Uncertainty in Algorithmic Mechanism Design

Anna R. Karlin

University of Washington

An Example

Classical Optimization Problem:

Maximum Weighted Matching

Input: Weighted Bipartite Graph

Output: Matching that maximizes the sum of matched edge weights.

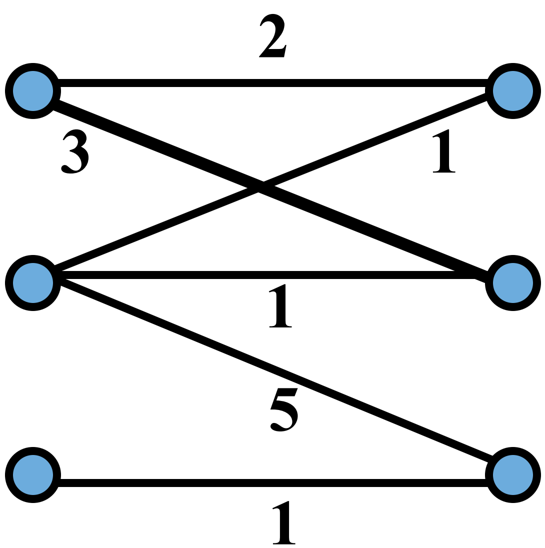

An Example

Classical Optimization Problem:

Maximum Weighted Matching

Input: Weighted Bipartite Graph

Output: Matching that maximizes the sum of matched edge weights.

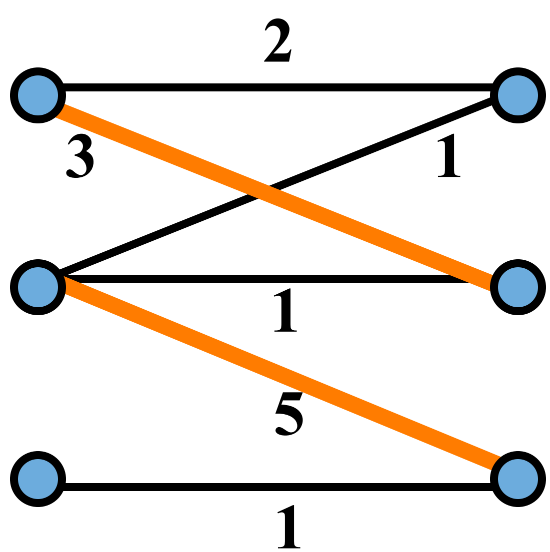

Example Application

Selling advertisement slots

- A search engine has advertising slots for sale
- Advertisers are willing to pay different amount to have their ad shown in a particular slot.

Suppose search engine wants to make as much money as possible.

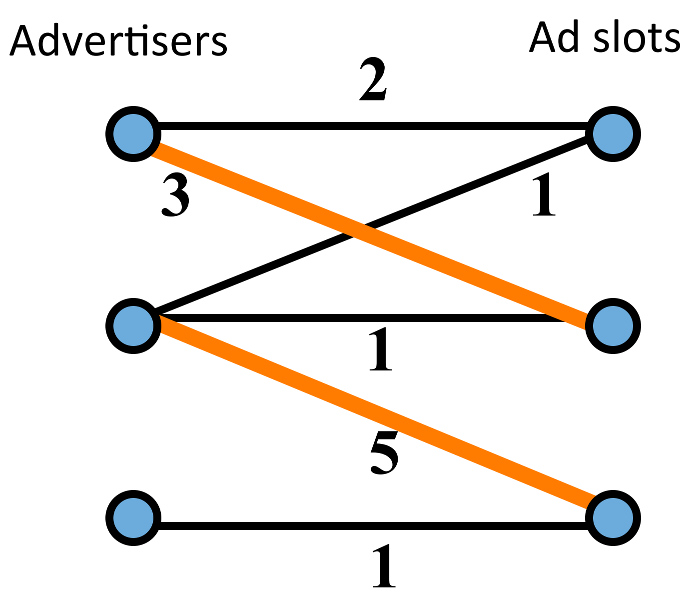

The values are private!

Algorithm must solicit values.

Advertisers may lie to get a better deal.

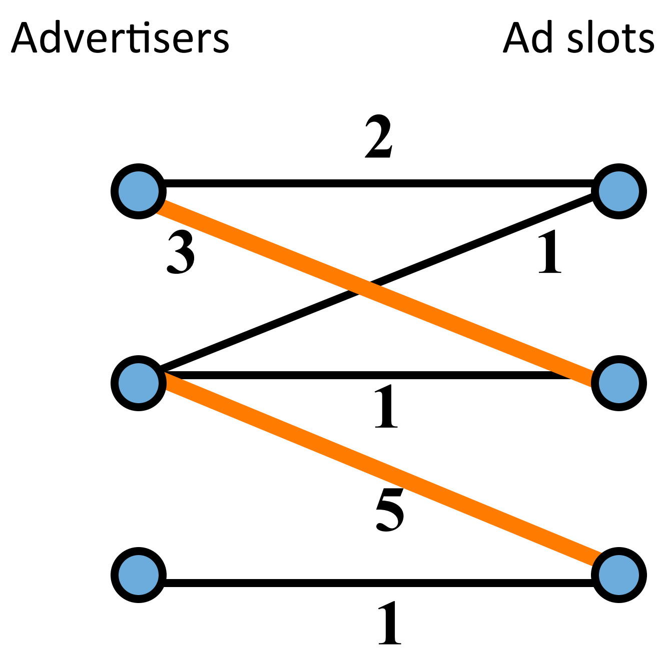

The values are private!

Algorithm must solicit values.

Advertisers may lie to get a better deal.

This is a game!

Google designs the game.

The advertisers play the game.

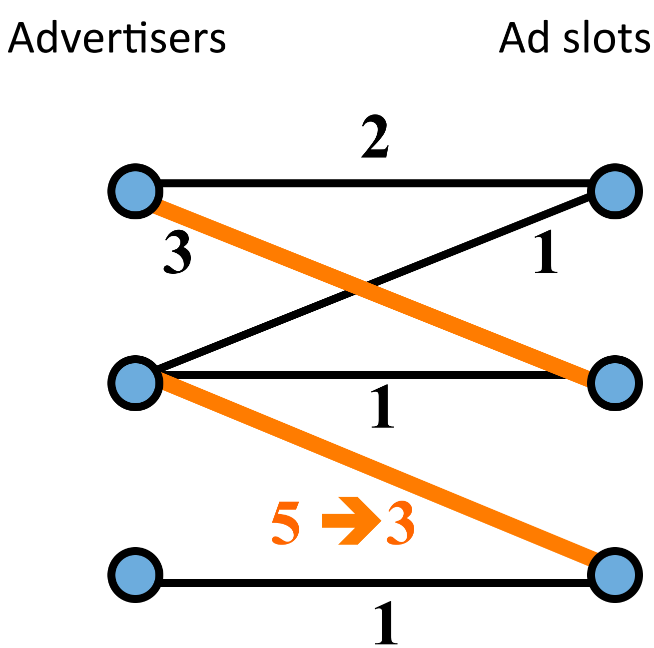

Big Picture

- Many problems where input is the private data of agents who will act selfishly to promote their own best interests
  - Resource allocation
  - Routing and congestion control
  - Electronic commerce

Mechanism design:
How do we optimize in a selfish world?

