Prior Free Profit Maximization: Random Sampling Auctions

Aaron Roth

University of Pennsylvania

April 13 2021

Overview

- We studied Myerson's optimal auctions for revenue maximization.

Overview

- We studied Myerson's optimal auctions for revenue maximization.
- And revenue-competitive pricings...

Overview

- We studied Myerson's optimal auctions for revenue maximization.
- And revenue-competitive pricings...
- But to use them, we needed to know the distribution D from which valuations are drawn.

Overview

- We studied Myerson's optimal auctions for revenue maximization.
- And revenue-competitive pricings...
- But to use them, we needed to know the distribution D from which valuations are drawn.
- To run the VCG mechanism, we didn't need to know anything at all.

Overview

- We studied Myerson's optimal auctions for revenue maximization.
- And revenue-competitive pricings...
- But to use them, we needed to know the distribution D from which valuations are drawn.
- To run the VCG mechanism, we didn't need to know anything at all.
- Can we think about revenue in a distribution independent way?

Overview

- We studied Myerson's optimal auctions for revenue maximization.
- And revenue-competitive pricings...
- But to use them, we needed to know the distribution D from which valuations are drawn.
- To run the VCG mechanism, we didn't need to know anything at all.
- Can we think about revenue in a distribution independent way?
- This lecture: A case study "digital goods auctions"

Digital Goods Auctions

- Digital goods auctions (*unlimited supply auctions*) models the sale of goods with zero marginal cost of production (e.g. software).

Digital Goods Auctions

- Digital goods auctions (*unlimited supply auctions*) models the sale of goods with zero marginal cost of production (e.g. software).
- Hence, there is no constraint on how many individuals can "win" the auction.

Digital Goods Auctions

- Digital goods auctions (*unlimited supply auctions*) models the sale of goods with zero marginal cost of production (e.g. software).
- Hence, there is no constraint on how many individuals can "win" the auction.

Definition

A digital goods auction is a single parameter domain with a set of alternatives $A = \{S \subseteq [n]\}$ – i.e. *any* set of bidders is a feasible outcome. For $a \in A$ we write $a_i = \begin{cases} 1, & \text{if } i \in S; \\ 0, & \text{otherwise.} \end{cases}$. Each bidder's valuation function is parameterized by $v_i \in \mathbb{R}_{\geq 0}$, and $v_i(a) := v_i \cdot a_i$.

Digital Goods Auctions

- Observe: Welfare and profit maximization are in conflict here.

Digital Goods Auctions

- Observe: Welfare and profit maximization are in conflict here.
- The VCG mechanism would allocate to everybody and charge nothing.

Digital Goods Auctions

- Observe: Welfare and profit maximization are in conflict here.
- The VCG mechanism would allocate to everybody and charge nothing.
- To maximize revenue, we'll need to artificially limit supply.

Digital Goods Auctions

- Observe: Welfare and profit maximization are in conflict here.
- The VCG mechanism would allocate to everybody and charge nothing.
- To maximize revenue, we'll need to artificially limit supply.
- But first, what should our benchmark be?

Revenue Benchmark

- When we had a prior distribution D, we could define the *optimal* revenue.

Revenue Benchmark

- When we had a prior distribution D, we could define the *optimal* revenue.
- But what is a reasonable benchmark?

Revenue Benchmark

- When we had a prior distribution D, we could define the *optimal* revenue.
- But what is a reasonable benchmark?
- If we knew D, the revenue optimal auction would correspond to a fixed price $p = \phi^{-1}(0)$.

Revenue Benchmark

- When we had a prior distribution D, we could define the *optimal* revenue.
- But what is a reasonable benchmark?
- If we knew D, the revenue optimal auction would correspond to a fixed price $p = \phi^{-1}(0)$.
- So if we could compete with the revenue of the *best* fixed price we'd be competing with the (unknown) Bayesian optimal benchmark.

Fixed Price Benchmarks

- At price p, everyone with value $v_i \geq p$ buys. We obtain revenue $p \cdot |\{i : v_i \geq p\}|$.

