Learning Auctions:
Reserve prices and Bidding Strategies

Vianney Perchet

GAMENET Training School
Krakow, September 2018

CMLA, ENS Paris-Saclay & Criteo AI Lab, Paris

Billions of auctions are run every day

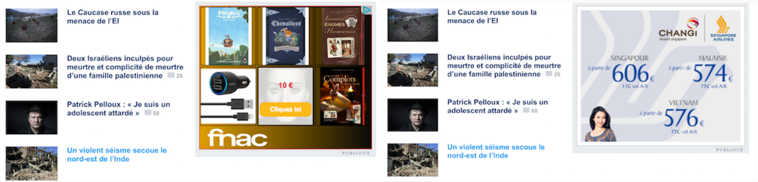

Ads slot sold by publisher through many **auctions**

- "First layer". between ad-exchange platforms
- "Second layer". Digital marketing companies (criteo, google)
- "Third layer". Advertising companies (fnac, singapore airlines)
- Each layer with different mechanism (allocation, payment)

Main Problem: "Data driven optimal auctions"

Learning of the the "optimal" **mechanisms & strategies**

Outline

1. Bayesian one-item auction principles
    - Mechanism design/auction theory (econ.)
2. Computing & learning the Optimal auctions
    - Approximation (c.s.)
    - Statistical learning (m.l.)
3. Earning while Learning
    - Online learning/ bandits (m.l.)
4. Manipulation of truthful auctions
    - Mechanism design/auction theory (econ.)

Bayesian one-item auction principles

General Model

- **One item** is sold to n buyers
- Buyer i has a **private valuation** $v_i \in [0, 1]$

  v_i is drawn from F_i, independently from other bidders
- Typical auction mechanism
  - Bidders place bids $b_i \in [0, 1]$, $b = (b_1, \ldots, b_n)$
  - Bidder $\star$ (or none of them) gets the item
  - Bidders pays $p_i(b) \in \mathbb{R}$ to the seller

> Bidder i expected utility: $\mathbb{E}\left[v_i \mathbb{1}\{i = \star\} - p_i(b)\right]$

- **Revenue of the seller:** $\mathbb{E} \sum_i p_i(b)$

$n = 1$ or Posted Price

Posted price, $n = 1$

Only **one** bidder.

Only auction: Posted price r (”take it or leave it”)

- **Truthful** ! bid $b_1 = v_1$
- $\max \mathbb{E}[\mathbb{1}\{v_1 \geq r\}r] = r\mathbb{P}(v_1 \geq r) = r(1 - F_1(r))$

  Maximum attained when

  $0 = (1 - F_1(r^*)) - r^* f_1(r^*) = -f_1(r^*)\left(r^* - \frac{1-F_1(r^*)}{f_1(r^*)}\right) = -f_1(r^*)\phi_1(r^*)$
- **Optimal price:** zero of virtual valuation function (if regular)

Regular distributions

Virtual valuation function $\phi(v) = v - \frac{1-F(v)}{f(v)}$ is **non-decreasing**

$n > 1$, First Price Auction

First Price Auction

- **Highest bidder** wins the auctions

  He pays the highest (=his) bid to the seller
- Non-**Truthful** auctions: Do not bid your valuations !
  - The utility would always be exactly zero.
- You should **Underbid !**
  - Nash equilibrium strategy of i depends on $F_1, \ldots, F_n$ and v_i
- The seller can put **reserve prices** r_i
  - Bidder i can only win if $b_i \geq r_i$

  Two different allocation rules
  - highest bidder clearing his reserve price (if any) gets the item
  - highest bidder gets the item if he clears his reserve price

Second Price Auction

Second Price Auction

- **Highest bidder** wins the auctions

  He pays the second highest bid to the seller

2nd price auctions are **truthful** ! $b_i = v_i$ is dominating strategy !