Real world mechanism design settings

- Sponsored search auctions; display advertising

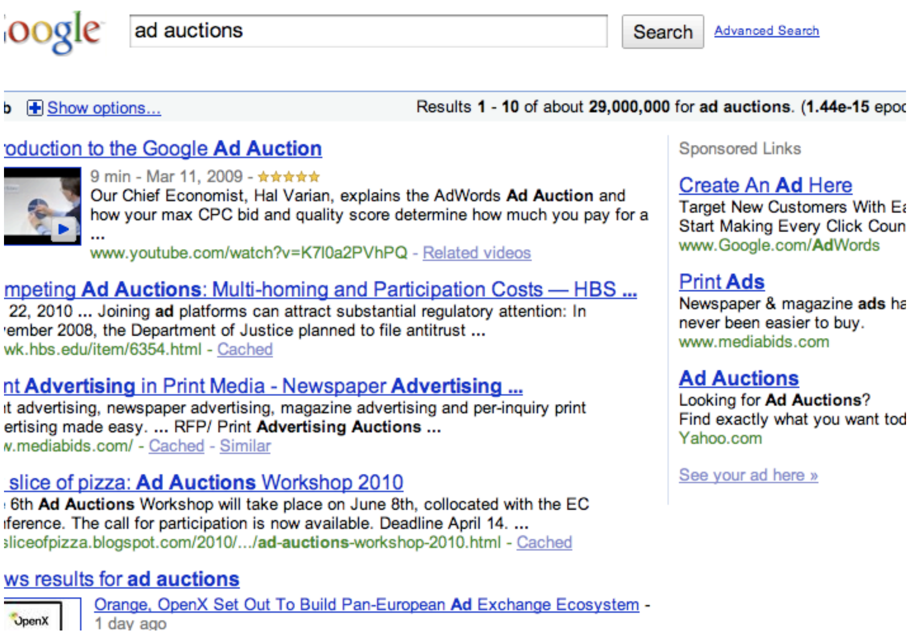

Google revenue in 2015 was approximately $74,500,000,000.

``What most people don't know is that all that money comes in pennies at a time."
Hal Varian, Google Chief Economist

Real world mechanism design settings

- Sponsored search auctions; display advertising
- FCC spectrum auctions
- Kidney exchange
- Healthcare systems
- Recommendation systems
- Routing on the Internet
- Resource allocation in the cloud
- Platform design for a sharing economy
- Energy and electricity markets
- Bitcoin
- Participatory democracy
- Crowdsourcing

What characterizes these problems?

- Many participants with
  - diverse incentives
  - private information of each agent unknown to designer and other agents (maybe even to themselves!)
  - varying atttitudes towards risk.
  - varying degrees of myopia
- Complex optimization problems
- **Dynamic and repeated interaction**

Plan for talk

- Meander…
  - Posted prices via prophet inequalities
  - Prior-independent and prior-free auctions, sample complexity
  - AGT and learning

Apologies: incomplete references.

Applications of prophet inequalities

Prophet inequalities, reminder

- Sequence of prizes $V_1 \sim F_1 \quad V_2 \sim F_2 \quad \ldots \quad V_n \sim F_n$
- You know all the priors.
- See them one at a time and make an irrevocable decision at that moment whether or not to keep it. Once done, game over.
- Compete with prophet who gets expected reward $\mathbb{E}(\max_i V_i)$

Version 1: Take the first prize that is above $\frac{1}{2}\mathbb{E}(\max_i V_i)$

Version 2: Choose threshold t such that

Pr(there is any value is higher than t) = ½

Take the first one above t.

Guarantee: expected value of prize selected $\frac{1}{2}\mathbb{E}(\max_i V_i)$

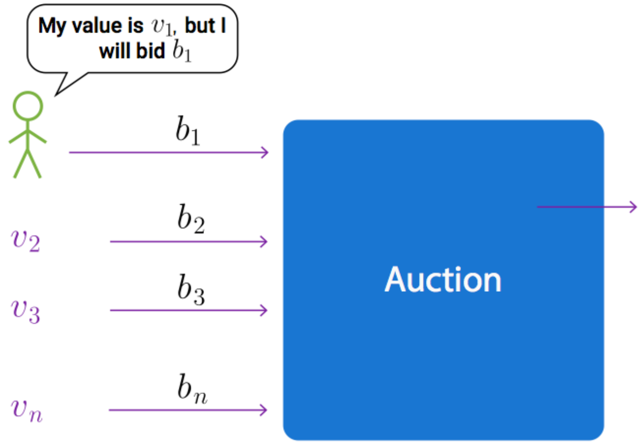

Win

v_1 = 100
v_2 = 80
Truthful bidding:
u_1 = 20
u_2= 0

With collusion, say
b_2= 10
Bidder 1 pays him $50
u_1 = 100-10 -50 = 40
u_2 = 50

Bidder's goal: maximize utility = value - payment

Maximize social surplus: allocate to bidde

Vickrey Second-price Auction: Allocate to highest bidder at second highest bid

Incentivizes truth-telling, i.e. $b_1 := v_1$ no matter what others bid

Issues with Vickrey Auction

- Collusion.
- Can be slow and inconvenient.
- May require more communication than we would like.
- All bidders needs to be present over the time the auction is run.

- But, if we are willing to settle for approximate optimality, and we have a prior…
- That is, we know that $V_i \sim F_i$ independently

Use posted price based on prophet inequality

- Choose threshold t such that Pr(there is any bidder whose value is higher than t) = ½
- Post a price of t; whoever grabs item first, gets it.
- This guarantees that the expected surplus of the outcome (value of winning bidder) is ½ $\mathbb{E}(\max_i v_i)$
- **Features:**
  - Very simple to implement: auction described by one number
  - Very simple for agents to "play".
  - Agent doesn't even need to know exact value.
  - Robust: it doesn't matter if bidders aren't all there at the same time, or come in some arbitrary or even worst-case order.
  - Resilient to collusion.
  - It doesn't matter if distributions change above or below the threshold!

Beautiful generalization
[Feldman, Gravin, Lucier]

Combinatorial auction:

m items, n bidders each with private submodular $v_i(\cdot)$

Goal: allocate items to maximize $\sum_i v_i(S_i)$

Combinatorial auctions, surplus maximization, submodular bidders

- Algorithmic version: 1-1/e approx.
  - [Dobzinski, Shapira][Vondrak][Feige]

- The best known truthful auction achieves $O(\sqrt{\log m})$

approx.
  - [Dobzinski]

- Simultaneous first-price auctions achieve 1-1/e approx. at every BNE
  - [Christodoulou, Kovacs, Schapira][Hassidim, Kaplan, Mansour, Nisan] [Syrgkanis, Tardos]

- [Feldman, Gravin, Lucier] constant approximation via posted prices, given priors.

