Multi-parameter Mechanism Design and Sequential Posted Pricing

Shuich Chawla Jason D. Hartline David L. Malec
Balasubramanian Sivan

May 1, 2023; Week #9

Overview

Overview

What is our problem?

Suppose that you need to reserve hotel rooms for the attendees of a conference. There are a number of rooms available with different features and attendees have preferences over the rooms. Given distributional knowledge on the preferences, How you can maximize your revenue?

Problem Design

1. n attendees
2. m attendee-room matchings
   - $J = [m]$; $\Pi = (J_1, \ldots, J_n)$ is a partition of J
3. Each attendee i has preference $v_j \sim F_j$ with density f_j for $j \in J_i$
4. Feasibility constraint $\mathcal{J} \subseteq 2^J$
   - Downward closed: $A \subsetneq B \in 2^J$ and $B \in \mathcal{J}$ implies $A \in \mathcal{J}$.
   - Unit-demand: $i \in [n]$, $S \in \mathcal{J}$, $|S \cap J_i| \leq 1$.

... hotel rooms for the attendees of a conference....

Problem Design

1. n attendees
2. m attendee-room matchings
   - $J = [m]$; $\Pi = (J_1, \ldots, J_n)$ is a partition of J
3. Each attendee i has preference $v_j \sim F_j$ with density f_j for $j \in J_i$
4. Feasibility constraint $\mathcal{J} \subseteq 2^J$
   - Downward closed: $A \subsetneq B \in 2^J$ and $B \in \mathcal{J}$ implies $A \in \mathcal{J}$.
   - Unit-demand: $i \in [n]$, $S \in \mathcal{J}$, $|S \cap J_i| \leq 1$.

... a number of rooms available ...

Problem Design

1. n attendees
2. m attendee-room matchings
   - $J = [m]$; $\Pi = (J_1, \ldots, J_n)$ is a partition of J
3. Each attendee i has preference $v_j \sim F_j$ with density f_j for $j \in J_i$
4. Feasibility constraint $\mathcal{J} \subseteq 2^J$
   - Downward closed: $A \subsetneq B \in 2^J$ and $B \in \mathcal{J}$ implies $A \in \mathcal{J}$.
   - Unit-demand: $i \in [n]$, $S \in \mathcal{J}$, $|S \cap J_i| \leq 1$.

... have preferences over the rooms. Given distributional knowledge on the preferences...

Problem Design

1. n attendees
2. m attendee-room matchings
   - $J = [m]$; $\Pi = (J_1, \ldots, J_n)$ is a partition of J
3. Each attendee i has preference $v_j \sim F_j$ with density f_j for $j \in J_i$
4. Feasibility constraint $\mathcal{J} \subseteq 2^J$
   - Downward closed: $A \subsetneq B \in 2^J$ and $B \in \mathcal{J}$ implies $A \in \mathcal{J}$.
   - Unit-demand: $i \in [n]$, $S \in \mathcal{J}$, $|S \cap J_i| \leq 1$.

You should not allocate a single room to several attendees.

Problem Design

1. n attendees
2. m attendee-room matchings
   - $J = [m]$; $\Pi = (J_1, \ldots, J_n)$ is a partition of J
3. Each attendee i has preference $v_j \sim F_j$ with density f_j for $j \in J_i$
4. Feasibility constraint $\mathcal{J} \subseteq 2^J$
   - Downward closed: $A \subsetneq B \in 2^J$ and $B \in \mathcal{J}$ implies $A \in \mathcal{J}$.
   - Unit-demand: $i \in [n]$, $S \in \mathcal{J}$, $|S \cap J_i| \leq 1$.

An attendee needs at most a single room

Problem Design

1. n attendees
2. m attendee-room matchings
   - $J = [m]$; $\Pi = (J_1, \ldots, J_n)$ is a partition of J
3. Each attendee i has preference $v_j \sim F_j$ with density f_j for $j \in J_i$
4. Feasibility constraint $\mathcal{J} \subseteq 2^J$
   - Downward closed: $A \subsetneq B \in 2^J$ and $B \in \mathcal{J}$ implies $A \in \mathcal{J}$.
   - Unit-demand: $i \in [n]$, $S \in \mathcal{J}$, $|S \cap J_i| \leq 1$.

If an allocation B is feasible, then its sub-allocation $A \subsetneq B$ should be feasible

Problem Design

1. n attendees
2. m attendee-room matchings
   - $J = [m]$; $\Pi = (J_1, \ldots, J_n)$ is a partition of J
3. Each attendee i has preference $v_j \sim F_j$ with density f_j for $j \in J_i$
4. Feasibility constraint $\mathcal{J} \subseteq 2^J$
   - Downward closed: $A \subsetneq B \in 2^J$ and $B \in \mathcal{J}$ implies $A \in \mathcal{J}$.
   - Unit-demand: $i \in [n]$, $S \in \mathcal{J}$, $|S \cap J_i| \leq 1$.

