# Single-parameter Mechanism design: Sponsored Search Auctions

G. Amanatidis

*Based on slides by A. Voudouris*

# Single-item auctions

- A seller with **one item for sale**
- $n$ **agents**
- Each agent $i$ has a **private value** $v_i$ for the item
  - This value represents the **willingness-to-pay** of the agent; that is, $v_i$ is the maximum amount of money that agent $i$ is willing to pay in order to buy the item
- The utility of each agent is **quasilinear in money**:
  - If agent $i$ loses the item, then her utility is $0$
  - If agent $i$ wins the item at price $p$, then her utility is $v_i - p$

# Single-item auctions

- General structure of an auction:
  - **Input:** every agent $i$ submits a bid $b_i$ (*agents* = *bidders*)
  - **Allocation rule:** decide the winner
  - **Payment rule:** decide a selling price

- Deciding the winner is easy: the highest bidder
- Deciding the selling price is more complicated
  - A selling price of $0$, creates a competition among the bidders as to who can think of the highest number
- We are interested in payment rules that incentivize the bidders to bid their true values
  - **Truthful auctions** that **maximize the social welfare**

# First-price auction

- **Allocation rule:** the winner is the highest bidder
- **Payment rule:** the winner pays her bid

- Is this a truthful auction?

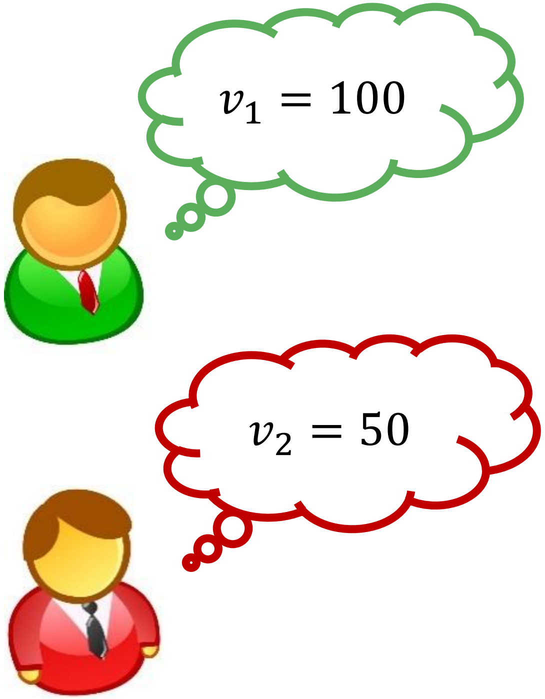

# First-price auction

- **Allocation rule:** the winner is the highest bidder
- **Payment rule:** the winner pays her bid

- Is this a truthful auction?

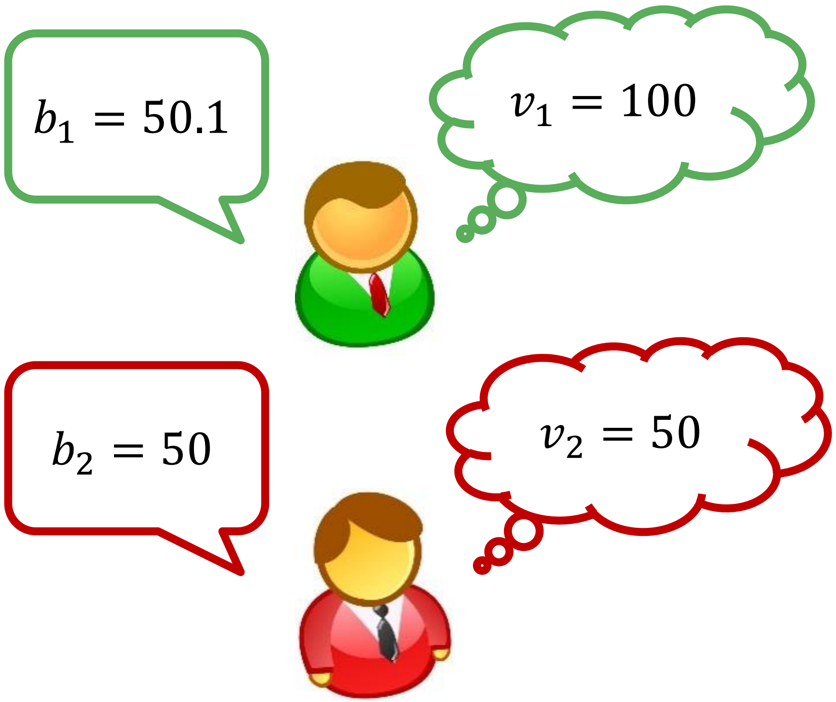

# Second-price auction

- **Allocation rule:** the winner is the highest bidder
- **Payment rule:** the winner pays the second highest bid

**Theorem** [Vickrey, 1961]
In a second-price auction
(a) it is a dominant strategy for every bidder $i$ to bid $b_i = v_i$, and
(b) every truthtelling bidder gets non-negative utility

- (b) is obvious:
  - the selling price is at most the winner's bid, and the bid of a truthtelling bidder is equal to her true value

# Second-price auction

- For (a), our goal is to show that the utility of bidder $i$ is maximized by bidding $v_i$, no matter what $v_i$ and the bids of the other bidders are
- Second highest bid: $B = \max_{j \neq i} b_j$
- The utility of bidder $i$ is either 0 if $b_i < B$, or $v_i - B$ otherwise

**Case I: $\boldsymbol{v_i < B}$**

- Maximum possible utility $= 0$
- Achieved by setting $b_i = v_i$

**Case II: $\boldsymbol{v_i \geq B}$**

- Maximum possible utility $= v_i - B$
- Bidder $i$ wins the item by setting $b_i = v_i$ $\square$

