i . The following lemma formalizes the connection between sincere bidding and Π -respecting orderings.

LEMMA 1. *For \mathbf{v}_{J_i} (the values of agent i) sincere bidding is a (weakly) dominant strategy for i in sequential posted pricing (\mathbf{p}, σ) if and only if σ is J_i -respecting.*

It is easy to see that if the condition on σ is met then an agent will have no reason not to respond sincerely. If the condition is not met, an agent may strategize in the following way. When offered item j' at a desirable price $p_{j'}$, the agent might reject the offer in hopes of later being offered item j for which the agent has even higher utility.

The condition of the lemma is met in three special cases of interest:

1. When the agent is single dimensional, i.e., $|J_i| = 1$.
2. When the agent has positive utility for at most one service, i.e., $|\{j \in J_i : v_j - p_j > 0\}| \leq 1$

3. When the agent can choose the relative order of σ on J_i .

Given the last point, for an OPM in the multi-parameter setting we assume that orderings of interest are Π -respecting and define the worst-case revenue of the mechanism accordingly:

$$\mathcal{R}_{(J,S,\Pi,F)}^{\mathbf{p}} = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\min_{\sigma: \sigma \text{ is } \Pi\text{-respecting}} \mathcal{R}_{(J,S,\Pi,F)}^{(\mathbf{p},\sigma)}(\mathbf{v}) \right]$$

In SPMs in multi-parameter settings, the ordering is not necessarily Π -respecting. We assume that all bidders for whom sincere bidding is a (weakly) dominant strategy indeed bid sincerely. We derive robust bounds on our mechanism performance in the presence of arbitrary manipulations of agents that do not have dominant strategies. This is a weaker form of implementation in dominant strategies, that we call *partial dominant strategy implementation*.

DEFINITION 1. *A mechanism is a partial dominant-strategy (PDS) implementation of a desired objective if that objective is met whenever every agent with a (weakly) dominant strategy plays that strategy (with other agents behaving arbitrarily).*

A mechanism is an α -approximation in partial dominant-strategy equilibrium (PDSE) of a desired objective if it obtains an α approximation to the objective in expectation whenever every agent with a (weakly) dominant strategy plays that strategy.

As is standard in mechanism design, we assume that the agents understand the mechanism. As our mechanisms are parameterized by prices \mathbf{p} , it is assumed that agents know these prices in advance. This assumption is only necessary for SPMs in multi-parameter settings where an agent i must know whether there is a future offer p_j for $j \in J_i$ such that $v_j - p_j > 0$.

2.4 Myerson’s optimal mechanism

Myerson [20] describes the revenue maximizing mechanism for the Bayesian single-parameter mechanism design problem. *Virtual valuations* are given by the formula $\phi(v) = v - \frac{1-F(v)}{f(v)}$. When the value distributions F_i are *regular*, i.e., virtual valuations are monotone in valuations, the optimal mechanism first computes 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 [20].

When the distributions F_i are irregular, that is, virtual valuations are not monotone in valuations, Myerson’s mechanism as described above will no longer be truthful. Myerson addressed this case by “ironing” the virtual valuation function and converting it into a monotone non-decreasing function called the ironed virtual value function denoted by $\bar{\phi}(v)$. We skip the description of this procedure; the reader is referred to [8, 9] for details.

We use $\mathcal{R}_{\mathcal{I}}^M$ to denote the expected revenue of Myerson’s mechanism on a single-parameter instance \mathcal{I} . For our analyses, we primarily require the following three characterizations of the expected revenue of any incentive compatible mechanism. See the appendix for a proof of Lemma 4.

PROPOSITION 2. [20] *When the distributions F_i are regular, the expected revenue of any incentive compatible single-parameter mechanism M is equal to its expected virtual surplus.*

PROPOSITION 3. [20] *The expected revenue of any incentive compatible 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.*

LEMMA 4. If F_i is regular for each i , the revenue of any incentive compatible mechanism M over the n agents 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 $\sum_i x_i \underline{p}_i \underline{q}_i + (1 - x_i) \overline{p}_i \overline{q}_i$.

2.5 Matroids and related set systems

Many of our techniques work for feasibility constraints \mathcal{S} that are matroids or close to matroids. We define these set systems here. The set system (X, \mathcal{S}) over a universe X is called a matroid if it satisfies the following conditions:

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

Sets in \mathcal{S} 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. The rank of a set $S \subseteq X$ is defined to be the size of any maximal independent subset of S . The *span* of a set $S \subseteq X$, $\text{sp}(S)$, is the maximal set $T \supseteq S$ with $\text{rank}(T) = \text{rank}(S)$.