You need to design a mechanism M that maps preferences $\mathbf{v}$ to an allocation $M(\mathbf{v}) \in \mathcal{J}$ and a pricing $\pi(\mathbf{v})$ that maximizes your revenue.

Overview

Bayesian Mechanism Design

Multi-parameter, Unit-demand

Given

1. n multi-parameter agents
2. a single seller providing m services
   - $J = [m]$; $\Pi = (J_1, \dots, J_n)$ is a partition of J.
3. Each agent i has value $v_j \sim F_j$ with density f_j for $j \in J_i$
4. Feasibility constraint $\mathcal{J} \subseteq 2^J$
   - Downward closed: $A \subsetneq B \in 2^J$ and $B \in \mathcal{J}$ implies $A \in \mathcal{J}$.
   - Unit-demand: $i \in [n]$, $S \in \mathcal{J}$, $|S \cap J_i| \leq 1$.

The Bayesian multi-parameter unit-demand mechanism design (BMUMD) problem is to design a mechanism M maps bids $\mathbf{v}$ to an allocation $M(\mathbf{v}) \in \mathcal{J}$ and a pricing $\pi(\mathbf{v})$.

Bayesian Mechanism Design

Single-parameter

Given

1. n single-parameter agents,
2. a single seller providing a service,
3. Each agent i has value $v_i \sim F_i$ with density f_i and
4. Feasibility constraint $\mathcal{J} \subseteq 2^{[n]}$,
   - Downward closed: $A \subsetneq B \in 2^{[n]}$ and $B \in \mathcal{J}$ implies $A \in \mathcal{J}$.

The Bayesian single-parameter mechanism design (BSMD) problem is to design a mechanism M that maps bids $\mathbf{v}$ to an allocation $M(\mathbf{v}) \in \mathcal{J}$ and a pricing $\pi(\mathbf{v})$.

Note

BSMD is a special case of BMUMD, where $n = m$ and $J_i = \{i\}$ for all $i \in [n]$.

Matroids

Definition

A set system $\mathcal{M} = (X, \mathcal{S})$ over X is a matroid if it satisfies the following conditions.

1. $\emptyset \in \mathcal{S}$,
2. (Downward-closed) If $A \in \mathcal{S}$ and $B \subsetneq A$, then $B \in \mathcal{S}$ and
3. (augmentation) If $A, B \in \mathcal{S}$ with $|A| > |B|$, then there exists $e \in A \setminus B$ such that $B \cup \{e\} \in \mathcal{S}$.

Matroids

Properties

For $S \subseteq X$ and $\mathcal{M} = (X, \mathcal{S})$,

- Rank $r(S) = \max_{A \subseteq S, A \in \mathcal{S}} |A|$
  (the cardinality of max. indep. set in S)
- Span $\mathrm{span}(S) = \{x \in X \mid r(S + x) = r(S)\} \supseteq S$
  (the max. superset T having the same rank)

Special Matroids

A matroid $\mathcal{M} = (X, \mathcal{S})$ is...

- k-Uniform matroid if $\mathcal{S} = \{S \in 2^X \mid |S| \leq k\}$.
- Partition matroid if it is a direct sum of uniform matroids
  - X partitioned into n sets $X_1, \ldots, X_n$.
  - $|X_i \cap S| \leq k_i$ for some k_i.

Sequential Posted-price Mechanism

Mechanism $\mathcal{S}$

This is a mechanism $\mathcal{S}$ for sequential posted pricing:

Require: ordering $\sigma : [m] \rightarrow [m]$ over services and prices $\mathbf{p} = \{p_j\}$.

$A \leftarrow \emptyset$
for $j \in [m]$ **do**
  if $A \cup \{\sigma(j)\} \in \mathcal{J}$ **then**
    Offer p_j for $\sigma(j)$; If accepted, $A \leftarrow A \cup \{\sigma(j)\}$
  end if
end for
Serve A

Note that p_j is the offering price for jth service; not the service j.
Let $\mathcal{R}^{(\mathbf{p},\sigma)}(\mathbf{v})$ be the expected revenue of this mechanism $\mathcal{S}$ on valuation profile $\mathbf{v}$.

