Prophet Inequalities

Part 1: Introduction

Paul Dütting, Google Research
duetting@google.com

ADFOCS 2024 Summer School

August 2024

Plan for Part 1

- What is a prophet inequality?
  - Statement and proof of the classic prophet inequality

- What's so exciting about prophet inequalities?
  - A powerful tool for mechanism design
  - A new ``beyond worst-case'' paradigm for online algorithms

- On the way: Sample / overview of research landscape

Outline Other Parts

Part 1: Introduction

Part 2: Online matching and contention resolution

Part 3: Online combinatorial auctions and balanced prices

Part 4: Data-driven prophet inequalities

Useful Resources

- WINE 2016 Tutorial “Posted-Price Mechanisms and Prophet Ineqaulities” by Brendan Lucier [website, slides]
- EC 2017 Tutorial ”Pricing in Combinatorial Markets” by Michal Feldman and Brendan Lucier [on request]
- IPCO 2017 Summer School “Prophets and Secretaries” by Anupam Gupta [lecture notes]
- EC 2021 Tutorial “Prophet Inequalities and Implications to Pricing and Online Algorithms” by Michal Feldman, Thomas Kesselheim, and Sahil Singla [website, slides-pt1, slides-pt2, slides-pt3]

(This course builds on these prior courses/tutorials, and re-uses some of the material)

Books and Surveys

- Survey “A Survey of Prophet Inequalities in Optimal Stopping” by Theodore Hill and Robert Kertz [pdf] (from 1992)
- Survey “An Economic View of Prophet Inequalities” by Brendan Lucier [pdf] (from 2017)
- Survey “Recent Developments in Prophet Inequalities” by Jose Correa, Patricio Foncea, Ruben Hoeksma, Tim Osterwijk, Tjark Vredeveld [pdf] (from 2018)
- Forthcoming book “Prophet Inequalities: Theory and Practice” by Jose Correa, Paul Dütting, Michal Feldman, Brendan Lucier, and Thomas Kesselheim (planned for 2025)

The Classic Prophet Inequality

The Problem

- Given known distributions $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_n$ over (non-negative) values:
  - A gambler gets to see realizations $v_i \sim \mathcal{D}_i$ one-by-one, and needs to immediately and irrevocable decide whether to accept v_i
  - The prophet sees the entire sequence of values $v_1, v_2, \dots, v_n$ at once, and can simply choose the maximum value
- **Question:** What's the worst-case gap between $\mathbb{E}$[value accepted by gambler] and $\mathbb{E}$[value accepted by prophet]?

$= \mathbb{E}[\max_i v_i]$

$=: \mathbb{E}[ALG]$

Let's Play

$\mathcal{D}_1 = U[0,1]$

$\mathcal{D}_2 = U[0,1]$

$\mathcal{D}_3 = U[0,1]$

$\mathcal{D}_4 = U[0,1]$

Let’s Play

$\mathcal{D}_1 = U[0,1]$

$\mathcal{D}_2 = U[0,1]$

$\mathcal{D}_3 = U[0,1]$

$\mathcal{D}_4 = U[0,1]$

0.4

Let’s Play

$\mathcal{D}_1 = U[0,1]$

0.4

reject

$\mathcal{D}_2 = U[0,1]$

$\mathcal{D}_3 = U[0,1]$

$\mathcal{D}_4 = U[0,1]$

Let’s Play

$\mathcal{D}_1 = U[0,1]$ $\mathcal{D}_2 = U[0,1]$ $\mathcal{D}_3 = U[0,1]$ $\mathcal{D}_4 = U[0,1]$

0.4 0.8

reject

Let’s Play

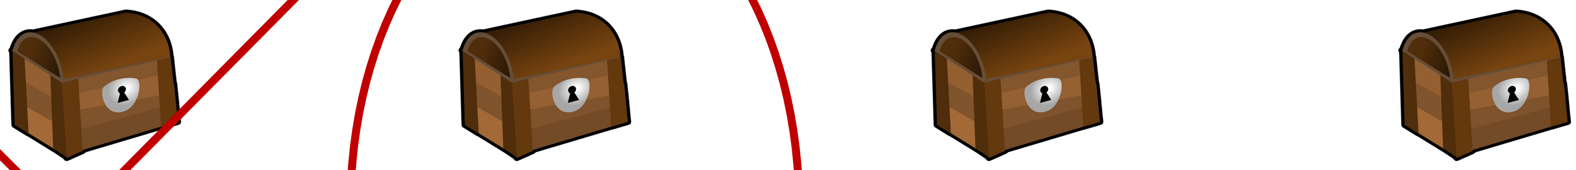

$\mathcal{D}_1 = U[0,1]$ $\mathcal{D}_2 = U[0,1]$ $\mathcal{D}_3 = U[0,1]$ $\mathcal{D}_4 = U[0,1]$

0.4 0.8

reject accept

Let's Play

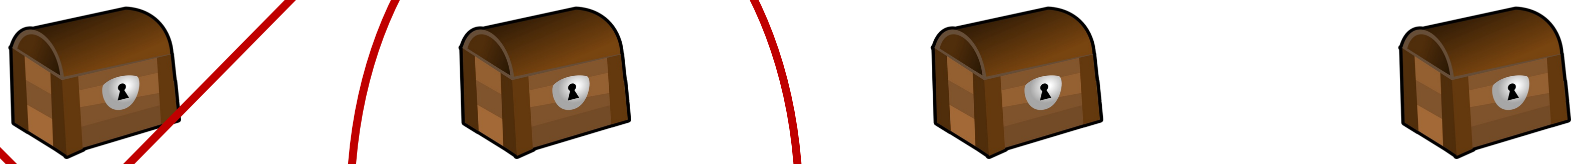

$\mathcal{D}_1 = U[0,1]$ | $\mathcal{D}_2 = U[0,1]$ | $\mathcal{D}_3 = U[0,1]$ | $\mathcal{D}_4 = U[0,1]$

0.4 | 0.8 | 0.1 | 0.9

reject | accept

Let's Play

$\mathcal{D}_1 = U[0,1]$ $\mathcal{D}_2 = U[0,1]$ $\mathcal{D}_3 = U[0,1]$ $\mathcal{D}_4 = U[0,1]$

0.4 0.8 0.1 0.9

reject accept

ALG = 0.8 vs. OPT = 0.9

Optimal Policy

For fixed distributions $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_n$, one can compute the optimal online algorithm by backward induction:

$$VAL_{n:n} := \mathbb{E}_{v_n \sim \mathcal{D}_n}[v_n]$$
$$VAL_{i:n} := \mathbb{E}_{v_i \sim \mathcal{D}_i, \dots, v_n \sim \mathcal{D}_n}[\max\{v_i, VAL_{i+1:n}\}]$$

$\Rightarrow$ Accept v_i if $v_i \geq VAL_{i+1:n}$

Competitive Ratio

Definition. The prophet inequality problem admits a competitive ratio of $\alpha \geq 1$ if, for all distributions $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_n$, there exists an online algorithm ALG such that

$$\mathbb{E}[ALG] \geq \frac{1}{\alpha} \cdot \mathbb{E}[\max_i v_i]$$

Prophet Inequality

Theorem [Krengel-Succheston '77+'78] (+ Garling)

For all distributions $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_n$, there is an online algorithm ALG such that $\mathbb{E}[ALG] \geq \frac{1}{2}\, \mathbb{E}[\max_i v_i]$.

Krengel and Succheston in Oberwolfach

Stronger Version

Theorem [Samuel-Cahn '84]

For all distributions $\mathcal{D}_1, \mathcal{D}_2, \ldots, \mathcal{D}_n$, there is a threshold algorithm ALG_τ such that $\mathbb{E}[ALG_\tau] \geq \frac{1}{2}\ \mathbb{E}[\max_i v_i]$.

Threshold algorithm: set threshold/price τ, accept first $v_i \geq \tau$

Samuel-Cahn (from Gil Kalai's Blog)

Tightness

The factor ½ cannot be improved upon:

Consider the following setting with $n = 2$ random variables:

$v_1 = 1$ w.p. 1, $v_2 = \frac{1}{\epsilon}$ w.p. ϵ and $v_2 = 0$ o.t.w.

a.k.a. "longshot"

Then $\mathbb{E}[ALG] \leq 1$, while

$$\mathbb{E}\left[\max_i v_i\right] = \epsilon \cdot \frac{1}{\epsilon} + (1-\epsilon) \cdot 1 = 2 - \epsilon$$

Sending $\epsilon \to 0$ shows the claim.

Re-Discovery in TCS

- Prophet inequalities are a powerful tool in mechanism design [Hajiaghayi, Kleinberg, Sandholm 2007]
- Prophet inequalities provide a new ``beyond the worst-case'' paradigm for online algorithms

This sparked a whole research field in (theoretical) computer science, exploring applications and extensions of the classic prophet inequality.

