

CSCI 357: Algorithmic Game Theory

Lecture 9: Revenue Maximization

Shikha Singh



Announcements and Logistics

- **Assignment 4** due tonight at 11 pm
 - Submit **code via Github**, latex answers and submit **PDF via Gradescope**
- Things to watch out for:
 - When targeting a slot, what if every slot gives negative utility? What is the balanced bidding condition in that case?
 - Related: make sure to never bid above value
 - Set $\alpha_{-1} = 2 \cdot \alpha_0$ in balanced bidding condition
- **VCG base case:** what should the occupant of the lowest slot pay?
- How to verify the outputs are reasonable?

Questions?

Empirical Analysis Takeaways

- Understand AGT theory vs practice
 - Harder to reason about asymmetric strategies in theory
- Humans reaching equilibrium vs auto-bidders
- Greedy balanced bidding heuristic is not perfect
 - Iteratively tries to converge to the theoretical equilibrium
- But, simple best-response dynamics get close even they do not converge
- Rules of mechanism matter
 - Simple mechanisms: complex bidding behavior
 - Complex mechanism: simple bidding behavior

Assignment 5 and Midterm 1

- **Assignment 5** will be released tomorrow
 - Shorter problem set, single-person assignment
 - Due a day early: Wed 11 pm instead of Thursday
 - To give time for returning graded feedback before midterm
- Self-scheduled midterm 1 on **March 12**
 - Pencil, paper exam, open book and notes
 - Pick up the exam anytime between **9am - 7pm**, will use a google form to coordinate
 - ~3 hour exam, TCL 202 room will be reserved for those who want to use it
- **Syllabus**: all topics covered until Monday March 7 and Assignment 5

Questions?

Last Time

- Computed the Bayes Nash equilibrium of first price auctions
- Showed **revenue equivalence** of first and second price auction
- Today: show how to use revenue equivalence to solve for BNE of any single-parameter **0/1** mechanism