Fixed Price Benchmarks

- At price p, everyone with value $v_i \geq p$ buys. We obtain revenue $p \cdot |\{i : v_i \geq p\}|$.
- The best fixed price in hindsight is always $p \in \{v_1, \ldots, v_n\}$. (why?)

Fixed Price Benchmarks

- At price p, everyone with value $v_i \geq p$ buys. We obtain revenue $p \cdot |\{i : v_i \geq p\}|$.
- The best fixed price in hindsight is always $p \in \{v_1, \ldots, v_n\}$. (why?)
- The revenue of the best fixed price is therefore:

$$\mathrm{OPT}(v) = \max_i v_i \cdot |\{j : v_j \geq v_i\}| = \max_i (i \cdot v_{(i)})$$

where $v_{(i)}$ is the i'th highest valuation in sorted order.

Fixed Price Benchmarks

- At price p, everyone with value $v_i \geq p$ buys. We obtain revenue $p \cdot |\{i : v_i \geq p\}|$.
- The best fixed price in hindsight is always $p \in \{v_1, \ldots, v_n\}$. (why?)
- The revenue of the best fixed price is therefore:

$$\mathrm{OPT}(v) = \max_i v_i \cdot |\{j : v_j \geq v_i\}| = \max_i (i \cdot v_{(i)})$$

  where $v_{(i)}$ is the i'th highest valuation in sorted order.
- ... But this isn't attainable by any truthful mechanism when $i = 1$. Consider the case of $n = 1$.

Fixed Price Benchmarks

- A slightly weaker benchmark: the revenue of the best fixed price *that sells to at least 2 people*.

$$\mathrm{OPT}^{\geq 2}(v) = \max_{i \geq 2} \left(i \cdot v_{(i)} \right)$$

Fixed Price Benchmarks

- A slightly weaker benchmark: the revenue of the best fixed price *that sells to at least 2 people*.

$$\mathrm{OPT}^{\geq 2}(v) = \max_{i \geq 2} \left(i \cdot v_{(i)} \right)$$

- We shouldn't think of this as a serious restriction in a large market...

Fixed Price Benchmarks

- A slightly weaker benchmark: the revenue of the best fixed price *that sells to at least 2 people*.

$$\mathrm{OPT}^{\geq 2}(v) = \max_{i \geq 2} \left(i \cdot v_{(i)} \right)$$

- We shouldn't think of this as a serious restriction in a large market...
- How should we obtain it?

Fixed Price Benchmarks

- A slightly weaker benchmark: the revenue of the best fixed price *that sells to at least 2 people*.

$$\mathrm{OPT}^{\geq 2}(v) = \max_{i \geq 2} \left(i \cdot v_{(i)} \right)$$

- We shouldn't think of this as a serious restriction in a large market...
- How should we obtain it?
- Attempt 1: Just compute the best fixed price v_j from the bids and use that. (Not truthful).

Fixed Price Benchmarks

- Attempt 2: Offer each i price p_i corresponding to $\mathrm{OPT}^{\geq 2}(v_{-i})$ – i.e. the best fixed price excluding agent i.

Fixed Price Benchmarks

- Attempt 2: Offer each i price p_i corresponding to $\mathrm{OPT}^{\geq 2}(v_{-i})$ – i.e. the best fixed price excluding agent i.
- This yields a truthful mechanism. How does it do with respect to the benchmark?

Fixed Price Benchmarks

- Attempt 2: Offer each i price p_i corresponding to $\mathrm{OPT}^{\geq 2}(v_{-i})$ – i.e. the best fixed price excluding agent i.
- This yields a truthful mechanism. How does it do with respect to the benchmark?

