Announcements

- HW1 grading will be out next Tuesday, and sample solution is out on Collab
- HW 2 is due next Tuesday

CS6501: Topics in Learning and Game Theory (Fall 2019)

Mechanism Design from Samples

Instructor: Haifeng Xu

Outline

- Optimal Auction and its Limitations
- The Sample Mechanism and its Revenue Guarantee

Recap: Optimal Auction for Single Item

Theorem. For single-item allocation with regular value distribution $v_i \sim f_i$ independently, the following auction is BIC and optimal:

1. Solicit buyer values $v_1, \cdots, v_n$
2. Transform v_i to "virtual value" $\phi_i(v_i)$ where $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$
3. If $\phi_i(v_i) < 0$ for all i, keep the item and no payments
4. Otherwise, allocate item to $i^* = \arg\max_{i \in [n]} \phi_i(v_i)$ and charge him the minimum bid needed to win, i.e., $\phi_i^{-1}\big(\max(\max_{j \neq i^*} \phi_j(v_j), 0) \big)$.

- Recall, "regular" means $\phi_i(v_i)$ is monotone non-decreasing
- Will always assume distributions are regular and "nice" henceforth

Recap: Optimal Auction for Single Item

An important special case: $v_i \sim F$ i.i.d.

- The second-price auction with reserve $\phi^{-1}(0)$ is optimal
  1. Solicit buyer values $v_1, \cdots, v_n$
  2. If $v_i < \phi^{-1}(0)$ for all i, keep the item and no payments
  3. Otherwise, allocate to $i^* = \arg\max_{i \in [n]} v_i$ and charge him the minimum bid needed to win, i.e., $\max(\max_{j \neq i^*} v_j , \phi^{-1}(0))$

Recap: Optimal Auction for Single Item

An important special case: $v_i \sim F$ i.i.d.

- The second-price auction with reserve $\phi^{-1}(0)$ is optimal
  1. Solicit buyer values $v_1, \cdots, v_n$
  2. If $v_i < \phi^{-1}(0)$ for all i, keep the item and no payments
  3. Otherwise, allocate to $i^* = \arg\max_{i \in [n]} v_i$ and charge him the minimum bid needed to win, i.e., $\max(\max_{j \neq i^*} v_j , \phi^{-1}(0))$

Intuitions about why *second-price auction with reserve* is good

- Incentive compatibility requires payment to not depend on bidder's own bid → second highest bid is pretty much the best choice
- Use the reserve to balance between "charging a higher price" and "disposing the item"

Recap: Optimal Auction for Single Item

Myerson's Lemma is central to the proof

Lemma. Consider any BIC mechanism M with interim allocation x and interim payment p, normalized to $p_i(0) = 0$. The expected revenue of M is equal to the expected virtual welfare served

$$\sum_{i=1}^{n} \mathbb{E}_{v_i \sim f_i}[\,\phi_i(v_i) x_i(v_i)]$$

Drawbacks of the Optimal Auction

1. Buyer's value v_i is assumed to be drawn from a distribution f_i
2. The precise distribution f_i is assumed to be known to seller

➢ In this lecture, we will keep Assumption 1, but relax Assumption 2

Drawbacks of the Optimal Auction

1. Buyer's value v_i is assumed to be drawn from a distribution f_i ✓
2. ~~The precise distribution f_i is assumed to be known to seller~~

- In this lecture, we will keep Assumption 1, but relax Assumption 2
- This is precisely the machine learning perspective
  - ML assumes data drawn from distributions
  - The precise distribution is unknown; instead samples are given

Task and Goal of This Lecture

- Will focus on setting with n buyer, i.i.d. values
- Buyer value v_i is drawn from regular distribution f, which is unknown to the seller

Goal: design an auction that has revenue close to the optimal revenue when knowing f

- Optimal auction is a second-price auction with reserve $\phi^{-1}(0)$
- "Closeness" will be measured by guaranteed approximation ratio

Task and Goal of This Lecture

- Will focus on setting with n buyer, i.i.d. values
- Buyer value v_i is drawn from regular distribution f, which is unknown to the seller

Goal: design an auction that has revenue close to the optimal revenue when knowing f

- Optimal auction is a second-price auction with reserve $\phi^{-1}(0)$
- "Closeness" will be measured by guaranteed approximation ratio

But wait . . . we cannot have any guarantee without assumptions on bidder values – is this a contradiction?