Proof:

- Utility of i is $\mathbb{1}\{b_i \geq m_i\}v_i - m_i$ with $m_i = \max_{j \neq i} b_j$
- Bidding $b_i = v_i$ gives $\mathbb{E}\left[\left(v_i - m_i\right)_+\right]$
- **Under-bidding** $b_i < v_i$ does not change payment m_i but lose profitable auctions $(b_i < m_i < v_i)$
- **Over-bidding** $b_i > v_i$ wins unprofitable auctions $(v_i < m_i < b_i)$

"Revenue Equivalence Principle" between 1st and 2nd price auctions

Reserve prices

- **Revenue of seller:** $\mathbb{E}[v^{(2)}]$ with $v^{(k)}$ is the kth max. comp. of v
- It can be improved with a unique reserve price r

$$\mathbb{E}\left[\mathbb{1}\{v^{(1)} - r \geq 0\}\max\{v^{(2)}, r\}\right]$$

- Or even with personalized reserve price r_i

$$\mathbb{E}\left[\mathbb{1}\{(v - r)^{(1)} \geq 0\}\max\{v^{(2)}, r_\star\}\right]$$

- The **optimal reserve prices** depend on $F_1, \dots, F_n$

2nd price auction with reserve prices still truthful

Myerson ”Optimal” Auction

Revenue Maximizing Auctions

- Myerson revenue maximising auctions
  - Compute i virtual value: $\phi_i(b_i) \quad (= \phi(v_i)$ because truthful)
  - The bidder with **highest non-negative virtual value** wins (if any)
  - He pays the **smallest winning bid** $\phi_i^{-1}(\max\{\phi^{(2)}(b), 0\})$
- Upsides:
  - "Simple" mechanism (if ϕ_i are "simple")
  - **Symmetric** bidder ($F_1 = F_2 = \ldots$):

    2nd price auctions with unique reserve price

    Same reserve price, no matter n
- Downsides:
  - Complicated to compute with asymmetric bidders.
  - Payments and allocations are "weird"

    The highest bidder/valuation might not win the auction

Computing & learning the Optimal auctions

Approximation of Myerson Auctions

2nd price auction with reserve price

- Myerson auction maximises revenue but complicated

  Good, simple approximation of it ?

2nd price auctions with **reserve prices** are 1/2-optimal !

- Which reserve prices ?
  - Find $t \geq 0$ such that $\mathbb{P}\left\{ \max_i \phi_i^+(v_i) \geq t \right\} = \frac{1}{2}$
  - Set $r_i = \phi_i^{-1}(t)$
- Direct consequence of **prophet inequality**
  - $\pi_1 \sim G_1, \ldots, \pi_n \sim G_n$. A prophet get $\mathbb{E} \max \pi_i$
  - See sequentially the quality (π_1 first, then π_2, ..).
  - Take it or leave it for ever. Can only take 1 item.
  - Threshold strategy: take if $\pi_i \geq \theta$ with $\mathbb{P}(\max \pi_j < \theta) = 1/2$

    $\frac{1}{2}$- **optimal** ! i.e, gets at least $\mathbb{E} \max \pi_i / 2$

M-Level Auctions

- Define **M** different **thresholds** $r_i^{(M)} \leq r_i^{(M-1)} \leq \ldots \leq r_i^{(1)}$
- After the bids b_i $(= v_i)$
  - Level $\ell_i \in \{1, ..., M, M+1\}$ is s.t. $r_i^{(\ell_i)} \leq v_i \leq r_i^{(\ell_i - 1)}$
  - **Winner:** bidder with lowest index smaller than M (if any)
    Ties broken lexicographically
  - **Payment:** smallest winning bid
- Discretization: Myerson auction = ∞-level auctions

Approximation

M-level auctions are $1 - O(\frac{\log(M) \log\log(M)}{M})$ optimal (actually even better)

Ratio of $1 - \varepsilon$, needs $M \sim 1/\varepsilon$ thresholds

Sample Complexity of Learning Auctions

Data-driven optimal auctions

- Myerson or optimal reserve prices depends on F_i **unknown!**
- Past i.i.d. bids/valuations $v_i^t \sim F_i$, $t = 1, \dots, T$
- Construct the "best" mechanism within a given class $\mathcal{A}$ (Myerson, 2nd price, M-levels...), **based on data**
- "**Empirical Revenue Maximizer**"

$$\widehat{A}_T = \arg\max_{A \in \mathcal{A}} \sum_{t=1}^{T} \text{revenue}[A](v^t)$$