A k -uniform matroid on the universe X is a matroid where every subset of X of size at most k is independent. A partition matroid (X, \mathcal{S}) is a union of two or more uniform matroids $\{(X_i, \mathcal{M}_i)\}_i$, where $\{X_i\}_i$ is a partition of X and $\mathcal{S} = \{\cup_i A_i : A_i \in \mathcal{M}_i \forall i\}$.

A set system (X, \mathcal{S}) is called a matroid intersection if there are two (or more) matroids (X, \mathcal{M}_1) and (X, \mathcal{M}_2) , such that $\mathcal{S} = \mathcal{M}_1 \cap \mathcal{M}_2$. An example of a matroid intersection is a matching in a bipartite graph.

Throughout this paper, the universe X is the set J of all services.

2.6 Computing the posted prices

For all but one of the approximately-optimal posted-price mechanisms that we present, prices and orderings can be computed efficiently in a computational model where we have black box access to the distribution \mathbf{F} . Please see the full version of this paper [11] for details.

3. REDUCING BMUMD TO BSMD

We now present a general reduction from the multi-parameter unit-demand mechanism design problem to the single-parameter problem. Using this reduction we can argue that if there exists an approximately optimal sequential posted-price mechanism in the single-parameter setting, there also exists one in the original multi-parameter setting. Understanding the properties of optimal mechanisms in multi-parameter settings is tricky so our approach is based on upper and lower bounds for single-parameter settings.

There are four main steps to give and instantiate our reduction. They are:

1. (Analogy) Define a single-parameter analog for any multi-parameter setting.
2. (Lower bound) Show that the revenue of the optimal single-dimensional analog is at least the revenue of the optimal multi-dimensional mechanism.
3. (Reduction) Show that if we had a sequential posted pricing for the single-dimensional analog, we can convert it into

a posted pricing for the multi-dimensional setting without much loss in performance.

4. (Instantiation) Show for a given multi-dimensional setting that there exist sequential pricings that approximate the optimal mechanism for the single-parameter analog.

We give the analogy and lower bound here, as well as a reduction for OPMs. Section 4 instantiates this reduction in various settings of interest. While we do not obtain a general purpose reduction from SPMs, as we do for OPMs, in Section 5 we describe SPMs that obtain approximate optimality in PDSE in many single as well as multi-parameter settings.

The analogy

The main concept behind our reduction is a single-parameter analog. Consider an instance $\mathcal{I} = (J, \mathcal{S}, \Pi, \mathbf{F})$ of the BMUMD with $n = |\Pi|$ agents and $m = |J|$ services. The *single-parameter analog* is the setting we get when we assume that each service is demanded by a distinct agent, i.e., $\mathcal{I}^{\text{reps}} = (J, \mathcal{S}, \mathbf{F})$. Formulaically, this analogy is trivial; intuitively, it replaces each agent i with $|J_i|$ distinct agents (called representatives or “reps” hereafter). Each rep is interested in a single service $j \in J_i$ and behaves independently of (and potentially to the detriment of) other reps. Notice that $\mathcal{I}^{\text{reps}}$ has $m = |J|$ agents and services.

Lower bound

Notice that $\mathcal{I}^{\text{reps}}$ is similar to \mathcal{I} except that it involves more competition (among different reps of the same multi-parameter agent). Therefore it is natural to expect that a seller can obtain more revenue in the instance $\mathcal{I}^{\text{reps}}$ than in \mathcal{I} . The following lemma formalizes this (see the appendix for a proof).

LEMMA 5. Let \mathcal{A} be any individually rational and incentive compatible 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 $\mathcal{I}^{\text{reps}}$.

A reduction for OPMs

The main advantage of a seller in the single-parameter analog is increased competition. Intuitively, if we can design mechanisms for the single-parameter setting that do not exploit competition, then it is reasonable to expect them to obtain similar performance in the multi-dimensional setting. Here, sequential posted pricings are exactly what is needed (see the appendix for a proof of the following).

THEOREM 6. If OPM \mathbf{p} is an α -approximation to the optimal mechanism for the single-parameter setting $\mathcal{I}^{\text{reps}}$ then it is an α -approximation in PDSE to the optimal mechanism for the multi-parameter setting \mathcal{I} .

4. APPROXIMATION THROUGH OPMs

In this section we instantiate the reduction described in Section 3 for several different settings. We begin by discussing the hotel rooms assignment setting described in Section 1 and giving a simple 6.75 approximation for this setting. We then describe a (tight) 2-approximation for a simpler “ k item auction” setting. We conclude this section by listing approximation factors for other settings of interest.