- No, we assumed $v_i \sim f$

A Natural First Attempt

- Since v_i's are all drawn from f, these n i.i.d. samples can be used to estimate f
- This results in the following "empirical Myerson" auction

Empirical Myerson Auction

1. Solicit buyer values $v_1, \cdots, v_n$
2. Use $v_1, \cdots, v_n$ to estimate an empirical distribution $\bar{f}$
3. Run second-price auction with reserve $\bar{\phi}^{-1}(0)$ where $\bar{\phi}$ is calculated using $\bar{f}$ instead

Q: does this mechanism work?

No, may fail in multiple ways

Issues of Empirical Myerson

Empirical Myerson Auction

1. $\boxed{\text{Solicit buyer values } v_1, \cdots, v_n}$ problematic
2. Use $v_1, \cdots, v_n$ to estimate an empirical distribution $\bar{f}$
3. Run second-price auction with reserve $\bar{\phi}^{-1}(0)$ where $\bar{\phi}$ is calculated using $\bar{f}$ instead

- Not incentive compatible – reserve depends on bidder's report
  - This is a crucial difference from standard machine learning tasks where samples are assumed to be correctly given

Issues of Empirical Myerson

Empirical Myerson Auction

1. Solicit buyer values $v_1, \cdots, v_n$ problematic
2. Use $v_1, \cdots, v_n$ to estimate an empirical distribution $\bar{f}$
3. Run second-price auction with reserve $\bar{\phi}^{-1}(0)$ where $\bar{\phi}$ is calculated using $\bar{f}$ instead

- Not incentive compatible – reserve depends on bidder's report
  - This is a crucial difference from standard machine learning tasks where samples are assumed to be correctly given
- Even bidders report true values, $\bar{f}$ may not be regular
- Even $\bar{f}$ is regular, $\bar{\phi}^{-1}(0)$ may not be close to $\phi^{-1}(0)$
  - Depend on how large is n, and shape of f

Outline

- Optimal Auction and its Limitations
- The Sample Mechanism and its Revenue Guarantee

The Basic Idea

- Want to use second-price auction with an estimated reserve
- Lesson from previous example – if a bidder's bid is used to estimate the reserve, we cannot use this reserve for him
- Main idea: pick a "reserve buyer" → use his bid to estimate the reserve but never sell to this buyer
  - I.e., we give up any revenue from the reserve buyer

The Basic Idea

- Want to use second-price auction with an estimated reserve
- Lesson from previous example – if a bidder's bid is used to estimate the reserve, we cannot use this reserve for him
- Main idea: pick a "reserve buyer" → use his bid to estimate the reserve but never sell to this buyer
  - I.e., we give up any revenue from the reserve buyer

Q: why only pick one reserve buyer, not two or more?

We have to give up revenue from reserve buyers, better not too many

The Basic Idea

- Want to use second-price auction with an estimated reserve
- Lesson from previous example – if a bidder's bid is used to estimate the reserve, we cannot use this reserve for him
- Main idea: pick a "reserve buyer" → use his bid to estimate the reserve but never sell to this buyer
  - I.e., we give up any revenue from the reserve buyer

Q: why only pick one reserve buyer, not two or more?

We have to give up revenue from reserve buyers, better not too many

Q: which buyer to choose as the reserve buyer?

A-priori, they are the same → pick one uniformly at random

The Basic Idea

- Want to use second-price auction with an estimated reserve
- Lesson from previous example – if a bidder's bid is used to estimate the reserve, we cannot use this reserve for him
- Main idea: pick a "reserve buyer" → use his bid to estimate the reserve but never sell to this buyer
  - I.e., we give up any revenue from the reserve buyer

Q: how to use a single buyer's value to estimate reserve?