Proof of the Classic Prophet Inequality

Prophet Inequality

Theorem [Samuel-Cahn '84]

For all distributions $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_n$, there is a threshold algorithm ALG_τ such that $\mathbb{E}[ALG_\tau] \geq \frac{1}{2}\,\mathbb{E}[\max_i v_i]$.

Threshold algorithm: set threshold/price τ, accept first $v_i \geq \tau$

Prophet Inequality

Theorem [Samuel-Cahn '84]

For all distributions $\mathcal{D}_1, \mathcal{D}_2, \ldots, \mathcal{D}_n$, there is a threshold algorithm ALG_τ such that $\mathbb{E}[ALG_\tau] \geq \frac{1}{2}\ \mathbb{E}[\max_i v_i]$.

Threshold algorithm: set threshold/price τ, accept first $v_i \geq \tau$

Let $p_\tau \coloneqq \Pr[\exists\ v_i \geq \tau]$

Prophet Inequality

Theorem [Samuel-Cahn '84]

For all distributions $\mathcal{D}_1, \mathcal{D}_2, \ldots, \mathcal{D}_n$, there is a threshold algorithm ALG_τ such that $\mathbb{E}[ALG_\tau] \geq \frac{1}{2}\,\mathbb{E}[\max_i v_i]$.

Threshold algorithm: set threshold/price τ, accept first $v_i \geq \tau$

Let $p_\tau := \Pr[\exists\, v_i \geq \tau]$

Actually, different rules work:

Median rule: Set τ such that $p_\tau = \frac{1}{2}$ [Samuel-Cahn '84]

Mean rule: Set $\tau = \frac{1}{2}\,\mathbb{E}\left[\max_i v_i\right]$ [Kleinberg-Weinberg '12]

Samuel-Cahn's Proof

Proof: Recall $p_\tau := \Pr[\exists\ v_i \geq \tau]$.

Samuel-Cahn's Proof

Proof: Recall $p_\tau \coloneqq \Pr[\exists\, v_i \geq \tau]$.

Note that $\max_i v_i \leq \tau + \sum_i (v_i - \tau)^+$ where $x^+ := \max\{x, 0\}$.

Samuel-Cahn's Proof

Proof: Recall $p_\tau \coloneqq \Pr[\exists\, v_i \geq \tau]$.

Note that $\max_i v_i \leq \tau + \sum_i (v_i - \tau)^+$ where $x^+ := \max\{x, 0\}$.

Using this, for any threshold rule,

$$\begin{aligned} \mathbb{E}[ALG_\tau] &= p_\tau \cdot \tau + \textstyle\sum_i \Pr\left[\forall_{j<i}\ v_j < \tau\right] \cdot \mathbb{E}[(v_i - \tau)^+] \\ &\geq p_\tau \cdot \tau + (1 - p_\tau) \cdot \textstyle\sum_i \mathbb{E}[(v_i - \tau)^+] \\ &\geq p_\tau \cdot \tau + (1 - p_\tau) \cdot \left(\mathbb{E}\left[\max_i v_i\right] - \tau\right). \end{aligned}$$

Samuel-Cahn's Proof

Proof: Recall $p_\tau := \Pr[\exists\, v_i \geq \tau]$.

Note that $\max_i v_i \leq \tau + \sum_i (v_i - \tau)^+$ where $x^+ := \max\{x, 0\}$.

Using this, for any threshold rule,

$$\begin{aligned} \mathbb{E}[ALG_\tau] &= p_\tau \cdot \tau + \textstyle\sum_i \Pr\left[\forall_{j<i}\ v_j < \tau\right] \cdot \mathbb{E}[(v_i - \tau)^+] \\ &\geq p_\tau \cdot \tau + (1 - p_\tau) \cdot \textstyle\sum_i \mathbb{E}[(v_i - \tau)^+] \\ &\geq p_\tau \cdot \tau + (1 - p_\tau) \cdot \left(\mathbb{E}\left[\max_i v_i\right] - \tau\right). \end{aligned}$$

For median rule $p_\tau = ½$, and so

$$\mathbb{E}[ALG_{\text{median}}] \geq \frac{1}{2} \cdot \tau + \frac{1}{2} \cdot \left(\mathbb{E}\left[\max_i v_i\right] - \tau\right) = \frac{1}{2}\,\mathbb{E}\left[\max_i v_i\right].$$

Q.E.D.

Proof for Mean Rule

Same proof works for mean rule:

Proof for Mean Rule

Same proof works for mean rule:

Recall: Previous proof showed that, for any threshold rule,

$$\mathbb{E}[ALG_\tau] \geq p_\tau \cdot \tau + (1 - p_\tau) \cdot \left(\mathbb{E}\left[\max_i v_i\right] - \tau\right).$$

Proof for Mean Rule

Same proof works for mean rule:

Recall: Previous proof showed that, for any threshold rule,

$$\mathbb{E}[ALG_\tau] \geq p_\tau \cdot \tau + (1 - p_\tau) \cdot \left(\mathbb{E}\left[\max_i v_i\right] - \tau\right).$$

For mean rule $\tau = \frac{1}{2} \cdot \mathbb{E}\left[\max_i v_i\right]$, and so

$$\mathbb{E}[ALG_{\text{mean}}] \geq \mathrm{p}_\tau \cdot \frac{1}{2}\mathbb{E}\left[\max_i v_i\right] + (1 - p_\tau) \cdot \frac{1}{2}\mathbb{E}\left[\max_i v_i\right] = \frac{1}{2}\mathbb{E}\left[\max_i v_i\right].$$

Q.E.D.

Several Alternative Proofs

- Induction [Hill Kertz '81]
- Stochastic dominance [Kleinberg Weinberg '12]
- Contention resolution [Feldman Svensson Zenklusen '16]
- Sample-based argument [Rubinstein Wang Weinberg '22]

Extensions to Richer Settings

- k-choice [Hajiaghayi Kleinberg Sandholm ’07, Alaei ‘12]
- Matroid and polymatroid constraints [Kleinberg Weinberg ’12, Dütting Kleinberg ‘15, Feldman Svensson Zenklusen ‘16]
- Downward-closed set systems [Rubinstein ’16, Singla Rubinstein ‘17]
- Matching constraints [Gravin Wang ’19, Ezra Feldman Gravin Tang ‘20]
- Combinatorial allocation [Feldman Gravin Lucier ‘15, Dütting Feldman Kesselheim Lucier ’17, Dütting Kesselheim Lucier ‘20, Correa Cristi ‘23]

Prophet Inequalities as a Tool in Mechanism Design

Single-Item Auction

Bidders with stochastic private values $v_i \sim \mathcal{D}_i$

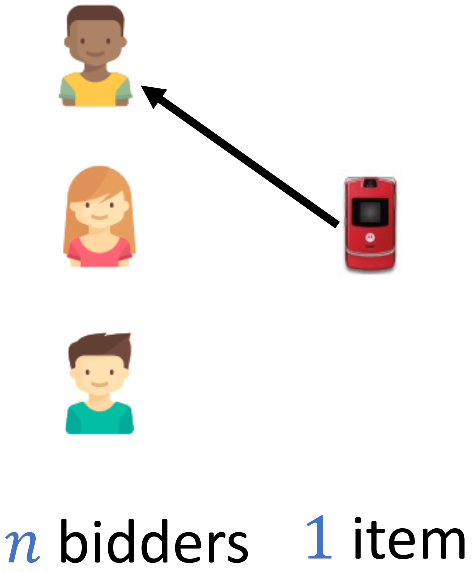

Single-Item Auction

Bidders with stochastic private values $v_i \sim \mathcal{D}_i$

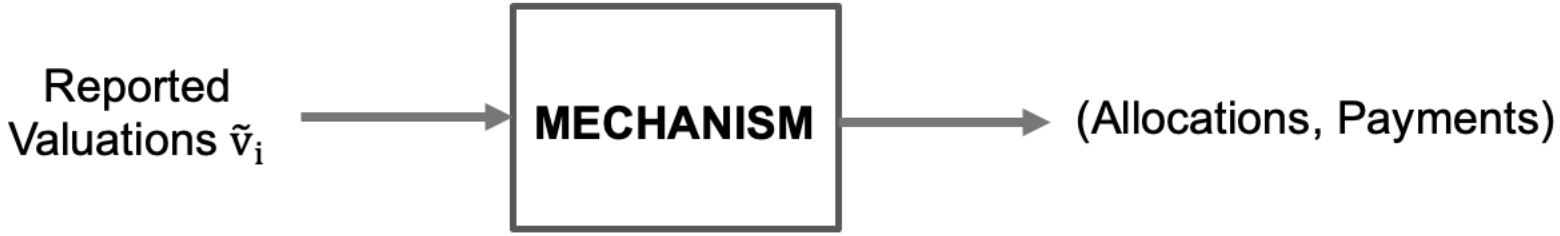

Strategic bidder maximizes utility $:= v_i \cdot 1_{i \text{ gets item}} - \text{payment}_i$