Example

Suppose we have 90 "low value" agents with $v_i = 1$, and 10 "high value" agents with $v_i = 10$. $\mathrm{OPT}^{\geq 2}(v) = 100$, achieved by charging either $p = 10$ or $p = 1$. But for $v_i = 1$, $\mathrm{OPT}^{\geq 2}(v_{-i}) \leftrightarrow p_i = 10$, and for $v_i = 10$, $\mathrm{OPT}^{\geq 2}(v_{-i}) \leftrightarrow p_i = 1$. So this auction gets profit only 10... (And the ratio to $OPT^{\geq 2}(v)$ can be made arbitrarily bad.)

Profit Extractors

- Lets start with an intermediate goal.

Profit Extractors

- Lets start with an intermediate goal.
- Given a target profit R, want a mechanism that will obtain profit R *if* $\mathrm{OPT}^{\geq 2}(v) \geq R$.

Profit Extractors

- Lets start with an intermediate goal.
- Given a target profit R, want a mechanism that will obtain profit R *if* $\mathrm{OPT}^{\geq 2}(v) \geq R$.
- Otherwise we won't require any revenue guarantee for the mechanism.

Profit Extractors

- Lets start with an intermediate goal.
- Given a target profit R, want a mechanism that will obtain profit R *if* $\mathrm{OPT}^{\geq 2}(v) \geq R$.
- Otherwise we won't require any revenue guarantee for the mechanism.

Definition

The digital goods profit extractor with target profit R ($\mathrm{Extract}(R, v)$) does the following: it finds the largest value k such that $v_{(k)} \geq R/k$, and then sells to the top k bidders at price R/k. If there is no such k, it sells to nobody.

Profit Extractors

- Lets start with an intermediate goal.
- Given a target profit R, want a mechanism that will obtain profit R *if* $\mathrm{OPT}^{\geq 2}(v) \geq R$.
- Otherwise we won't require any revenue guarantee for the mechanism.

Definition

The digital goods profit extractor with target profit R (Extract(R, v)) does the following: it finds the largest value k such that $v_{(k)} \geq R/k$, and then sells to the top k bidders at price R/k. If there is no such k, it sells to nobody.

Lemma

Extract(R, v) is dominant strategy truthful.

Profit Extractors are Dominant Strategy IC

- View the profit extractor as running the following process:

Profit Extractors are Dominant Strategy IC

- View the profit extractor as running the following process:
  1. Start with $k = n$, and offer a price of $p = R/k$ to the bidders.
  2. If any bidder *rejects* the offer (i.e. $v_{(k)} < R_i$), remove her from the auction, set $k \leftarrow k - 1$ and repeat the offer of $p = R/k$ (now a higher offer, to 1 fewer bidders).
  3. If all k bidders *accept* the offer, then they (the top k) bidders receive the good and pay the last offer price.

Profit Extractors are Dominant Strategy IC

- View the profit extractor as running the following process:
  1. Start with $k = n$, and offer a price of $p = R/k$ to the bidders.
  2. If any bidder *rejects* the offer (i.e. $v_{(k)} < R_i$), remove her from the auction, set $k \leftarrow k - 1$ and repeat the offer of $p = R/k$ (now a higher offer, to 1 fewer bidders).
  3. If all k bidders *accept* the offer, then they (the top k) bidders receive the good and pay the last offer price.
- Note that if any bidder rejects the offer, she can never win in any future round.

Profit Extractors are Dominant Strategy IC

- View the profit extractor as running the following process:
  1. Start with $k = n$, and offer a price of $p = R/k$ to the bidders.
  2. If any bidder *rejects* the offer (i.e. $v_{(k)} < R_i$), remove her from the auction, set $k \leftarrow k - 1$ and repeat the offer of $p = R/k$ (now a higher offer, to 1 fewer bidders).
  3. If all k bidders *accept* the offer, then they (the top k) bidders receive the good and pay the last offer price.
- Note that if any bidder rejects the offer, she can never win in any future round.
- So rejecting any offer of $p < v_i$ is a dominated strategy.