Not much we can do . . . just use his value as reserve

The Mechanism

Second-Price auction with Random Reserve (SP-RR)

1. Solicit buyer values $v_1, \cdots, v_n$
2. Pick $j \in [n]$ uniformly at random as the reserve buyer
3. Run second-price auction with reserve v_j but only among bidders in $[n] \setminus \{j\}$.

The Mechanism

Second-Price auction with Random Reserve (SP-RR)

1. Solicit buyer values $v_1, \cdots, v_n$
2. Pick $j \in [n]$ uniformly at random as the reserve buyer
3. Run second-price auction with reserve v_j but only among bidders in $[n] \setminus \{j\}$.

Claim. SP-RR is dominant-strategy incentive compatible.

For any bidder i

- If i is picked as reserve, his bid does not matter to him, so truthful bidding is an optimal strategy
- If i is not picked, he faces a second-price auction with reserve. Again, truthful bidding is optimal

The Mechanism

Theorem. Suppose F is regular. In expectation, SP-RR achieves at least $\frac{1}{2} \cdot \frac{n-1}{n}$ fraction of the optimal expected revenue.

Remarks

- $\frac{1}{2} \cdot \frac{n-1}{n}$ is a worst-case guarantee
- The first time we use approximation as a lens to analyze algorithms in this class
- It is possible to have a good auction even without knowing F
  - But we still assumed $v_i \sim F$ i.i.d.

The Mechanism

Theorem. Suppose F is regular. In expectation, SP-RR achieves at least $\frac{1}{2} \cdot \frac{n-1}{n}$ fraction of the optimal expected revenue.

- Equivalently, SP-RR is a second-price auction for $(n-1)$ i.i.d. bidders, with a reserve r drawn from F.
- To prove its revenue guarantee, we have to argue
  1. Discarding one buyer does not hurt revenue much (the $\frac{n-1}{n}$ term)
  2. Using a random $v \sim F$ as an estimated reserve is still good (the $\frac{1}{2}$ term)

The Mechanism

Theorem. Suppose F is regular. In expectation, SP-RR achieves at least $\frac{1}{2} \cdot \frac{n-1}{n}$ fraction of the optimal expected revenue.

Next, we will give a formal proof

Step 1: discarding a buyer does not hurt revenue much

Lemma 1. The expected optimal revenue for an environment with $(n-1)$ buyers is at least $\frac{n-1}{n}$ fraction of the optimal expected revenue for n buyers.

Proof: use Myerson's Lemma

- Expected revenue for n buyers is $\sum_{i=1}^{n} \mathbb{E}_{v_i \sim f_i} \left[\phi_i(v_i) x_i^{(n)}(v_i) \right]$
  - $x_i^{(n)}$ = interim allocation of the optimal auction for n buyers
- By symmetry of the auction and buyer values, each buyer's interim allocation must be the same, i.e., $x_i^{(n)}(v) = x^{(n)}(v)$ for some $x^{(n)}$
- Thus, optimal revenue with n bidders is

$$R(n) = \sum_{i=1}^{n} \mathbb{E}_{v_i \sim f_i} \left[\phi_i(v_i) x_i^{(n)}(v_i) \right]$$

$$= n \cdot \mathbb{E}_{v \sim f}\, \phi(v) x^{(n)}(v)$$

Step 1: discarding a buyer does not hurt revenue much

Lemma 1. The expected optimal revenue for an environment with $(n-1)$ buyers is at least $\frac{n-1}{n}$ fraction of the optimal expected revenue for n buyers.