- **Control the error** $\mathbb{E}\left[\text{revenue}[\widehat{A}_T](v)\right] - \mathbb{E}\left[\text{revenue}[A^*](v)\right]$

Flavor of Statistical Learning

Binary classification

- **Binary Classification.**
  - joint distribution $(X, Y) \in \mathbb{R}^k \times \{0, 1\} \sim \mathcal{D}$,
  - family of classifiers $\mathcal{F} \subset \{f : \mathbb{R}^k \to \{0, 1\}$
  - I.i.d. data $(X_1, Y_1), (X_2, Y_2), \ldots (X_T, Y_T) \sim \mathcal{D}$
  - Find the **risk minimizer** $\text{arg min}\, \mathbb{E}\Big[\mathbb{1}\{f(X) \neq Y\}\Big] = R(f)$
- **Empirical Risk Minimization**
  - $\widehat{f}_T = \text{arg min}_{f \in \mathcal{F}} \frac{1}{T} \sum_{t=1}^{T} \mathbb{1}\{f(X_y) \neq Y_t\} = \widehat{R}_T(f)$
- **Example.**
  - X are pictures of cat/dog $(= 0/1)$
  - Dataset of labelled pictures and a possible class of classifier
  - Find the best automatic labelling of cats/dogs

Controlling the error - easy cases

- Hoeffding inequality For a fixed f
  - Exponential decay $\mathbb{P}\left\{\left|\widehat{R}_T(f) - R(f)\right| \geq \varepsilon\right\} \leq e^{-2T\varepsilon^2}$
  - Need $\frac{\log(1/\delta)}{2\varepsilon^2}$ samples for ε-error with proba $1-\delta$
- Union bound if $\mathcal{F}$ is finite
  - Need $\frac{\log(|\mathcal{F}|/\delta)}{2\varepsilon^2}$ samples for ε-error with proba $1-\delta$
- **What if $\mathcal{F}$ is infinite ?**
  - It is not $|\mathcal{F}|$ that really matters, but **the number of possible different labelings of T data points that $\mathcal{F}$ can generate**

VC-dimension + Sauer-Shelah lemma

There exists $d \in \mathbb{N} \cup +\infty$ such that the number of different labelings grows as T^d (instead of 2^T)

Example If $\mathcal{F}$ are linear classifier, $d = k + 1$

A small illustration (taken from Wikipedia)

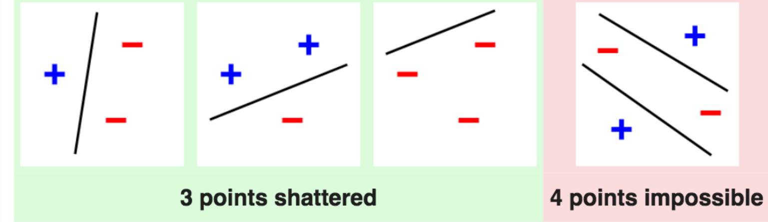

- For linear classifiers, in dimension 2
  - If $T \leq 3$, then number of possible labelings can be $2^3 = 8$
  - If $T = 4$, then number of possible labelings can never be 2^4 !
  - If $T \geq 4$, number of possible labelings can never be 2^T but at most $\mathcal{O}\left((\frac{T}{3})^3\right)$!
- The $T \leq 3$ and $(\frac{T}{3})^3$ is **not a coincidence**, but the VC-dimension