Profit Extractors are Dominant Strategy IC

- View the profit extractor as running the following process:
  1. Start with $k = n$, and offer a price of $p = R/k$ to the bidders.
  2. If any bidder *rejects* the offer (i.e. $v_{(k)} < R_i$), remove her from the auction, set $k \leftarrow k - 1$ and repeat the offer of $p = R/k$ (now a higher offer, to 1 fewer bidders).
  3. If all k bidders *accept* the offer, then they (the top k) bidders receive the good and pay the last offer price.
- Note that if any bidder rejects the offer, she can never win in any future round.
- So rejecting any offer of $p < v_i$ is a dominated strategy.
- Similarly, accepting an offer of $p > v_i$ is a dominated strategy since prices only rise.

Profit Extractors are Dominant Strategy IC

- View the profit extractor as running the following process:
  1. Start with $k = n$, and offer a price of $p = R/k$ to the bidders.
  2. If any bidder *rejects* the offer (i.e. $v_{(k)} < R_i$), remove her from the auction, set $k \leftarrow k - 1$ and repeat the offer of $p = R/k$ (now a higher offer, to 1 fewer bidders).
  3. If all k bidders *accept* the offer, then they (the top k) bidders receive the good and pay the last offer price.
- Note that if any bidder rejects the offer, she can never win in any future round.
- So rejecting any offer of $p < v_i$ is a dominated strategy.
- Similarly, accepting an offer of $p > v_i$ is a dominated strategy since prices only rise.
- Hence the dominant strategy for every bidder i is to report their true value.

Profit Extractors

Lemma

Extract(R, v) *obtains revenue* R *if* $\mathrm{OPT}^{\geq 2}(v) \geq R$, *and otherwise obtains revenue* 0.

Profit Extractors

Lemma

Extract(R, v) *obtains revenue* R *if* $\mathrm{OPT}^{\geq 2}(v) \geq R$*, and otherwise obtains revenue* 0*.*

Proof.

- Recall: $\mathrm{OPT}^{\geq 2}(v) = k \cdot v_{(k)}$ for some $k \in \{2, \ldots, n\}$.

$\square$

Profit Extractors

Lemma

Extract(R, v) *obtains revenue* R *if* $\mathrm{OPT}^{\geq 2}(v) \geq R$*, and otherwise obtains revenue* 0*.*

Proof.

- Recall: $\mathrm{OPT}^{\geq 2}(v) = k \cdot v_{(k)}$ for some $k \in \{2, \ldots, n\}$.
- If $\mathrm{OPT}^{\geq 2}(v) \geq R$ then $v_{(k)} \geq \frac{R}{k}$.

$\square$

Profit Extractors

Lemma

Extract(R, v) *obtains revenue* R *if* $\mathrm{OPT}^{\geq 2}(v) \geq R$*, and otherwise obtains revenue* 0*.*

Proof.

- Recall: $\mathrm{OPT}^{\geq 2}(v) = k \cdot v_{(k)}$ for some $k \in \{2, \ldots, n\}$.
- If $\mathrm{OPT}^{\geq 2}(v) \geq R$ then $v_{(k)} \geq \frac{R}{k}$.
- Hence, the profit extractor finds some $k' \geq k$ such that $v_{(k')} \geq R/k'$, and obtains profit $k' \cdot R/k' = R$.

$\square$

Profit Extractors

Lemma

Extract(R, v) obtains revenue R if $\mathrm{OPT}^{\geq 2}(v) \geq R$, and otherwise obtains revenue 0.

Proof.

- Recall: $\mathrm{OPT}^{\geq 2}(v) = k \cdot v_{(k)}$ for some $k \in \{2, \ldots, n\}$.
- If $\mathrm{OPT}^{\geq 2}(v) \geq R$ then $v_{(k)} \geq \frac{R}{k}$.
- Hence, the profit extractor finds some $k' \geq k$ such that $v_{(k')} \geq R/k'$, and obtains profit $k' \cdot R/k' = R$.
- If $R > \mathrm{OPT}^{(2)}(v) = \max_k k \cdot v_{(k)}$, then there is no k such that $v_{(k)} \geq R/k$. So the mechanism halts without selling to anybody.