Sequential Posted-price Mechanism

Sequential Posted-price Mechanism

A sequential posted-price mechanism (SPM) has an expected revenue of

$$\mathcal{R}^{(\mathbf{p},\sigma)} = \mathrm{E}_{\mathbf{v}\sim\mathbf{F}}[\mathcal{R}^{(\mathbf{p},\sigma)}(\mathbf{v})] = \sum_j c_j p_j q_j,$$

where c_j is the probability with which the mechanism offers to $\sigma(j)$ (at price p_j) and $q_j = 1 - F_j(p_j) = \Pr[v_j \geq p_j]$.
We need to maximize $\mathcal{R}^{(\mathbf{p},\sigma)}$ by choosing **p** and σ.

Sequential Posted-price Mechanism

Example

Given two attendees 1 and 2 and a single room; their values are i.i.d. uniformly between \$100 and \$200.

- The optimal mechanism (Vickery/Myerson) has an expected revenue of \$133 ($\mathrm{E}[\min\{v_1, v_2\}]$).
- The optimal SPM is to offer 1 at \$150, and 2 at \$100; its expected revenue is \$125.

Obviously, $133 > 125$. Why should we use SPM?

Sequential Posted-price Mechanism

Why do we use SPM?

SPM is easily extensible to multi-parameter settings.

Sequential Posted-price Mechanism

Caveats

- $\mathcal{S}$ requires two parameters: ordering σ and prices $\mathbf{p}$.
- As a seller, we choose $\mathbf{p}$ to offer.
- What if we cannot choose σ on our own?

Sequential Posted-price Mechanism

Order-oblivious pricing

An order-oblivious posted-pricing mechanism (OPM) has an expected (worst) revenue of

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

Myerson's Auction

Regularity

Definition

A value distribution F is *regular* if the revenue function $R(q) = F^{-1}(1-q) \cdot q = v \cdot (1 - F(v))$ is concave. Equivalently, $\phi(v) = \frac{dR}{dq}$ is monotone non-decreasing.

This talk will only consider *regular* distributions.

Myerson's Auction

Mechanism M

Require: Agents' valuation **v**; $\phi_j(v_j) = v_j - \frac{1-F_j(v_j)}{f_j(v_j)}$

Choose $A \in \mathcal{J}$ that maximizes $\sum_{j \in A} \phi_j(v_j)$.

Serve A

$\mathcal{R}^M = \mathrm{E}_{\mathbf{v} \sim \mathbf{F}}[\sum_{j \in A} \phi_j(v_j)]$ is the expected revenue of Myerson's auction.

Myerson's Auction

Optimality

Proposition

For any incentive-compatible mechanism A with its expected revenue $\mathcal{R}^A$, $\mathcal{R}^M \geq \mathcal{R}^A$.

Overview

Reducing BMUMD to BSMD

Given a BMUMD instance $I = (J = [m], \mathcal{J}, \Pi, \mathbf{F})$, construct a BSMD instance I' by replacing i into $|J_i|$ distinct representatives $j \in J_i$ with value distribution F_j for $j \in J_i$. Each representative j is interested in j.

Reducing BMUMD to BSMD

Lower bound property

Since I' involves more competition than I, the following holds.

Lemma

Let $\mathcal{A}$ be any IR and IC deterministic mechanism for I. Then, its expected revenue $\mathcal{R}_I^{\mathcal{A}}$ is no more than the expected revenue $\mathcal{R}_{I'}^{M}$ of Myerson's auction for I'.

Reducing BMUMD to BSMD

Reduction for OPM

Theorem

If an OPM with prices **p** *is an α-approximation to the optimal mechanism for BSMD I', then it is an α-approximation to the optimal mechanism for BMUMD I.*

Proof.

An ordering σ is good if $v_{\sigma(a)} - p_a \geq v_{\sigma(b)} - p_b$ for all $a < b$, $a, b \in J_i$.
I.e., an agent i should always take the first offer to maximize his surplus.
Fix a good ordering σ on I. Then,

$$\mathcal{R}_I^{(\mathbf{p},\sigma)}(\mathbf{v}) = \mathcal{R}_{I'}^{(\mathbf{p},\sigma)}(\mathbf{v}) \geq \mathcal{R}_{I'}^{\mathbf{p}}(\mathbf{v})$$

□

Overview

Summary

Feasibility const.	Base Mechanism	Bound
General matroid	SPM	$\sqrt{\pi/2}$ – 2
	OPM	2 – $O(\log k)$
	VCG	2
k-uniform, partition	SPM	1.25 – $e/(e-1) \approx 1.58$
	OPM	2
Graphical	OPM	2 – 3
Intersection matroid	SPM	1.25 – 3
Intersection of part. mat.	OPM	2 – 6.75
Non-matroid	SPM, OPM	$\Omega(\log n/\log\log n)$ – ?