# Sponsored search auctions

CTR

0.36

0.29

0.18

0.5

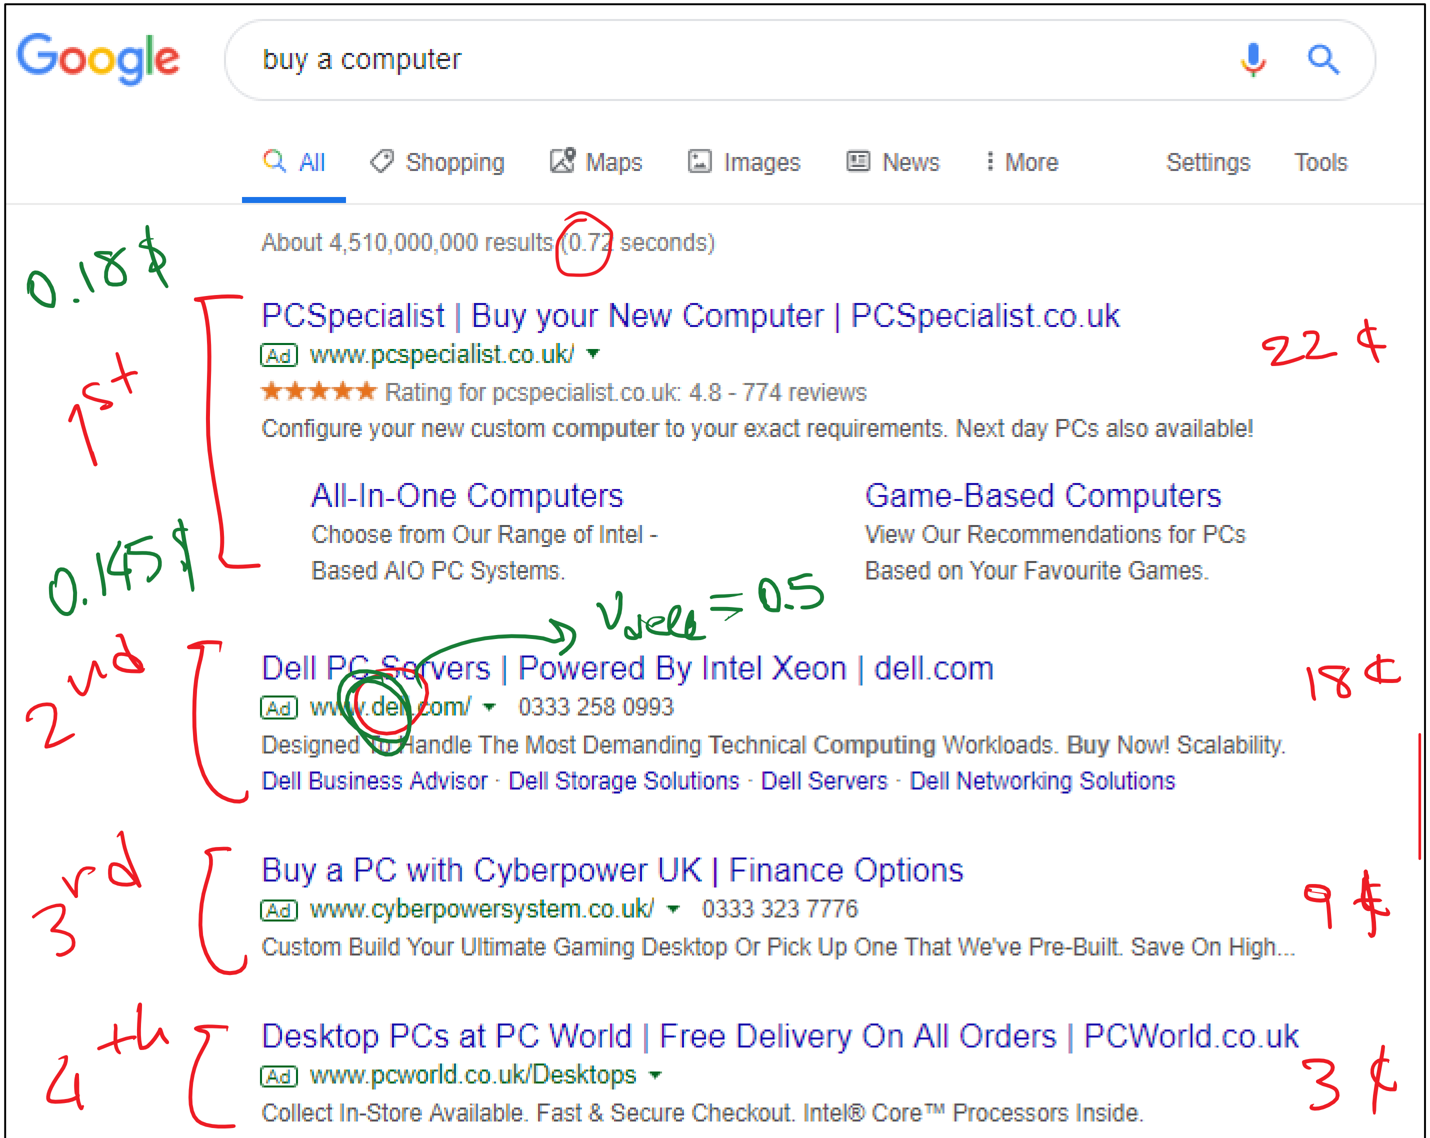

# Sponsored search auctions

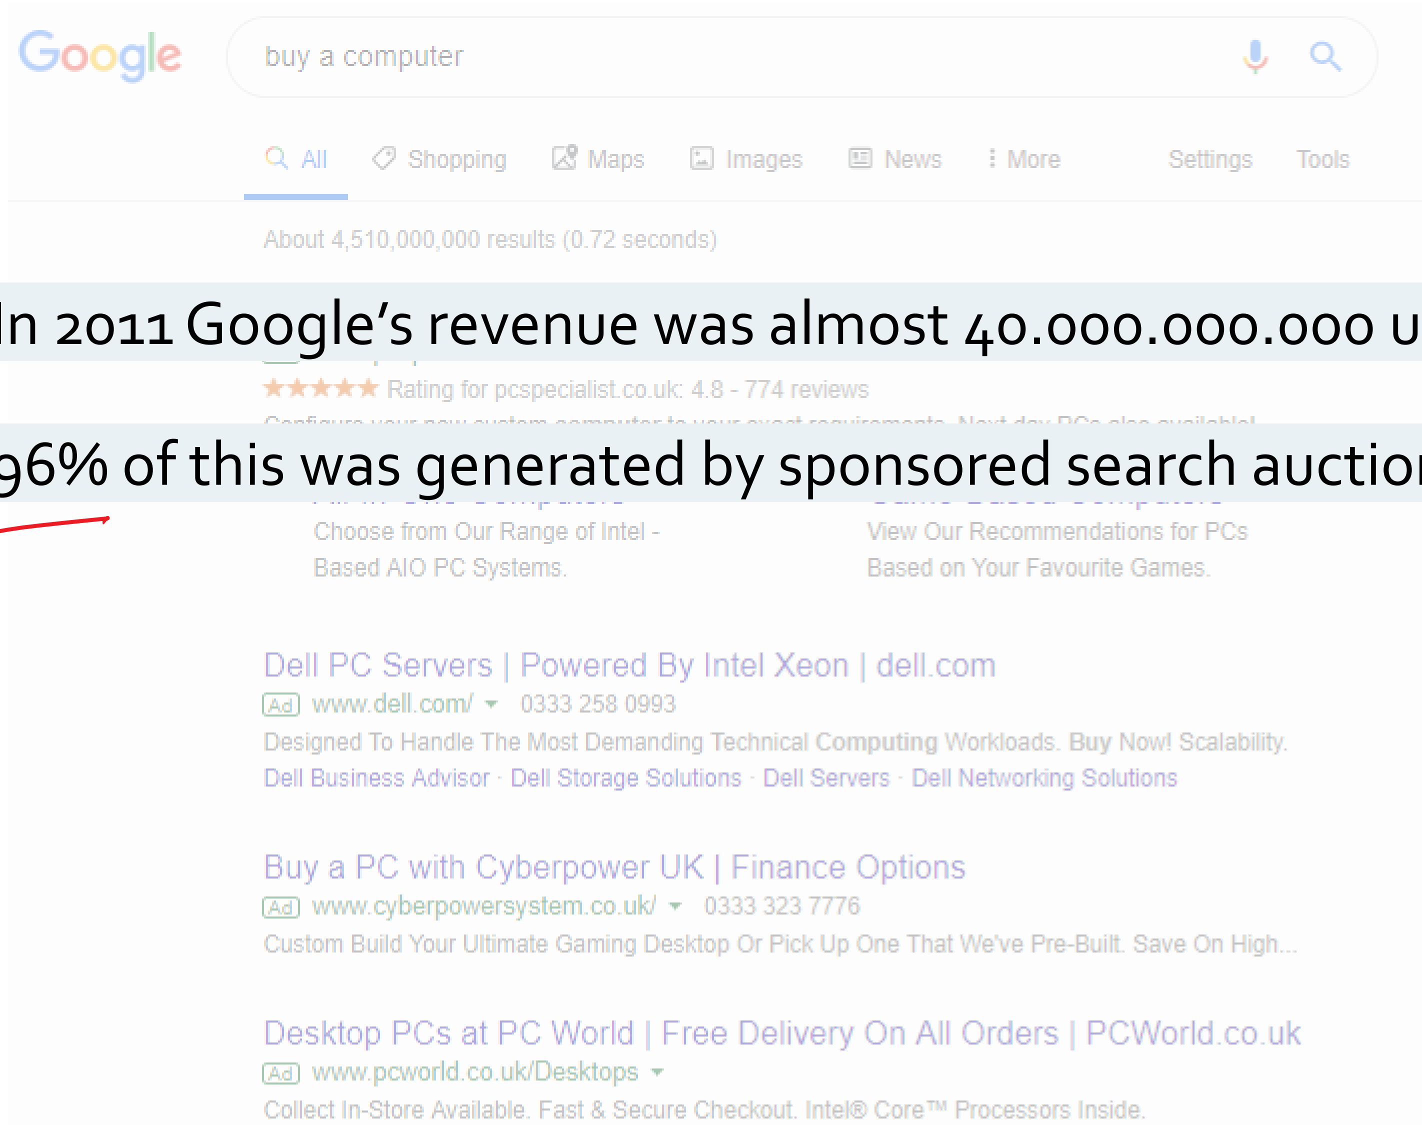


- In 2011 Google's revenue was almost 40.000.000.000 usd
- 96% of this was generated by sponsored search auctions

# Sponsored search auctions

- $k$ advertising **slots**
- $n$ **bidders** (advertisers) who aim to occupy a slot
- Slot $j$ has a **click-through-rate** (CTR) $a_j$
  - The CTR of a slot represents the **probability that the ad placed at this slot will be clicked on**
  - Assumption: the CTRs are independent of the ads that occupy the slots
- The slots are ranked so that $a_1 \geq \cdots \geq a_k$
- Each bidder $i$ has a **private value** $v_i$ **per click**
  - Bidder $i$ derives utility $\boxed{a_j \cdot v_i}$ from slot $j$

# Sponsored search auctions: goals

50 32 29 20 15 9

1
2
3
4

- **Truthfulness:** It is a dominant strategy for each bidder to bid her true value
- **Social welfare maximization:** $\sum_i v_i \cdot x_i$
  - $x_i$ is the CTR of the slot that bidder $i$ is assigned to, or $0$ otherwise
- **Poly-time execution:** running the auction should be quick

$\Theta(n)$