$\square$

Using Profit Extractors

- We now have a useful tool.

Using Profit Extractors

- We now have a useful tool.
- We can obtain revenue R *if* we know that it is possible to obtain revenue R with a fixed price.

Using Profit Extractors

- We now have a useful tool.
- We can obtain revenue R *if* we know that it is possible to obtain revenue R with a fixed price.
- But we're not done, since we don't know R.

Using Profit Extractors

- We now have a useful tool.
- We can obtain revenue R *if* we know that it is possible to obtain revenue R with a fixed price.
- But we're not done, since we don't know R.
- We've reduced our problem to finding a good *estimate* of the true optimal revenue R^*.

Using Profit Extractors

- We now have a useful tool.
- We can obtain revenue R *if* we know that it is possible to obtain revenue R with a fixed price.
- But we're not done, since we don't know R.
- We've reduced our problem to finding a good *estimate* of the true optimal revenue R^*.
- For truthfulness, it is important that R is defined independently of the bidders we run the profit extractor on.

The Random Sampling Auction

Idea: Try and estimate R^* from a random sample of the bidders, and then run the profit extractor on the remaining bidders.

The Random Sampling Auction

Idea: Try and estimate R^* from a random sample of the bidders, and then run the profit extractor on the remaining bidders.

RS(v):

Randomly partition the agents by assigning each agent uniformly at random to one of two sets: S' or S''.
Calculate $R' = \mathrm{OPT}^{\geq 2}(v_{S'})$ and $R'' = \mathrm{OPT}^{\geq 2}(v_{S''})$.
Run Extract$(R', v_{S''})$ on S'' and Extract$(R'', v_{S'})$ on S'.

The Random Sampling Auction

Theorem

The random sampling auction is dominant strategy truthful.

The Random Sampling Auction

Theorem

The random sampling auction is dominant strategy truthful.

Proof.

Extract(R, v) is truthful whenever it is run with a value R computed independently of the bidders it is run on. $\square$

The Random Sampling Auction

Lemma

The revenue of the random sampling auction is at least $\min(R', R'')$.

The Random Sampling Auction

Lemma

The revenue of the random sampling auction is at least $\min(R', R'')$*.*

Proof.

Either $R' \geq R''$ or $R'' \geq R'$ (or possibly both). So at least one copy of Extract succeeds. $\square$

The Random Sampling Auction

Lemma

The revenue of the random sampling auction is at least $\min(R', R'')$*.*

Proof.

Either $R' \geq R''$ or $R'' \geq R'$ (or possibly both). So at least one copy of Extract succeeds. $\square$

So it remains to understand $\min(R', R'')$ as a function of $R := \mathrm{OPT}^{\geq 2}(v)$.

Some Simple Math

Theorem

If we flip $k \geq 2$ coins, then $\mathbb{E}[\min(\#heads, \#tails)] \geq k/4$.

Proof.

$\square$

Some Simple Math

Theorem

If we flip $k \geq 2$ coins, then $\mathbb{E}[\min(\#heads, \#tails)] \geq k/4$.

Proof.

- Let M_i be $\mathbb{E}[\min(\#\text{heads}, \#\text{tails})]$ after i coin flips.

$\square$

Some Simple Math

Theorem

If we flip $k \geq 2$ coins, then $\mathbb{E}[\min(\#heads, \#tails)] \geq k/4$.

Proof.

- Let M_i be $\mathbb{E}[\min(\#\text{heads}, \#\text{tails})]$ after i coin flips.
- Some direct calculations show: $M_1 = 0$ and $M_2 = 1/2$.

$\square$

Some Simple Math

Theorem

If we flip $k \geq 2$ coins, then $\mathbb{E}[\min(\#heads, \#tails)] \geq k/4$.

Proof.

- Let M_i be $\mathbb{E}[\min(\#\mathrm{heads}, \#\mathrm{tails})]$ after i coin flips.
- Some direct calculations show: $M_1 = 0$ and $M_2 = 1/2$.
- Now define $X_i := M_i - M_{i-1}$, the expected change to $\min(\#\mathrm{heads}, \#\mathrm{tails})$ after we flip the i'th coin.