Given priors, compute posted prices

Utility of bidder = value of bundle - price

$50

$39

$99

$114

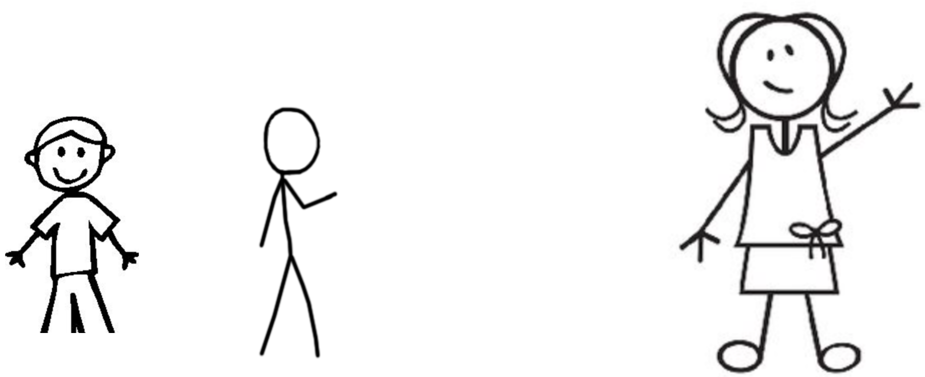

Prices guarantee constant fraction of the optimal expected surplus!

utility of  $= v(\text{kindle} + \text{tap}) - (99 + 114)$

Note

- Here we don't even know how to design a truthful mechanism that gets close to the same approximation ratio as we can get in the algorithmic setting.
- One of the big open questions in mechanism design:
  - to what extent is computability with incentive compatibility harder than without. **[Nisan, Ronen]**

Profit Maximization: prophet inequalities very useful here too!

What if we want to design an auction to maximize auctioneer revenue?

assuming **known** priors
completely solved by **[Myerson]**

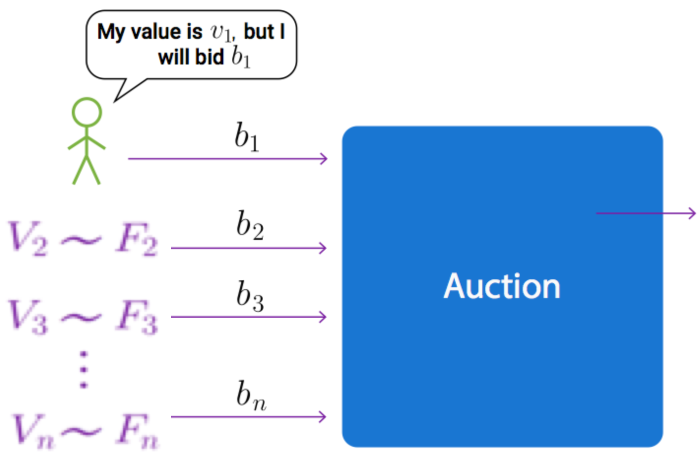

Winner & price

Simplest case: 1 bidder

$$V \sim F$$

Optimal auction: posted price

Best price to offer: $p^* = \text{argmax}_p \quad p(1 - F(p))$

Called monopoly price for the distribution.

Maximizing revenue

Bayesian setting: $V_i \sim F$

Under ***i.i.d. assumption***

[Myerson]

- The optimal auction is just Vickrey with a monopoly reserve price.
- Reserve price r: like an opening bid on eBay
- Truthful auction

Vickrey auction:

- Winner = highest bidder above r, if any
- Price = maximum of r and 2nd highest bid

Beyond i.i.d.

Bayesian setting: $V_i \sim F_i$ independently

[Myerson]

- Ask bidders to report values
- For each bidder, compute virtual value

$$\psi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$$

E(revenue of an auction) = E(virtual value of winner)

Optimal auction*: Allocate to bidder with highest $\psi(v_i)$ (if positive) breaking ties by value.

* distributions regular

Example

- Bidder 1 has value drawn exp(1)
- Bidder 2 has value drawn U[0,1]
- $(v_1, v_2) = (1.5, 0.8)$ =>

  bidder 2 wins (and pays 0.75)

Indep, but not identical →

optimal auction complicated, unintuitive and depends on detailed distributional information!

Is there a simpler, more practical, more robust way, if we are willing to sacrifice optimality?

Prophet inequalities to the rescue!

[Chawla, Hartline, Kleinberg] [Chawla, Hartline, Malec, Sivan]

- Myerson tells us that our revenue will be the expected virtual value of winning bidder.
- To apply prophet inequality, think of $\varphi_i(V_i)^+$

as i^{th} prize.

- Choose t so that $\mathbb{P}(\max_i \varphi_i(V_i)^+ \geq t) = \frac{1}{2}$
- For each bidder this induces a threshold price. $\varphi_i(p_i) = t$
- Sell to first bidder willing to pay his personal reserve price.
- Gives 2 approximation!

Observations

- Although constant virtual price t gives bidder-specific posted prices, it's still way simpler than optimal auction
- All the nice properties we saw earlier.
- Called **oblivious posted pricing** because don't care what order the agents show up.
- If the designer is allowed to consider the agents in the order of his choosing, better approximations possible. Called **sequential posted pricing.**

Sequential Posted Pricing

$V_i \sim F_i$

[Chawla, Hartline, Kleinberg]
[Chawla, Hartline, Malec, Sivan] [Alaei] [Yan]

- Relax the problem by considering the **optimal ex-ante relaxation** i.e., agent i wins with probability q_i

$$\text{maximize} \sum_i R_i(q_i)$$

$$\text{subject to} \sum_i q_i \leq 1$$

$$q_i \geq 0$$

- If distributions nice (regular), then optimal ex-ante pricing is a posted pricing with solution say

$$\sum_i R_i(q_i) = \sum_i q_i p_i$$

$$q_i = 1 - F_i(p_i)$$

Sequential Posted Pricing

[Alaei] [Yan][Agrawal, Ding, Saberi, Y

$F_i(p_i)$

- Example: if distributions regul
  ante pricing is a posted pricin

$$\sum_i R_i(q_i) = \sum_i q_i p_i$$

- Suggests: order by decreasing p_i.
- Then $$\text{Rev} = \sum_i q_i p_i \prod_{j<i} (1 -$$

$$\text{Example: } p_i = 1, q_i = 1/n, \text{ ex-ante opt: } 1$$

$$\text{Rev} = \sum_i \frac{1}{n} 1 (1 - 1/n)^{i-1} \to 1 - 1/e$$

[ADSY]

$$f : 2^N \to \mathbb{R}$$

submodular

$\mathcal{D}$ distn over subsets S of N with marginals q_i

$$\text{Corr gap} = \frac{\mathbb{E}_{S \sim I(D)}(f(S))}{\mathbb{E}_{S \sim D}(f(S))}$$

$$\geq 1 - 1/e$$

Revenue maximization beyond single-item auctions

- All of this even more interesting beyond single item,.e.g. many agents can be served, with combinatorial feasibility constraint, e.g. matroid.
- See **[Chawla, Hartline, Malec, Sivan]...**

Summary

- Many results on posted price auctions (often based on prophet inequalities). Many extensions beyond single item auctions.
- In mechanism design, we love posted prices! Simple, strategyproof, robust, collusion-resistant, less information revealed, don't need everyone present at once, they are what we see in the real world, etc.

- Take-away: Auctions can be approximately optimal without being complex! "simple vs optimal" **[Hartline, Roughgarden]**
- These results depend on deep understanding of optimal auction for single item auctions.
- Take-away: The optimal (or approx optimal) auction serves as benchmark for evaluating more practical designs.
- Huge on-going quest to understand optimal and approximately optimal auctions in ever-more complex settings.

- All of the above depended on the fact that we had an accurate prior in our hands.

Where does the prior come from?

The prior…

- Where does it come from?
  - Previous experience in the market
  - On the fly market analysis
- What if the prior is not accurate? Results potentially sensitive to small errors in the prior.
- What if the prior is changing over time?
- Even if we can get our hands on it, we may not want to redesign the mechanism every time the prior changes

Prior independence
[Dhangwatnotai, Roughgarden, Yan]

- **Prior-independent** mech design; **unknown *F***
  - *Assume values come from some prior, design single auction (with no knowledge of F) so that, no matter what F is*

$$\mathbb{E}_{\mathbf{v}\sim F}(A(\mathbf{v})) \geq c \cdot \mathbb{E}(\mathrm{opt}_F(\mathbf{v}))$$

Not possible without any assumptions on F. However, benign assumptions suffice!