# Sponsored search auctions: goals

- **Truthfulness:** It is a dominant strategy for each bidder to bid her true value
- **Social welfare maximization:** $\sum_i v_i \cdot x_i$
  - $x_i$ is the CTR of the slot that bidder $i$ is assigned to, or $0$ otherwise
- **Poly-time execution:** running the auction should be quick

  $\Theta(n \log n)$
- If the bidders are truthful, then maximizing the social welfare is easy: sort the bidders in decreasing order of their bids
- So, the problem is to incentivize them to be truthful, again
- Can we extend the ideas we exploited for single-item auctions?

# Generalized second-price auction

- **Allocation rule:** sort the bidders in decreasing order of their bids and rename them so that $b_1 \geq \cdots \geq b_n$
- **Payment rule:** every bidder $i \leq k$ (who is assigned at slot $i$) pays the next highest bid $b_{i+1}$ per click, and every bidder $i > k$ pays $0$

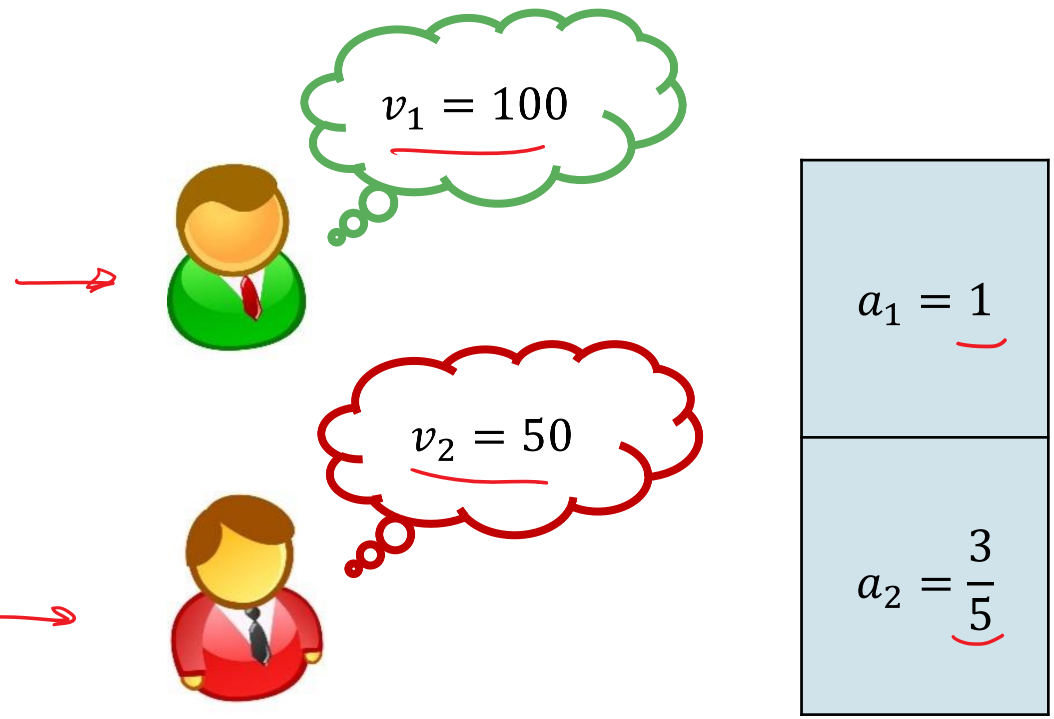

# Generalized second-price auction

- **Allocation rule:** sort the bidders in decreasing order of their bids and rename them so that $b_1 \geq \cdots \geq b_n$
- **Payment rule:** every bidder $i \leq k$ (who is assigned at slot $i$) pays the next highest bid $b_{i+1}$ per click, and every bidder $i > k$ pays $0$