4.1 Intersections of partition matroids

Consider the instance $(J, \mathcal{S}, \mathbf{F})$ of the BSMD, where \mathcal{S} is an intersection of two partition matroids $(\{X_a\}, \mathcal{S}_1)$ and $(\{Y_b\}, \mathcal{S}_2)$,

with $\{X_a\}$ and $\{Y_b\}$ being partitions of the ground set $[m]$. Note that such an instance arises as the single-parameter analog $\mathcal{I}^{\text{reps}}$ of the hotel rooms assignment example discussed in the introduction; here the matroid \mathcal{S}_1 corresponds to the unit-demand constraint over the agents (i.e. any feasible allocation must contain at most one element in J_i for all i) and the matroid \mathcal{S}_2 corresponds to the supply constraints over hotel rooms (i.e. any feasible allocation must allocate any hotel room at most once).

THEOREM 7. *Let $\mathcal{I} = (J, \mathcal{S}, \mathbf{F})$ be an instance of the BSMD with \mathcal{S} being an intersection of two partition matroids. Then, there exists a set of prices \mathbf{p} such that $\mathcal{R}_{\mathcal{I}}^{\mathbf{p}} \geq \frac{1}{6.75} \mathcal{R}_{\mathcal{I}}^M$.*

PROOF. We first consider the setting where the distributions F_j are regular. Recall from Lemma 4 that in this case the revenue of the optimal truthful mechanism for \mathcal{I} is bounded by $\sum_j p_j^M q_j^M$, where q_j^M is the probability with which the optimal mechanism allocates service j and $p_j^M = F_j^{-1}(1 - q_j^M)$. We now describe the prices that achieve the claimed approximation. Let $q_j = q_j^M/3$. Our approximately optimal pricing is given by \mathbf{p} , where $p_i = F_j^{-1}(1 - q_j)$. Note that $p_j \geq p_j^M$.

Consider any ordering σ over the services, and let c_j denote the probability that service j gets offered under this ordering. We claim that $c_j \geq 4/9$ for all $j \in [m]$. Then, the expected revenue of the OPM \mathbf{p} is given by

$$\mathcal{R}_{\mathcal{I}}^{\mathbf{p}} = \sum_j c_j p_j q_j \geq \sum_j (4/9) p_j^M (q_j^M/3) \geq \sum_j (1/6.75) p_j^M q_j^M.$$

This along with Lemma 4 implies the theorem.

To prove the claim, let \mathcal{S}_1 and \mathcal{S}_2 be the two partition matroids with $\mathcal{S} = \mathcal{S}_1 \cap \mathcal{S}_2$. Let j belong to the partition $X_a \subset [m]$ in \mathcal{S}_1 and the partition $Y_b \subset [m]$ in \mathcal{S}_2 . Let k_a (resp. k_b) denote the maximum number of elements of X_a (resp. Y_b) that can be present in an independent set in \mathcal{S}_1 (resp. \mathcal{S}_2). Then, service j is offered if and only if the number of services in X_a that are offered prior to j in σ and are accepted is at most $k_a - 1$ and the number of services in Y_b that are offered prior to j in σ and are accepted is at most $k_b - 1$.

Note that by the feasibility of the allocation made by the optimal mechanism \mathcal{M} , we have

$$\begin{aligned} \sum_{j \in X_a} q_j &= 1/3 \cdot \sum_{j \in X_a} q_j^M \leq k_a/3 \\ \sum_{j \in Y_b} q_j &= 1/3 \cdot \sum_{j \in Y_b} q_j^M \leq k_b/3 \end{aligned}$$

Therefore, using Markov's inequality, with probability at least $2/3$, at most $k_a - 1$ services in X_a are accepted if offered, and the same for Y_b . Using Baye's rule and noting that the two events are positively correlated, the probability that both the events happen is at least $4/9$, and so the claim holds.

When the distributions F_j are non-regular, we pick the prices p_j randomly as suggested by Lemma 4, such that the probability that service j is accepted if offered a randomized price is exactly $q_j^M/3$, and $\mathcal{R}_{\mathcal{I}}^M$ is bounded from above by $3 \sum_j \mathbb{E} p_j q_j$. To bound the expected revenue that the OPM obtains from allocating service j , we note that in any instantiation of the prices (and corresponding "worst" ordering over services), we can pessimistically defer offering service j until all other services have been offered. Then, following the above analysis, the probability that j is offered is at least $c_j \geq 4/9$. Then, the expected revenue from j is at least $4/9 \mathbb{E} p_j q_j$, which is $4/27$ -th of the revenue that this service contributes to $\mathcal{R}_{\mathcal{I}}^M$. Therefore, the claimed approximation holds. \square

4.2 k -uniform matroids

Next we present an improved (and tight) 2-approximation in the case where \mathcal{S} is a k -uniform matroid. Our analysis follows an approach developed in the context of prophet inequalities in [23].