Single item, i.i.d. setting (regular distributions)

- **[Bulow, Klemperer]**

$$\underset{\text{with } n+1 \text{ bidders}}{\mathbb{E}(\text{Rev Vickrey})} \geq \underset{\text{with } n \text{ bidders}}{\mathbb{E}(\text{Rev opt}_F)}$$

- Interpretation:
  - A little more competition is more important than precise knowledge of prior.
  - High value from one extra sample
  - Random price from the distribution almost as good as the optimal reserve price.

Prior-independent results inspired by [BK]

[Dhangwatnotai, Roughgarden, Yan]

- Pick one random bidder and use his bid as the reserve for others – "market research on the fly"
- "single sample mechanism" gets 2 approximation to optimal mechanism tailored to the distributions, no matter what the distribution*
- Many extensions.
- Beginning of flurry of activity on prior-independent auctions.

[Roughgarden, Talgam-Cohen, Yan][Devanur, Hartline, K, Nguyen][Roughgarden, Talgam-Cohen] [Goldner, K]….

* Under regularity assumption, +..

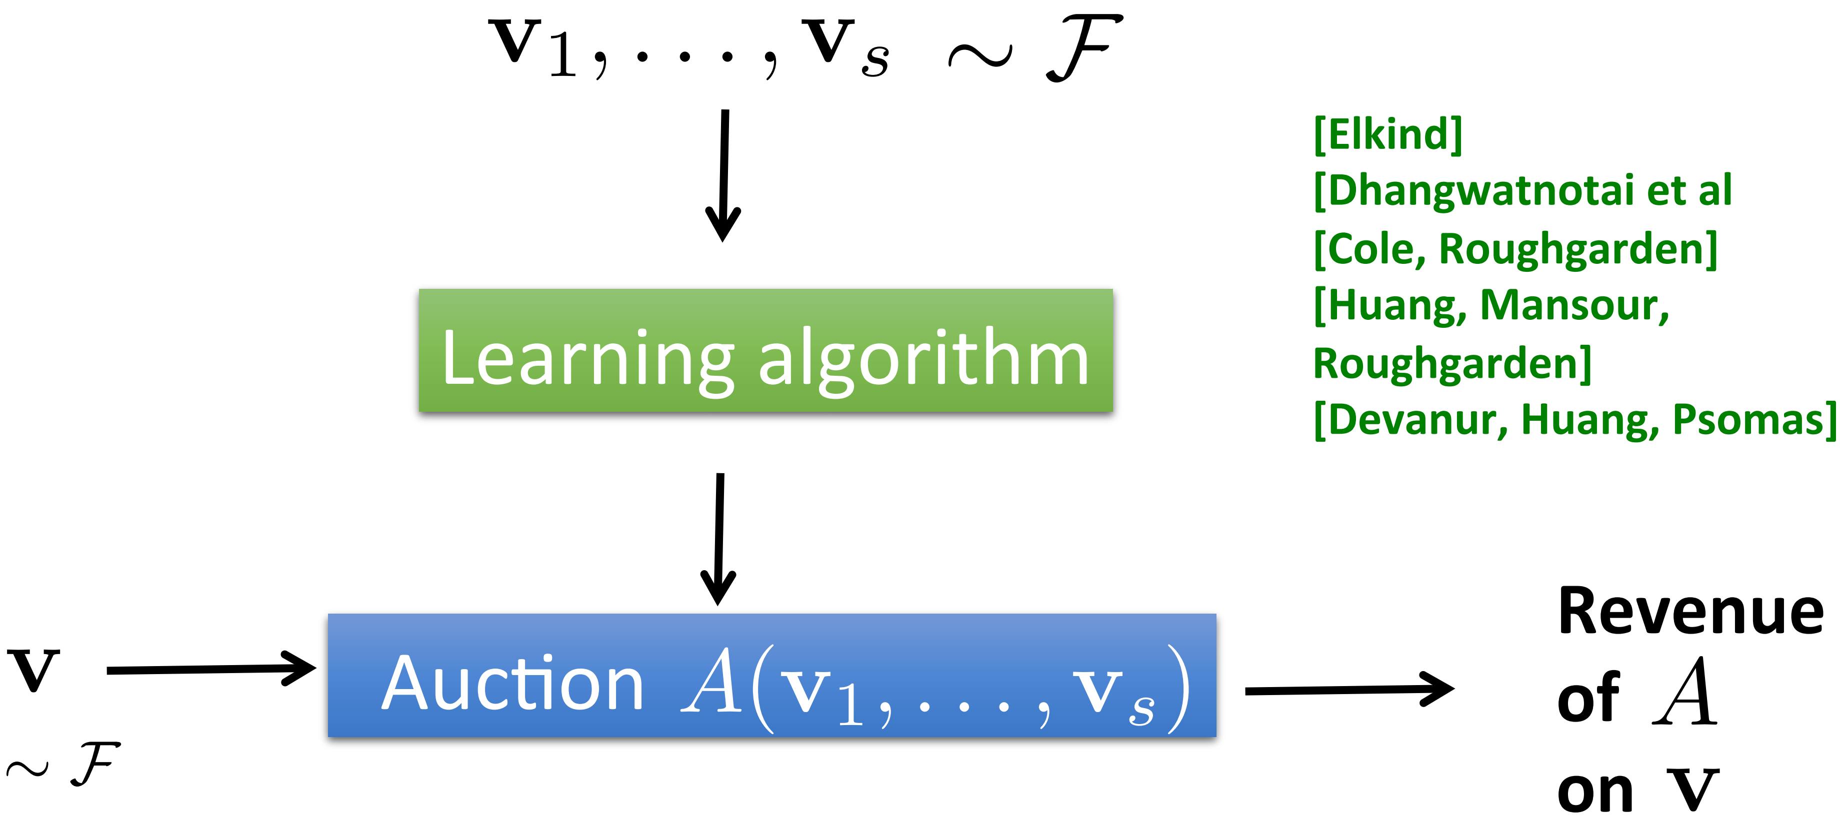

Question: How many samples are necessary and sufficient to get a $1 - \epsilon$ approximation to the optimal expected revenue?

