Prophet Inequalities

A Crash Course

Brendan Lucier, Microsoft Research

EC18: ACM Conference on Economics and Computation

Mentoring Workshop, June 18, 2018

Profit

From Wikipedia, the free encyclopedia

Not to be confused with Prophet.

Prophet

From Wikipedia, the free encyclopedia

Not to be confused with Profit.

The Plan

1. Introduction to Prophet Inequalities

2. Connections to Pricing and Mechanism Design

Prophet Inequality

The gambler's problem:

D_1 D_2 D_3 D_4 D_5

Prophet Inequality

The gambler's problem:

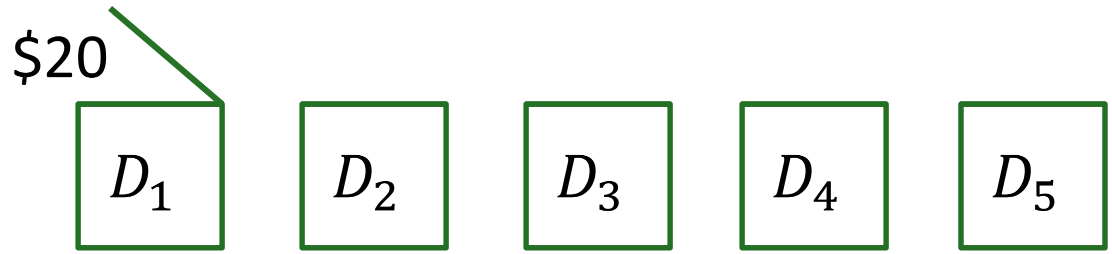

Keep: win $20, game stops.
Discard: prize is lost, game continues with next box.

Let's Play…

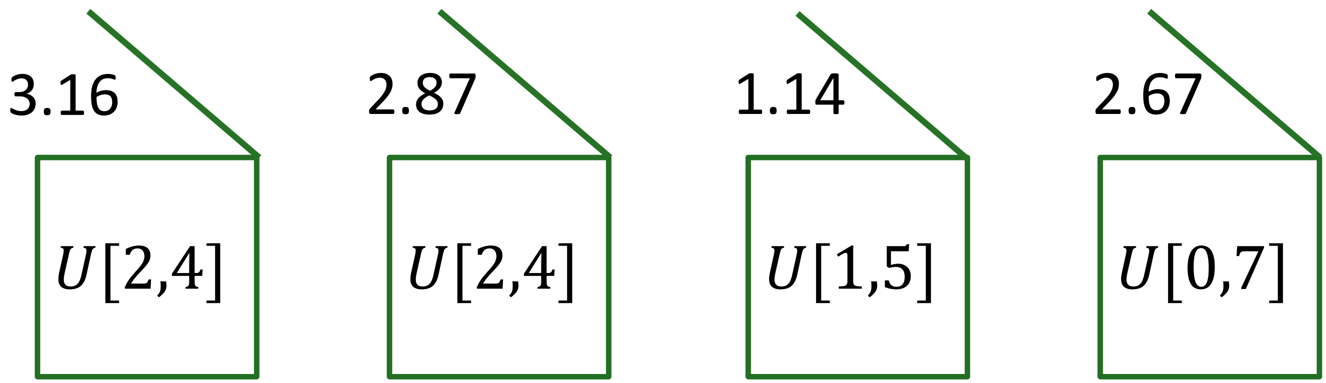

Prophet Inequality

Theorem: [Krengel, Sucheston, Garling '77]

There exists a strategy for the gambler such that

$$E[prize] \geq \frac{1}{2} E\left[\max_i v_i\right]$$

and the factor 2 is tight.

[Samuel-Cahn '84] … a fixed threshold strategy:
choose a single threshold t, accept first prize $\geq t$.

Lower Bound: 2 is Tight

$\mathbb{E}[\max_i \pi_i] = 1(1-\epsilon) + \frac{1}{\epsilon} \cdot \epsilon$

$\sim 2-\epsilon$

1

1

$\frac{1}{\epsilon}$ w.p. ϵ

0 otherwise

Theorem: [Samuel-Cahn '84]

Given distributions $G_1, \ldots, G_n$ where $\pi_i \sim G_i$, there exists a fixed threshold strategy (accept first prize $\geq t$) such that

$$E_\pi[prize] \geq \frac{1}{2} E_\pi \left[\max_i \pi_i\right]$$

Proof:

Application: Posted Pricing

A mechanism design problem:

1 item to sell, n buyers, independent values $v_i \sim D_i$.