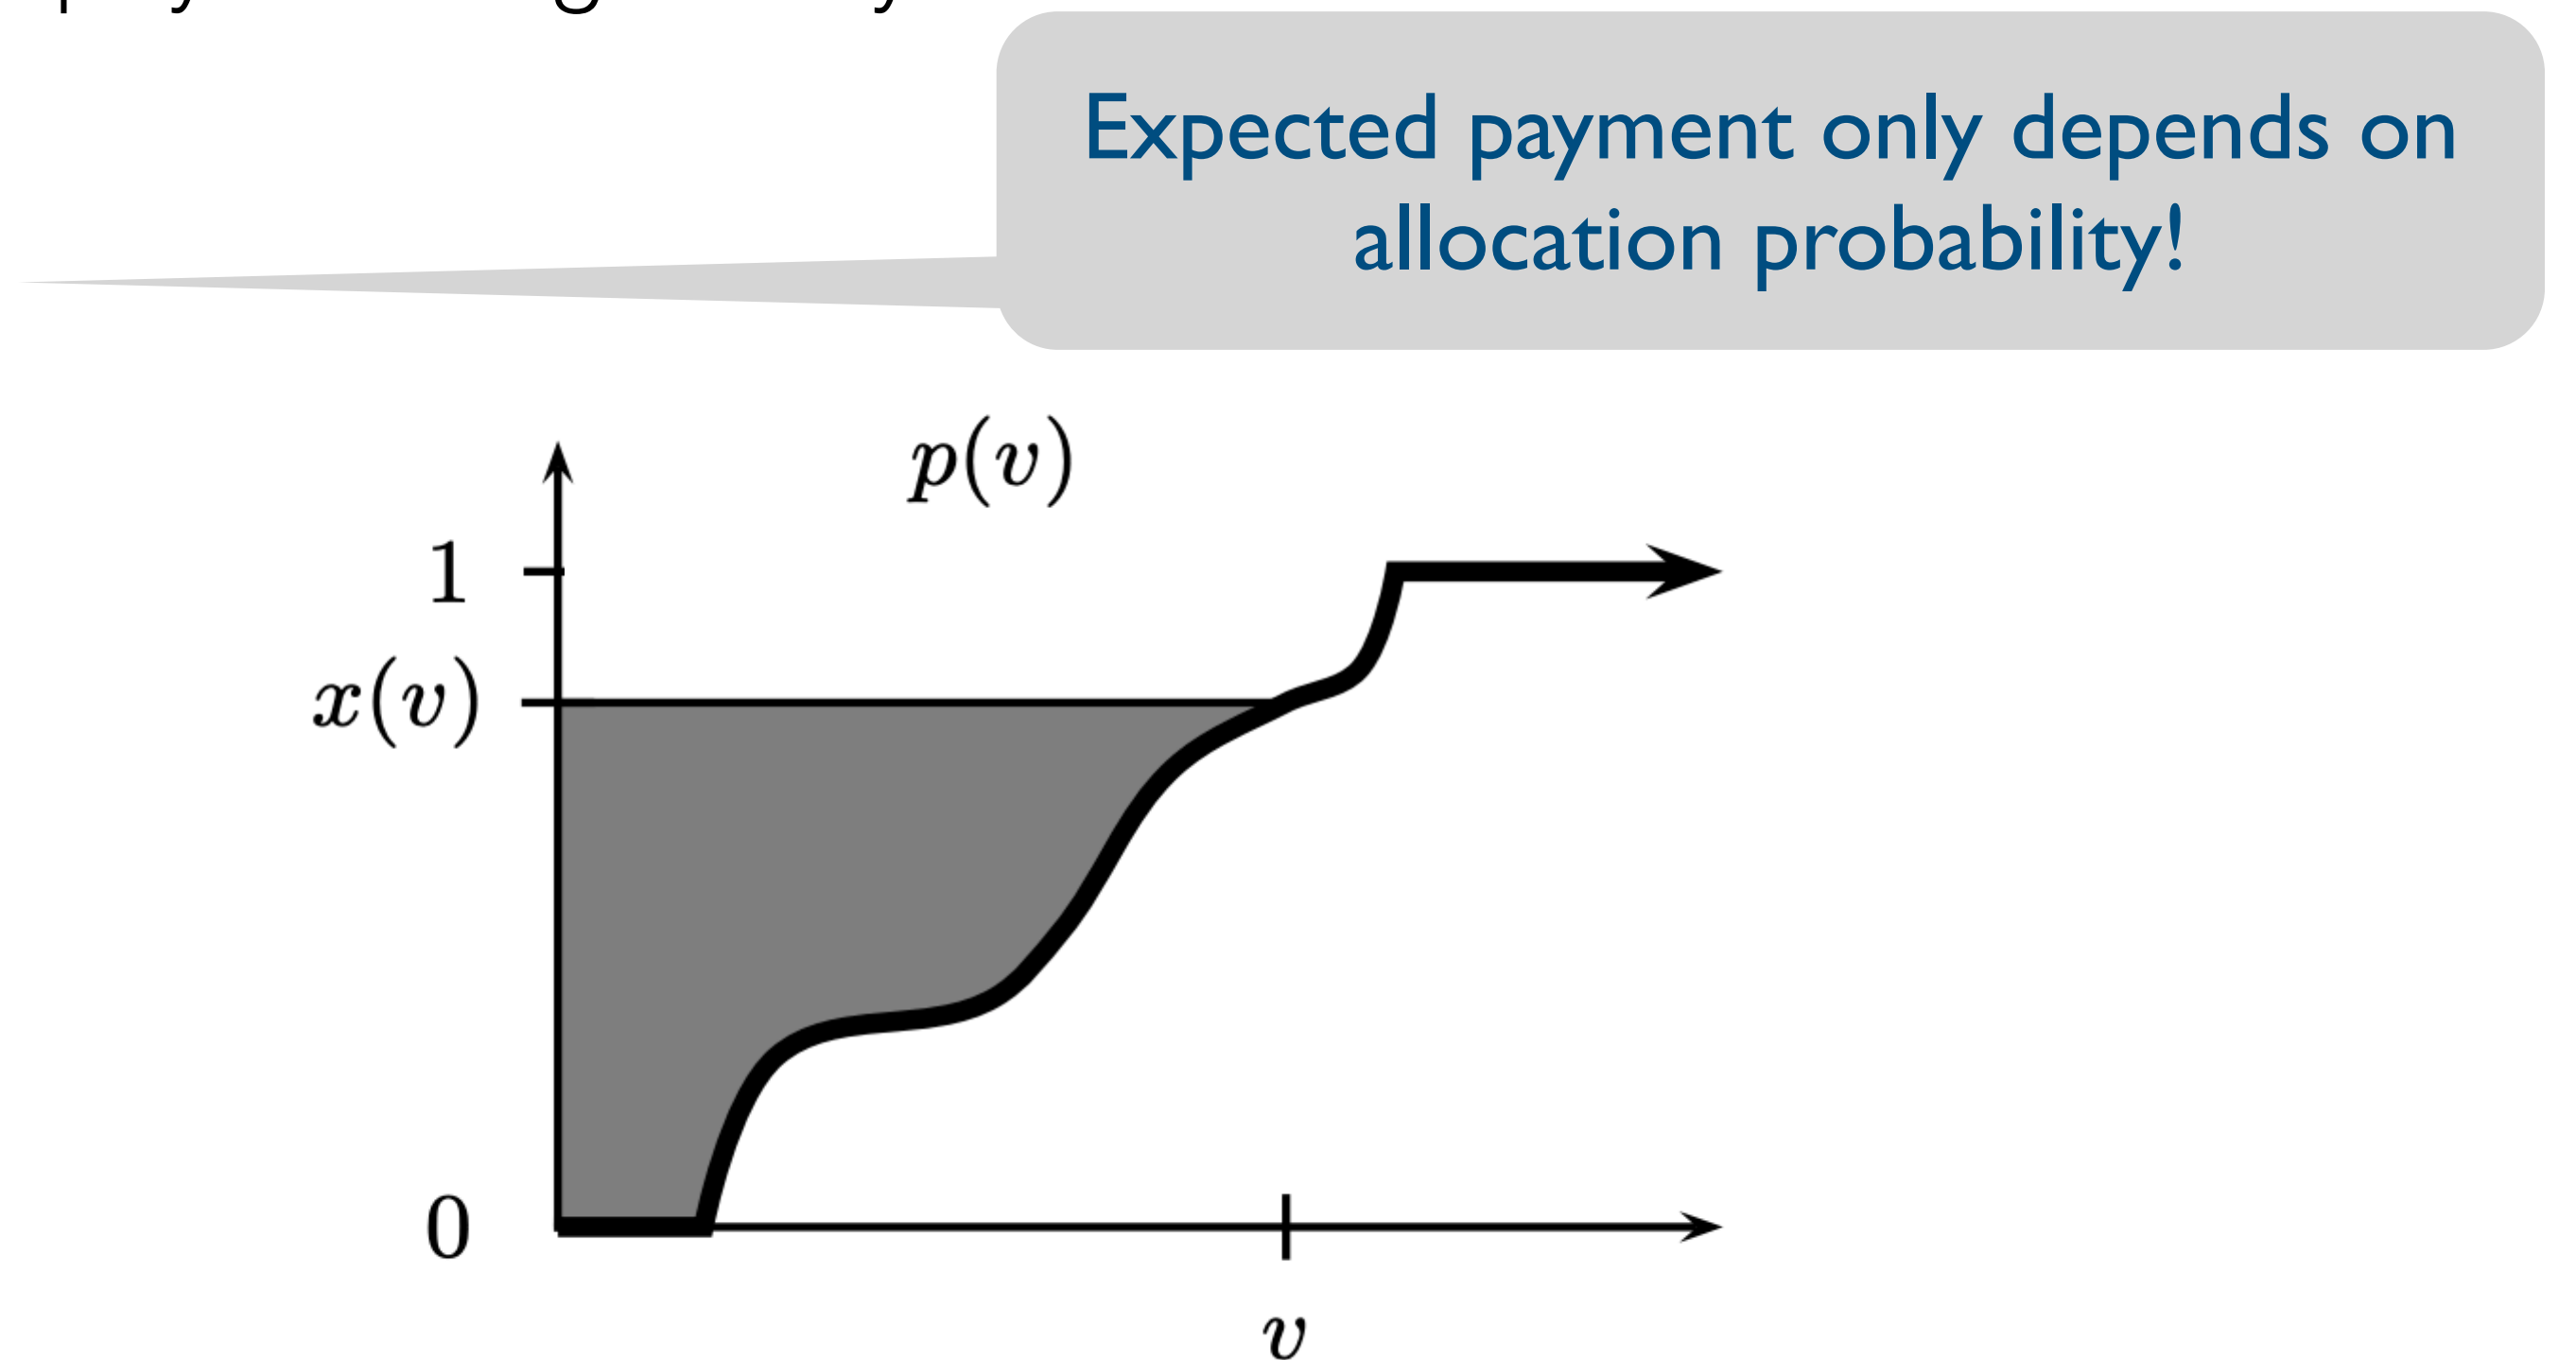
Myerson's Lemma: BNE

- Consider any single-parameter 0/1 allocation with values drawn i.i.d. from distribution G
- A strategy profile s is a **Bayes' Nash equilibrium** in (\mathbf{x}, \mathbf{p}) **if and only if** for all i
 - (a) **(monotonicity)** the allocation probability $x_i(v_i)$ is monotone non decreasing
 - (b) **(payment identity)** agent i 's expected payment is given by:

$$p_i(v_i) = v_i \cdot x_i(v_i) - \int_0^{v_i} x_i(z) dz$$

Assuming that $p_i(0) = 0$.

Proof is analogous to the DSE case.



Revenue Equivalence

- Most significant observation in auction theory
- A mechanism with the same allocation in DSE (BNE) have the same (expected) revenue!
 - In fact, each agent has the same expected payment in each mechanism
- Direct corollary of Myerson's lemma
 - The interim expected payments depend only on the allocation probability!
- **Corollary (Revenue equivalence).**
 - For any two mechanisms in 0-1 single-parameter setting, if the mechanism have the same BNE allocation, then they have the same expected revenue (assuming 0-valued agents pay nothing)

If we want to increase the (expected) revenue, **changing payments or charging more won't do it!** You need to change how you allocate!

Solving for BNE of an Auction

- Myerson's lemma tells us what outcomes are possible in BNE, but not how to solve for BNE strategies
- Solving for BNE is important: if you are a bidder in an auction, you want to know how to bid!