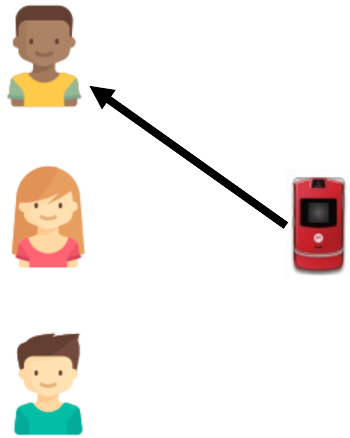

n bidders 1 item

Single-Item Auction

Bidders with stochastic private values $v_i \sim \mathcal{D}_i$

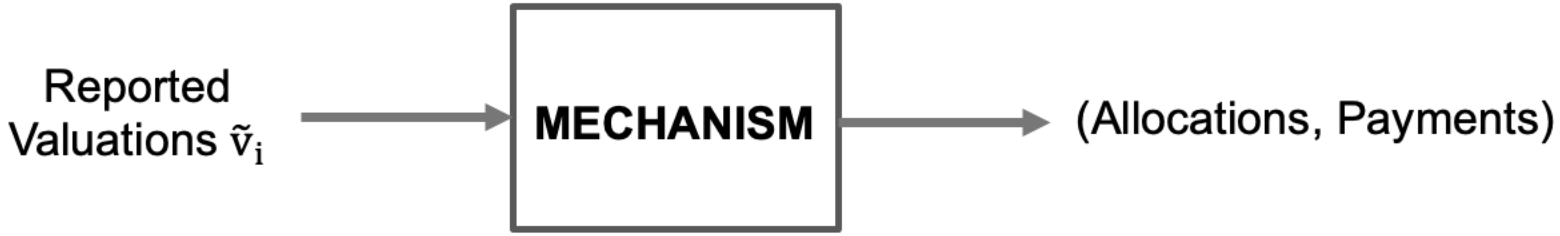

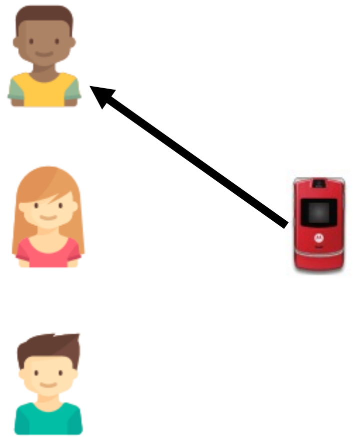

n bidders 1 item

Strategic bidder maximizes utility $:= v_i \cdot 1_{i \text{ gets item}} - \text{payment}_i$

Seek truthful mechanism that

1. Maximizes welfare $:= \mathbb{E}[\sum_i v_i \cdot 1_{i \text{ gets item}}]$
2. Maximizes revenue $:= \mathbb{E}[\sum_i \text{payment}_i]$

Single-Item Auction

Bidders with stochastic private values $v_i \sim \mathcal{D}_i$

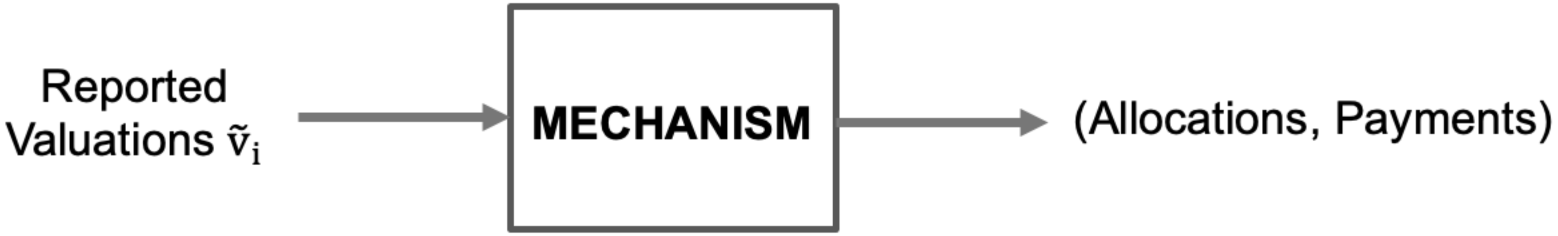

Strategic bidder maximizes utility $:= v_i \cdot 1_{i \text{ gets item}} - \text{payment}_i$

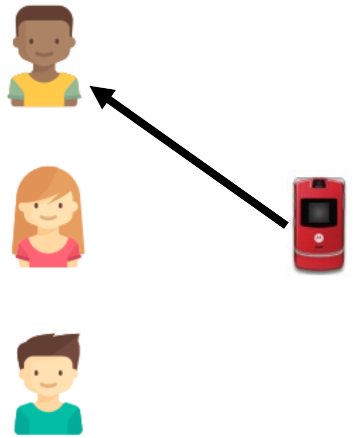

n bidders 1 item

No incentive to misreport

Seek truthful mechanism that

1. Maximizes welfare $:= \mathbb{E}[\sum_i v_i \cdot 1_{i \text{ gets item}}]$
2. Maximizes revenue $:= \mathbb{E}[\sum_i \text{payment}_i]$

Single-Item Auction

Bidders with stochastic private values $v_i \sim \mathcal{D}_i$

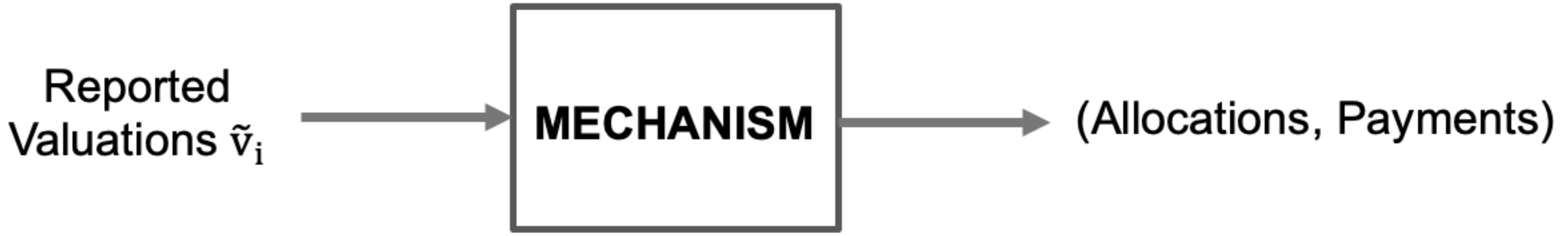

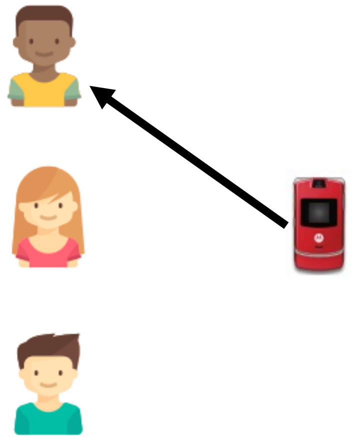

Strategic bidder maximizes utility $:= v_i \cdot 1_{i \text{ gets item}} - \text{payment}_i$

n bidders 1 item

Seek truthful mechanism that

No incentive to misreport

1. Maximizes welfare $:= \mathbb{E}[\sum_i v_i \cdot 1_{i \text{ gets item}}]$
2. Maximizes revenue $:= \mathbb{E}[\sum_i \text{payment}_i]$

1st Price Auction:	Under-report
2nd Price Auction:	Truthful

Single Item: Welfare

- **2nd-Price Auction:** Is truthful and maximizes welfare, but
  - Bidder payments “less natural’’
  - Bidders need to find their values: expensive/impossible
  - Assumes bidders don’t collude
  - …

“The Lovely but Lonely Vickrey Auction”
[Ausubel Milgrom ‘06]

Single Item: Welfare

- **2nd-Price Auction:** Is truthful and maximizes welfare, but
  - Bidder payments "less natural"
  - Bidders need to find their values: expensive/impossible
  - Assumes bidders don't collude
  - …

"The Lovely but Lonely Vickrey Auction"
[Ausubel Milgrom '06]

simpler

- **Posted-price mechanism (PPM):** …
  - Bidders come in arbitrary order
  - Offer them a take-it-or-leave-it price τ
  - Sell to first bidder that is willing to pay price

Single Item: Welfare

- **2nd-Price Auction:** Is truthful and maximizes welfare, but
  - Bidder payments "less natural''
  - Bidders need to find their values: expensive/impossible
  - Assumes bidders don't collude
  - …

"The Lovely but Lonely Vickrey Auction"
[Ausubel Milgrom '06]

simpler

- **Posted-price mechanism (PPM):** …
  - Bidders come in arbitrary order
  - Offer them a take-it-or-leave-it price τ
  - Sell to first bidder that is willing to pay price

Corollary. [Hajiaghayi et al. 2007]
Prophet inequality implies PPM gives
welfare $\geq \frac{1}{2}\mathbb{E}\left[\max_i v_i\right]$.

