CS 396: Online Markets

Lecture 2: Online Allocation

Last Time:

- ride sharing problem
- algorithms, online algorithms, mechanisms
- first price auction
- ascending auction
- second price auction

Today:

- course philosophy
- second price auction (cont)
- online allocation: prophet inequality
- online mechanisms: sequential pricing

Exercise: Uniform Expectation

Calculate the expected value of X for $X \sim U[0, 1]$

Course Philosophy

- online markets combines:
  - online algorithm design
  - mechanism design
- lectures:
  - algorithms/mechanisms and proofs
  - goal: learn **why** markets work
- projects:
  - hands on experience.
  - goal: learn how theory applies.

Mech 3: second-price auction

1. ask buyers to report values.
2. winner is highest bidder.
3. winner pays second-highest bid.

Thm: bid = value is dominant strategy in second-price auction

Proof:

Consider bidder i:

- $\hat{v}_i = \max_{j \neq i} b_j$
- if $b_i > \hat{v}_i$:
  - i wins and pays $\hat{v}_i$
  - utility: $u_i = v_i - \hat{v}_i$
- if $b_i < \hat{v}_i$:
  - i loses and pays 0
  - utility: $u_i = 0$
- consider cases
  - $v_i > \hat{v}_i$
  - $v_i < \hat{v}_i$

  [PICTURE]
- $b_i = v_i$ is "best" in both cases.
- thus, dominant strategy.

Cor: second-price auction maximizes social welfare

Exercise: Second-price Performance

Setup:

- two buyers, one item
- values $U[0, 3]$
- second price auction

Question: Calculate the seller's expected revenue and the mechanism's expected social welfare (the expected value of the winner)

Online Algorithms

"buyers show up one at a time, must make decision for each buyer before next buyer arrives"

Note: no good algorithm without more information.

Example:

- $\mathbf{v} = (1, 0, \dots, 0)$
- $\mathbf{v}' = (1, 0, \dots, 100)$

Assumption: values are drawn from known probability distributions

"$v_1 \sim F_1, \dots, v_n \sim F_n$"

Example: $F_1 = \ldots = F_n = U[0, 1]$

Online Gambling:

- online gambler faces n games.
- game i has prize $v_i \sim F_i$.
- gambler plays games in order:
  - realized prize $v_i \sim F_i$
  - either keeps prize and quits.
  - or discards prize and continues.

Q: what is gambler's optimal strategy?

A: backwards induction

- on day n, take prize

  $\Rightarrow \mathbf{E}[v_n] = \hat{v}_{n-1}$

- on day $n-1$, take prize if $v_{n-1} \geq \hat{v}_{n-1}$

  $\Rightarrow \mathbf{E}[\max(v_{n-1}, \hat{v}_{n-1})] = \hat{v}_{n-2}$

- on day $n-2$, take prize if $v_{n-2} \geq \hat{v}_{n-2}$

- ...

$\Rightarrow$ "decreasing sequence of thresholds"

Exercise: Two-day Gamble

Setup:

- $n = 2$ prizes
- uniformly distributed $F_1 = F_2 = U[0, 1]$
- realize first price
- claim it or discard and realize and claim second price

Question: Find optimal strategy. What is its expected payoff? What is its probability of claiming first price?

Comparing online algorithm to optimal offline algorithm

Defn: an online algorithm ALG is a β-approximation to the optimal offline algorithm OPT if

$$\mathbf{E}[\mathrm{ALG}(\mathbf{v})] \geq \frac{1}{\beta}\mathbf{E}[\mathrm{OPT}(\mathbf{v})]$$

A.k.a. "competitive analysis"

Thm: gambler can used single threshold $\hat{v}$ and obtain expected payoff at least half of expected optimal offline payoff, i.e., 2-approximation.

A.k.a. prophet inequality

"optimal offline payoff" is prophet's payoff (who can see future)"

Proof:

0. Notation

   - $\hat{v}$ = threshold
   - $(\cdot)^+ = \max(\cdot, 0)$
   - $\chi = \prod_i \mathbf{P}r[v_i < \hat{v}] = \mathbf{P}r[\text{no prize}]$
   - event $\mathcal{E}_i$ = "$\forall j \neq i,\ v_j < \hat{v}$"

1. upperbound on prophet

$\mathbf{E}[\mathrm{OPT}(\mathbf{v})]$

$$
\begin{aligned}
&= \mathbf{E}[\max_i v_i] \\
&= \hat{v} + \mathbf{E}[\max_i (v_i - \hat{v})] \\
&\leq \hat{v} + \mathbf{E}[\max_i (v_i - \hat{v})^+] \\
&\leq \hat{v} + \mathbf{E}[\sum_i (v_i - \hat{v})^+] \\
&= \hat{v} + \sum_i \mathbf{E}[(v_i - \hat{v})^+]
\end{aligned}
$$

2. lowerbound on gambler

   $\mathbf{E}[\mathrm{ALG}(\mathbf{v})]$

$$
\begin{aligned}
&= (1-\chi)\,\hat{v} + \mathbf{E}[\text{amount selected prize exceeds } \hat{v}] \\
&\geq (1-\chi)\,\hat{v} + \sum_i \mathbf{E}[(v_i - \hat{v})^+ \,|\, \mathcal{E}_i]\,\mathbf{P}r[\mathcal{E}_i] \\
&\geq (1-\chi)\,\hat{v} + \chi \sum_i \mathbf{E}[(v_i - \hat{v})^+]
\end{aligned}
$$

3. choose $\chi = 1/2$ and combine.

   $\mathbf{E}[\mathrm{ALG}(\mathbf{v})] \geq \frac{1}{2}\mathbf{E}[\mathrm{OPT}(\mathbf{v})]$

Online Mechanisms

"buyers arrive one at a time and are strategic"

"the eBay 'buy it now' problem"

Idea: use thresholds as prices

Defn: in a **sequential pricing** buyers arrive one at a time and the mechanism offers take-it-or-leave-it while-supplies-last prices

Defn: in a **uniform pricing** is a sequential pricing where all buyers are offered the same price.

Thm: bid = value is dominant strategy in sequential pricing.

Proof:

- each bidder i faces threshold $\hat{v}_i$
- same analysis as second-price auction follows.

Thm: the sequential pricing with prices equal to thresholds from backwards induction maximizes social welfare (among online mechanisms)

Proof: by optimality of thresholds.

Thm: uniform pricing is 2-approximation

Proof: direct corollary of prophet inequality

Summary

Four paradigms for allocating a single-item

- algorithms.
- mechanisms.
- online algorithms.
- online mechanisms.