- Can use revenue equivalence to solve for BNE strategies in symmetric environments
 - If two mechanisms have the same allocation rule, they must have the same expected payment
- To solve for BNE strategies in a mechanism M , then the approach is:
 - Express agents expected payment in terms of their strategy and value
 - Set it equal to the expected payment in a strategically-simpler revenue equivalent mechanism M' (usually a "second-price implementation")

Solving for BNE: Steps

- **Step 1.** *Guess* what the allocation might be in a Bayes-Nash equilibrium (usually a surplus-maximizing one)
 - E.g., in a single-item auction, guess that highest bidder wins is a BNE
- **Step 2.** *Calculate* the interim expected payment of an agent in a strategically-simpler auction (usually the DSIC version, e.g. the second-price version)
 - E.g., consider a 2-bidder second price auction with values i.i.d. in $U(0,1)$:

$$\begin{aligned}
 E[p_1(v_1)] &= E[p_1(v_1) \mid 1 \text{ wins}] \cdot \Pr[1 \text{ wins}] + \overset{\text{Loser pays zero}}{E[p_1(v_1) \mid 1 \text{ loses}] \cdot \Pr[1 \text{ loses}]} \\
 &= E[v_2 \mid 1 \text{ wins}] \cdot \Pr[1 \text{ wins}] \\
 &= E[v_2 \mid v_2 \leq v_1] \cdot \Pr[1 \text{ wins}] \\
 &= \frac{v_1}{2} \cdot \Pr[1 \text{ wins}]
 \end{aligned}$$

v_2 is a uniform random variable that evenly divides the interval it is on, so if $v_2 \in [0, v_1]$ then $E[v_2] = v_1/2$

If we draw n samples i.i.d. from $U(0, a)$ then the expected value of k th highest draw is $\frac{n - (k - 1)}{n + 1} \cdot a$