$\square$

Some Simple Math

Theorem

If we flip $k \geq 2$ coins, then $\mathbb{E}[\min(\#heads, \#tails)] \geq k/4$.

Proof.

- Let M_i be $\mathbb{E}[\min(\#\text{heads}, \#\text{tails})]$ after i coin flips.
- Some direct calculations show: $M_1 = 0$ and $M_2 = 1/2$.
- Now define $X_i := M_i - M_{i-1}$, the expected change to $\min(\#\text{heads}, \#\text{tails})$ after we flip the i'th coin.
- By linearity of expectation:

$$M_k = \sum_{i=1}^{k} X_i$$

so we are done if we can compute X_i for all i.

□

Some Simple Math

There are two cases:

- **Case 1: i is even.** $i - 1$ is odd, and so we have $\#\text{heads} \neq \#\text{tails}$ after $i - 1$ coin flips.

Some Simple Math

There are two cases:

- **Case 1: i is even.** $i - 1$ is odd, and so we have $\#\text{heads} \neq \#\text{tails}$ after $i - 1$ coin flips.
- Hence $X_i = 1/2$, since with probability $1/2$, the coin flip contributes to the smaller of the two quantities.

Some Simple Math

There are two cases:

- **Case 1: i is even.** $i - 1$ is odd, and so we have $\#\text{heads} \neq \#\text{tails}$ after $i - 1$ coin flips.
- Hence $X_i = 1/2$, since with probability $1/2$, the coin flip contributes to the smaller of the two quantities.
- **Case 2: i is odd.** $X_i \geq 0$.

Some Simple Math

There are two cases:

- **Case 1: i is even**. $i - 1$ is odd, and so we have $\#\text{heads} \neq \#\text{tails}$ after $i - 1$ coin flips.
- Hence $X_i = 1/2$, since with probability $1/2$, the coin flip contributes to the smaller of the two quantities.
- **Case 2: i is odd**. $X_i \geq 0$.

So:

$$M_k = \sum_{i=1}^{k} X_k \geq \frac{k}{2} \cdot \frac{1}{2} = \frac{k}{4}$$

Some Simple Math

There are two cases:

- **Case 1: i is even**. $i-1$ is odd, and so we have $\#\text{heads} \neq \#\text{tails}$ after $i-1$ coin flips.
- Hence $X_i = 1/2$, since with probability $1/2$, the coin flip contributes to the smaller of the two quantities.
- **Case 2: i is odd**. $X_i \geq 0$.

So:

$$M_k = \sum_{i=1}^{k} X_k \geq \frac{k}{2} \cdot \frac{1}{2} = \frac{k}{4}$$

(Actually, we were a little sloppy... we only showed that $M_k \geq \lfloor \frac{k}{2} \rfloor \cdot \frac{1}{2}$, which might be a little less than $k/4$. To be fully rigorous, we have to directly verify that $X_3 = 1/4$ which makes up the difference).

The Random Sampling Auction

Theorem

Let Rev be the expected revenue of the random sampling auction. Then:

$$Rev \geq \frac{\mathrm{OPT}^{\geq 2}(v)}{4}.$$

The Random Sampling Auction

- Recall:

$$Rev \geq \mathbb{E}[\min(R', R'')]$$

The Random Sampling Auction

- Recall:

$$Rev \geq \mathbb{E}[\min(R', R'')]$$

- We know that $\mathrm{OPT}^{\geq 2}(v) = k \cdot p$ for some $k \geq 2$ and some price p.