Recall from Section 2.4 that the expected revenue of any mechanism is equal to the expected virtual surplus from its allocation. Therefore, the optimal mechanism in this case allocates the k services with the largest virtual values. Let \mathcal{O} denote this (random) set of services.

For a random variable X , let $(X)^+$ denote the positive portion of X , i.e. $(X)^+ = \max(0, X)$.

Our approximately optimal OPM chooses a virtual value r and sets $p_i = \phi_i^{-1}(r)$. The value r is chosen as the unique solution to $kr = \mathbb{E}[\sum_{i \in \mathcal{O}} (\phi_i(v_i) - r)^+]$. When the distributions F_j are non-regular, we use the same approach, except with ironed virtual valuations instead of virtual valuations. For the rest of this section, we assume that the distributions are regular.

THEOREM 8. *For any instance $\mathcal{I} = (J, \mathcal{S}, \mathbf{F})$ with a k -uniform matroid feasibility constraint, the OPM \mathbf{p} described above achieves a 2-approximation to the optimal revenue, i.e., $\mathcal{R}_{\mathcal{I}}^{\mathbf{p}} \geq 1/2 \mathcal{R}_{\mathcal{I}}^M$.*

PROOF. Let σ be the worst-case ordering for the pricing \mathbf{p} . In particular, this corresponds to the order of increasing prices. Without loss of generality, we assume that $\sigma(i) = i$. Let S denote the (random) set of services allocated.

We upper bound the optimal revenue as

$$\begin{aligned} \mathbb{E} \left[\sum_{i \in \mathcal{O}} \phi_i \right] &\leq \mathbb{E} \left[\sum_{i \in \mathcal{O}} r + (\phi_i - r)^+ \right] \\ &\leq kr + \mathbb{E} \left[\sum_{i \in \mathcal{O}} (\phi_i - r)^+ \right] \\ &= 2kr, \end{aligned}$$

and lower bound the revenue of the OPM as

$$\mathbb{E} \left[\sum_{i \in S} \phi_i \right] = kr \Pr[|S| = k] + \mathbb{E} \left[\sum_{i \in S} \phi_i - r \right].$$

Let χ_i be the indicator random variable associated with the event $|S \cap \{1, 2, \dots, i\}| < k$. We may then lower bound the second term above as

$$\begin{aligned} \mathbb{E} \left[\sum_{i \in S} \phi_i - r \right] &= \sum_{i=1}^n \mathbb{E}[(\phi_i - r)^+ \cdot \chi_{i-1}] \\ &\geq \sum_{i=1}^n \mathbb{E}[(\phi_i - r)^+] \cdot \Pr[|S| < k] \\ &\geq kr \cdot \Pr[|S| < k]. \end{aligned}$$

Thus, we get

$$\mathbb{E} \left[\sum_{i \in S} \phi_i \right] \geq kr \cdot \Pr[|S| = k] + kr \cdot \Pr[|S| < k] = kr,$$

which implies the desired result. \square

A lower bound 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

ϵ 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 ϵ . In either case, the expected revenue of the mechanism is 1.

4.3 Graphical matroids

While we do not know how to obtain constant factor approximations through OPMs for general matroid feasibility constraints, we now demonstrate that OPMs are constant factor optimal for a large class of matroids, namely graphical matroids. The ground set for a graphical matroid is the set of all edges of an undirected graph; A subset of edges is independent if it forms a forest (that is, it contains no cycles).

In order to obtain an approximation, however, we need to extend our definition of OPMs to allow the mechanism to be more restrictive in enforcing feasibility. Specifically, a **constrained order-oblivious posted-price mechanism** (COPM, for short) is given by the tuple $(\mathbf{p}, \mathcal{S}')$ where $\mathcal{S}' \subseteq \mathcal{S}$, and (as for OPMs) we allow the order σ over the agents to be picked adversarially, after the valuations of the agents are drawn. The selling protocol for a COPM offers a service j if the service along with previously allocated services is feasible in the set system \mathcal{S}' , and not merely in \mathcal{S} .

For graphical matroids, Babaioff et al. [3] and Korula and Pál [17] develop approaches for reducing this case to a partition matroid that in our setting yield a 8-approximation to the optimal revenue; we use a similar approach but exploit the connection between prophet inequalities and partition matroids to obtain a 3-approximation.