Solving for BNE: Steps

- **Step 3.** Write the expression for the interim expected payment in terms of the strategy in the auction you are trying to solve for the BNE
 - e.g., consider a 2-bidder first price auction

Loser pays zero

$$E[p_1(v_1)] = E[p_1(v_1) \mid 1 \text{ wins}] \cdot \Pr[1 \text{ wins}] + \cancel{E[p_1(v_1) \mid 1 \text{ loses}] \cdot \Pr[1 \text{ loses}]}$$
$$= s_1(v_1) \cdot \Pr[1 \text{ wins}]$$

- **Step 4.** Solve for the BNE strategy by equating the expected payments between the two auctions

$$\frac{v_1}{2} \cdot \Pr[1 \text{ wins}] = s_1(v_1) \cdot \Pr[1 \text{ wins}]$$

$$s_1(v_1) = v_1/2$$

Prob of winning at a particular value is the same in both auctions!

Solving for BNE: Steps

- **Step 5.** Finally, verify that the initial guess was correct and the strategy is a symmetric BNE of the mechanism !
 - We have already verified for first-price auction
 - This is important since the other steps are just a way to guess the equilibrium strategy
- This process can be used for **any single-parameter 0/1 auction** (not just single-item auction)
- Let us apply this to an all-pay auction

All Pay Auctions

- Revenue equivalence governs all kinds of auctions
- Consider a **single-item all pay sealed-bid auction**, where the highest bidder wins but everyone pays their bid
 - Thus, utility of participating can be negative in this case
 - Participating has a cost!
- Can you think of examples of this setting in practice?

Solving for BNE: All Pay

- **Example:** All-pay auction with 2 bidders i.i.d. $U(0,1)$
- Step 1. Assume highest bidder wins
- Step 2. Interim expected payment in second-price auction is $\frac{v_1}{2} \cdot \Pr[1 \text{ wins}]$
- Step 3. Interim expected payment in all-pay auction?

Regardless of winning or losing you pay your bid

$$\begin{aligned} E[p_1(v_1)] &= E[p_1(v_1) \mid 1 \text{ wins}] \cdot \Pr[1 \text{ wins}] + E[p_1(v_1) \mid 1 \text{ loses}] \cdot \Pr[1 \text{ loses}] \\ &= s_1(v_1) \cdot \Pr[1 \text{ wins}] + s_1(v_1) \cdot \Pr[1 \text{ loses}] \\ &= s_1(v_1) \end{aligned}$$

- Step 4. Set them equal: here we need to compute the probability of winning

$$E^{\text{SP}}[p_1(v_1)] = \frac{v_1}{2} \cdot \Pr[v_2 \leq v_1] = \frac{v_1^2}{2}$$

Solving for BNE: All Pay

- **Example:** All-pay auction with 2 bidders i.i.d. $U(0,1)$
- **Step 4.** Set the interim expected payment equal to find the BNE strategy

$$s_1(v_1) = \frac{v_1^2}{2}$$

- **Step 5.** Verify that this a symmetric Bayes Nash of the auction
- Guess for n -bidder symmetric BNE?

- $\left(\frac{n-1}{n}\right)v_i^n$

- HW problem: verify that this is a symmetric BNE of all-pay auctions
 - Same way as we did for first-price auctions

Bayesian Analysis Challenges

- So far we considered symmetric settings:
 - Values are drawn i.i.d. from the same distribution
 - For these settings, Bayes Nash equilibrium are pretty well understood
 - They also lead to efficient outcomes: surplus maximizing outcomes
- For asymmetric settings:
 - Information asymmetry: e.g. different companies do different levels of "market research"
 - Can still study BNE but often more complicated: no closed form solutions and multiple (some inefficient) equilibrium
 - Near CS-driven approach: study approximation bounds (price of anarchy) of these auctions

Price of Anarchy: Auctions

- 2017 Survey

The Price of Anarchy in Auctions

Tim Roughgarden

*Computer Science Department, Stanford University,
Stanford, CA USA*

TIM@CS.STANFORD.EDU

Vasilis Syrgkanis

*Microsoft Research, 1 Memorial Drive,
Cambridge, MA USA*

VASY@MICROSOFT.COM

Éva Tardos

*Computer Science Department, Cornell University,
Ithaca, NY USA*

EVA@CS.CORNELL.EDU

Abstract

This survey outlines a general and modular theory for proving approximation guarantees for equilibria of auctions in complex settings. This theory complements traditional economic techniques, which generally focus on exact and optimal solutions and are accordingly limited to relatively stylized settings.

We highlight three user-friendly analytical tools: smoothness-type inequalities, which immediately yield approximation guarantees for many auction formats of interest in the special case of complete information and deterministic strategies; extension theorems, which extend such guarantees to randomized strategies, no-regret learning outcomes, and incomplete-information settings; and composition theorems, which extend such guarantees from simpler to more complex auctions. Combining these tools yields tight worst-case approximation guarantees for the equilibria of many widely-used auction formats.