Back to ERM of auctions

- Same results for classes of functions (not 0/1 valued):
  There exists $d \in \mathbb{N}$, such that $\frac{1}{\varepsilon^2}\left(d \log(\frac{1}{\varepsilon}) + \log(\frac{1}{\delta})\right)$ samples are needed for ε-error with proba $1-\delta$
- **VC-(pseudo)dimension of classes of mechanisms**
  - M-levels auctions. $d = \mathcal{O}(nM \log(nM))$
  - Myerson auctions $d = \infty$ (since $M \to \infty$)
  - 2nd price with reserve prices $d = \mathcal{O}(n \log(n))$

Sample Complexity of optimal auctions

For ε-error with proba $1-\delta$, need $\frac{n}{\varepsilon^3}$ (up to log)

- M-levels auctions with $M \sim \frac{1}{\varepsilon}$

Earning while Learning

Learning the Reserve Price

Regret while learning auctions

- Items sequentially sold at second price, with reserve prices r_t
  - n bidders, same distribution of valuation v_t (i.e., $F_i = F$).
  - Seller only sees the revenue $\mathbb{1}\{v_t^{(1)} \geq r_t\} \max\{v_t^{(2)}, r_t\}$
- Optimal reserve price $r^* = \phi^{-1}(0)$ to be learned.
- **Minimization of Regret:**

$$T.\mathbb{E}\left[\mathbb{1}\{v^{(1)} \geq r^*\} \max\{v^{(2)}, r^*\}\right] - \sum_{t=1}^{T} \mathbb{E}\left[\mathbb{1}\{v_t^{(1)} \geq r_t\} \max\{v_t^{(2)}, r_t\}\right]$$

Difference between revenue of optimal auctions vs of algorithm

A flavor on Multi-Armed Bandits

K-Armed Stochastic Bandit Problems

- K actions/arms $i \in [K]$, reward $X_t^i \in \mathbb{R}$ sub-Gaussian/bounded

$$X_1^i, X_2^i, \dots, \sim \mathcal{N}(\mu^i, 1) \quad \textbf{i.i.d.}$$

- Non-Anticipative Policy: $\pi_t\left(X_1^{\pi_1}, X_2^{\pi_2}, \dots, X_{t-1}^{\pi_{t-1}}\right) \in [K]$
- Goal: Maximize expected reward $\sum_{t=1}^T \mathbb{E} X_t^{\pi_t} = \sum_{t=1}^T \mu^{\pi_t}$
- Performance: Cumulative Regret

$$\max_{i \in [K]} \sum_{t=1}^T \mu^i - \sum_{t=1}^T \mu^{\pi_t} = \sum_i \Delta_i \sum_{t=1}^T \mathbb{1}\{\pi_t = i\} = \sum_i \Delta_i N_T^i$$

with $\Delta_i = \mu^\star - \mu^i$, the “gap” or cost of error i.

Upper Confidence Bound [Auer, Cesa-Bianchi, Fisher, '02]

- Estimate μ^i with $\overline{X}_t^i = \frac{1}{N_t^i}\sum_{s=1}^{N_t^i} X_s^i$? negative bias

  Positive bias: add an error term $2\sqrt{\frac{\log(t)}{N_t^i}}$

$$\text{UCB: } \pi_{t+1} = \arg\max_i \left\{ \overline{X}_t^i + 2\sqrt{\frac{\log(t)}{N_t^i}} \right\}$$

- Regret bounds
  - Distribution dependent

$$\mathbb{E}\, R_T \lesssim \sum_i \frac{\log(T)}{\Delta_i}$$

  - Minimax

$$\mathbb{E}\, R_T \lesssim \sqrt{KT\log(T)} \quad \text{and even } \sqrt{\log(K)T} \text{ if all } X_s^i \text{ are observed}$$

Ideas of proof $\pi_{t+1} = \arg\max_i \left\{ \overline{X}_t^i + 2\sqrt{\frac{\log(t)}{N_t^i}} \right\}$

- 2-lines proof:

$$\pi_{t+1} = i \neq \star \iff \overline{X}_t^\star + 2\sqrt{\frac{\log(t)}{N_t^\star}} \leq \overline{X}_t^i + 2\sqrt{\frac{\log(t)}{N_t^i}}$$

$$\text{“} \Longrightarrow \text{”} \Delta_i \leq 2\sqrt{\frac{\log(t)}{N_t^i}} \Longrightarrow N_t^i \lesssim \frac{\log(t)}{\Delta_i^2}$$

- Number of mistakes grows as $\frac{\log(t)}{\Delta_i^2}$; each mistake costs Δ_i.