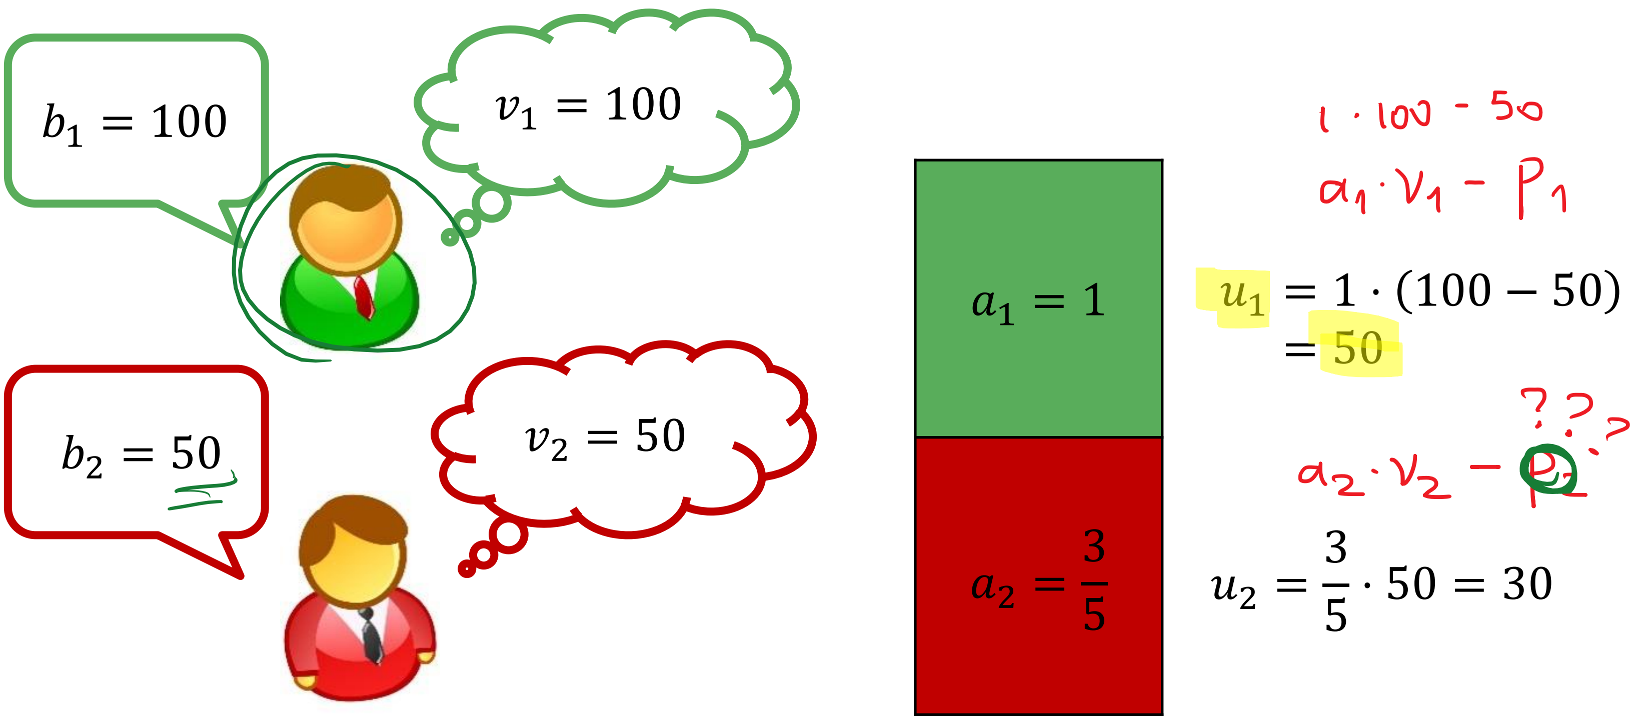

# Generalized second-price auction

- **Allocation rule:** sort the bidders in decreasing order of their bids and rename them so that $b_1 \geq \cdots \geq b_n$
- **Payment rule:** every bidder $i \leq k$ (who is assigned at slot $i$) pays the next highest bid $b_{i+1}$ per click, and every bidder $i > k$ pays $0$

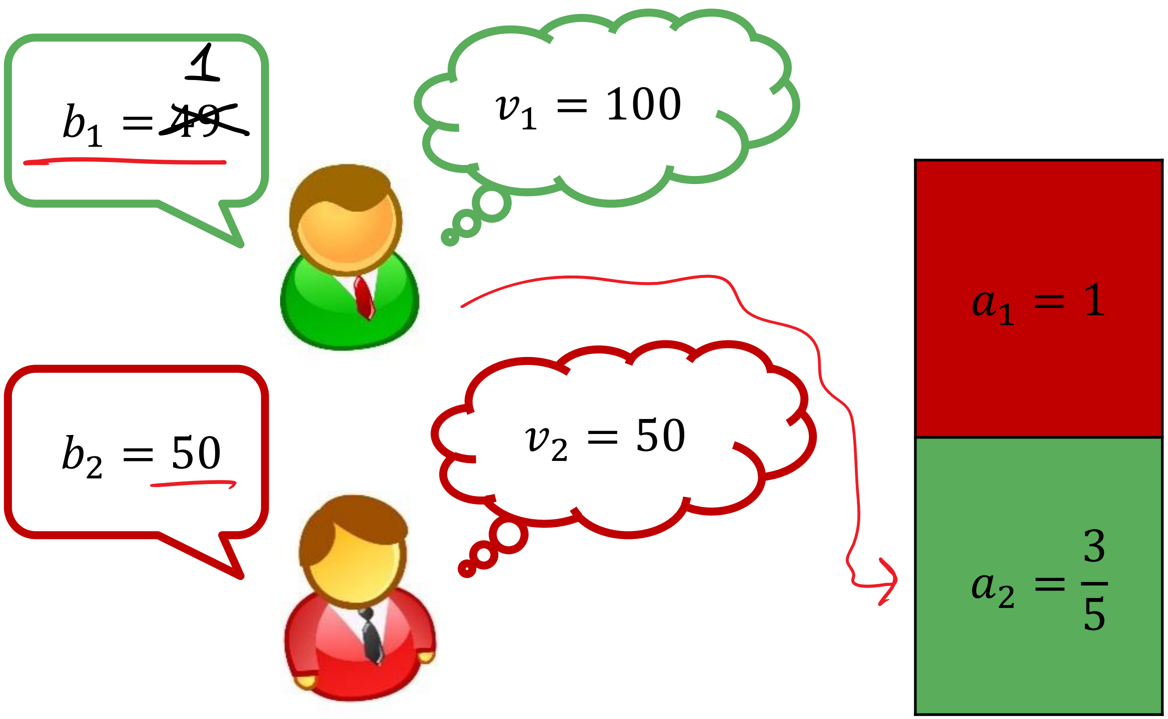


$$u_2 = 1 \cdot (50 - \overset{1}{\cancel{49}})$$
$$= 1$$

$$u_1 = \frac{3}{5} \cdot 100 = 60$$

$a_2 \cdot v_1 - 0 = 60$

# Myerson's Lemma

- That didn't work for sponsored search auctions, so what now?
- Let's try to see how the optimal truthful auction should look like, for any single parameter environment

- Input by bidders: $\boldsymbol{b} = (b_1, \ldots, b_n)$
- Allocation rule: $\boldsymbol{x}(\boldsymbol{b}) = (x_1(\boldsymbol{b}), \ldots, x_n(\boldsymbol{b}))$
- Payment rule: $\boldsymbol{p}(\boldsymbol{b}) = (p_1(\boldsymbol{b}), \ldots, p_n(\boldsymbol{b}))$
- The utility of bidder $i$ is $u_i(\boldsymbol{b}) = v_i \cdot x_i(\boldsymbol{b}) - p_i(\boldsymbol{b})$

$$x_5(b) = \begin{cases} \text{3rd slot} \\ \text{5th slot} \\ \text{no slot} \end{cases}$$

- Focus on payment rules such that $p_i(\boldsymbol{b}) \in [0, b_i \cdot x_i(\boldsymbol{b})]$
  - $p_i(\boldsymbol{b}) \geq 0$ ensures that the seller does not pay the bidders
  - $p_i(\boldsymbol{b}) \leq b_i \cdot x_i(\boldsymbol{b})$ ensures non-negative utility for truthful bidders

# Myerson's Lemma

- An allocation rule $\boldsymbol{x}$ is **implementable** if there exists a payment rule $\boldsymbol{p}$ such that $(\boldsymbol{x}, \boldsymbol{p})$ is a truthful auction
- An allocation rule $\boldsymbol{x}$ is **monotone** if for every bidder $i$ and bid vector $\boldsymbol{b}_{-i}$, the allocation $x_i(z, \boldsymbol{b}_{-i})$ is non-decreasing in the bid $z$ of bidder $i$

can give a truthful mech.