THEOREM 9. *Let \mathcal{I} be an instance of the BSMD with a graphical matroid feasibility constraint. Then there is a COPM $(\mathbf{p}, \mathcal{S}')$, where \mathcal{S}' is a partition matroid, that 3-approximates \mathcal{R}^M for \mathcal{I} .*

PROOF. 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 $G = (V, J)$ be the graph defining our matroid constraint, where J is the set of services/edges. As before, let q_j^M denote the probability with which edge (service) j is allocated by the optimal mechanism. Let $\delta(v)$ denote the set of edges incident on a vertex v , and for each $v \in V$ define $q_v = \sum_{j \in \delta(v)} q_j^M$. Now, we can see that

$$\sum_{v \in V} q_v = \sum_{j \in E} 2q_j^M \leq 2(|V| - 1),$$

This implies that there exists a v for which $q_v \leq 2$; Let $\delta(v)$ be one of our partitions. Note that the edge set $\delta(v)$ forms a cut in G , and so given an independent set of edges from $J \setminus \delta(v)$ we may add any single edge from $\delta(v)$ while retaining independence. We apply this argument recursively to $(V \setminus \{v\}, J \setminus \delta(v))$ to form the rest of our partition. At the end, we have a partition of J 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 first note how to obtain a simple 8-approximation and then describe the changes needed to obtain a 3-approximation. We define \mathcal{S}' to be the union of 1-uniform matroids, each over the different parts of J defined above. The prices \mathbf{p} for the 8-approximation are defined as follows: $q_j = q_j^M/4$ and $p_j = F_j^{-1}(1 - q_j)$. Then,

the optimal revenue is at most $4 \sum_j p_j q_j$, whereas, our mechanism offers each service with probability at least $1/2$, and therefore, obtains a revenue of $1/2 \sum_j p_j q_j$.

To obtain a 3-approximation, we use the same constraint \mathcal{S}' as before, but modify the prices employing the approach developed in the proof of Theorem 8. Let us focus on a single part, say $X \subseteq J$, in the partition of J defined above. Recall that $\sum_{j \in X} q_j \leq 2$. Let $p_j^M = F_j^{-1}(1 - q_j^M)$, and r be defined such that $r = \sum_{j \in X} q_j^M (p_j^M - r)^+$. Let $p_j = p_j^M$ if $p_j^M \geq r$ and ∞ otherwise. Our COPM is given by $(\mathbf{p}, \mathcal{S}')$.

Then the optimal revenue over X can be bounded as

$$\sum_{j \in X} p_j^M q_j^M \leq \sum_{j \in X} q_j^M (r + (p_j^M - r)^+) = r(1 + \sum_{j \in X} q_j^M) \leq 3r.$$

On the other hand, following the proof of Theorem 8, the revenue of the COPM can be bounded as

$$\begin{aligned} & r \Pr[\text{at least one service in } X \text{ is allocated}] \\ & + \left(\sum_{j \in X} q_j^M (p_j^M - r)^+ \right) \Pr[\text{no service in } X \text{ is allocated}] \\ & = r \end{aligned}$$

□

4.4 Non-matroids constraints

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 OPM is $\Omega(\log n / \log \log n)$ where n is the number of agents. In fact the same lower bounds holds even for SPMs, that is, when we are able to choose the best ordering over offers.

EXAMPLE 1. *For a given k , set $n = k^{k+1}$. Partition $[n]$ into k^k groups G_1, \dots, G_{k^k} of size k each, with $G_i \cap G_j = \emptyset$ for all $i \neq j$. The set system \mathcal{S} contains all subsets of the groups G_i , that is, $\mathcal{S} = \cup_i 2^{G_i}$. Each agent has a value of 1 with probability $1 - 1/k$ and k with probability $1/k$.*

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

Next consider any SPM. Once the mechanism commits to serving an agent, it can only serve agents within the same group in the future. These have a total expected value less than $2k$. Therefore, the revenue of any SPM is at most k from the first agent it serves and $2k$ in expectation from subsequent agents, for a total of $3k$. We get a gap of $\Omega(k) = \Omega(\log n / \log \log n)$.

4.5 Other feasibility constraints

Table 1 outlines the approximation factors we can achieve with OPMs, as well as lower bounds on achievable factors. Please see the full version of this paper [11] for details.

5. APPROXIMATION THROUGH SPMS

In this section we consider settings for the BSMD where we don't know of good approximations through OPMs. We show that we can achieve good approximations by picking the right ordering σ over offers, that is, through sequenced posted-price mechanisms.

Feasibility constraint \mathcal{S}	Gap from optimal	
	upper bound	lower bound
Uniform matroid, Partition matroid	2	2
Graphical matroid	3	2
Intersection of two part. matroids	6.75	2
Graphical matroid \cap partition matroid	10.66	2
General matroid	$O(\log k)$	2
Non-matroid downward closed	-	$\Omega(\frac{\log n}{\log \log n})$

Table 1: A summary of approximation factors for the BSMD achievable through OPMs. Here k is the rank of \mathcal{S} .