Price of Anarchy: Auctions

- Even with two bidders and asymmetric distributions, the BNE can be very complicated:

EXAMPLE 2.2. (Two bidders with uniform $[0, 1]$ and uniform $[0, 2]$ distributions (Vickrey, 1961)) One can verify that the following bidding functions constitute an equilibrium in this example (see also Krishna, 2002):

$$s_1(v_1) = \frac{4}{3v_1} \left(1 - \sqrt{1 - \frac{3v_1^2}{4}} \right)$$
$$s_2(v_2) = \frac{4}{3v_2} \left(\sqrt{1 + \frac{3v_2^2}{4}} - 1 \right).$$

- Instead of trying to find the Bayes Nash, the literature quantifies the approximation ratio of any Bayes Nash equilibrium using price of anarchy
- **Theorem** (Price of Anarchy of First-Price Single-Item Auctions). The price of anarchy of the first-price single-item auction format is at least $1 - 1/e \approx 0.63$.
 - justifies conventional wisdom about such auctions

Revenue Maximization

Let's Talk about Revenue

- So, far revenue is incidental: payment was necessary to impose truthful behavior
- Start with **one bidder** with private value v and **one item**
- What is the unique DSIC surplus maximizing auction for this setting?
 - Allocate the item to the the bidder
 - Charge critical bid: zero
- Thus, our DSIC surplus-maximizing auction has zero revenue
- Any ideas on how we can improve the revenue?
 - Post a minimum price
 - Ignore bids below it: don't sell
 - If bid is above it, charge posted price

v



Single Bidder Single Item

- Suppose we knew the bidder's value v (maximum willingness to pay)
 - What should the posted price r be?
 - $r = v$ (also called reservation/**reserve price** or the **monopoly price**)
- Unfortunately we don't know v
- What are the tradeoffs of setting r too high or too low?
 - Set it too high, might not sell the item
 - Set it too low, might get less revenue than is possible
- What if the seller knew the distribution F from which v is drawn?

v



Posted Price for One Bidder

- Suppose we set the reserve price to r and the value v of the bidder is drawn from a distribution with CDF F
 - If $v < r$: no sale
 - Otherwise we sell the item at price r
- What is the expected revenue?

$$\begin{aligned} E[R] &= r \cdot \Pr(\text{sale}) + 0 \cdot \Pr(\text{no sale}) \\ &= r \cdot (1 - F(r)) \end{aligned}$$

- If $v \sim$ i.i.d. $U(0,1)$, then $F(r) = r$

$$E[R] = r(1 - r), \text{ maximized at } r = 1/2$$

- Achieving an expected revenue of $1/4$

Notice that we **sometimes don't sell the item, i.e. (this is not surplus maximizing)**: revenue equivalence says we must allocate item differently to generate more revenue

v




Second Price with Reserve

- Suppose now we have two bidders: suppose both their values is drawn uniform i.i.d. from $U(0,1)$ and no reserve price
- We calculated the revenue of second price auction without reserve last lecture



$$E[R_2] = \frac{n-1}{n+1} = \frac{1}{3}$$

- Can we improve this revenue if we have a reserve price?
- Suppose $r = 1/2$





Second Price with Reserve

- Suppose now we have two bidders: suppose both their values is drawn uniform i.i.d. from $U(0,1)$ and $r = 1/2$
- Probability that both bidder values are below $1/2$
 - Probability that two uniformly randomly thrown balls fall into the first half (when both are thrown independently)
 - $1/2 \cdot 1/2 = 1/4$
 - Expected revenue in this case?
 - 0





Second Price with Reserve

- Suppose now we have two bidders: suppose both their values is drawn uniform i.i.d. from $U(0,1)$ and $r = 1/2$
- Probability that one bidder value is above $1/2$, other below
 - Probability that at exactly one ball (thrown uniformly randomly and independently) lands in the first half
 - $1/2 \cdot 1/2 + 1/2 \cdot 1/2 = 1/2$
- Expected revenue in this case?
 - Reserve price $r = 1/2$



Second Price with Reserve

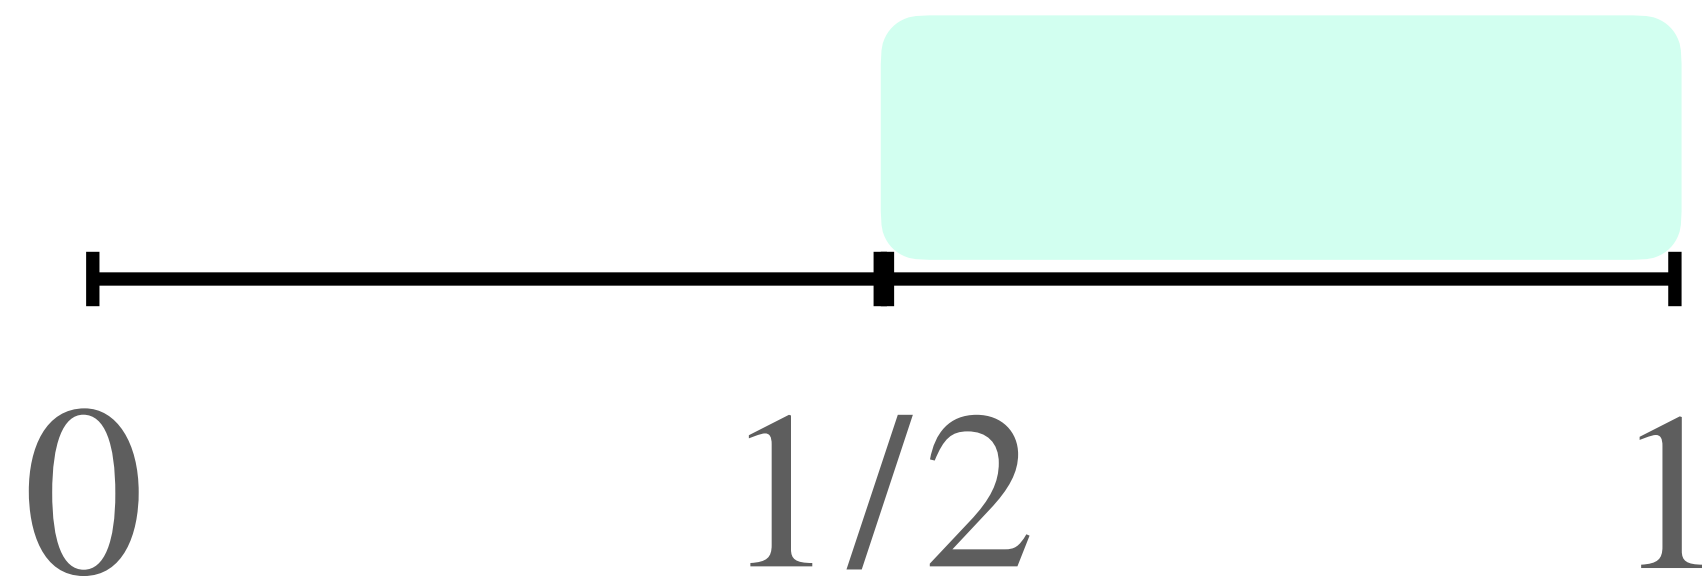
- Suppose now we have two bidders: suppose both their values is drawn uniform i.i.d. from $U(0,1)$ and $r = 1/2$
- Probability that both bidder values are above $1/2$
 - Probability that two uniformly randomly thrown balls fall into the second half (when both are thrown independently)
 - $1/2 \cdot 1/2 = 1/4$
- Expected revenue in this case?
 - Expected value of second-highest sample when two samples are drawn iid from $U(0.5, 1)$