Buyers arrive sequentially, in an arbitrary order.

For each buyer: interact according to some protocol, decide whether or not to trade, and at what price.

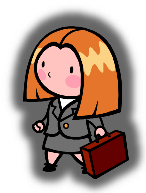

$v_1 \sim D_1$

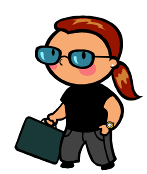

$v_2 \sim D_2$

$v_3 \sim D_3$

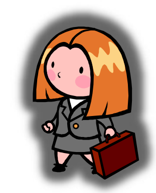

$v_4 \sim D_4$

s.w. $\geq \frac{1}{2} \mathbb{E}_{v_i \sim D_i}\left[\max_i v_i\right]$

Corollary of Prophet Inequality:

Posting an appropriate take-it-or-leave-it price t yields at least half of the expected optimal social welfare.

[Hajiaghayi Kleinberg Sandholm '07]

Applications

What about revenue?

[Chawla Hartline Malec Sivan '10]: Can apply prophet inequality to *virtual values* to achieve half of optimal revenue.

$$E[Rev] = E_v\left[\sum_i p_i(v)\right] = E_v\left[\sum_i \phi_i(v_i) x_i(v)\right]$$

(for single item) $$= E_v[\max_i \phi_i(v_i)^+]$$

r_i

Auction w/ E[Rev]$\geq \frac{1}{2} OPT$

1. *Distribution G_i on $\phi_i(v_i)^+$ using F_i on v_i*
2. *Compute t s.t.* $\Pr\left[\max_i \phi_i(v_i)^+ \geq t\right] = 1/2$ (t s.t. Prob. Of selling is ½)
3. *Give to an agent with $\phi_i(v_i)^+ \geq t$*
   - *With highest value*
4. *Payment = max{$\phi_i^{-1}(t)$, second highest bid}*

Alternate Pricing

Multiple choices of p that achieve the 2-approx of total value. Here's one due to [Kleinberg Weinberg 12]:

Theorem (prophet inequality): for one item, setting threshold $p = \frac{1}{2} E\left[\max_i v_i\right]$ yields expected **welfare** $\geq \frac{1}{2} E\left[\max_i v_i\right]$.

Example:

1 or 6	0 or 8	2 or 10

(each box: prizes equally likely)

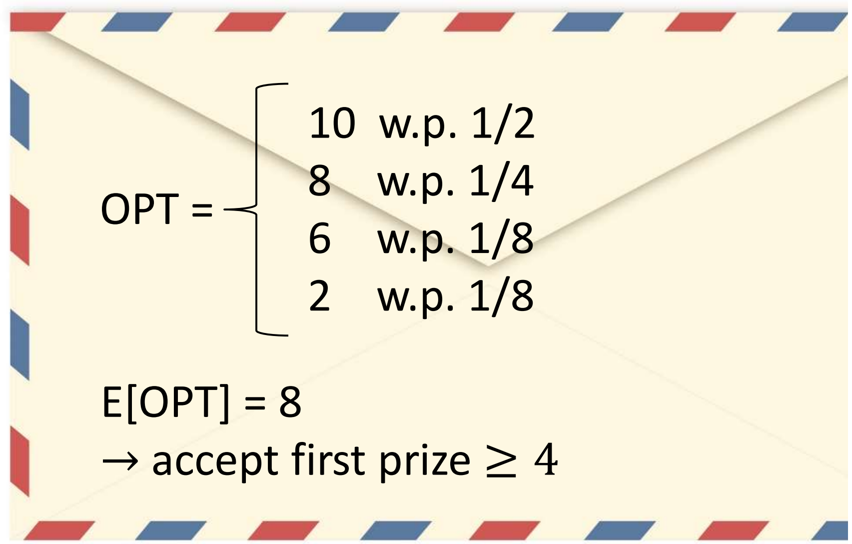

Prophet Inequality: Proof

Theorem (prophet inequality): for one item, setting threshold $\mathrm{p} = \frac{1}{2} E\left[\max_i v_i\right]$ yields expected value $\geq \frac{1}{2} E\left[\max_i v_i\right]$.

What can go wrong?

If threshold is

- Too low: we might accept a small prize, preventing us from taking a larger prize in a later round.
- Too high: we don't accept *any* prize.