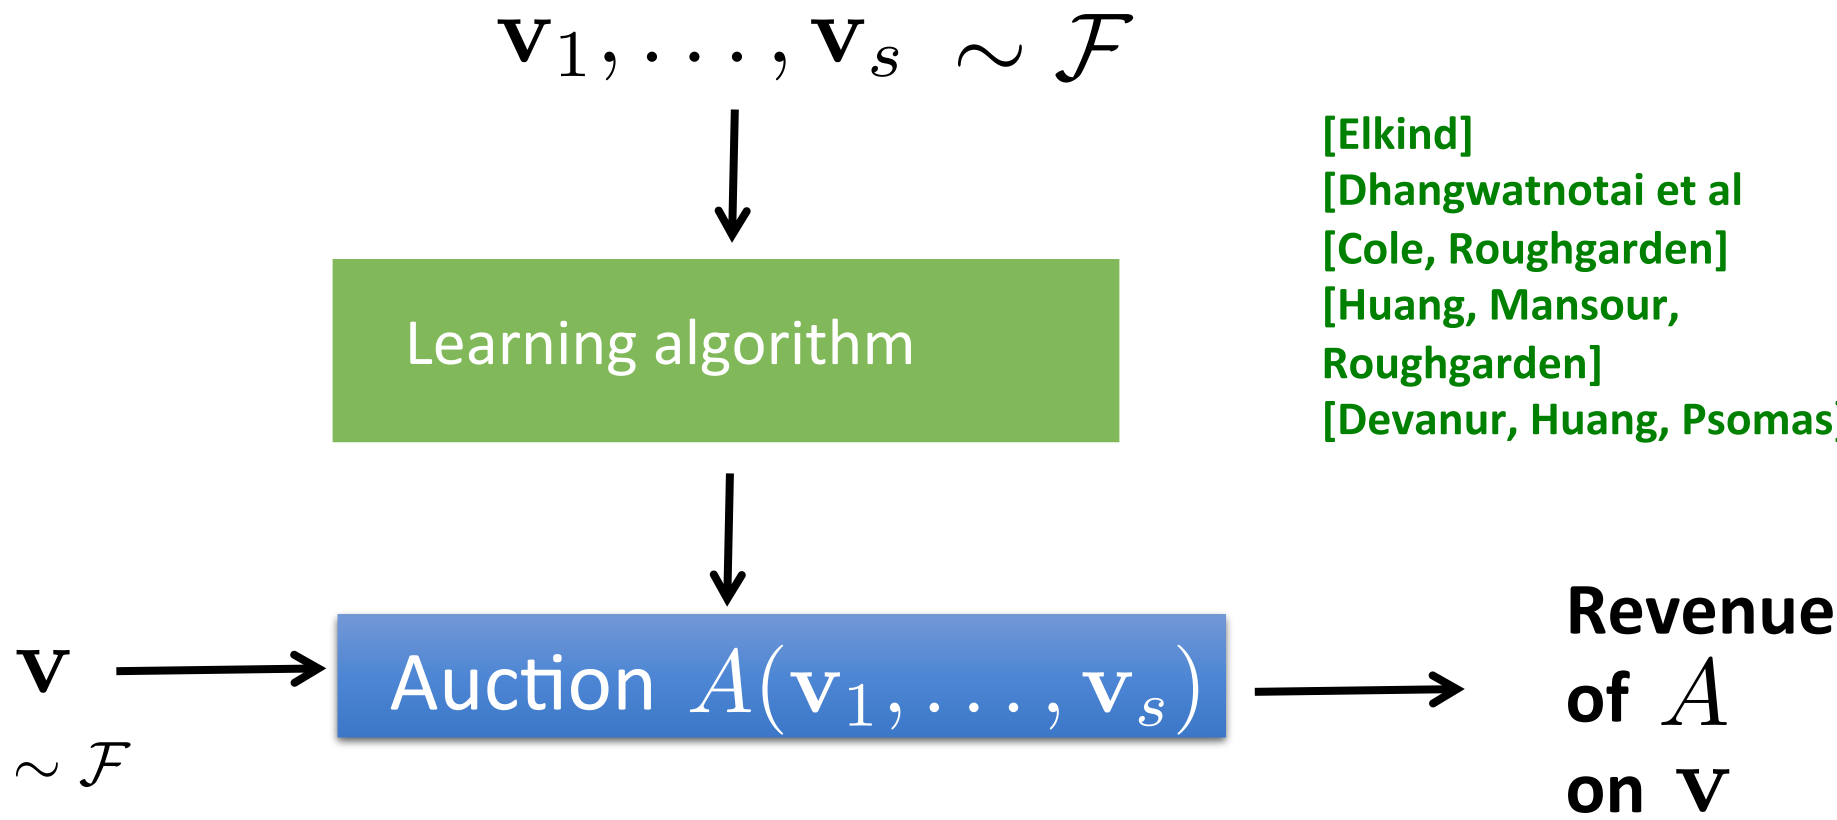

Approach:

- Discretize the value space losing only O(ϵ) fraction of revenue.
- Bounds number of mechanisms need to consider.
- Chernoff bounds => best mechanism in class is $1 - \epsilon$ approximation on new random sample.