Single Item: Welfare

- **2^{nd}-Price Auction:** Is truthful and maximizes welfare, but
  - Bidder payments "less natural"
  - Bidders need to find their values: expensive/impossible
  - Assumes bidders don't collude
  - …

"The Lovely but Lonely Vickrey Auction"
[Ausubel Milgrom '06]

simpler

- **Posted-price mechanism (PPM):** …
  - Bidders come in arbitrary order
  - Offer them a take-it-or-leave-it price τ
  - Sell to first bidder that is willing to pay price

Corollary. [Hajiaghayi et al. 2007]
Prophet inequality implies PPM gives
welfare $\geq \frac{1}{2}\mathbb{E}\left[\max_i v_i\right]$.

Truthful and much simpler, but approximation & a stochastic assumption

Single Item: Revenue

Revenue maximization:

- Stochastic private values $v_i \sim \mathcal{D}_i$ (assume regular)
- Optimal mechanism:

  "2nd Price Auction" on virtual value $\hat{v}_i \coloneqq v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$

Theorem. [Myerson 1983]

Opt revenue = $\mathbb{E}\left[\max_i \hat{v}_i^+\right]$

Single Item: Revenue

Revenue maximization:

- Stochastic private values $v_i \sim \mathcal{D}_i$ (assume regular)
- Optimal mechanism:
  "2nd Price Auction" on virtual value $\hat{v}_i := v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$

Disadvantages:

- Highest bidder may loose
- Payments complicated functions of distributions

Theorem. [Myerson 1983]
Opt revenue = $\mathbb{E}\left[\max_i \hat{v}_i^+\right]$

"Simple versus Optimal Mechanisms"
[Hartline Roughgarden '09]

Single Item: Revenue

Revenue maximization:

- Stochastic private values $v_i \sim \mathcal{D}_i$ (assume regular)
- Optimal mechanism:
  ”2nd Price Auction” on virtual value $\hat{v}_i := v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$

Disadvantages:

- Highest bidder may loose
- Payments complicated functions of distributions

2nd Price Auction with Personalized Reserves:

- Set bidder specific reserves
- Sell to highest bidder above reserve
- Payment is max of reserve and highest competing bid

Theorem. [Myerson 1983]
Opt revenue = $\mathbb{E}\left[\max_i \hat{v}_i^+\right]$

“Simple versus Optimal Mechanisms”
[Hartline Roughgarden ‘09]

Theorem. Prophet ineq. implies “simple” auction achieves revenue $\geq \frac{1}{2}\, \mathbb{E}\left[\max_i \hat{v}_i^+\right]$

See [Roughgarden ‘16] book.

Combinatorial Auctions

Stochastic private values $v_i \sim \mathcal{D}_i$; v_i: subset of items $\to \mathbb{R}_{\geq 0}$

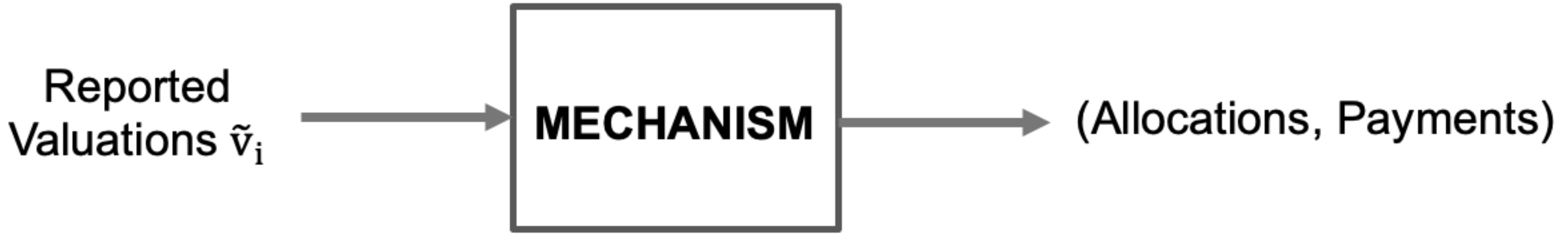

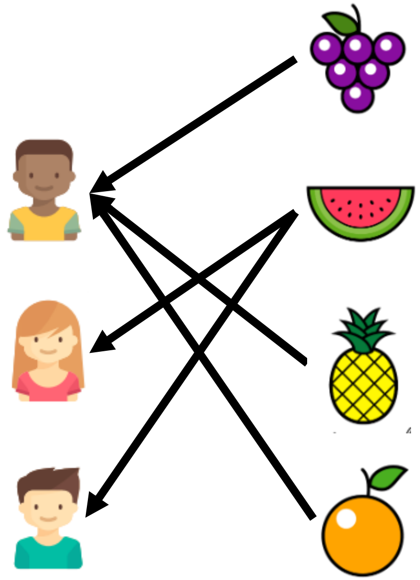

n bidders m items

Strategic bidder maximizes utility $:= v_i(S_i) - \mathrm{payment}_i$

Seek truthful mechanism that

No incentive to misreport

1. Maximizes welfare $:= \mathbb{E}[\sum_\mathrm{i} v_i(S_i)]$
2. Maximizes revenue $:= \mathbb{E}[\sum_i \mathrm{payment}_i]$

Multiple Items: Welfare

- **VCG Mechanism:**

  generalizes 2nd price auction

  - Truthful and maximizes welfare [Vickrey '61, Clarke '71, Groves '73]
  - Not poly-time beyond "simple" classes of values
    - Additive: $v_i(A \cup B) = v_i(A) + v_i(B)$
    - Subadditive: $v_i(A \cup B) \leq v_i(A) + v_i(B)$

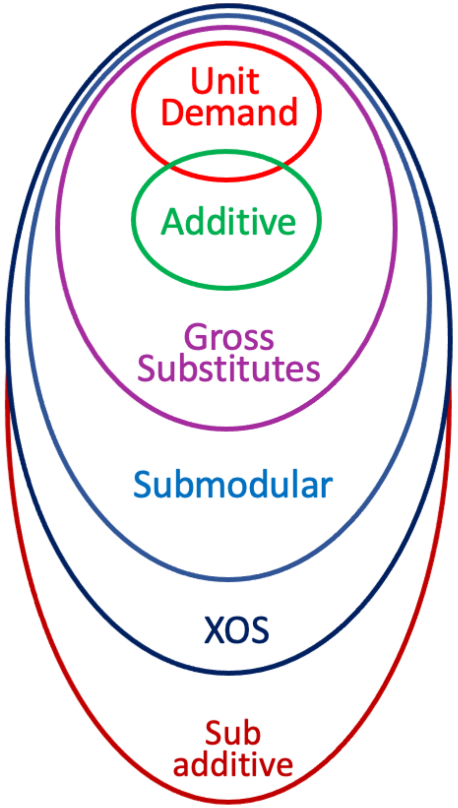

Multiple Items: Welfare

- **VCG Mechanism:** generalizes 2nd price auction
  - Truthful and maximizes welfare [Vickrey ‘61, Clarke ‘71, Groves ‘73]
  - Not poly-time beyond “simple” classes of values
    - Additive: $v_i(A \cup B) = v_i(A) + v_i(B)$
    - Subadditive: $v_i(A \cup B) \leq v_i(A) + v_i(B)$

- **Posted-Price Mechanism (PPM):** truthful and poly-time
  - Set fixed prices $\boldsymbol{p} \in \mathbb{R}^m_{\geq 0}$
  - Buyers come in arbitrary order
  - Select best subset of remaining items:

    $\mathrm{argmax}_{S \subseteq \text{remaining items}} \{v_i(S) - \sum_{j \in S} p_j\}$