While there is no direct reduction from SPMs in multi-parameter settings to SPMs in single-parameter settings analogous to Theorem 6, we show that for matroid and matroid intersection settings our results carry over in an approximation preserving way to multi-parameter instances as well.

We begin with a 2 approximation to single-parameter instances with a general matroid feasibility constraint, and then describe an 8 approximation in PDSE to multi-parameter instances with a general matroid intersection constraint. We conclude this section with a table of approximation factors achievable through SPMs.

5.1 A 2-approximation for matroids

Consider the instance $\mathcal{I} = (J, \mathcal{S}, \mathbf{F})$ where \mathcal{S} is a matroid with rank k . Assume first that all the distributions F_j are regular. Let $q_j = q_j^M$ be the probability with which the optimal truthful mechanism (Myerson’s mechanism) allocates service j . Let $p_j = p_j^M = F_j^{-1}(1 - q_j^M)$. Let σ be the order of decreasing prices p_j over the services. Our approximately optimal SPM is given by (\mathbf{p}, σ) . When the distributions F_j are non-regular, we define the prices randomly as suggested by Lemma 4, and for each instantiation of the prices, pick the greedy ordering over services in order of decreasing prices.

THEOREM 10. *Let \mathcal{I} be an instance of the BSMD with a matroid feasibility constraint. Then, the mechanism (\mathbf{p}, σ) described above 2-approximates the revenue of Myerson’s mechanism for \mathcal{I} .*

PROOF. We show that the mechanism (\mathbf{p}, σ) achieves an expected revenue of at least $\frac{1}{2} \sum_i p_i q_i$. Once again we start with the assumption that all the distributions F_j are regular. Note that if the mechanism ignored the feasibility constraint \mathcal{S} , 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 $B_j = (\text{sp}(S_j) \setminus \text{sp}(S_{j-1})) \cap \{i : i > i_j\}$. Note that the sets B_j partition the set of agents blocked by those previously served. Moreover, $p_i \leq p_{i_j}$ for all $i \in B_j$, since $B_j \subseteq \{i : i > i_j\}$.

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

$$\begin{aligned} & \sum_{j=1}^{\ell} \sum_{i \in B_j} p_i q_i \\ & \leq p^1 \left(\sum_{i \in \text{sp}(S_1)} q_i \right) + \sum_{j=2}^{\ell} p^j \left(\sum_{i \in \text{sp}(S_j)} q_i - \sum_{i \in \text{sp}(S_{j-1})} q_i \right) \end{aligned}$$

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

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

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

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

Next consider the case of non-regular distributions. As mentioned earlier, we pick prices randomly as suggested by Lemma 4. Let p_j be the average price offered for service j if j is the first service offered. Consider a hypothetical posted-price mechanism that orders the services in decreasing order of p_j , and then defers the instantiation of the prices to be offered to just before the service is offered. Then, the acceptance probabilities for the services are exactly q_j , and the previous analysis continues to work in this case. Our SPM, that picks the greedy ordering for every instantiation of prices performs no worse than this hypothetical mechanism and we obtain the same approximation factor as before. \square

We note that this approximation factor is not known to be tight. Blumrosen et al. [6] 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.

5.2 BMUMD with a matroid intersection constraint

Recall that Theorems 6 and 7 together show that we can achieve a constant factor approximation through OPMs to instances of the BMUMD with a feasibility constraint that is an intersection of two partition matroids. We now extend this result to general matroid intersections albeit through SPMs in a slightly weaker solution concept — partial dominant-strategy implementation.

The prices and ordering in our approximately optimal SPM is picked in a manner similar to the one employed in Section 5.1 for matroids. Specifically, let $\mathcal{I} = (J, \mathcal{S}, \Pi, \mathbf{F})$ be the instance of BMUMD that we are interested in. Assume, to begin with, that all the distributions F_j are regular. Consider the instance $\mathcal{I}^{\text{reps}} = (J, \mathcal{S}, \mathbf{F})$, and let q_j^M be the probability with which the optimal truthful mechanism (Myerson’s mechanism) allocates service j in that setting. Let $p_j^M = F_j^{-1}(1 - q_j^M)$. We define q_j to be $q_j^M/2$ and $p_j = F_j^{-1}(1 - q_j)$. Let σ be the order of decreasing prices p_j over the services. Our approximately optimal SPM for \mathcal{I} is (\mathbf{p}, σ) .