Proof: use Myerson's Lemma

- Due to less competition, we have $x^{(n-1)}(v) \geq x^{(n)}(v)$
  - They face the same reserve $\phi^{-1}(0)$, but with $n-1$ buyers, bidder i has more chance to win
- Therefore,

$$
\begin{aligned}
R(n-1) &= (n-1) \cdot \mathbb{E}_{v\sim f}\, \phi(v) x^{(n-1)}(v) \\
&\geq (n-1) \cdot \mathbb{E}_{v\sim f}\, \phi(v) x^{(n)}(v) \\
&\geq \frac{n-1}{n} R(n)
\end{aligned}
$$

Step 2: using random reserve is not bad

Consider the following two auctions for i.i.d. bidders with $v_i \sim F$

- SP-OR: second price auction with optimal reserve $r^* = \phi^{-1}(0)$
- SP-RR: second price auction with random reserve $r \sim F$

Lemma 2. Rev(SP-RR) $\geq \frac{1}{2}$ Rev(SP-OR) for any n and regular F.

Note: this completes our proof of the theorem

Proof of Lemma 2

Lemma 2. Rev(SP-RR) $\geq \frac{1}{2}$ Rev(SP-OR) for any n and regular F.

Proof of Lemma 2

Lemma 2. Rev(SP-RR) $\geq \frac{1}{2}$ Rev(SP-OR) for any n and regular F.

Step 1: characterize how much revenue i contribute in each auction

Let us focus on SP-OR first

- Fix v_{-i}, buyer i contributes to revenue only when he wins
- Whenever i wins, he pays $p = \max(t, r^*)$ where $t = \max[v_{-i}]$ and $r^* = \phi^{-1}(0)$
- Conditioning on v_{-i}, i contributes the following amount to revenue

$$p\big(1 - \mathrm{F(p)}\big) = \hat{R}(p) = \hat{R}(\max(t, r^*))$$

- In expectation, i contributes $\mathbb{E}_{v_{-i}}[\hat{R}(\max(t, r^*))]$

Proof of Lemma 2

Lemma 2. Rev(SP-RR) $\geq \frac{1}{2}$ Rev(SP-OR) for any n and regular F.

Step 1: characterize how much revenue i contribute in each auction

In expectation, i contributes $\mathbb{E}_{v_{-i}}[\hat{R}(\max(t, r^*))]$ in SP-OR

Proof of Lemma 2

Lemma 2. Rev(SP-RR) $\geq \frac{1}{2}$ Rev(SP-OR) for any n and regular F.

Step 1: characterize how much revenue i contribute in each auction

What about SP-RR?

- Similar argument, but use a random reserve r instead
- In expectation, i contributes $\mathbb{E}_{r\sim F}\mathbb{E}_{v_{-i}}[\hat{R}(\max(t,r))]$

In expectation, i contributes $\mathbb{E}_{v_{-i}}[\hat{R}(\max(t,r^*))]$ in SP-OR

Proof of Lemma 2

Lemma 2. Rev(SP-RR) $\geq \frac{1}{2}$ Rev(SP-OR) for any n and regular F.

Step 1: characterize how much revenue i contribute in each auction

What about SP-RR?

- Similar argument, but use a random reserve r instead
- In expectation, i contributes $\mathbb{E}_{r\sim F}\mathbb{E}_{v_{-i}}[\hat{R}(\max(t, r))]$

In expectation, i contributes $\mathbb{E}_{v_{-i}}[\hat{R}(\max(t, r^*))]$ in SP-OR

Proof of Lemma 2

Lemma 2. Rev(SP-RR) $\geq \frac{1}{2}$ Rev(SP-OR) for any n and regular F.

Step 1: characterize how much revenue i contribute in each auction

What about SP-RR?

- Similar argument, but use a random reserve r instead
- In expectation, i contributes $\mathbb{E}_{r\sim F}\mathbb{E}_{v_{-i}}[\hat{R}(\max(t, r))]$

In expectation, i contributes $\mathbb{E}_{v_{-i}}[\hat{R}(\max(t, r^*))]$ in SP-OR

Step 2: prove $\mathbb{E}_{r\sim F}[\hat{R}(\max(t, r)) \geq \frac{1}{2}\hat{R}(\max(t, r^*))$ for any t

This proves Lemma 2

Claim. $\mathbb{E}_{r\sim F}[\hat{R}(\max(t,r))] \geq \frac{1}{2}\hat{R}(\max(t,r^*))$ for any t.