Summary

Feasibility const.	Base Mechanism	Bound
General matroid	SPM	$\sqrt{\pi/2}$ – 2
	OPM	2 – $O(\log k)$
	VCG	2
k-uniform, partition	SPM	1.25 – $e/(e-1) \approx 1.58$
	OPM	2
Graphical	OPM	2 – 3
Intersection matroid	SPM	1.25 – 3
Intersection of part. mat.	OPM	2 – 6.75
Non-matroid	SPM, OPM	$\Omega(\log n/\log\log n)$ – ?

2-approx. for general matroids

Proof (1)

Theorem

For a BSMD instance I, there exist prices $\boldsymbol{p}$ and ordering σ such that $\mathcal{R}^{(\boldsymbol{p},\sigma)} = \mathcal{S}$ 2-approx. $\mathcal{R}^M$ for I.

Proof.

Note that, without the feasibility constraints, we can archive revenue of $\sum_i p_i q_i$. Let $S = \{i_1 < i_2 < \cdots < i_l\}$ be the set of agents served and $S_j = \{i_1 < \ldots < i_j\}$. Let $B_j = \mathrm{span}(S_j) \setminus \mathrm{span}(S_{j-1}) \subseteq \{i \mid i \geq i_j\}$, then B_j is the set of agents blocked by i_{j+1}.

2-approx. for general matroids

Proof (2)

Proof.

Then, the lost revenue given that S is served is:

$$\begin{aligned}\sum_{1\le j\le l}\sum_{i\in B_j} p_i q_i &\le p_1\left(\sum_{i\in \mathrm{span}(S_1)} q_i\right) + \sum_{1<j\le l} p_j\left(\sum_{i\in B_j} q_i\right)\\ &= \sum_{1\le j<l}\left((p_j - p_{j+1})\sum_{i\in \mathrm{span}(S_j)} q_i\right) + p_l\left(\sum_{i\in \mathrm{span}(S_l)} q_i\right)\\ &\le \sum_{1\le j<l}(p_j - p_{j+1})\cdot j + p_l\cdot l \le \sum_{1\le j<l} p_j.\end{aligned}$$

2-approx. for general matroids

Proof (3)

Proof.

Thus,

$$\mathrm{E}[\text{revenue lost}] = \sum_{S} \sum_{j \in S} p_j \cdot \Pr[S \text{ served}] = \mathcal{R}^{(\boldsymbol{p},\sigma)}.$$

which follows $\mathcal{R}^M \leq \sum_j p_j q_j \leq 2\mathcal{R}^{(\boldsymbol{p},\sigma)}$ □

$O(\log k)$-approx. for general matroids

Proof (1)

Theorem

For a BSMD instance I, there exist prices $\boldsymbol{p}$ such that $\mathcal{R}^{\boldsymbol{p}}$ $O(\log k)$-approx. $\mathcal{R}^M$ for I, where $k = \max_{S \in \mathcal{J}} r(S)$ is the maximum rank of independent sets of the matroid.

Proof.

Note that the worst allocation is when agents arrive in the order of increasing prices; let σ be that order. Note that $\mathcal{R}^{(\boldsymbol{p},\sigma)} = \sum_i c_i p_i q_i \geq \frac{1}{2} \sum_i p_i q_i$.

$O(\log k)$-approx. for general matroids

Proof (2)

Proof.

Consider $\boldsymbol{p} = 1$; then $\sum_i c_i q_i \geq \frac{1}{2} \sum_i q_i$. Then, we have

$$\frac{1}{2} \sum_i q_i \leq \sum_i c_i q_i \leq \frac{1}{4} \sum_i q_i + \frac{3}{4} \sum_{i:c_i \geq 1/4} q_i,$$

and thus,

$$\sum_{i:c_i \geq 1/4} q_i \geq \frac{1}{3} \sum_i q_i.$$

$O(\log k)$-approx. for general matroids

Proof (3)

Proof.

Then, Let $G = \{i \mid c_i \geq \frac{1}{4}\}$, the revenue by G is

$$\sum_{i \in G} c_i p_i q_i \geq \frac{1}{4} \sum_{i \in G} p_i^M q_i^M .$$

Note that $|G| \geq \frac{1}{3} n$.

$O(\log k)$-approx. for general matroids

Proof (4)

Proof.

By setting $l = \lceil 1 + \log_{3/2} k \rceil$, we can partition $[n]$ into l sets with total revenue of at least $\mathcal{R}^M/4$. We can conclude that there exists a set whose revenue is at least $1/4l \cdot \mathcal{R}^M = \Omega(1/\log k) \cdot \mathcal{R}^M$.