THEOREM 11. *Given any instance $\mathcal{I} = (J, \mathcal{S}, \Pi, \mathbf{F})$ of the BMUMD, if \mathcal{S} is a matroid intersection set system, then the SPM (\mathbf{p}, σ) described above is an 8-approximation in PDSE to the revenue of the optimal incentive compatible mechanism for \mathcal{I} .*

PROOF. Let the pricing \mathbf{p} and ordering σ be as defined above. Recall that Lemmas 4 and 5 together imply that the revenue of any incentive compatible mechanism for \mathcal{I} is bounded above by $\sum_j p_j^M q_j^M \leq 2 \sum_j p_j q_j$.

Now consider the SPM (\mathbf{p}, σ) . We say that an agent i desires a service $j \in J_i$ if $v_j > p_j$, and i uniquely desires j if j is the only service in J_i with that property. As noted in Section 2.3, for an agent that uniquely desires a service, sincere bidding is a (weakly) dominant strategy. We first note that for every agent i , with probability $1/2$, i bids sincerely. This follows from Markov’s

Feasibility constraint \mathcal{S}	Gap from optimal	
	upper bound	lower bound
General matroid	2	$\sqrt{\pi/2} \approx 1.25$
Uniform matroid, Partition matroid	$e/(e-1) \approx 1.58$	1.25
Intersection of two matroids (BSMD)	3	1.25
Intersection of two matroids (BMUMD)	8	1.25
Non-matroid downward closed	-	$\Omega(\log n / \log \log n)$

Table 2: A summary of approximation factors for the BSMD and the BMUMD achievable through SPMs.

inequality by noting that $\sum_{j \in J_i} \Pr[i \text{ desires } j] = \sum_{j \in J_i} q_j = 1/2 \sum_{j \in J_i} q_j^M \leq 1/2$.

Now 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 . Note that these sets S , B , and U are random variables depending on the instantiation of agents’ values. Then our observation above 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, every desired service is uniquely desired with probability at least $1/2$; therefore, the expected total price in U is at most half the total price of all desired services, that is, at most $1/2 \sum_j p_j q_j$. Finally, following the proof of Theorem 10, 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. \square

5.3 Other feasibility constraints

Table 2 outlines the approximation factors we can achieve with SPMs, as well as lower bounds on achievable factors.

6. DISCUSSION

We presented constant factor approximations to revenue for several classes of multi-parameter mechanism design problems by designing approximately-optimal posted price mechanisms for single-parameter settings. The approximation factors we obtain in both the single-parameter and multi-parameter settings depend on the kind of feasibility constraint that the seller faces. The exact constants are summarized in Tables 1 and 2.

While these approximation factors are with respect to the optimal deterministic incentive compatible mechanism, in [10] we show that (slightly worse) constant-factor approximation guarantees can be obtained against the optimal randomized incentive compatible mechanism as well.

Our 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-parameter settings that do not exploit competition among agents. Such mechanisms may lead to approximately-optimal multi-parameter 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 single-parameter setting with representatives 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. A partial result along these lines was obtained by [10]—they show that for a class of positively-correlated distributions, called

the common base value model, the optimal revenue of the single-parameter analog still bounds from above the optimal revenue of the multi-parameter problem within a small constant factor.

7. 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] A. Blum and J. Hartline. Near-optimal online auctions. In *Proc. 16th ACM Symp. on Discrete Algorithms*. ACM/SIAM, 2005.
- [5] A. Blum, V. Kumar, A. Rudra, and F. Wu. Online learning in online auctions. *Theoretical Computer Science*, 324:137–146, 2004.
- [6] L. Blumrosen and T. Holenstein. Posted prices vs. negotiations: an asymptotic analysis. In *Proc. 10th ACM Conf. on Electronic Commerce*, 2008.
- [7] P. Briest. Towards hardness of envy-free pricing. Technical Report TR06-150, ECCC, 2006.
- [8] J. Bulow and J. Roberts. The simple economics of optimal auctions. *The Journal of Political Economy*, 97:1060–90, 1989.
- [9] S. Chawla, J. Hartline, and R. Kleinberg. Algorithmic pricing via virtual valuations. In *Proc. 9th ACM Conf. on Electronic Commerce*, pages 243–251, 2007.
- [10] S. Chawla, D. Malec, and B. Sivan. The power of randomness in Bayesian optimal mechanism design. In *Proc. 12th ACM Conf. on Electronic Commerce*, 2010.
- [11] Shuchi Chawla, Jason D. Hartline, David Malec, and Balu Sivan. Multi-parameter mechanism design and sequential posted pricing. *CoRR*, abs/0907.2435, 2009.
- [12] A. Goldberg and J. Hartline. Collusion-resistant mechanisms for single-parameter agents. In *Proc. 16th ACM Symp. on Discrete Algorithms*, 2005.
- [13] 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.
- [14] J. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *Proc. 11th ACM Conf. on Electronic Commerce*, 2009.