**<u>Lemma</u>** [Myerson, 1981]

(a) An allocation rule $\boldsymbol{x}$ is implementable if and only if it is monotone

monotone

(b) For every allocation rule $\boldsymbol{x}$, there exists a unique payment rule $\boldsymbol{p}$ such that $(\boldsymbol{x}, \boldsymbol{p})$ is a truthful auction

# Proof of Myerson's Lemma

- Fix a bidder $i$, and the bids $\boldsymbol{b}_{-i}$ of the other bidders
- Given that these quantities are now fixed, we simplify our notation:
  - $x(z) = x_i(z, \boldsymbol{b}_{-i})$
  - $p(z) = p_i(z, \boldsymbol{b}_{-i})$
  - $u(z) = u_i(z, \boldsymbol{b}_{-i})$

- The idea:
  - assuming $(\boldsymbol{x}, \boldsymbol{p})$ is a truthful auction, the bidder has no incentive to unilaterally deviate to any other bid
  - This will give us a relation between $\boldsymbol{x}$ and $\boldsymbol{p}$, which we can use to derive an explicit formula for $\boldsymbol{p}$ as a function of $\boldsymbol{x}$

# Proof of Myerson's Lemma

- Consider two bids $0 \le z < y$ and assume $\boldsymbol{x}$ is implementable by $\boldsymbol{p}$
- True value $= z$, deviating bid $= y$:

$$u(z) \ge u(y) \Leftrightarrow z \cdot x(z) - p(z) \ge z \cdot x(y) - p(y)$$
$$\Leftrightarrow p(y) - p(z) \ge z \cdot \big(x(y) - x(z)\big)$$

- True value $= y$, deviating bid $= z$:

$$u(y) \ge u(z) \Leftrightarrow y \cdot x(y) - p(y) \ge y \cdot x(z) - p(z)$$
$$\Leftrightarrow p(y) - p(z) \le y \cdot \big(x(y) - x(z)\big)$$

# Proof of Myerson's Lemma

- Combining these two, we get:

$$z \cdot \big(x(y) - x(z)\big) \le p(y) - p(z) \le y \cdot \big(x(y) - x(z)\big)$$

- This also implies that

$$(y - z) \cdot \big(x(y) - x(z)\big) \ge 0$$

- Since $0 \le z < y$, this is possible if and only if $\boldsymbol{x}$ is monotone so that $y - z > 0$ and $x(y) - x(z) > 0$

  ⇨ (a) is now proved

# Proof of Myerson's Lemma

3 slots

- We can now assume that $\boldsymbol{x}$ is monotone
- Assume $\boldsymbol{x}$ is piecewise constant, like in sponsored search auctions