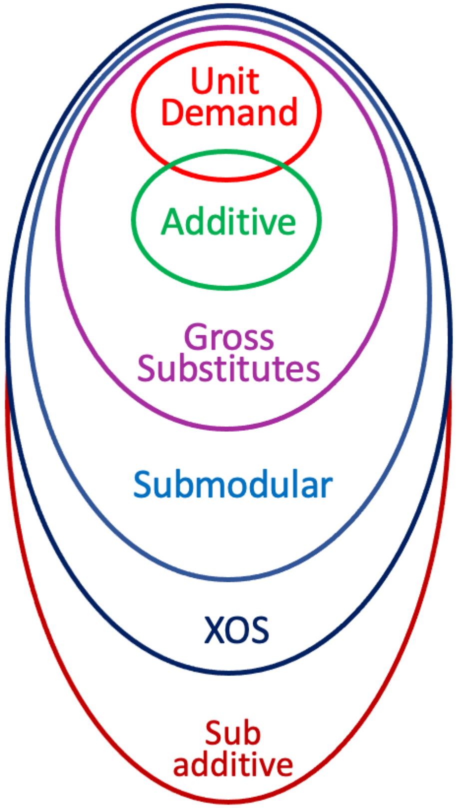

Multiple Items: Welfare

- **VCG Mechanism:**

generalizes 2nd price auction

  - Truthful and maximizes welfare [Vickrey '61, Clarke '71, Groves '73]
  - Not poly-time beyond "simple" classes of values
    - Additive: $v_i(A \cup B) = v_i(A) + v_i(B)$
    - Subadditive: $v_i(A \cup B) \leq v_i(A) + v_i(B)$

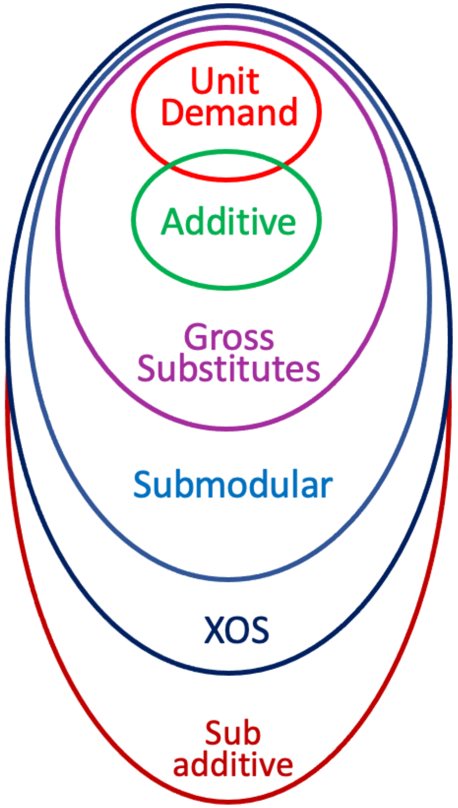

- **Posted-Price Mechanism (PPM):**

truthful and poly-time

  - Set fixed prices $\boldsymbol{p} \in \mathbb{R}^{m}_{\geq 0}$
  - Buyers come in arbitrary order
  - Select best subset of remaining items:

$$\text{argmax}_{S \subseteq \text{remaining items}} \{v_i(S) - \sum_{j \in S} p_j\}$$

Theorem: For welfare max, generalized prophet inequalities imply
- 2 approx for submodular/XOS
- $O(\log\log m)$ approx for subadditive

[Feldman Gravin Lucier '15, Dütting Feldman Kesselheim Lucier '17, Dütting Kesselheim Lucier '20]

Multiple Items: Revenue

- Myerson's mechanism does not work in multi-dimensional settings
  - Single bidder, and multiple items
  - Multiple bidders, and multiple items
  - Multiple combinatorial bidders, and multiple items

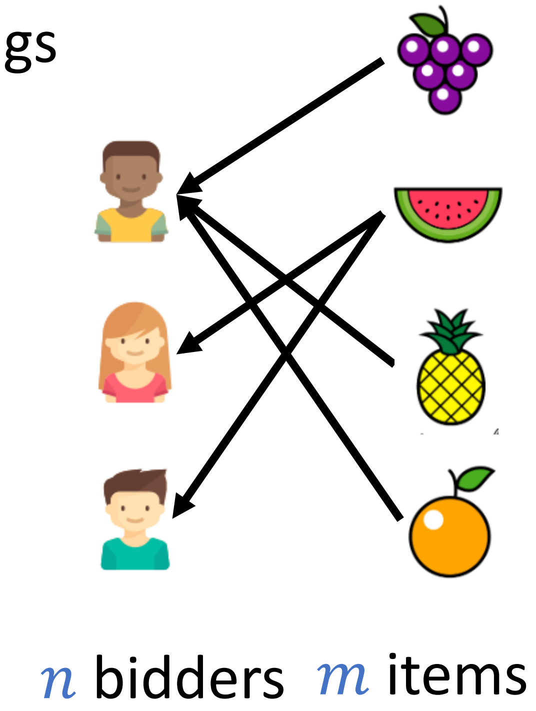

Theorem: For revenue max, generalized prophet inequalities used to get
- 2 approx for submodular/XOS
- $O(\log\log m)$ approx for subadditive

[Chawla Hartline Malec Sivan '10, Chawla Miller'16, Cai Zhao'17, Dütting Kesselheim Lucier'20]

Take Aways

What did we gain?

- Simple, (often) poly-time mechanisms
- Work for both welfare and revenue maximization
- Work for combinatorial auctions
  (& also for other combinatorial feasibility constraints)

What did we lose?

- Stochastic assumption on bidders for welfare (necessary for revenue)
- Approximation algorithms (necessary for combinatorial auctions)

Implications for Online Algorithms

Online Maximization Problems

Inputs arrive one-by-one and must decide immediately and irrevocably

Online Maximization Problems

Inputs arrive one-by-one and must decide immediately and irrevocably

Example: Selecting a large item / bidder

- Max in an online sequence $v_1, \ldots, v_n$

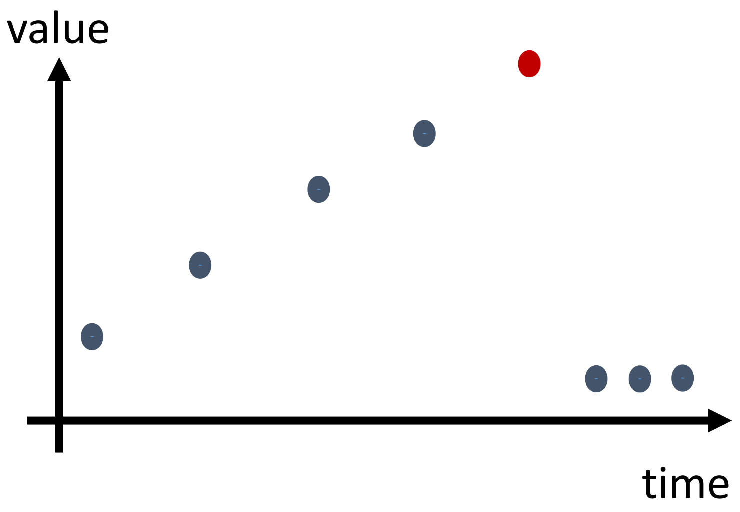

Online Maximization Problems

Inputs arrive one-by-one and must decide immediately and irrevocably

Example: Selecting a large item / bidder

- Max in an online sequence $v_1, \ldots, v_n$

Worst-case arrivals:

- Values determined by adversary
- Best algo selects at random: $\mathbb{E}[ALG] \geq \frac{1}{n} \max_i v_i$

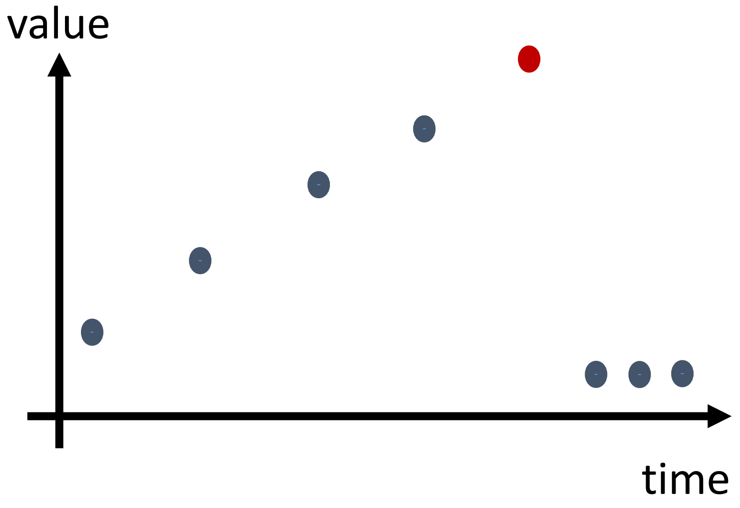

Online Maximization Problems

Inputs arrive one-by-one and must decide immediately and irrevocably

Example: Selecting a large item / bidder

- Max in an online sequence $v_1, \dots, v_n$

Worst-case arrivals:

- Values determined by adversary
- Best algo selects at random: $\mathbb{E}[ALG] \geq \frac{1}{n} \max_i v_i$

Prophet model: Beyond the worst case

- Values from known, non-identical distributions: $v_i \sim \mathcal{D}_i$
- Prophet ineq. gives: $\mathbb{E}[ALG] \geq \frac{1}{2} \mathbb{E}[\max_i v_i]$

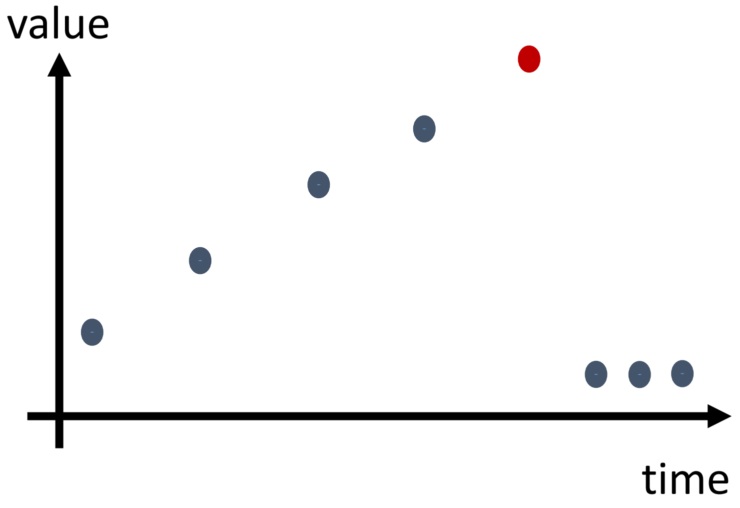

a semi-random model,
stronger than i.i.d.

Overview: Maximization Problems

- k-choice: $1 + o(1)$
  [Hajiaghayi Kleinberg Sandholm '07, Alaei '12]
- Matroid and polymatroid constraints: $O(1)$
  [Kleinberg Weinberg '12, Dütting Kleinberg '15, Feldman Svensson Zenklusen '16]
- General downward-closed: $O(\log n)$ resp. $O(\log n \cdot \log r)$
  [Rubinstein '16, Rubinstein Singla '17]
- Matching constraints: $O(1)$
  [Gravin Wang '19, Ezra Feldman Gravin Tang '20]
- Combinatorial allocation: $O(1)$ (all the way up to subadditive)
  [Feldman Gravin Lucier '15, Dütting Feldman Kesselheim Lucier '17, Dütting Kesselheim Lucier '20, Correa Cristi '23]

Online Minimization Problems

Requirements arrive one-by-one, and must be met while minimizing total cost

Online Minimization Problems

Requirements arrive one-by-one, and must be met while minimizing total cost

Example: Online Steiner Tree

- Given a graph $G = (U, E)$ with edge costs $c_e \geq 0$ and a root $r \in U$
- Vertices $u_1, \ldots, u_n \in U$ arriving online
- Immediately purchase edges to connect u_i to the root r
- Minimize sum of purchased edge costs

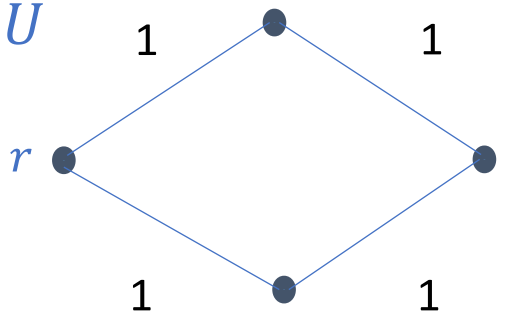

competitive ratio: algorithm's cost to optimal hindsight cost

Online Minimization Problems

Requirements arrive one-by-one, and must be met while minimizing total cost

Example: Online Steiner Tree

- Given a graph $G = (U, E)$ with edge costs $c_e \geq 0$ and a root $r \in U$
- Vertices $u_1, \ldots, u_n \in U$ arriving online
- Immediately purchase edges to connect u_i to the root r
- Minimize sum of purchased edge costs

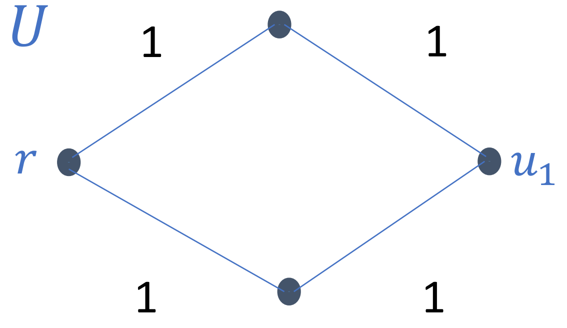

competitive ratio: algorithm's cost to optimal hindsight cost

Online Minimization Problems

Requirements arrive one-by-one, and must be met while minimizing total cost

Example: Online Steiner Tree

- Given a graph $G = (U, E)$ with edge costs $c_e \geq 0$ and a root $r \in U$
- Vertices $u_1, \ldots, u_n \in U$ arriving online
- Immediately purchase edges to connect u_i to the root r
- Minimize sum of purchased edge costs

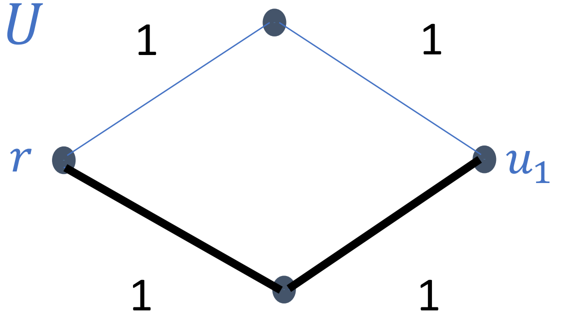

competitive ratio: algorithm's cost to optimal hindsight cost

Online Minimization Problems

Requirements arrive one-by-one, and must be met while minimizing total cost

Example: Online Steiner Tree

- Given a graph $G = (U, E)$ with edge costs $c_e \geq 0$ and a root $r \in U$
- Vertices $u_1, \ldots, u_n \in U$ arriving online
- Immediately purchase edges to connect u_i to the root r
- Minimize sum of purchased edge costs

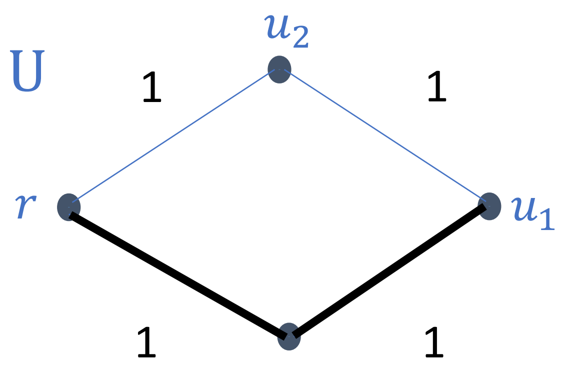

competitive ratio: algorithm's cost to optimal hindsight cost

Online Minimization Problems

Requirements arrive one-by-one, and must be met while minimizing total cost

Example: Online Steiner Tree

- Given a graph $G = (U, E)$ with edge costs $c_e \geq 0$ and a root $r \in U$
- Vertices $u_1, \dots, u_n \in U$ arriving online
- Immediately purchase edges to connect u_i to the root r
- Minimize sum of purchased edge costs

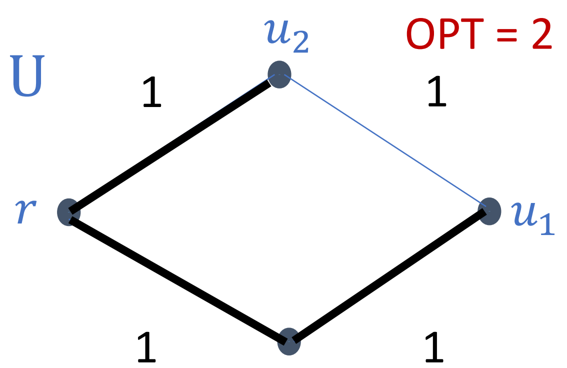

competitive ratio: algorithm's cost to optimal hindsight cost

Online Minimization Problems

Requirements arrive one-by-one, and must be met while minimizing total cost

Example: Online Steiner Tree

- Given a graph $G = (U, E)$ with edge costs $c_e \geq 0$ and a root $r \in U$
- Vertices $u_1, \ldots, u_n \in U$ arriving online
- Immediately purchase edges to connect u_i to the root r
- Minimize sum of purchased edge costs