Regret at stage T $\lesssim \sum_i \frac{\log(T)}{\Delta_i^2} \times \Delta_i \approx \sum_i \frac{\log(T)}{\Delta_i}$

- “ $\Longrightarrow$ ” actually happens with overwhelming proba
- “optimal”: no algo with regret always smaller than $\sum_i \frac{\log(T)}{\Delta_i}$

Back to learning second price

- First technique ? **Brute force**
  - ε-approximations, with $1/\varepsilon$ possible reserve prices
  - UCB on those prices
  - Regret of $\sqrt{\frac{1}{\varepsilon}T} + \varepsilon T.$
  - Optimal choice $\varepsilon = T^{-1/3}$, regret in $T^{2/3}$
- We can do **better** ! Regret in $\frac{\sqrt{T}}{\text{revenue}[r*]}$
  - Trivial if all bids are seen $\sqrt{\log(\frac{1}{\varepsilon})T} + \varepsilon T.$
  - Play by block of length $T^{1-2^{-k}}$ (i.e., $T^{1/2}, T^{3/4}, T^{15/16}$...)

On **block** k

- At all stages, same reserve prive $r^k \leq r^*$
- From data, empirical approximation of F
- Remove all reserve prices clearly sub-optimal (optimistic revenue lower than highest pessimistic revenue, with error terms as in UCB)
- r^{k+1} is the smallest remaining reserve price.

Learning the Correct Posted Price

Regret while learning auctions

- Items sequentially sold at posted price p_t
  - 1 bidders, valuation $v_t \sim F$. Accept the price if $v_t \geq p_t$
  - Seller only sees the revenue $p_t \mathbb{1}\{v_t \geq p_t\}$
- Optimal posted price $p^* = \phi^{-1}(0)$ to be learned.
- **Minimization of Regret:**

$$T.\mathbb{E}\Big[\mathbb{1}\{v \geq p^*\}p^*\Big] - \sum_{t=1}^{T} \mathbb{E}\Big[\mathbb{1}\{v_t \geq p_t\}p_t\Big]$$

- ε-discretization + UCB, Regret bounded as $\sqrt{\frac{T}{\varepsilon}} + \varepsilon T \leq T^{2/3}$
- Can be improved to $\sqrt{T}$ under strong smoothness of $F(\cdot)$

Posted price with only 1 valuation

- Assume $v_t = v^*$ but **unknown**. Regret of finding v^* ?
- $p_t \leq v^*$ accepted, $p_t > v^*$ rejected... Binary search ? (up to $1/T$)
- **Bad idea !!** if $v^* = \frac{1}{2}$ then regret $\frac{\log(T)}{2}$

Better idea, "cautious binary search"

- Cautious binary search proceed by blocks $k \in \mathbb{N}$
  - At block k, valid interval $[a_k, b_k]$
  - Post price $a_k + \varepsilon_k$, $a_k + 2\varepsilon_k$... with $\varepsilon_k = \frac{1}{2^{2^k}}$
  - a_{k+1} last accepted price, b_{k+1} first rejected price
- **Regret of Cautious binary search:** $\log\log(T)$

Regret increases by 1 (small errors) +1 (overshoot) at each block

Posted price with K valuations

- $v_t \in \{v^{(1)}, v^{(2)}, \dots, v^{(K)}\}$ i.i.d. with $\mathbb{P}(v_t = v^{(k)}) = \gamma_k$

$$\Delta_k = v^* \mathbb{P}(v_t \geq v^*) - v^{(k)} \mathbb{P}(v_t \geq v^{(k)})$$

- A **lower bound** $\gamma \leq \gamma_k$ is known

Achievable regret of $\mathcal{O}\left(\sum_k \frac{\log(T)}{\Delta_k} + \frac{\log(T)}{\gamma^2} \log\log(T)\right)$

- Post each price $\frac{\log(T)}{\gamma^2}$.
  - Drop in demand **always observed** w.h.p.
  - Again because of Hoeffding !
- Create and maintain one cautious binary search for each v_k
- Choose which cautious bin. search to update via UCB type algo

Special case $K = 2$. Regret of $\mathcal{O}(\frac{\log(T)}{\Delta})$ without prior !