6.75-approx. for intersection of partition matroids

Proof (1)

Theorem

Let I be a BSMD instance with a feasibility constraints given by the intersection of two partition matroids. Then, there exists a set of prices $\boldsymbol{p}$ such that $\mathcal{R}^{\boldsymbol{p}}$ 6.75-approximates $\mathcal{R}^M$ for I.

Proof.

Let $q_i = q_i^M/3$ and $p_i = F_i^{-1}(1 - q_i)$. This mechanism serves agents in any arbitrary order (hence OPM), but offers a price p_i for agent i. We now prove that $c_i \geq 4/9 = 1/6.75$, then $\mathcal{R}^{\boldsymbol{p}} = \sum_i c_i p_i q_i \geq 4/9 \sum_i p_i^M q_i^M/3$.

6.75-approx. for intersection of partition matroids

Proof (2)

Proof.

Let $\mathcal{M}_1$ and $\mathcal{M}_2$ be two partition matroids. for $j = 1, 2$, let agent i be in partition P_j of $\mathcal{M}_j$ and $k_x = r_{\mathcal{M}_x}(P_x)$. Then, the expected number of agents in P_j desiring service is

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

6.75-approx. for intersection of partition matroids

Proof (3)

Proof.

Let E_j be the event that at most $k_j - 1$ agents from P_j desire service. Then i is always considered when both E_1 and E_2 happen. Thus,

$$c_i \geq \Pr[E_1 \cap E_2] \geq 2/3 \cdot 2/3.$$

□

Overview

Summary

Feasibility const.	Concept / Mechanism	Bound
Multi-unit, multi-item, unit-demand	DSIC / OPM	6.75
Graphical w/ unit-demand	DSIC / OPM	$32/3$
Intersection matroid	undominated / SPM	8
Comb. auction w/ small bundles	undominated / SPM	8

Summary

Feasibility const.	Concept / Mechanism	Bound
Multi-unit, multi-item, unit-demand	DSIC / OPM	6.75
Graphical w/ unit-demand	DSIC / OPM	$32/3$
Intersection matroid	undominated / SPM	8
Comb. auction w/ small bundles	undominated / SPM	8

Multi-unit, Multi-item, unit-demand

Theorem

Consider an instance of BMUMD where the seller has multiple copies of n items on sale, and agents are unit-demand. then, there exists an 6.75-approximate OPM for this instance.

Proof.

From Theorem 4 (α-approximation for BMUMD) and Theorem 13 (6.75-approx. for BSMD w/ two part. mat.). □

Overview

Assumptions

- An algorithm that computes the optimal price p_i to charge to a single-parameter agent given by F_i. (Note that, with given x, we can use this algorithm to find an optimal price in $[x, \infty)$.)
- An oracle that, given a value v, returns $F_i(v)$ and $f_i(v)$.
- An oracle that, given a probability α, returns $F_i^{-1}(\alpha)$.
- An algorithm to maximize social welfare over the given feasibility constraint (Myerson's).

Algorithm

1. Let $\epsilon = 1/3n$. Sample $N = 4n^4 \log n/\epsilon^2$ value profiles from $\boldsymbol{F}$.
2. Estimate q_i^M using the samples; call $\widehat{q_i^M}$.
3. If $\widehat{q_i^M} < 1/n^2$, set $\widehat{q_i} = 1/n^2$. Else, set $\widehat{q_i} = \widehat{q_i^M}/(1-\epsilon)$.
4. Compute $\widehat{p_i} = F_i^{-1}(1-\widehat{q_i})$.
5. Find the optimal price p_i in $[\widehat{p_i}, \infty)$; let $q_i = 1 - F_i(p_i)$.
6. Output p_i's, and order of agents in decreasing prices.

Proof

Lemma

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

Proof.

$\Pr[|\widehat{q_i^M} - q_i^M| \geq \epsilon q_i^M] \leq 2/n^2$ by Chernoff bounds. □

With $p_i^M \in [\widehat{p_i}, \infty)$, we have $p_i^M q_i^M \leq p_i q_i$.

Proof

Let $S = \{i \mid \widehat{q_i^M} < 1/n^2\}$. Then, the probability of a mechanism offer to anyone in S is at most $1/n$. Suppose not, then, by the prob. of $1 - 1/n$, our revenue from i is $p_i q_i > p_i^M q_i^M$. Thus, conditioned on the lemma (with probability of $1 - 2/n$), we get a $(1 - o(1))$ approx. to the optimal mechanism.