Order Statistics: Uniform

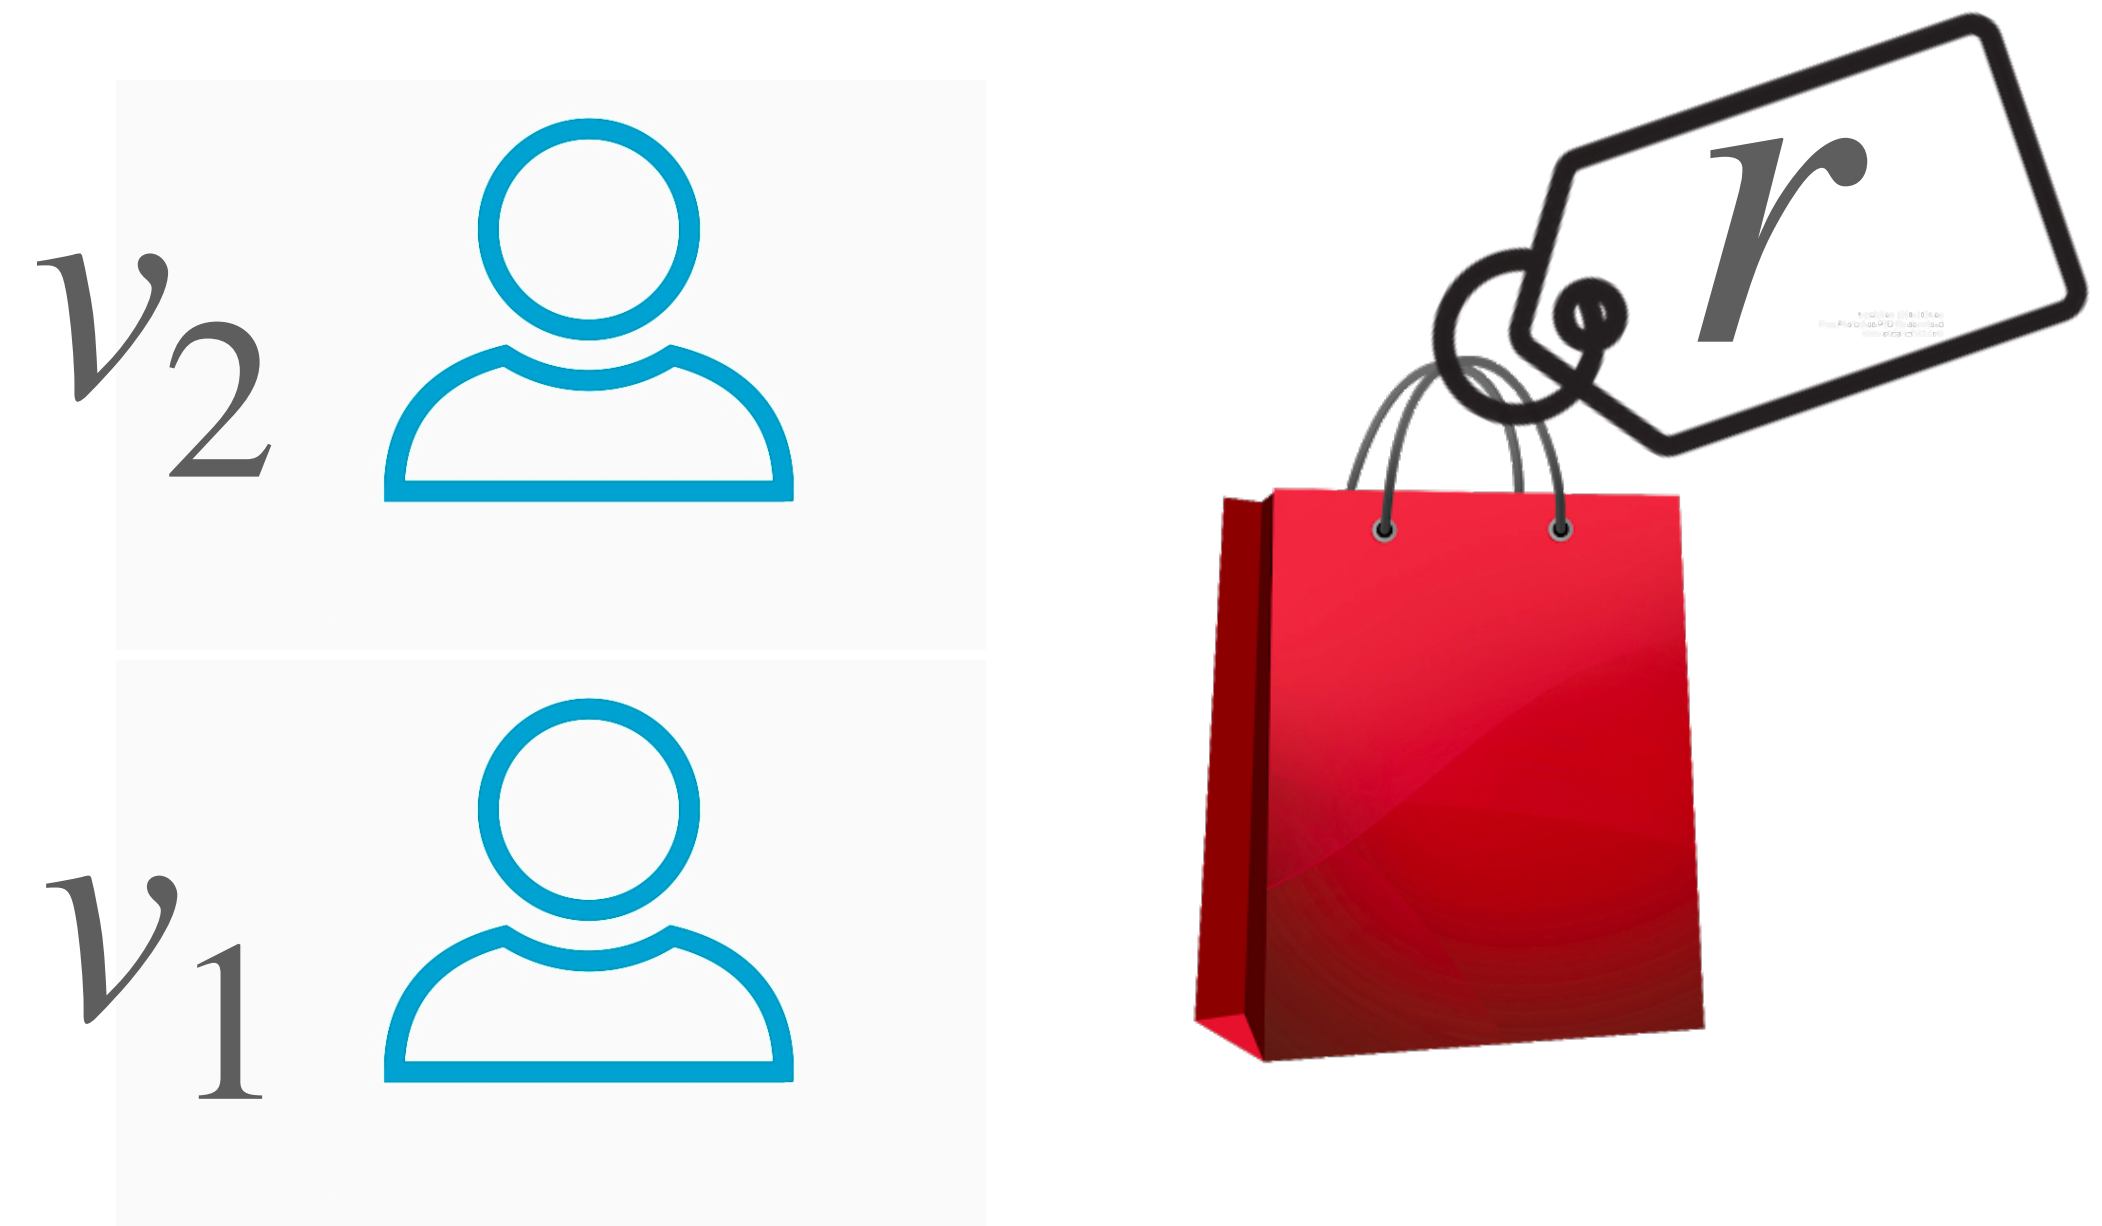

- We use the fact that a uniform random variable evenly divides the interval its over
- In this case the interval is $[a, b]$ where $a = 0.5$, $b = 1$
- Expected value of k th order statistic for n samples drawn iid from the range $[a, b]$ is

$$a + \frac{n - (k - 1)}{n + 1} \cdot (b - a)$$



Second Price with Reserve

- Suppose now we have two bidders: suppose both their values is drawn uniform i.i.d. from $U(0,1)$ and $r = 1/2$
- Probability that both bidder values are above $1/2$
 - $1/2 \cdot 1/2 = 1/4$
- Expected revenue in this case?
 - Expected value of second-highest sample when two samples are drawn iid from $U(0.5, 1)$ is
$$\frac{1}{2} + \frac{2 - (2 - 1)}{2 + 1} \cdot \frac{1}{2} = \frac{2}{3}$$