A Proof for Full Information

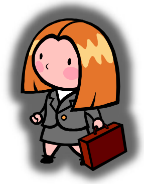

$v_1 = 10$

$v_2 = 50$

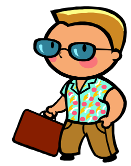

$v_3 = 80$

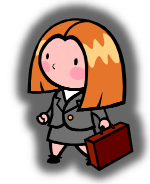

$v_4 = 15$

Idea: price $\frac{1}{2}\max_i v_i$ is "balanced"

Let $v_{i^*} = \max_i v_i$.

Case 1: Somebody $i < i^*$ buys the item.

$\Rightarrow$ revenue $\geq \frac{1}{2} v_{i^*}$

Case 2: Nobody $i < i^*$ buys the item.

$\Rightarrow$ utility of $i^* \geq v_{i^*} - \frac{1}{2} v_{i^*} = \frac{1}{2} v_{i^*}$

In either case: welfare = revenue + buyer utilities $\geq \frac{1}{2} v_{i^*}$ ∎

Extending to Stochastic Setting

Thm: setting price $p = \frac{1}{2} E\left[\max_i v_i\right]$ yields value $\geq \frac{1}{2} E\left[\max_i v_i\right]$.

v^*

Proof. Random variable: $v^* = \max_i v_i = OPT$

1. $\text{REVENUE} = p \cdot \Pr[\text{item is sold}] = \frac{1}{2}\mathrm{E}[v^*] \cdot \Pr[\text{item is sold}]$
2. $\text{SURPLUS} = \sum_i \mathrm{E}[\text{utility of buyer } i]$

$$\geq \sum_i \mathrm{E}\left[(v_i - \mathrm{p})^+ \cdot \mathbf{1}[i \text{ sees item}]\right]$$
$$= \sum_i \mathrm{E}[(v_i - \mathrm{p})^+] \cdot \Pr[i \text{ sees item}]$$
$$\geq \sum_i \mathrm{E}[(v_i - \mathrm{p})^+] \cdot \Pr[\text{item not sold}]$$
$$\geq \mathrm{E}\left[\max_i (v_i - \mathrm{p})\right] \cdot \Pr[\text{item not sold}]$$
$$\geq \frac{1}{2}\mathrm{E}[v^*] \cdot \Pr[\text{item not sold}]$$

3. $\text{Total Value} = \text{REVENUE} + \text{SURPLUS} \geq \frac{1}{2}\mathrm{E}[\mathrm{v}^*]$. ∎

Prophet Inequality: Proof

Thm: for one item, price $\mathrm{p} = \frac{1}{2}E[OPT]$ yields value $\geq \frac{1}{2}E[OPT]$.

Summary:

- Price is high enough that expected revenue offsets the opportunity cost of selling the item.
- Price is low enough that expected buyer surplus offsets the value left on the table due to the item going unsold.

Secretaries and Prophet Secretaries

A Variation

Prophet Inequality:

Prizes drawn from distributions, order is arbitrary

A Related Problem:

Prizes are arbitrary, order is uniformly random

Let’s Play…

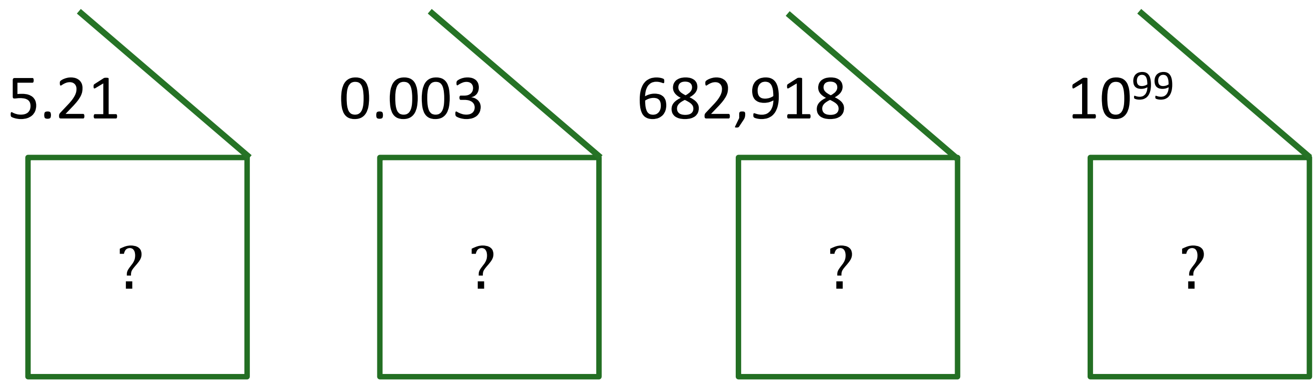

The game of googol [Gardner ‘60]

Secretary Problem

Theorem: [Lindley ’61, Dynkin ‘63, Gilbert and Mosteller ‘66]

There exists a strategy for the secretary problem such that

$$Pr[select\ largest] \geq \frac{1}{e}$$

and the factor e is tight as n grows large.