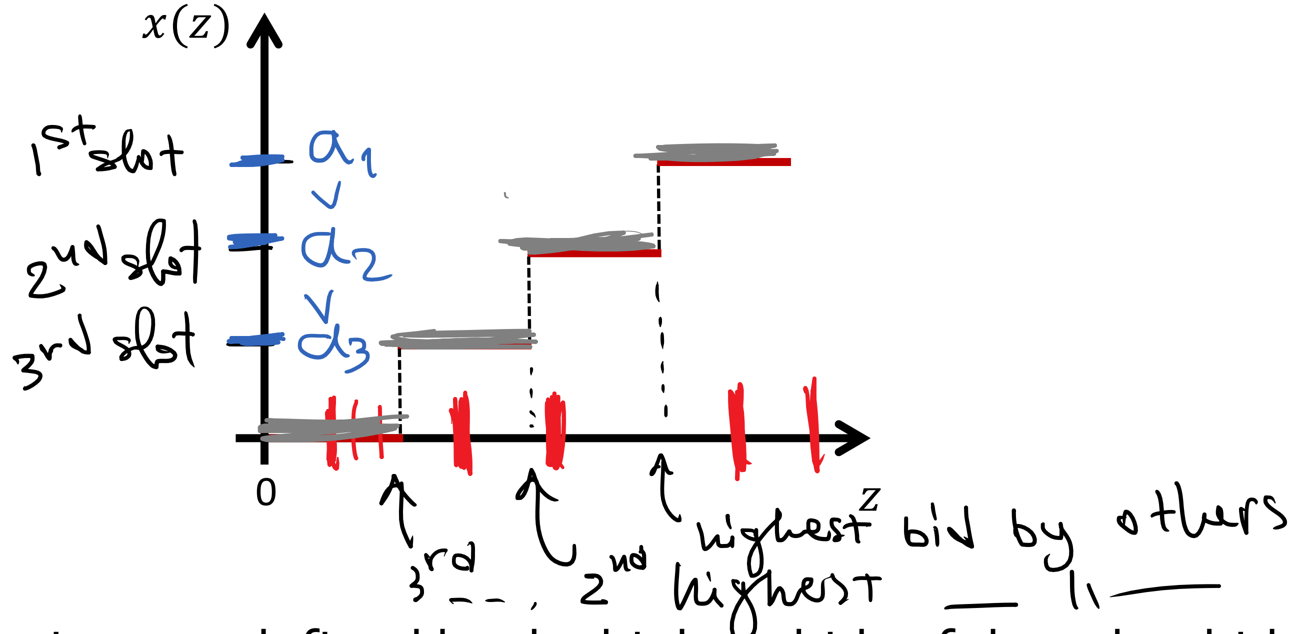


- The break points are defined by the highest bids of the other bidders

# Proof of Myerson's Lemma

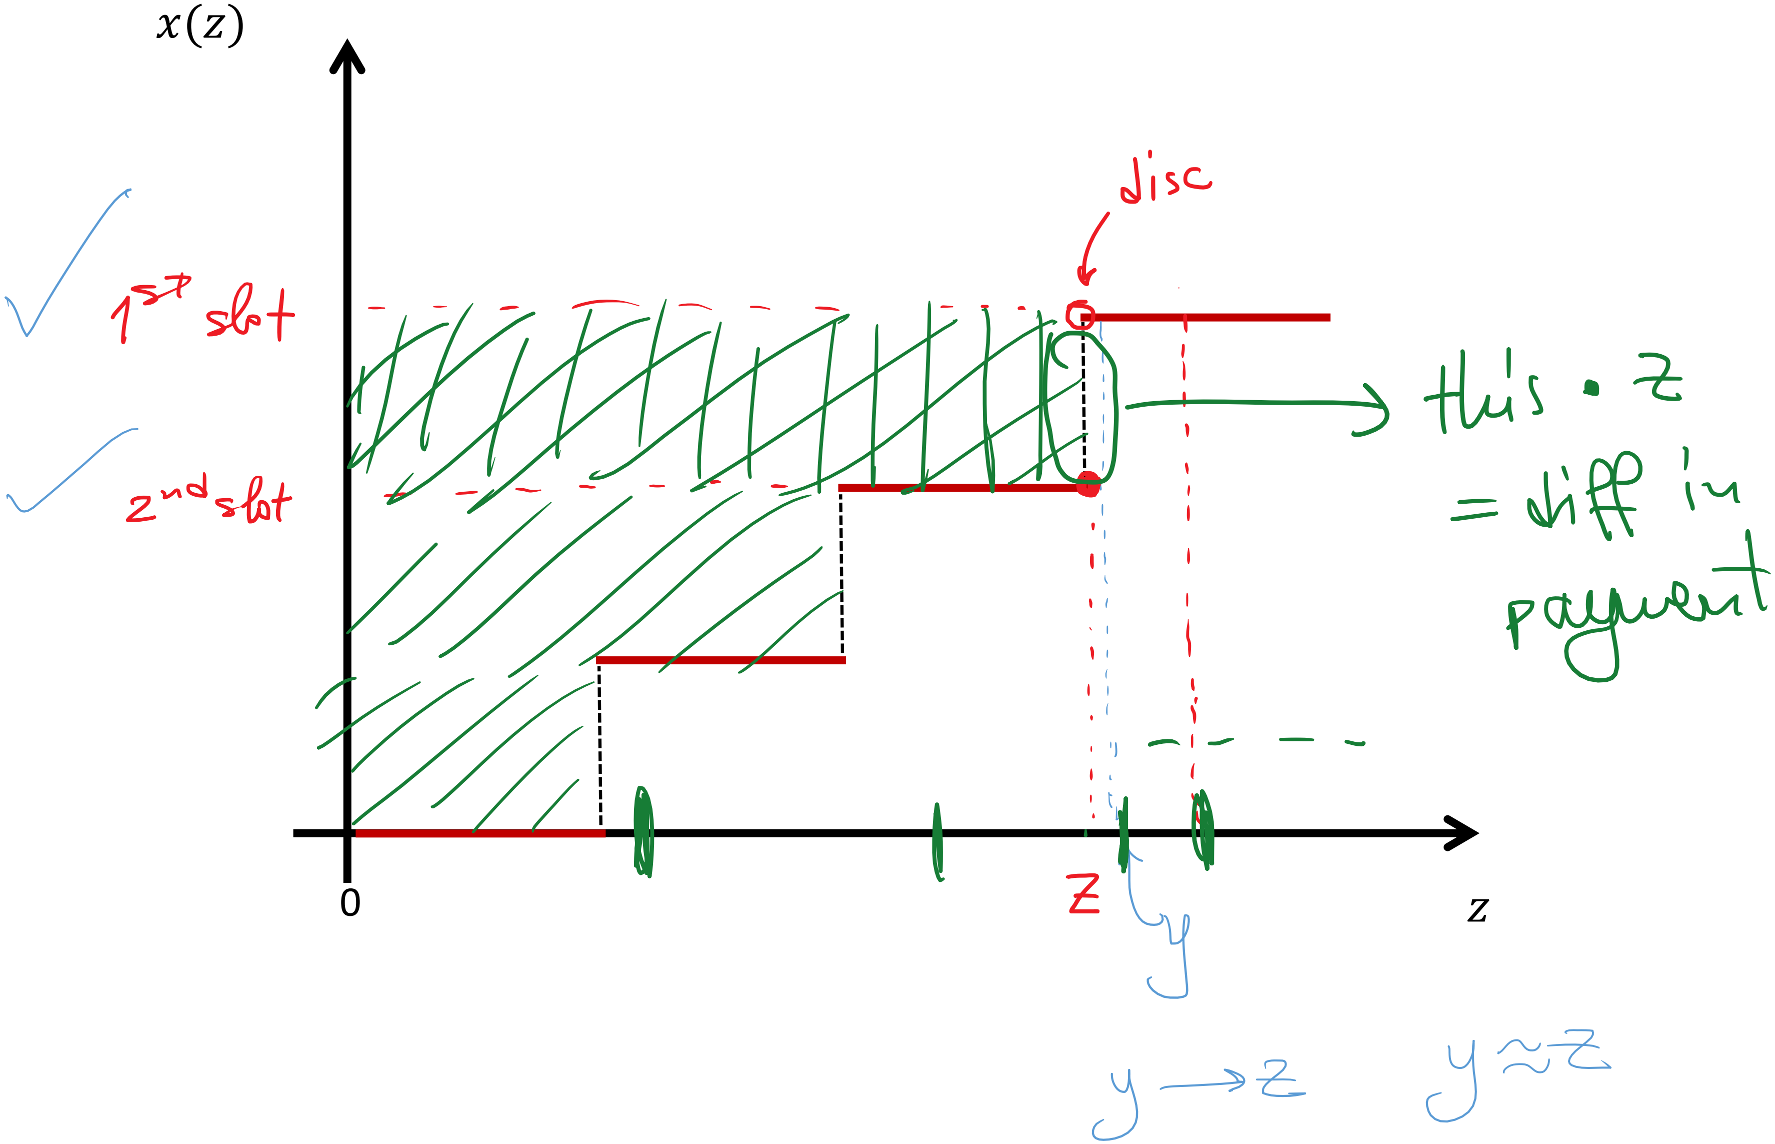

# Proof of Myerson's Lemma

If $z$ is at a "jump":

$$z \cdot \big(\underbrace{x(y) - x(z)}\big) \le p(y) - p(z) \le \overset{z}{\overset{\uparrow}{y}} \cdot \big(\underbrace{x(y) - x(z)}\big)$$

- By fixing $z$ and taking the limit as $y$ tends to $z$, we have that

$$\text{jump of } p \text{ at } z = z \cdot (\text{jump of } x \text{ at } z)$$

$$p(z+\varepsilon) - p(z) = z\big(x(z+\varepsilon) - x(z)\big)$$

# Proof of Myerson's Lemma

$$z \cdot \big(x(y) - x(z)\big) \leq p(y) - p(z) \leq y \cdot \big(x(y) - x(z)\big)$$

- By fixing $z$ and taking the limit as $y$ tends to $z$, we have that

$$\text{jump of } p \text{ at } z = z \cdot (\text{jump of } x \text{ at } z)$$

- Therefore, we can define the payment of the bidder as

$$p(b) = \sum_{y \in [0,b]} y \cdot (\text{jump of } x \text{ at } y)$$

where $y$ enumerates all break points of $x$ in $[0, b]$

# Proof of Myerson's Lemma

- Example:

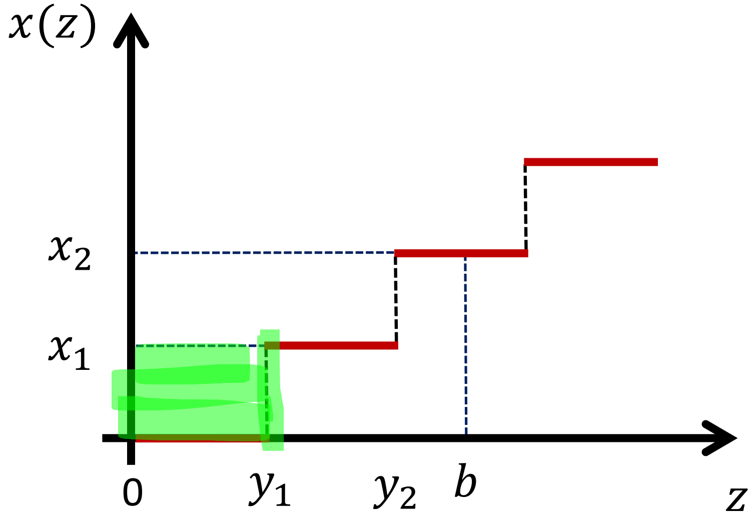


$p(b) = \sum_{y \in [0,b]} y \cdot (\text{jump of } x \text{ at } y)$

# Proof of Myerson's Lemma

- Example:

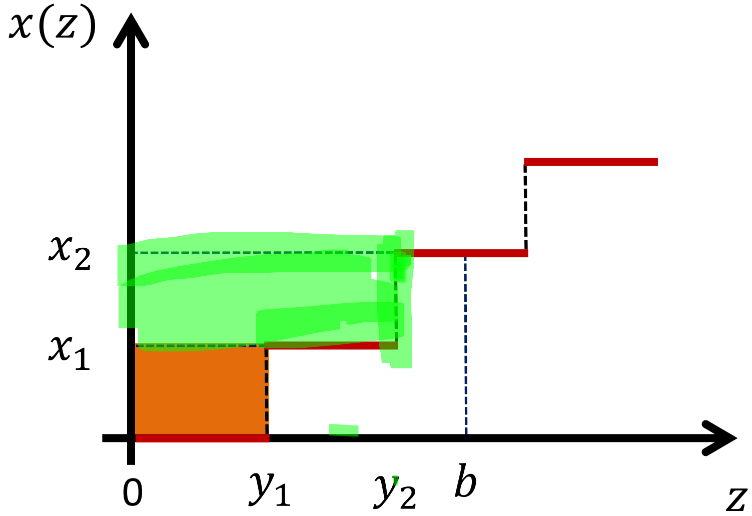


$p(b) = \sum_{y \in [0,b]} y \cdot (\text{jump of } x \text{ at } y) = y_1 \cdot x_1$

# Proof of Myerson's Lemma

- Example:

orange area = payment for any bid between $y_2$ and $y_3$

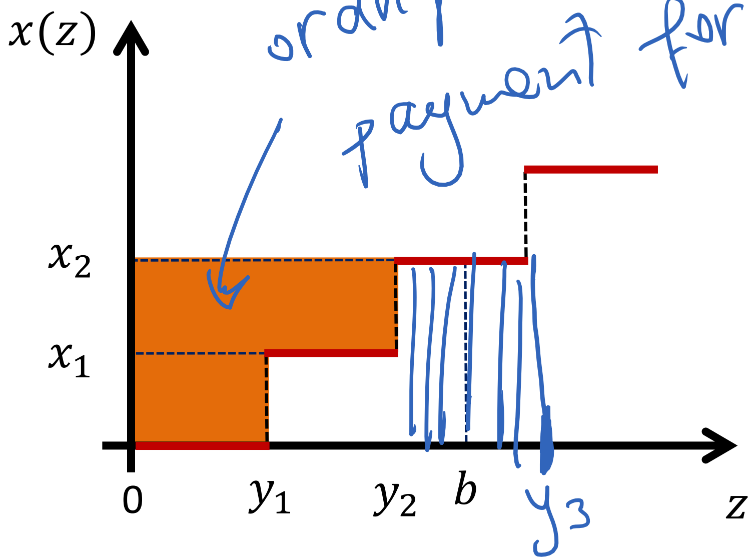


$$p(b) = \sum_{y \in [0,b]} y \cdot (\text{jump of } x \text{ at } y) = y_1 \cdot x_1 + y_2 \cdot (x_2 - x_1)$$