Learning the Bidding Strategies

Second Price Auctions - Learning Bids

Online Marketing: Bidders do not know their own valuation

- Online Marketing company paid if user clicks on the ad !
- At round $t = 1, \ldots, T$:

  bidder bids $b_t \in [0, 1]$

  if $b_t > m_t$ (maximum other bids & reserve price)

  win good, observe value $v_t \in [0, 1]$ (or $v_t \in \{0, 1\}$ for clicks)
- Total utility: $\sum_{t=1}^{T}(v_t - m_t)\mathbb{1}\{b_t > m_t\}$
- Total **regret**:

$$\max_{b \in [0,1]} \sum_{t=1}^{T}(v_t - m_t)\mathbb{1}\{b > m_t\} - \sum_{t=1}^{T}(v_t - m_t)\mathbb{1}\{b_t > m_t\}$$

Data Assumptions - Stochastic vs Adversarial

- **Stochastic:** v_t i.i.d. $\mathbb{E}[v_t] = v \in [0, 1]$

  m_t stochastic (i.i.d. $\mathbb{E}[m_t] = m$), indpt. of v_t

  m_t adversarial (no assumptions), indpt. of v_t

  In both cases, **expected** regret attained at v.

$$\sum_{t=1}^{T}(v - m_t)\mathbb{1}\{v > m_t\} - \sum_{t=1}^{T}(v - m_t)\mathbb{1}\{b_t > m_t\}$$

- **Adversarial:** no assumptions at all on v_t and m_t

Our policy: UCBid

- Auctions: infinite action space, but with a special structure.
- Round $t+1$ bid

$$b_{t+1} = \min\left(\bar{v}_{\omega_t} + \sqrt{\frac{3\log(t)}{2\omega_t}}, 1\right)$$

where ω_t number of auctions won.

Theorem - Stochastic auctions

UCBid yields a regret bound of

$$\mathbb{E}R_T \leq 3 + 12\frac{\log(T)}{\Delta} \wedge 5\sqrt{T\log(T)}$$

where Δ is such that no bid m_t is in the interval $(v, v+\Delta)$

Fully stochastic case: UCBid

- If $m_t \sim \mu$ satisfies margin condition, parameter α (unknown):

Definition - margin condition

$$\forall u > 0,\ \mu\{(v, v+u)\} \leq Cu^{\alpha} \text{ for some constant } C.$$

The bigger α, the easier.

Theorem - Fully stochastic case

$$\mathbb{E}R_T \leq \begin{cases} c_\alpha T^{\frac{1-\alpha}{2}} \log^{\frac{1+\alpha}{2}}(T) & \text{if } \alpha < 1 \\ c_\alpha \log^2(T) & \text{if } \alpha = 1 \\ c_\alpha \log(T) & \text{if } \alpha > 1 \end{cases}$$

- Almost matching lower bound

$$\mathbb{E}R_T \geq \begin{cases} c_\alpha T^{\frac{1-\alpha}{2}} & \text{if } \alpha < 1 \\ c_\alpha \log(T) & \text{if } \alpha \geq 1 \end{cases}$$

Manipulation of truthful auctions

Truthful or not ?

- All the learning results holds if $b_t = v_t$, **truthful bidding**
- OK since Myerson/2nd price are **truthful mechanisms....**

> Repeated truthful mechanisms are **not truthful**

- Possible to manipulate **repeated** truthful auctions !

  because mechanisms learning from past data.