- Note: this is really the fundamental reason for why using uniform reserve is not bad
- Proof is based on an elegant geometric argument

Claim. $\mathbb{E}_{r\sim F}[\hat{R}(\max(t,r))] \geq \frac{1}{2}\hat{R}(\max(t,r^*))$ for any t.

- Note: this is really the fundamental reason for why using uniform reserve is not bad
- Proof is based on an elegant geometric argument
- Recall $\hat{R}(p) = p \cdot (1 - F(p))$. The (not so) magic step: change variable for function $\hat{R}(p)$
  - Define new variable $q = 1 - F(p)$, so $p = F^{-1}(1-q)$
  - Define $R(q) = q \cdot F^{-1}(1-q)$
  - Note: value of $R(q)$ equals value of $\hat{R}(p)$ (when $q = 1 - F(p)$)
- It turns out that $R(q)$ is concave if and only if F is regular
  - This is also the intrinsic interpretation of the regularity assumption

Claim. $\mathbb{E}_{r\sim F}[\hat{R}(\max(t,r))] \geq \frac{1}{2}\hat{R}(\max(t,r^*))$ for any t.

Calculating derivative of $R(q) = q \cdot F^{-1}(1-q)$:

$$\frac{d\,R(q)}{d\,q} = F^{-1}(1-q) + q \cdot \frac{d\,F^{-1}(1-q)}{d\,q}$$
$$= F^{-1}(1-q) - q \cdot \frac{1}{f(F^{-1}(1-q))}$$

Derive on the board

Claim. $\mathbb{E}_{r\sim F}[\hat{R}(\max(t,r))] \geq \frac{1}{2}\hat{R}(\max(t,r^*))$ for any t.

Calculating derivative of $R(q) = q \cdot F^{-1}(1-q)$:

$$\frac{d\,R(q)}{d\,q} = F^{-1}(1-q) + q \cdot \frac{d\,F^{-1}(1-q)}{d\,q}$$

$$= F^{-1}(1-q) - q \cdot \frac{1}{f(F^{-1}(1-q))}$$

$$= p - (1-F(p)) \cdot \frac{1}{f(p)}$$

Use the equation $1 - F(p) = q$

Claim. $\mathbb{E}_{r\sim F}[\hat{R}(\max(t,r))] \geq \frac{1}{2}\hat{R}(\max(t,r^*))$ for any t.

Calculating derivative of $R(q) = q \cdot F^{-1}(1-q)$:

$$\begin{aligned}
\frac{d\,R(q)}{d\,q} &= F^{-1}(1-q) + q \cdot \frac{d\,F^{-1}(1-q)}{d\,q} \\
&= F^{-1}(1-q) - q \cdot \frac{1}{f(F^{-1}(1-q))} \\
&= p - (1-F(p)) \cdot \frac{1}{f(p)} \\
&= \phi(p)
\end{aligned}$$

Use the equation $1 - F(p) = q$

- Regularity means $\phi(p)$ is increasing in p
- Moreover, p is decreasing in q, so $R'(q)$ is decreasing in q
- This implies $R(q)$ is concave

Claim. $\mathbb{E}_{r\sim F}[\hat{R}(\max(t,r))] \geq \frac{1}{2}\hat{R}(\max(t,r^*))$ for any t.