The Random Sampling Auction

- Recall:

$$Rev \geq \mathbb{E}[\min(R', R'')]$$

- We know that $\mathrm{OPT}^{\geq 2}(v) = k \cdot p$ for some $k \geq 2$ and some price p.
- Of the k winners when using price p, let k' be the number in S' and k'' be the number in S''. Observe that $R' \geq k' \cdot p$ and $R'' \geq k'' \cdot p$

The Random Sampling Auction

- Recall:

$$Rev \geq \mathbb{E}[\min(R', R'')]$$

- We know that $\mathrm{OPT}^{\geq 2}(v) = k \cdot p$ for some $k \geq 2$ and some price p.
- Of the k winners when using price p, let k' be the number in S' and k'' be the number in S''. Observe that $R' \geq k' \cdot p$ and $R'' \geq k'' \cdot p$
- Hence:

$$\frac{Rev}{\mathrm{OPT}^{\geq 2}(v)} \geq \frac{\mathbb{E}[\min(R', R'')]}{k \cdot p}$$

The Random Sampling Auction

- Recall:
$$Rev \geq \mathbb{E}[\min(R', R'')]$$
- We know that $\mathrm{OPT}^{\geq 2}(v) = k \cdot p$ for some $k \geq 2$ and some price p.
- Of the k winners when using price p, let k' be the number in S' and k'' be the number in S''. Observe that $R' \geq k' \cdot p$ and $R'' \geq k'' \cdot p$
- Hence:
$$\begin{aligned} \frac{Rev}{\mathrm{OPT}^{\geq 2}(v)} &\geq \frac{\mathbb{E}[\min(R', R'')]}{k \cdot p} \\ &\geq \frac{\mathbb{E}[\min(k' \cdot p, k'' \cdot p)]}{k \cdot p} \end{aligned}$$

The Random Sampling Auction

- Recall:
$$Rev \geq \mathbb{E}[\min(R', R'')]$$
- We know that $\mathrm{OPT}^{\geq 2}(v) = k \cdot p$ for some $k \geq 2$ and some price p.
- Of the k winners when using price p, let k' be the number in S' and k'' be the number in S''. Observe that $R' \geq k' \cdot p$ and $R'' \geq k'' \cdot p$
- Hence:
$$\begin{aligned} \frac{Rev}{\mathrm{OPT}^{\geq 2}(v)} &\geq \frac{\mathbb{E}[\min(R', R'')]}{k \cdot p} \\ &\geq \frac{\mathbb{E}[\min(k' \cdot p, k'' \cdot p)]}{k \cdot p} \\ &\geq \frac{\mathbb{E}[\min(k', k'')]}{k} \end{aligned}$$

The Random Sampling Auction

- Recall:
$$Rev \geq \mathbb{E}[\min(R', R'')]$$
- We know that $\mathrm{OPT}^{\geq 2}(v) = k \cdot p$ for some $k \geq 2$ and some price p.
- Of the k winners when using price p, let k' be the number in S' and k'' be the number in S''. Observe that $R' \geq k' \cdot p$ and $R'' \geq k'' \cdot p$
- Hence:
$$\begin{aligned} \frac{Rev}{\mathrm{OPT}^{\geq 2}(v)} &\geq \frac{\mathbb{E}[\min(R', R'')]}{k \cdot p} \\ &\geq \frac{\mathbb{E}[\min(k' \cdot p, k'' \cdot p)]}{k \cdot p} \\ &\geq \frac{\mathbb{E}[\min(k', k'')]}{k} \\ &\geq \frac{1}{4} \end{aligned}$$

Summary

- So we can approximate the revenue of the optimal auction without knowing D.

Summary

- So we can approximate the revenue of the optimal auction without knowing D.
- We got a 4 approximation, but...

Summary

- So we can approximate the revenue of the optimal auction without knowing D.
- We got a 4 approximation, but...
- This was only because we needed to handle the case in which the optimal auction sold to only 2 people.

Summary

- So we can approximate the revenue of the optimal auction without knowing D.
- We got a 4 approximation, but...
- This was only because we needed to handle the case in which the optimal auction sold to only 2 people.
- Similar ideas lead to a $(1+\epsilon)$ approximation of $\mathrm{OPT}^{\geq k}(v)$ as k becomes large.

Thanks!

See you next class — stay healthy, and wear a mask!