Strategy: observe the first n/e values, then accept the next value that is larger than all previous.

Prophets vs Secretaries

Prophet Inequality:

Prizes drawn from distributions, order is arbitrary

Secretary Problem / Game of Googol:

Prizes are arbitrary, order is uniformly random

Prophet Secretary:

Prizes drawn from [known] distributions, order is uniformly random [and revealed online]

[Esfandiari, Hajiaghayi, Liaghat, Monemizadeh ‘15]

Recall:

$U[2,4]$ $U[2,4]$ $U[1,5]$ $U[0,7]$

Recall:

$U[0,7]$ $U[1,5]$ $U[2,4]$ $U[2,4]$

Prophet Secretary

Theorem: [Esfandiari, Hajiaghayi, Liaghat, Monemizadeh '15]

There exists a strategy for the gambler such that

$$E[prize] \geq \left(1 - \frac{1}{e}\right) E\left[\max_i v_i\right].$$

[Azar, Chiplunkar, Kaplan EC'18]: A strategy for the gambler that beats $\left(1 - \frac{1}{e}\right)$.

Prophet Secretary

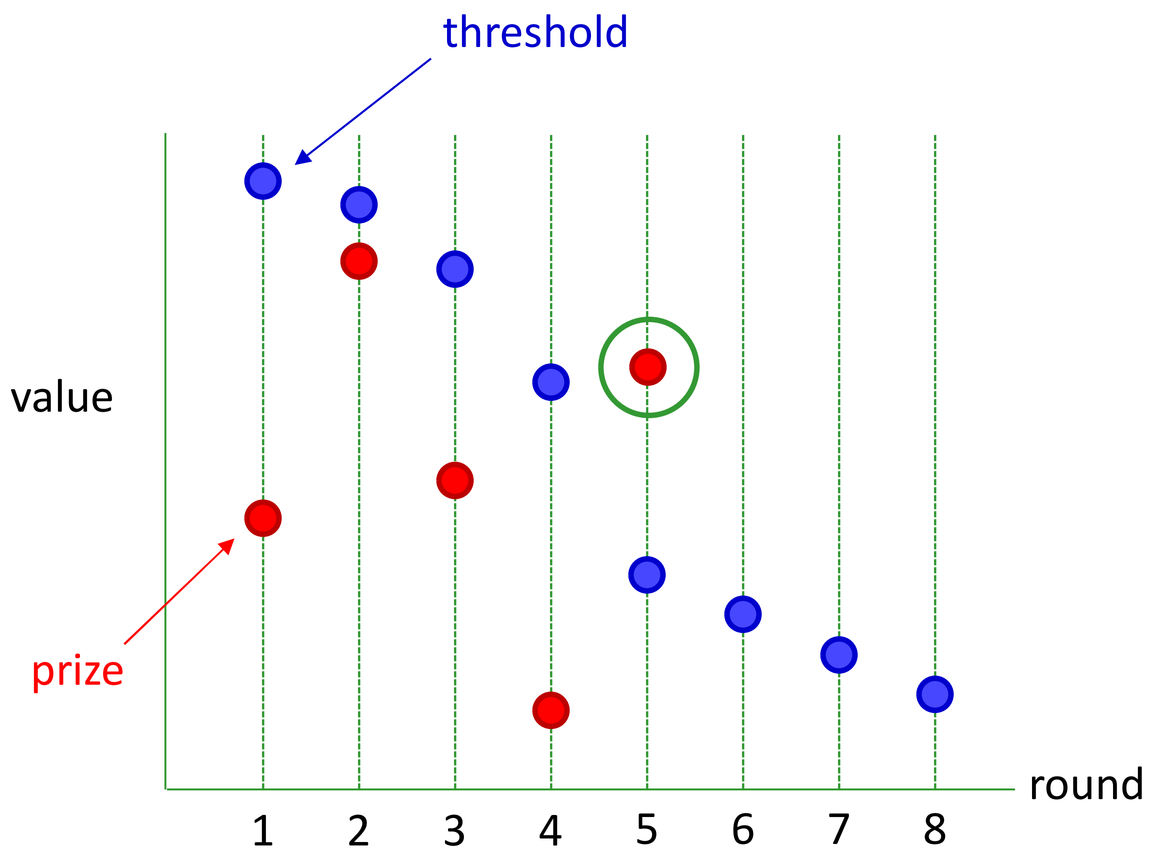

Prophet Secretary

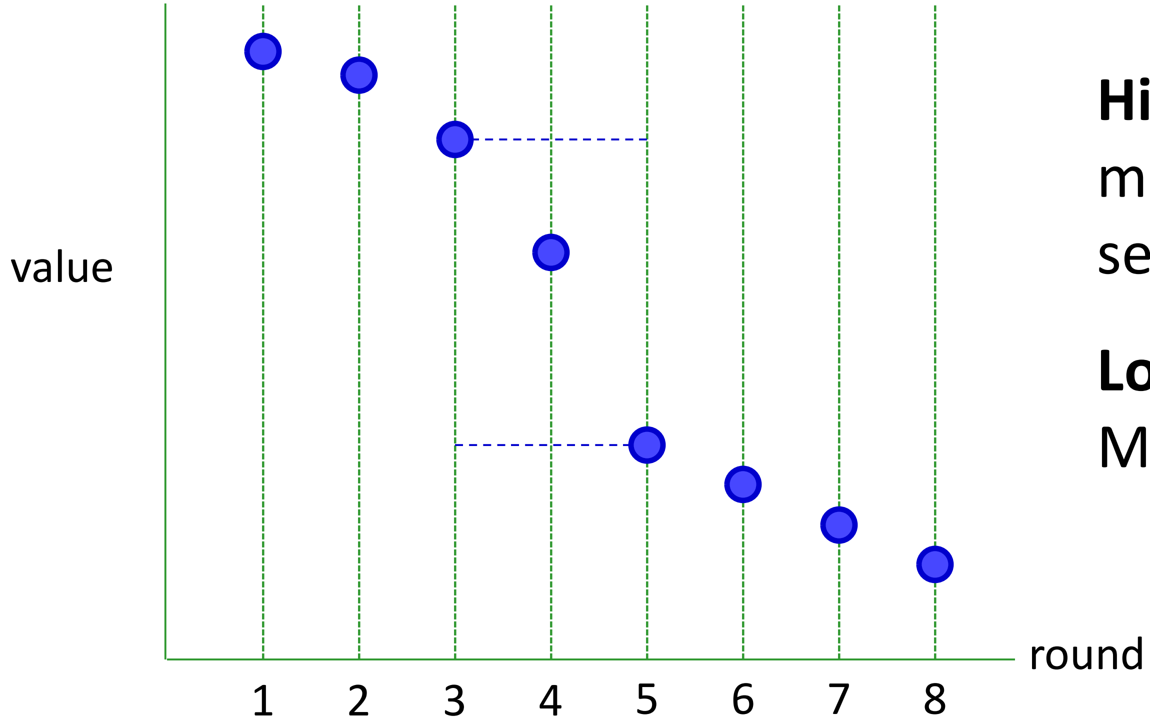

Higher threshold:
more revenue when we sell the item to this buyer.

Lower threshold:
More surplus for this buyer.

Extension: Multiple Prizes

Multiple-Prize Prophet Inequality

Prophet inequality, but gambler can keep up to k prizes

$k = 1$: original prophet inequality: 2-approx

$k \geq 1$: [Hajiaghayi, Kleinberg, Sandholm '07]

There is a threshold p such that picking the first k values $\geq p$ gives a $1 + O(\sqrt{\log k / k})$ approximation.

Idea: choose p s.t. expected # of prizes taken is $k - \sqrt{2k \log k}$.

Then w.h.p. # prizes taken lies between $k - \sqrt{4k \log k}$ and k.

[Alaei '11] [Alaei Hajiaghayi Liaghat '12] Can be improved to $1 + O\left(\frac{1}{\sqrt{k}}\right)$ using a randomized strategy, and this is tight.