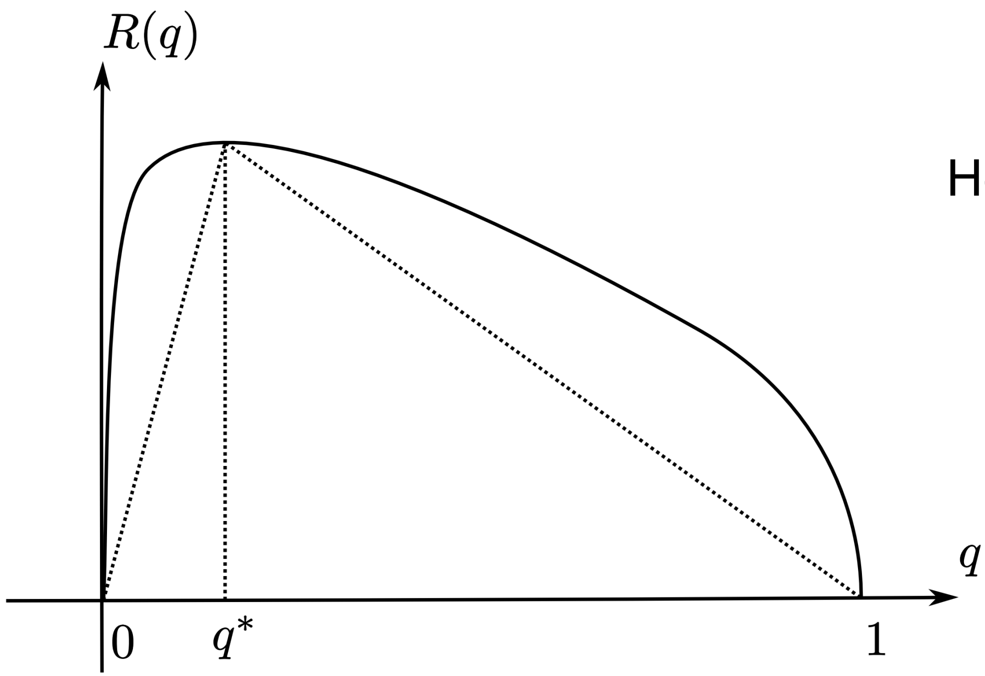

Here: $1 - F(r^*) = q^*$

r^* satisfies $\phi(r^*) = 0$, i.e., the point where derivative of $R(q)$ is 0

Claim. $\mathbb{E}_{r\sim F}[\hat{R}(\max(t,r))] \geq \frac{1}{2}\hat{R}(\max(t,r^*))$ for any t.

First, prove the $t = 0$ case.

Claim (when $t = 0$). $\mathbb{E}_{r\sim F}[\hat{R}(r)] \geq \frac{1}{2}\hat{R}(r^*)$.

Proof

- $\mathbb{E}_{r\sim F}[\hat{R}(r)] = \mathbb{E}_{q\sim U[0,1]}[R(q)]$ by variable change $q = 1 - F(r)$
  - If $r \sim f$, then $F(r) \sim U[0,1]$
- $\mathbb{E}_{q\sim U[0,1]}[R(q)]$ is precisely the area under the $R(q)$ curve
- $\hat{R}(r^*) = R(q^*)$ is precisely the area of the rectangle
- By geometry, $\mathbb{E}_{r\sim F}[\hat{R}(r)] \geq \frac{1}{2}\hat{R}(r^*)$

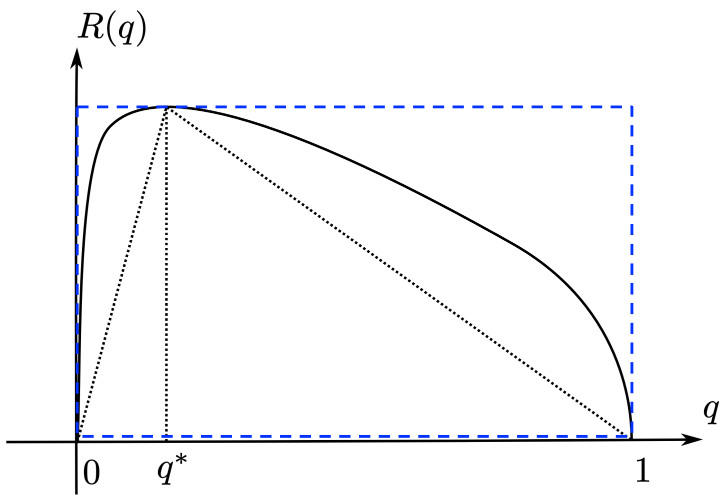

Claim. $\mathbb{E}_{r\sim F}[\hat{R}(\max(t,r))] \geq \frac{1}{2}\hat{R}(\max(t,r^*))$ for any t.