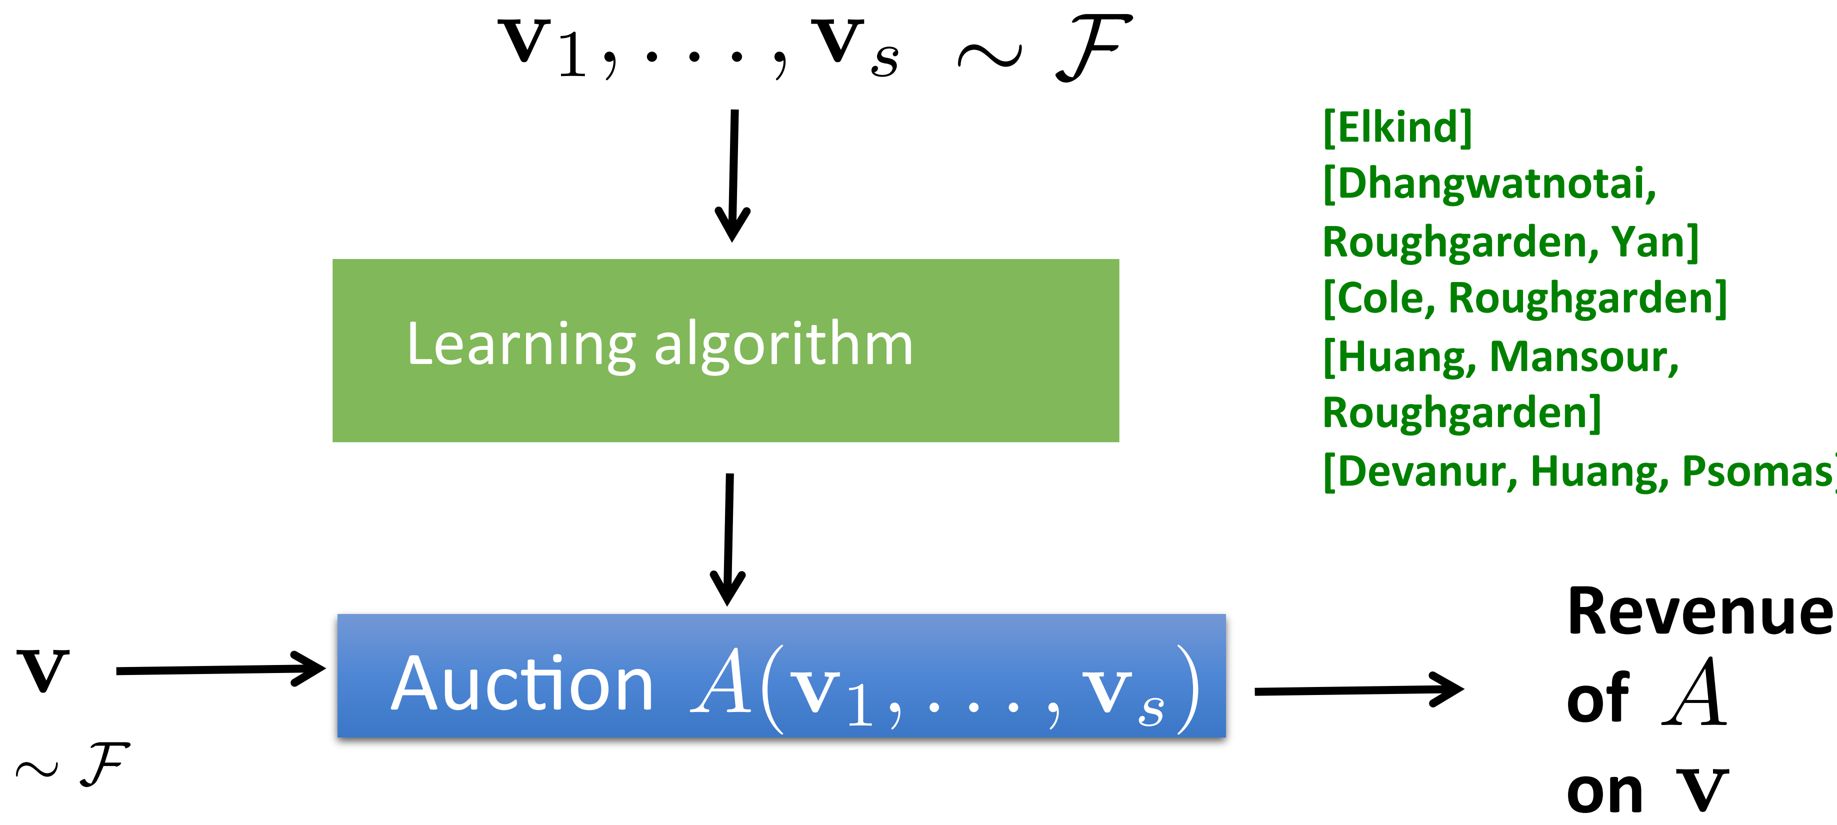

[Elkind]
[Dhangwatnotai, Roughgarden, Yan]
[Cole, Roughgarden]
[Huang, Mansour, Roughgarden]
[Devanur, Huang, Psomas]

Results:

- Without any assumptions on distribution, can't do anything.
- In i.i.d. setting, $\mathrm{poly}(\epsilon^{-1})$ samples suffices.
- In non-i.i.d. setting $n\ \mathrm{poly}(\epsilon^{-1})$ samples suffice.

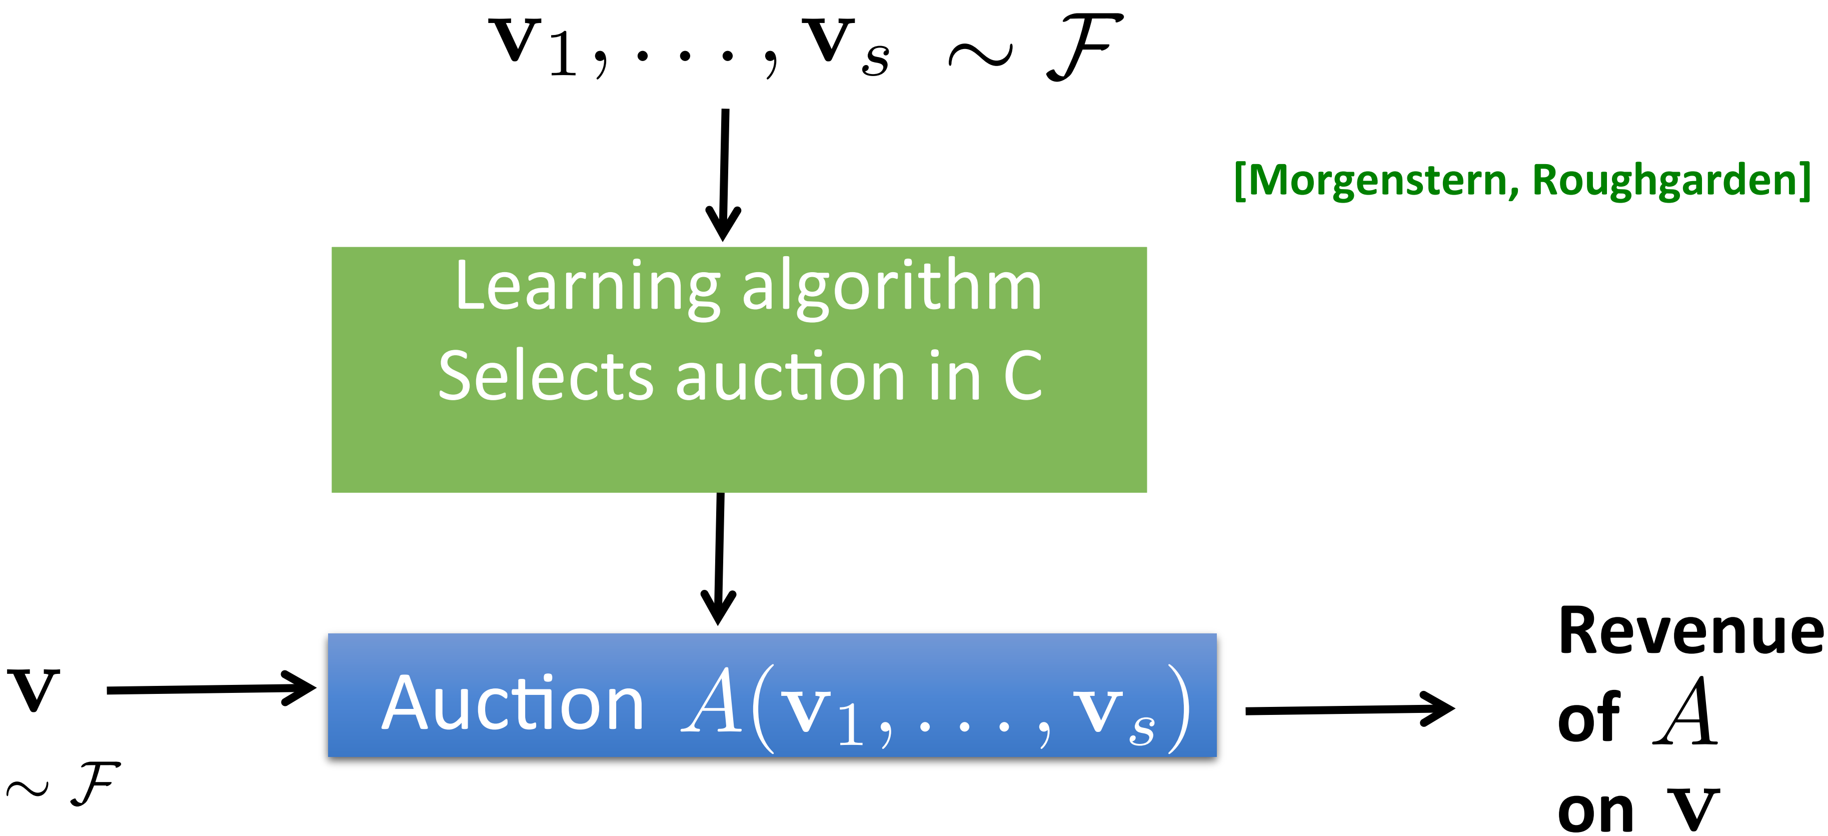

- Apply notion of ``**pseudodimension**'' from learning theory.
- pseudodimension d(C) matches intuitive notion of simplicity/complexity of family of auctions C.
- Implies good sample complexity bounds. Roughly $H^2\epsilon^{-2}d$ samples suffice.
- Results imply that when it's possible to compute a near-optimal simple auction with a known prior, can do so with unknown prior (with polynomial number of samples).