- [15] Catherine Holahan. Auctions on eBay: A Dying Breed. http://www.businessweek.com/technology/content/jun2008/tc2008062_112762.htm, June 3 2008.
- [16] 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.
- [17] 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.
- [18] 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.
- [19] R.P. McAfee and J. McMillan. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory*, 46(2):335–354, 1988.
- [20] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [21] N. Nisan and A. Ronen. Algorithmic mechanism design. In *Proc. 31st ACM Symp. on Theory of Computing*, pages 129–140. ACM Press, 1999.
- [22] Jean-Charles Rochet and Philippe Chone. Ironing, Sweeping, and Multidimensional Screening. *Econometrica*, 66(4):783–826, 1998.
- [23] Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *The Annals of Probability*, 12(4):1213–1216, 1984.
- [24] 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.
- [25] D. Vincent and A. Manelli. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory*, 137(1):153–185, 2007.
- [26] Robert B. Wilson. *Nonlinear Pricing*. Oxford University Press, 1997.

APPENDIX

In this section we prove Lemmas 4 and 5, as well as Theorem 6.

Proof of Lemma 4. 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 2, 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 \bar{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 $\bar{q}_i = 1 - F_i(\bar{p}_i)$. Then, $\bar{q}_i \leq q_i^M \leq \underline{q}_i$, and there exists an x_i such that $x_i \underline{q}_i + (1 - x_i) \bar{q}_i = q_i^M$. Now an easy consequence of Proposition 3 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 \bar{p}_i , and with probability 1 if the value is above \bar{p}_i . The revenue of this mechanism is exactly $x_i \underline{q}_i p_i + (1 - x_i) \bar{q}_i \bar{p}_i$.

Proof of Lemma 5. Truthful mechanisms in multi-parameter settings satisfy the weak monotonicity condition defined below. For a Mechanism M , and a value vector \mathbf{v} , let $M_j(\mathbf{v})$ denote the probability with which M provides service j at value vector \mathbf{v} .

DEFINITION 2. A mechanism M satisfies weak monotonicity if for any agent i and any two types \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:

$$\sum_{j \in J_i} (M_j(\mathbf{v}^1) v_j^1 + M_j(\mathbf{v}^2) v_j^2) \geq \sum_{j \in J_i} (M_j(\mathbf{v}^2) v_j^1 + M_j(\mathbf{v}^1) v_j^2)$$

We show that we can construct a truthful mechanism $\mathcal{A}^{\text{reps}}$ for the instance $\mathcal{I}^{\text{reps}}$ 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 $\mathcal{A}^{\text{reps}}$ allocates to the set that \mathcal{A} allocates to in \mathcal{I} for the same value vector \mathbf{v} . We first claim that the allocation rule of $\mathcal{A}^{\text{reps}}$ 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 α_x and α_y denote the probabilities of serving agent i with service j under the two value vectors respectively, and let β_x and β_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 2) of \mathcal{A} implies that

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

This is equivalent to $(x - y)(\alpha_x - \alpha_y) \geq 0$ and so the claim holds.

It remains to prove that the expected revenue of $\mathcal{A}^{\text{reps}}$ given $\mathcal{I}^{\text{reps}}$ 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 [26]. 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 $\mathcal{A}^{\text{reps}}$ 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 6. Let \mathbf{p} be an α -approximate OPM for $\mathcal{I}^{\text{reps}}$. Consider its performance on I . For a fixed instantiation \mathbf{v} of values let σ be any Π -respecting ordering that minimizes the revenue of the mechanism. Note that whenever the mechanism (\mathbf{p}, σ) offers a service to agent i it is a dominant strategy for the agent to accept the service if and only if the agent gets non-negative value from the service. This is because any future offers that the agent gets can only bring him lower utility. Therefore, the sequence of offers and outcome of (\mathbf{p}, σ) is identical under \mathcal{I} and $\mathcal{I}^{\text{reps}}$, and, $\mathcal{R}_{\mathcal{I}}^{(\mathbf{p}, \sigma)}(\mathbf{v}) = \mathcal{R}_{\mathcal{I}^{\text{reps}}}^{(\mathbf{p}, \sigma)}(\mathbf{v}) \geq \mathcal{R}_{\mathcal{I}^{\text{reps}}}^{(\mathbf{p})}(\mathbf{v})$. Therefore, the revenue of \mathbf{p} in the multi-parameter setting is no less than its revenue in the single-parameter setting. Then the result follows from Lemma 5.