For general t

- If $t \leq r^*$, left-hand side increases, right-hand side no change ✓
- If $t > r^*$, $\hat{R}(\max(t,r^*)) = \hat{R}(t)$

$$\begin{aligned}\mathbb{E}_{r\sim F}\left[\hat{R}(\max(t,r))\right] &= \Pr(r \leq t)\cdot \hat{R}(t) + \Pr(r > t)\cdot \mathbb{E}_{r\sim F|\mathrm{r}\geq t}\hat{R}(r)\\ &\geq \Pr(r \leq t)\cdot \hat{R}(t) + \Pr(r > t)\cdot \frac{1}{2}\hat{R}(t)\\ &\geq \frac{1}{2}\hat{R}(t) \quad ✓\end{aligned}$$

Similar geometric argument shows $\mathbb{E}_{r\sim F|\mathrm{r}\geq t}\hat{R}(r) \geq \frac{1}{2}\hat{R}(t)$

Remarks

- Approximation ratio can be improved to $\frac{1}{2}$ (i.e. without the $\frac{n-1}{n}$ term)
  - Idea: don't discard the reserve buyer; instead randomly choose another buyer's bid as the reserve for him
- $\frac{1}{2}$ approximation is the best possible guarantee for SP-RR
  - The worst case is precisely when $R(q)$ curve is a triangle

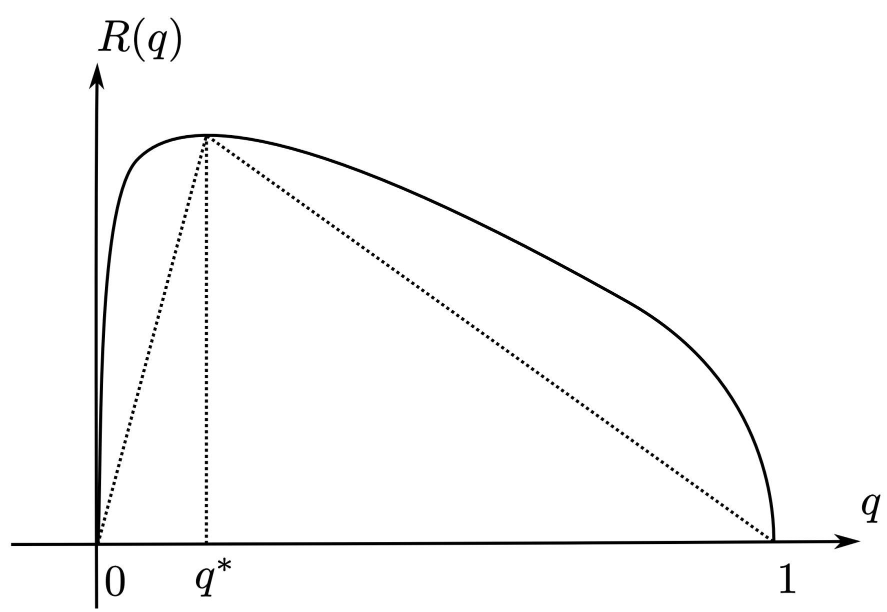

Remarks

- If we have sufficiently many bidders (more than $\Theta(\epsilon^{-4} \ln \epsilon^{-1})$ many), can obtain ϵ-optimal auction
  - Idea: pick many reserve bidders and use their values to estimate a better reserve
  - The estimation is tricky, not simply using the empirical distribution of the reserve bidders' values
- These results can all be generalized to "single-parameter" settings
  - E.g., selling k identical copies of items to n buyers
- Many open questions in this broad field of learning optimal auctions

Thank You

Haifeng Xu

University of Virginia

hx4ad@virginia.edu