- **Bayesian optimal** mechanism design
  - *Given priors, design mechanism that maximizes or approximates the optimal expected revenue*

- **Prior-independent** opt mechanism design – **unknown F**
  - *Knowing that values are drawn from some large class of distributions, design single auction so that*

$$\mathbb{E}_{\mathbf{v}\sim F}(A(\mathbf{v})) \geq c \cdot \mathbb{E}(\mathrm{opt}_F(\mathbf{v}))$$

- **Prior-free** optimal mechanism design
  - *Design truthful auctions so that for every input*

$$A(\mathbf{v}) = \Omega(\mathrm{B}(\mathbf{v}))$$

where $\mathrm{B}(\mathbf{v})$ *is some profit benchmark.*

We'd like B(v) to be Opt(v), but in the setting of truthful auction design, there is no such thing as an optimal auction!

- **Prior-free** optimal mechanism design
  - *Design auction so that for every input v*

$$A(\mathbf{v}) = \Omega(\mathrm{B}(\mathbf{v}))$$

where $\mathrm{B}(\mathbf{v})$ *is some profit benchmark.*

Example: digital good auctions

[Goldberg, Hartline, Wright] [GHKWS]

- Auctioneer has unlimited supply of items.
- n agents, each has private value v_i for getting one item
- Auction as before takes as input set of bids, and chooses as output, a subset of winners.

- Design a truthful auction that obtains good revenue pointwise.

$$A(\mathbf{v}) = \Omega(\mathrm{B}(\mathbf{v}))$$

- What should benchmark *B(v)* be?
- As I said, there is no optimal truthful auction.

- Good benchmark : optimal fixed price profit : order values $$v_1 \geq v_2 \geq \ldots \geq v_n$$

$$B(v) = \max_i i v_i$$

- Nice thing about this benchmark:
  - If you did have a prior, then this quantity is at least as large as $\mathrm{opt}_F(\mathbf{v})$: sell at price $\max_p p(1 - F(p))$
  - If we can compete with this benchmark, then **simultaneously competitive** with all Bayesian optimal auctions.
- Question: Can we construct a truthful auction that gets revenue *c B(**v**)* on every input?

Competing with best fixed price

- **Truthful mechanism:** price an agent charged can't depend on their own value.
- Suggests: offer best fixed price from rest of values.
- Doesn't work:
  - if bidder high, right price looks low (little revenue)
  - If bidder low, right price too high (rejects)
- **[Goldberg, Hartline, Wright]** No deterministic auction that treats bidders symmetrically can get any constant competitive ratio.

Constant competitive auction*
[GHW][GHKSW]

- Use random sampling:
  - Partition the bidders at random into two sets, S and S'
  - Compute the best fixed price p for S and best fixed price p' for S'
  - Offer price p to bidders in S' and price p' to bidders in S
- Many other results.
- Tight bound of 2.42 on competitive ratio is now known.
- **[Goldberg, Hartline, K, Saks] [Chen, Gravin, Lu]**

* assumes no dominant bidder

AGT and Learning

Issues

- Interactions in the marketplace and beyond are highly dynamic and/or repetitive.
- Agents know that information they reveal today may be used against them tomorrow.
- Lots and lots of data, but it may be strategic.

Repeated interactions very tricky

Example: Fishmonger
sells a fish each day to a buyer.
The buyer's value V is U[0,1]
but doesn't change from day
to day.

[Hart, Tirole][Schmidt]
[Devanur, Peres, Sivan]
[Mohri,Medina]…

Single sale: price ½
Expected revenue: ¼

Buyer's value on day 1 = value on day 2
Seller knows it's a U[0,1] draw

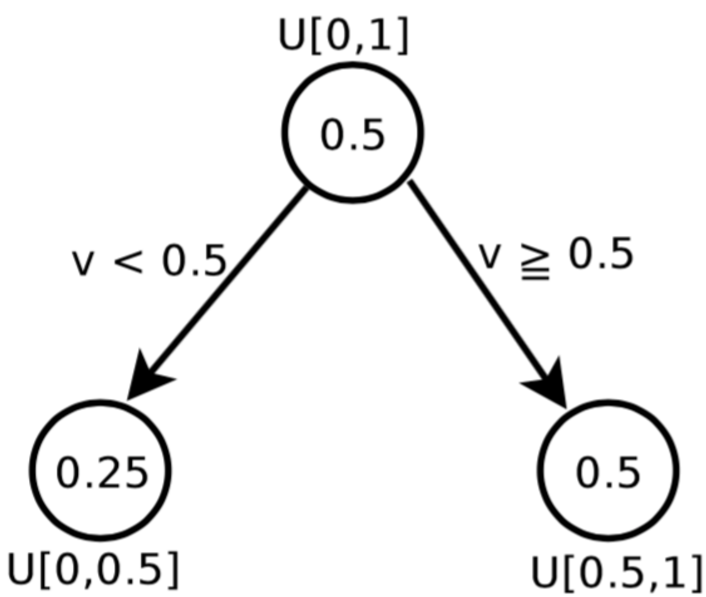

e.g. Buyer with value 0.6

Accepts both days

=> utility = 0.2

Rejects 1st, accepts 2nd

=> utility = 0.35

But if seller could commit…

If not, exp revenue at most 0.45

Remember: buyer utility is value - price

Without commitment, the seller's revenue in n days is o(n)!

[Devanur, Peres, Sivan]

Learning – other questions

- How do you incentivize providers of data to put the effort in to give you high quality data?
- How do you learn from data when the data is both noisy and strategically presented? E.g., if data providers want to influence the outcome of the learning algorithm.

- Huge body of work emerging from various communities

e.g. [Cai, Daskalakis, Papadimitriou]

Incentivizing exploration
[Mansour, Slivkins, Syrgkanis, Wu]
[Frazier, Kempe, Kleinberg, Kleinberg]…

- Waze recommends routes for drivers, but relies on the drivers to do the discovery.
- Retailers like amazon want products to be explored and reviewed, but rely on their users to do this.

- Principal's goal is to collect information about many alternatives: necessitates exploring!

- Users' goal is to select the best alternative for them right now: exploit!!

- How can the principal incentivize the users so that long term learning is as good as it can be? And how good can it be?

[Frazier, Kempe, Kleinberg, Kleinberg] model

- K alternatives ("arms"), each with type that governs reward distribution when selected