ALG = 3
OPT = 2

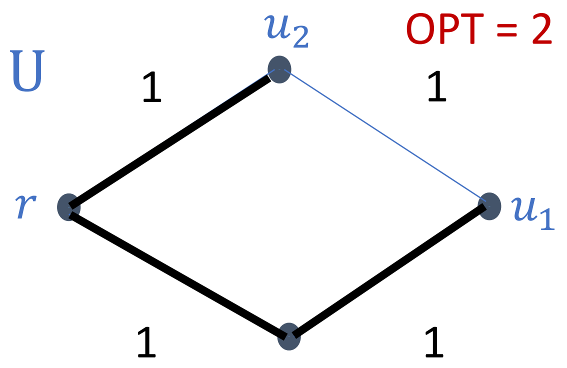

Offline already hard: MST gives 2 approx.

competitive ratio: algorithm's cost to optimal hindsight cost

Online for worst-case arrivals: [Imase Waxman '91]

- No algorithm can be better than $\Omega(\log n)$
- Greedy achieves $O(\log n)$

Online Minimization Problems

Requirements arrive one-by-one, and must be met while minimizing total cost

Example: Online Steiner Tree

- Given a graph $G = (U, E)$ with edge costs $c_e \geq 0$ and a root $r \in U$
- Vertices $u_1, \ldots, u_n \in U$ arriving online
- Immediately purchase edges to connect u_i to the root r
- Minimize sum of purchased edge costs

ALG = 3
OPT = 2

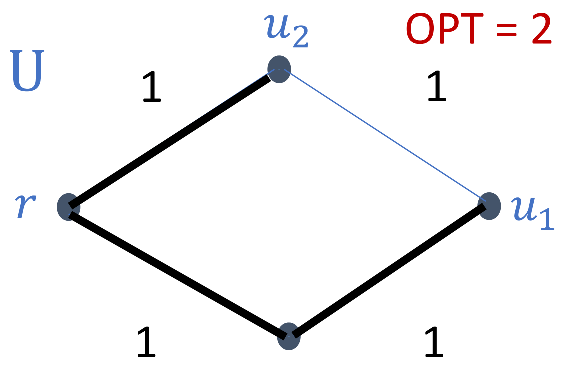

Offline already hard: MST gives 2 approx.

competitive ratio: algorithm's cost to optimal hindsight cost

Online for worst-case arrivals: [Imase Waxman '91]

- No algorithm can be better than $\Omega(\log n)$
- Greedy achieves $O(\log n)$

Can we do better?

Prophet Model

Prophet Steiner Tree:

- Given a graph $G = (U, E)$ with edge costs $c_e \geq 0$ and a root $r \in U$
- Vertices $u_1, \ldots, u_n \in U$ arriving online
- Each vertex $u_i \sim \mathcal{D}$ (known distribution over vertices)
- Immediately purchase edges to connect u_i to the root
- Minimize sum of purchased edge costs

competitive ratio: algorithm's expected cost to expected optimal hindsight cost

Prophet Model

Prophet Steiner Tree:

- Given a graph $G = (U, E)$ with edge costs $c_e \geq 0$ and a root $r \in U$
- Vertices $u_1, \dots, u_n \in U$ arriving online
- Each vertex $u_i \sim \mathcal{D}$ (known distribution over vertices)
- Immediately purchase edges to connect u_i to the root
- Minimize sum of purchased edge costs

competitive ratio: algorithm's expected cost to expected optimal hindsight cost

Other minimization problems:

- Facility location
- Vertex cover

Prophet Model

Prophet Steiner Tree:

- Given a graph $G = (U, E)$ with edge costs $c_e \geq 0$ and a root $r \in U$
- Vertices $u_1, \dots, u_n \in U$ arriving online
- Each vertex $u_i \sim \mathcal{D}$ (known distribution over vertices)
- Immediately purchase edges to connect u_i to the root
- Minimize sum of purchased edge costs

competitive ratio: algorithm's expected cost to expected optimal hindsight cost

Other minimization problems:

- Facility location
- Vertex cover

Theorem: In the prophet model, online Steiner Tree/Facility Location/Vertex Cover admit $O(1)$ competitive ratio.

[Garg Gupta Leonardi Sankowski '08]

Algorithm and Analysis

Algorithm:

1. Take n fresh samples $\hat{v}_1, \hat{v}_2, \dots, \hat{v}_n$, where $\hat{v}_i \sim \mathcal{D}$
2. Construct MST on samples and the root
3. When requirement $v_i \sim \mathcal{D}$ arrives, connect it greedily to MST

Recall:
$\mathbb{E}[\text{MST cost}] \leq 2 \cdot \mathbb{E}[\text{OPT}]$

Algorithm and Analysis

Algorithm:

1. Take n fresh samples $\hat{v}_1, \hat{v}_2, \dots, \hat{v}_n$, where $\hat{v}_i \sim \mathcal{D}$
2. Construct MST on samples and the root
3. When requirement $v_i \sim \mathcal{D}$ arrives, connect it greedily to MST

Recall:
$\mathbb{E}[\text{MST cost}] \leq 2 \cdot \mathbb{E}[\text{OPT}]$

Proof idea:

- Main observation: In expectation,

  greedy cost of connecting v_i to MST $\leq$ cost of connecting $\hat{v}_i$ to closest other vertex in MST

- Summing over i: $\mathbb{E}[\text{total augmentation cost}] \leq \mathbb{E}[\text{MST cost}]$
- $\mathbb{E}[\text{ALG}] = \mathbb{E}[\text{MST cost}] + \mathbb{E}[\text{total augmentation cost}] \leq 2 \cdot \mathbb{E}[\text{MST cost}] \leq 4 \cdot \mathbb{E}[\text{OPT}]$

Q.E.D.

Minimization is Harder

Prophet problem (minimization variant):

- costs $c_i \sim \mathcal{D}_i$ (known distributions),
- need to accept at least one
- **Goal:** minimize expected cost
- **Benchmark:** $\mathbb{E}[\min_i c_i]$

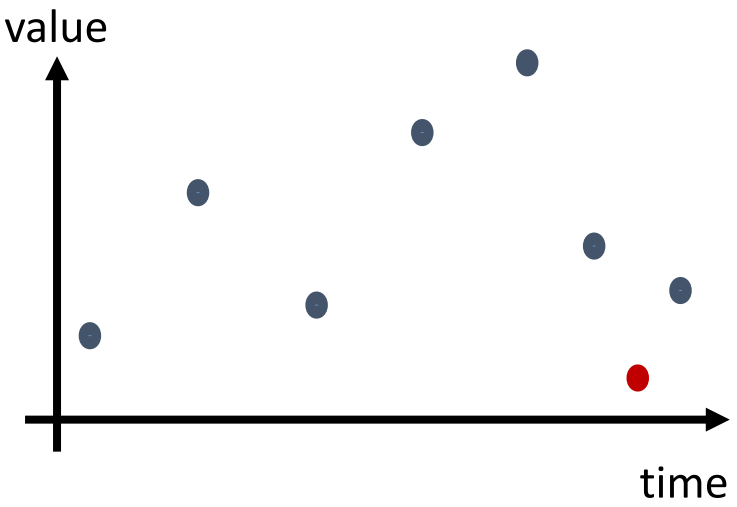

Minimization is Harder

Prophet problem (minimization variant):

- costs $c_i \sim \mathcal{D}_i$ (known distributions),
- need to accept at least one
- **Goal:** minimize expected cost
- **Benchmark:** $\mathbb{E}[\min_i c_i]$

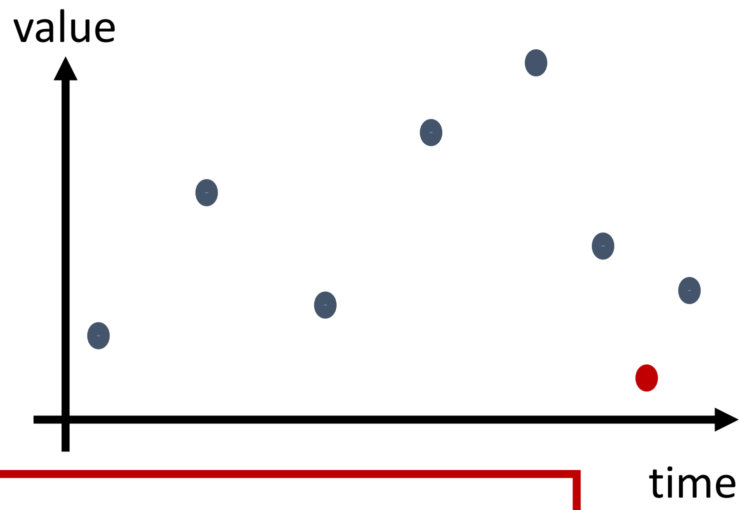

No algorithm can achieve a bounded ratio:

≈ same instance that is hard for max. problem

- $c_1 = 1$ w.p. 1, $c_2 = L$ w.p. $\frac{1}{L}$ and $c_2 = 0$ o.t.w.
- Then $\mathbb{E}[ALG] \geq 1$, while

$$\mathbb{E}\left[\min_i c_i\right] = \frac{1}{L} \cdot 1 + \left(1 - \frac{1}{L}\right) \cdot 0 = \frac{1}{L}.$$

Positive results for i.i.d. costs: [Livanos Mehta ’24]

Take Aways

New "beyond-worst-case" paradigm for online algorithms:

- Many positive results for maximization problems
- To a lesser extent also for minimization problems

Suggestions for your online problems:

- May allow you to go beyond the worst case
- New way of thinking, e.g., when you don't know how to design better worst-case online algorithms

Further Directions

The I.I.D. Case

Theorem [Hill Kertz '82, Correrea-Foncea-Hoeksma-Oosterwijk-Vredeveld '17]

For every distribution $\mathcal{D}$, and n draws $v_i \sim \mathcal{D}$ there exists an algorithm ALG such that

$$\mathbb{E}[ALG] \geq 0.745 \cdot \mathbb{E}[\max_i v_i],$$

and this is best possible.

- There is a sequence of increasing "quantiles" $q_1 \leq q_2 \leq \ldots \leq q_n$ (independent of the distribution)
- The algorithm sets a sequence of decreasing thresholds $\tau_1 \geq \tau_2 \geq \cdots \geq \tau_n$ where $\Pr[v_i \geq \tau_i] = q_i$, and accepts the first $v_i \geq \tau_i$

Alternative Arrival Orders

Given distributions $\mathcal{D}_1, \dots, \mathcal{D}_n$, what if the arrival order is not adversarial?

Free-Order Prophet Inequality:

- Algorithm chooses the arrival order
- Connections to Stochastic Probing

Prophet Secretary:

- Arrival order chosen uniformly at random
- Connections to Secretary Problem

	Lower bound	**Upper bound**
Free-Order	≥ 1.342 [Correa et al. '17]	≤ 1.495 [Correa Saona Zilliotto '19] ≤ 1.379 [Peng Tang '22] ≤ 1.3778 [Bubna Chiplunkar '23]
Prophet Secretary	≥ 1.342 [Correa et al. '17] ≥ 1.366 [Correa Saona Zilliotto '19] ≥ 1.3785 [Bubna Chiplunkar '23]	≤ 1.581 [Esfandiari et al.'15] ≤ 1.495 [Correa Saona Zilliotto '19]

Open question: i.i.d. worst case for free order?

Sample Access to Distributions

What if we only have sample access to distributions?

Single-Sample Prophet Inequality:

- Tight 2 approx. for single item [Rubinstein-Wang-Weinberg '20]
- $O(1)$ approx. for simple-matroids and matching [Azar Kleinberg Weinberg '14, Caramanis et al. '22]
- $O(1)$ approx. for XOS combinatorial auctions [Dütting Kesselheim Lucier Reiffenhäuser Singla '24]

Open questions:

Single-sample $O(1)$ approx for general matroids? For subadditive combinatorial auctions?

For i.i.d. model: Tradeoff between # samples and approx.:

- e for $o(n)$ samples, $\geq \frac{e}{e-1}$ for n samples [Correa Dütting, Fischer, Schewior '19]
- $1.342 + O(\epsilon)$ for $O(n \cdot poly(1/\epsilon))$ samples [Rubinstein Wang Weinberg '20]

Competing w/ Optimal Online Policy

Often the optimal online algorithm via backward induction is computationally infeasible.

Question: What is the best approximation achievable by a poly-time online algorithm, when evaluated against the optimal online policy?

"philosopher inequality"
[Braverman et al 24+]

For example:

- PTAS for "simple" laminar matroids [Anari Niazadeh Saberi Shameli '19]
- 1.96 approximation for online matching [Papadimitriou Pollner Saberi Wajc '21] (and lots of follow up work)
- PTAS for Prophet Secretary [Dütting et al. '23]

Open questions: PTAS for general matroids? Better than 2 approx. for XOS combinatorial auctions?

Summary

- What is a prophet inequality?
  - Statement and proof of the classic prophet inequality

- What's so exciting about prophet inequalities?
  - A powerful tool for mechanism design
  - A new ``beyond worst-case'' paradigm for online algorithms

- On the way: Sample / overview of research landscape

Thanks! Coffee!

Additional Slides

Overview: Online Matching Against Optimal Online Policy

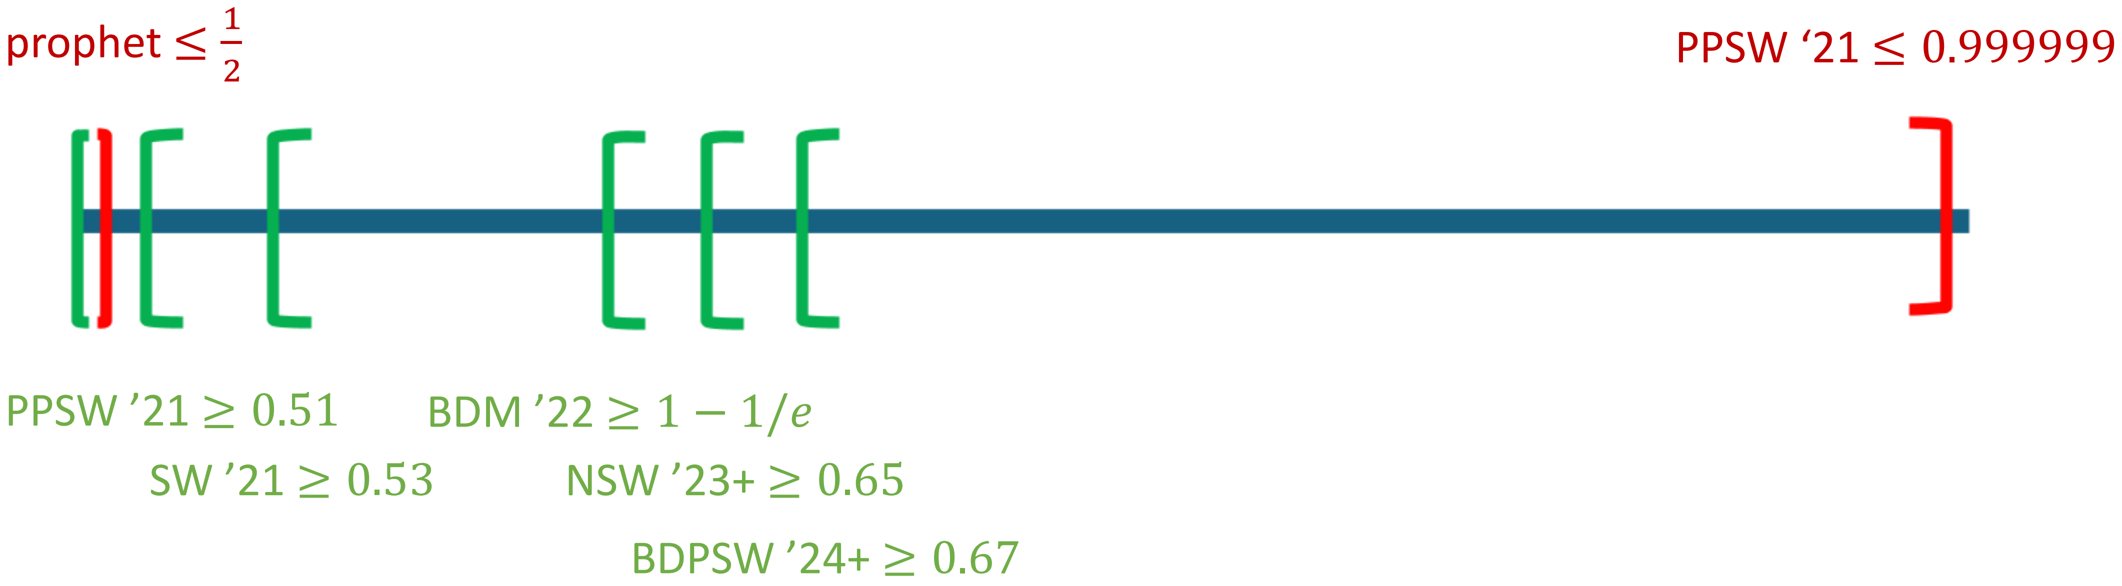

Figure: Fraction of $\mathbb{E}[\text{optimal online policy}]$ achievable with poly-time algorithm.
(Figure due to David Wajc)