Aside: Beyond Cardinality

Constraint	Upper Bound	Lower Bound
Single item	2	2
k items	$1 + O\left({}^{1}/_{\sqrt{k}}\right)$	$1 + \Omega\left({}^{1}/_{\sqrt{k}}\right)$
Matroid	2 [Kleinberg Weinberg '12]	2
k matroids	$e \cdot (k+1)$ [Feldman Svensson Zenklusen '15]	$\sqrt{k} + 1$ [Kleinberg Weinberg '12]
Knapsack	5 [Duetting Feldman Kesselheim L. '17]	2
Downward-closed, max set size $\leq r$	$O(\log n \log r)$ [Rubinstein '16]	$\Omega\left({}^{\log n}/_{\log \log n}\right)$ [Babaioff Immorlica Kleinberg '07]

Directly imply posted-price mechanisms for welfare, revenue

Multiple-Prize Prophet Inequality

A different variation on cardinality:

- The gambler can choose up to $k \geq 1$ prizes
- Afterward, gambler can keep the *largest* of the prizes chosen

Theorem [Assaf, Samuel-Cahn '00]: There is a strategy for the gambler such that $E[prize] \geq \left(1 - \frac{1}{k+1}\right) E\left[\max_i v_i\right]$

[Ezra, Feldman, Nehama EC'18]: An extension to settings where gambler can choose up to k prizes and keep up to ℓ. Includes an improved bound for $\ell = 1$!

Combinatorial Variants

More general valuation functions:

Reward for accepting a set of prizes S is a function $f(S)$.
Example: arbitrary submodular. [Rubinstein, Singla '17]

Multiple prizes per round:

Multiple boxes arrive each round.
Revealed in round i: valuation function $f_i(S)$ for accepting set of prizes S_i on round i. (Note: possible correlation!)

Application: posted-price mechanisms for selling many goods
[Alaei, Hajiaghayi, Liaghat '12], [Feldman Gravin L '13], [Duetting Feldman Kesselheim L '17]

Summary

- Prophet Inequalities: analyzing the power of sequential decision-making, vs an offline benchmark.
- Recent connections to pricing and mechanism design
- MANY variations! A very active area of research

Open Challenge: Best-Order Prophet Inequality

Suppose the gambler can choose which order to open boxes.

- What fraction of $E\left[\max_i v_i\right]$ can the gambler guarantee?
- Can the best order be computed efficiently?

Thanks!

Bonus: Multi-Dimensional Prophets

A General Model

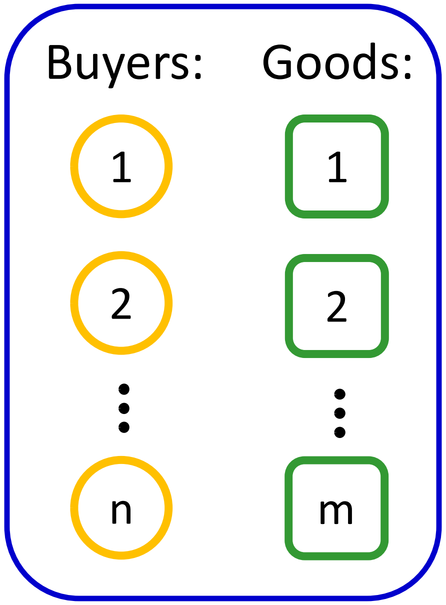

Combinatorial allocation

- Set M of m resources (goods)
- n buyers, arrive sequentially online
- Buyer i has valuation function $v_i: 2^M \to R_{\geq 0}$
- Each v_i is drawn indep. from a known distribution D_i
- Allocation: $\boldsymbol{x} = (x_1, \dots, x_n)$.
  There is a downward-closed set F of feasible allocations.

Goal: feasible allocation maximizing $\sum_i v_i(x_i)$

Posted Price Mechanism

1. For each bidder in some order π:
2. Seller chooses prices $p_i(x_i)$
3. Bidder i's valuation is realized: $v_i \sim F_i$
4. i chooses some $x_i \in \arg\max\{v_i(x_i) - p_i(x_i)\}$

Notes:

- "Obviously" strategy proof [Li 2015]
- Tie-breaking can be arbitrary
- Prices: static vs dynamic, item vs. bundle
- Special case: oblivious posted-price mechanism (OPM) prices chosen in advance, arbitrary arrival order

Applications

Problem	Approx.	Price Model
Combinatorial auction, XOS valuations	2	Static item prices
Bounded complements (MPH-k) [Feige et al. 2014]	$4k - 2$	Static item prices
Submodular valuations, matroid constraints	2 (existential) 4 (polytime)	Dynamic prices
Knapsack constraints	5	Static prices
d-sparse Packing Integer Programs	8d	Static prices

[Feldman Gravin L '13], [Duetting Feldman Kesselheim L '17]