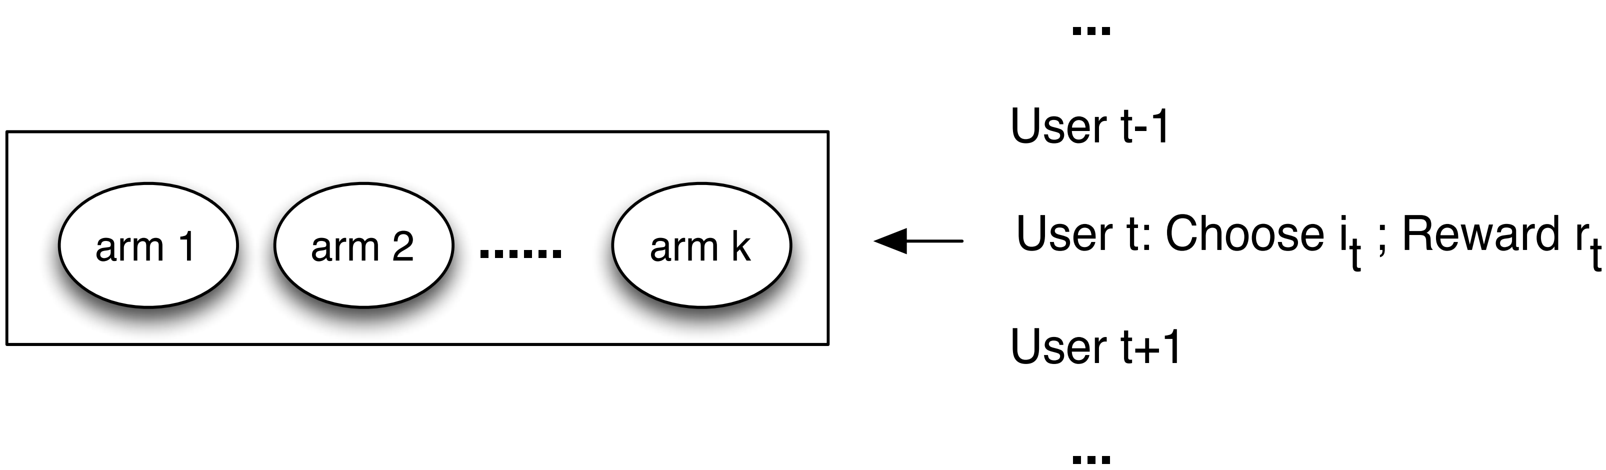

- Users observe all past rewards before making their selection.

(Bayesian) multi-arm bandits

- K alternatives ("arms") (Users observe all past rewards before making their selection.)

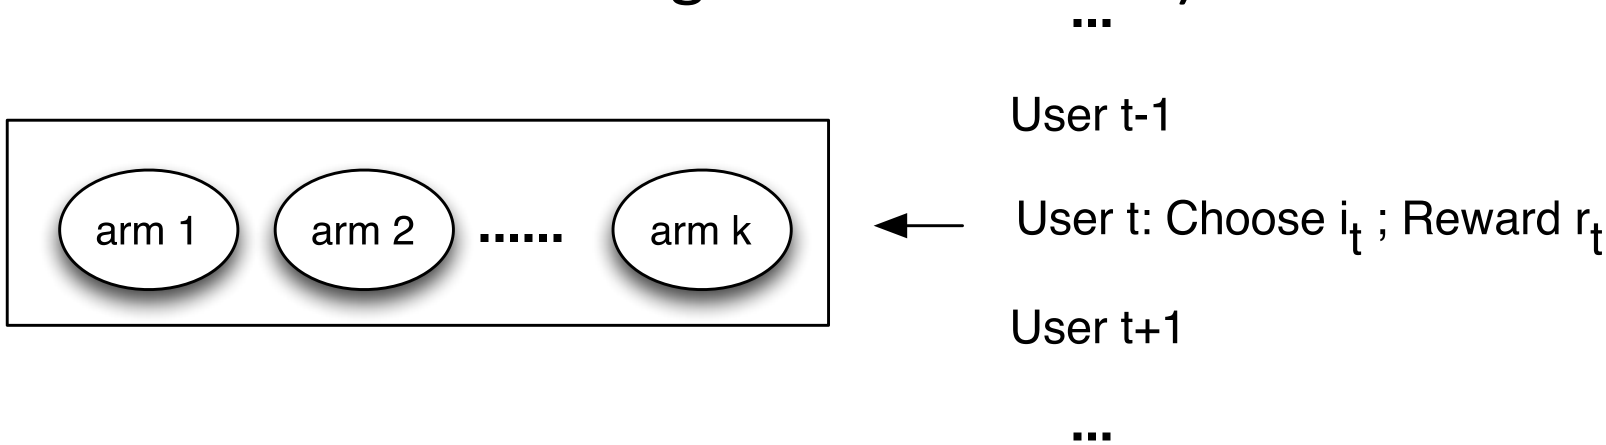

Principal's goal:

$$\text{maximize} \sum_{t=0}^{\infty} \gamma^t r_t$$

User's goal:

$$\text{maximize } r_t$$

(Bayesian) multi-arm bandits

- K alternatives ("arms")

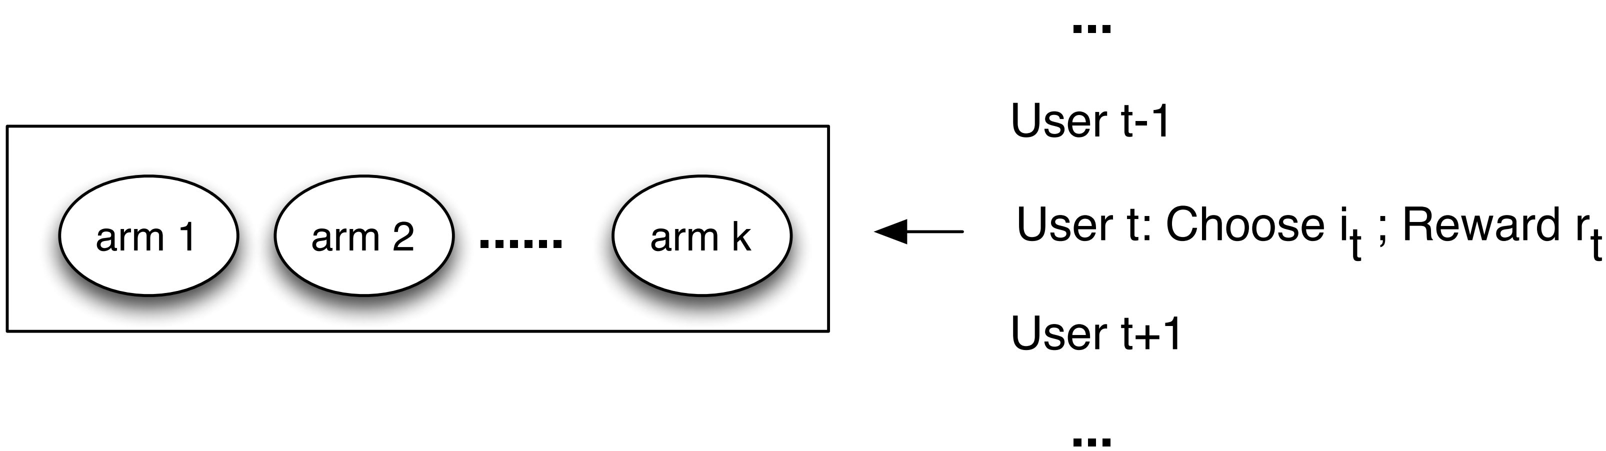

Optimal policy (Gittins index)

$$\text{maximize} \sum_{t=0}^{\infty} \gamma^t r_t$$

Myopic policy

$$\text{maximize } r_t$$

Incentive payments

At time t, announce bonus $c_{t,i} \geq 0$ for each arm i.
User now chooses i to maximize $\mathbb{E}[r_{i,t}] + c_{i,t}$.

Paper precisely characterizes the tradeoff between the incentive payments and the opportunity cost (what you lose in rewards from not playing the optimal policy).

Conclusions

- Many missing topics including:
  - modeling of agents (e.g. value distributions that aren't independent)
  - Mechanism design in more complex settings (multi-parameter)
  - complexity of equilibria – beyond worst case?
  - dynamic mechanism design.

- Many exciting applications that we don't understand!

- For more see
  - Jason Hartline's book
  - Tim Roughgarden's lecture notes and videos
  - Simons Institute Economics and Computation semester workshop videos!

Thank you!!!!!