Asymptotic, steady-state scenario

- Bidder i, distribution of valuations F_i.
- Steady state strategy, mapping $\sigma_i : [0,1] \to [0,1]$.
- Mechanism see the distribution $\widetilde{F}_i = \sigma_i \sharp F_i$

  Optimal Myerson/reserve prices w.r.t. $\widetilde{F}_i$ (compute $\widetilde{\phi}_i$ or $\widetilde{r}_i$)

Utility of bidder i: $\mathbb{E}_F \left[\text{revenue}[A(\widetilde{F})] \big(\sigma_i(v_i), \sigma_{-i}(v_{-i}) \big) \right]$

- This meta-mechanism defined on distributions **is not truthful**

Example: symmetric bidders

- **Learning mechanism**: Myerson auction w.r.t. $\sigma_i \sharp F$
- Symmetric equilibrium of bidders ?
  - There is a threshold θ
  - **Truthful above θ:** $\sigma(v) = v$
  - **Overbid below θ:** $\sigma(v) = \theta \frac{1-F(\theta)}{1-F(v)}$
- Consequences
  - Bidders overbids low valuation : **Reserve prices decrease** !
  - Previously won auctions: payment **smaller**
  - New auctions won: Overbid, so possibly pay more than value
  - First effect dominates the second one
- For the seller (and bidders actually):
  - Same revenue than in 2nd price auctions without reserve prices

illustrations

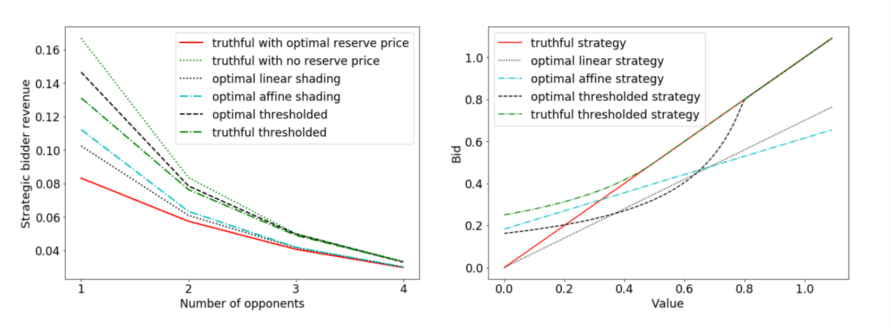

Conclusions

Conclusions

- Billions of auctions run every day. Huge market
- No prior on the valuations distributions. Must be learned
- **Batch method:** Empirical Revenue Maximization

  Sample complexity (number of data required)

  Depends on the complexity of auctions

  (2nd price, Myerson, t-levels)
- **Online method:** Save money when acquiring data

  Multi-armed bandit techniques

  Control of Regret $\mathcal{O}(\log(T), \sqrt{T}...)$
- Be **careful** ! Repeated truthful mechanism can **be manipulated**

References

- Great book on auctions & algo GT: T. Roughgarden, *20 lectures on Algorithmic Game theory*
- ERM, ML and more: P. Rigollet, *Mathematics of Machine Learning*
- Multi-armed bandits: S. Bubeck & N. Cesa-Bianchi, *Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems*
- ERM for auctions J. Morgenstern & T. Roughgarden, *The Pseudo-Dimension of Near-Optimal Auctions*
- Learning 2nd price auctions N. Cesa-Bianchi, C. Gentile & Y. Mansour, *Regret Minimization for Reserve Prices in Second-Price Auctions*
- Cautious Binary Search R. Kleinberg & T. Leighton. *The value of knowing a demand curve: Bounds on regret for online posted-price auctions.*
- Learning Posted Price N. Cesa-Bianchi, T. R. Cesari, & V. Perchet, *Dynamic Pricing with Finitely Many Unknown Valuations*
- Learning Bidding Strategy J. Weed, V. Perchet, P. Rigollet, *Online learning in repeated auctions.*
- Manipulation of auctions M. Abeille, C. Calauzenes, N. El Karoui, T. Nedelec & V. Perchet *Explicit shading strategies for repeated truthful*