Second Price with Reserve

- Suppose now we have two bidders: suppose both their values is drawn uniform i.i.d. from $U(0,1)$ and $r = 1/2$
- Putting it all together:
 - $\frac{1}{4} \cdot 0 + \frac{1}{2} \cdot \frac{1}{2} + \frac{1}{4} \cdot \frac{2}{3} = \frac{5}{12}$
- Expected revenue increased!
 - Without reserve $1/3$, with a reserve of $1/2$ it is $5/12$
- **Question.** Can we do better?
 - By using a different reserve price or using a totally different auction format?
- **(Side question.)** Is this auction still DSIC?
 - Pretend there is an another bidder with bid r

Revenue Optimal Auctions

- **(Regular distributions).** A distribution F with density f from which value x is drawn is regular if $x - \frac{1 - F(x)}{f(x)}$ is **strictly increasing**
- Uniform distributions, exponential distributions and lognormal distributions are regular
- Irregular distribution examples: bi-model or multi model distributions and distributions with sufficiently heavy tails
- Myerson's proved the following foundational result on such class of distributions
- **Theorem.** If bidders values are drawn IID from a regular distribution F , then the Vickrey (second-price) auction with a reserve price equal to the "monopoly price" (i.g., $\operatorname{argmax}_r r \cdot (1 - F(r))$) is DSIC and maximizes expected revenue.
- The above generalizes to all single-parameter settings!

Takeaways

- Going from surplus maximization to revenue is not a big jump
 - We just need to add a suitable reserve price!
- Revenue optimal auctions are simple
 - This holds for all the single-parameter auctions we have seen, including sponsored ad auctions
- Surprisingly do not depend on number of bidders n
- Myerson received the 2007 Nobel Prize in part for this work

Myerson's theory can be generalized to non IID and non regular distributions as well!

Theory vs Practice

- **eBay Auctions** are essentially second price with a suitable "opening bid" (a reserve price)
- Thus the theory we developed argues that eBay auctions are the best possible for revenue and are strategyproof!



Reserve Prices in Yahoo!

- Does our optimal auction theory apply well in practice?
- Ostrovsky and Schwarz did a field experiment in 2008 exploring the affect of reserve price in Yahoo! Keyword auctions
- Before 2008, Yahoo had been using relatively small reserve prices: around 1 or 5 cents and the same reserve price for all keywords

Reserve Prices in Internet Advertising Auctions:
A Field Experiment*

Michael Ostrovsky[†]

Michael Schwarz[‡]

YAHOO!

Reserve Prices in Yahoo!

- Does our optimal auction theory apply to practice?
- Ostrovsky and Schwarz did a field experiment in 2008 exploring the affect of reserve price in Yahoo! Keyword auctions
- Before 2008, Yahoo had been using relatively small reserve prices: around 1 or 5 cents and the same reserve price for all keywords
- How did Yahoo fare when reserve prices were updated close to the theoretical optimal?
 - Under some reasonable assumptions, the theory said the reserve price should be around 30-40 cents
- Instead of trying out the new prices, they used a conservative approach by averaging the old and new theoretically optimal prices

YAHOO!

Reserve Prices in Yahoo!

- **And it worked!**
- Yahoo's revenue went up several percent (of a huge number!)
- The change was especially effective in "thin" markets: not as competitive (less than 6 bidders)

On the [revenue per search] front I mentioned we grew 11% year-over-year in the quarter [...], so that's north of a 20% gap search growth rate in the US and that is a factor of, attributed to rolling out a number of the product upgrades we've been doing. [Market Reserve Pricing] was probably the most significant in terms of its impact in the quarter. We had a full quarter impact of that in Q3, but we still have the benefit of rolling that around the world.

Sue Decker, President, Yahoo! Inc. Q3 2008 Earnings Call.^{18,19}

Prior-Free Auctions

- Something about our optimal auction is not satisfying:
 - the auctioneer needs to know the distribution from which bidder values are drawn a priori
- Can we generate good revenue if we did not know bidder values?
- That is, is there an auction that does not use the distribution **in its design**, but we may use it to analyze the revenue
- There is a beautiful result from classical auction theory:

[Bulow Klemperer] Vickrey auction (with no reserve) with $n + 1$ bidders generates just as much expected revenue as the revenue-optimal auction with n bidders!

The Power of One Additional Bidder

- Implies Vickrey auction always generates $(n - 1)/n$ times the revenue of the optimal auction!
- One additional bidder has the same power as if we knew the distribution of bidder values ahead of time!
- This result that has anecdotal support in practice:
 - Extra competition is more important than getting the details of the auction just right!
 - More useful to generate interest from more participants, than learning more about their preferences!

[Bulow Klemperer] Vickrey auction (with no reserve) with $n + 1$ bidders generates just as much expected revenue as the revenue-optimal auction with n bidders!

Bulow Klemperer: Proof Idea

- Define a fictitious auction A that does the following:
 - Simulate the revenue-optimal auction on n bidders
 - If the item is not allocated, then give it to bidder $n + 1$ for free
- This auction has two useful properties:
 - Its expected revenue with $n + 1$ bidders is exactly that of the optimal with n bidders
 - It always allocates the item
- We can finish the proof by showing that:
- **Claim.** Vickrey auction obtains at least as much expected revenue as any auction that is guaranteed to allocate the item (and thus A)
- $R(\text{Vickrey}, (n + 1)) \geq R(A) \geq \text{OPT}(n)$