# Proof of Myerson's Lemma

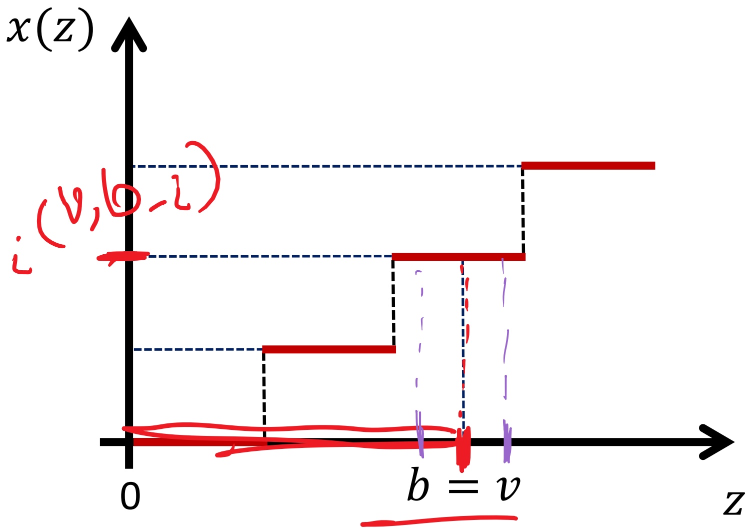


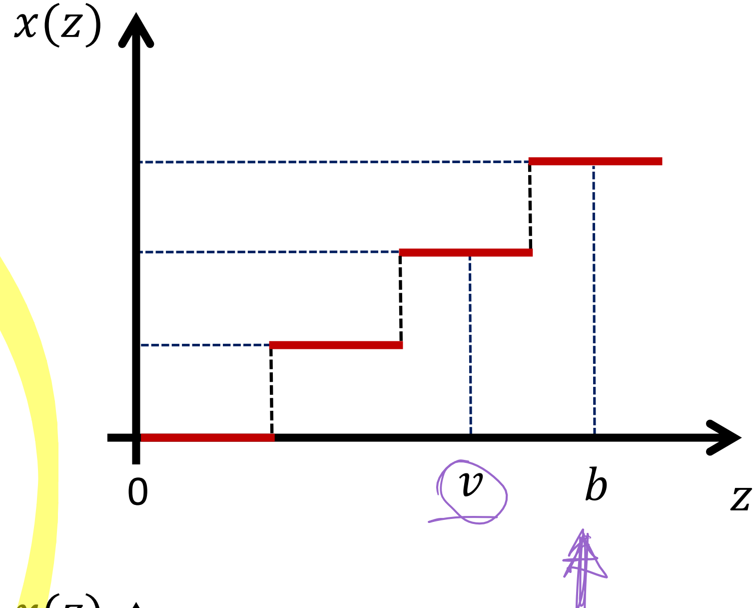


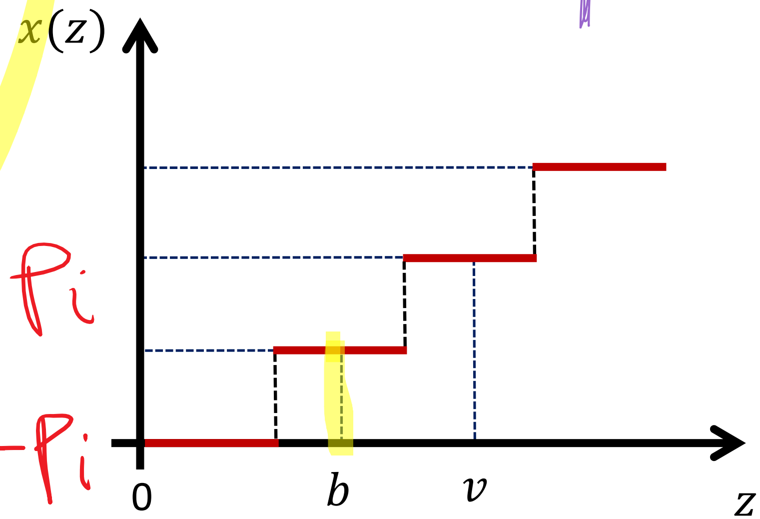


$a_j \cdot v - p_i$

$u_i(v, b_{-i}) = \boxed{x_i(v, b_{-i}) \cdot v} - p_i$

# Proof of Myerson's Lemma

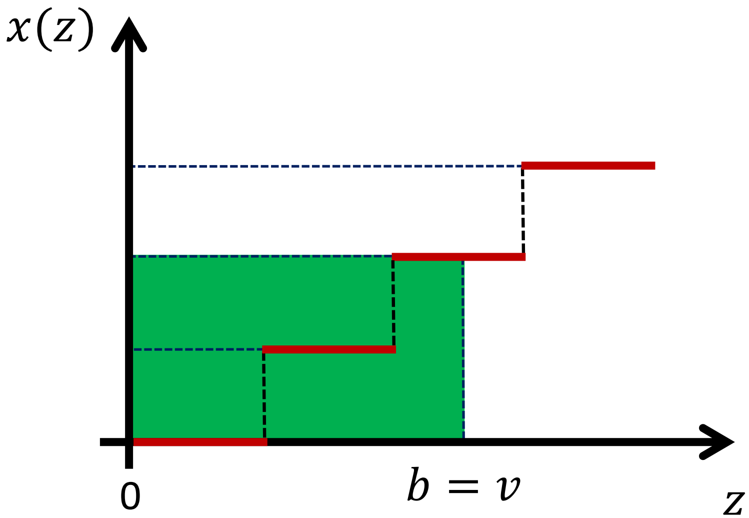


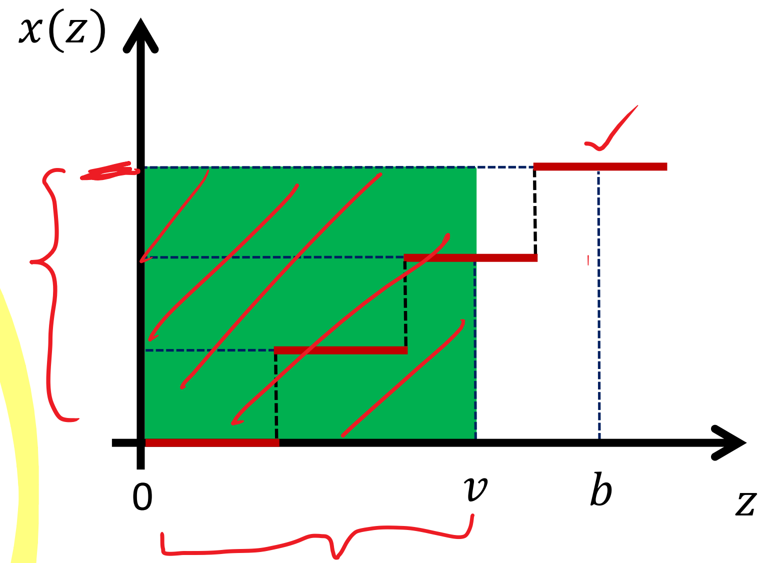


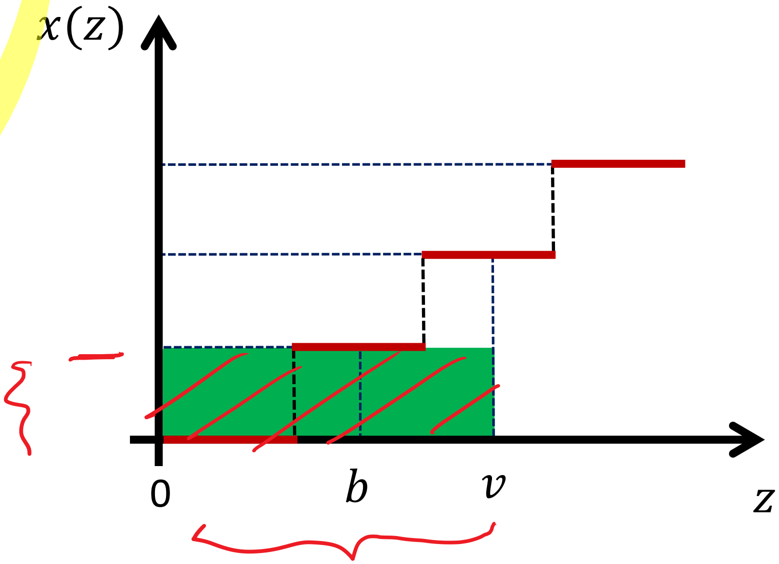

# Proof of Myerson's Lemma

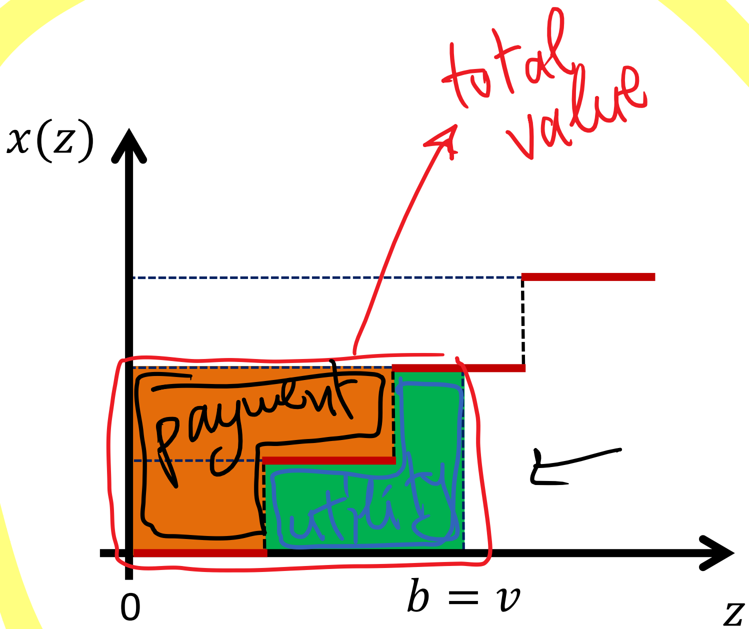


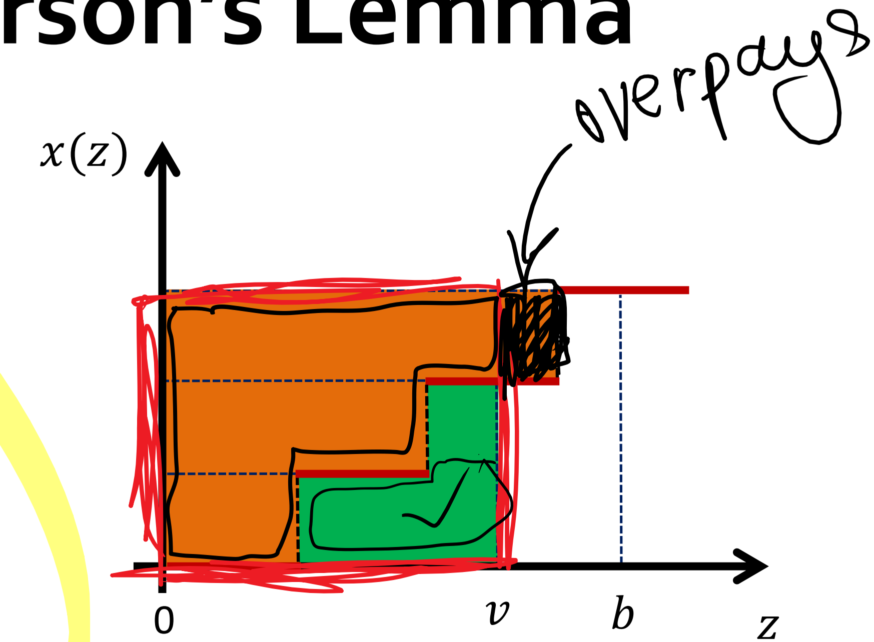


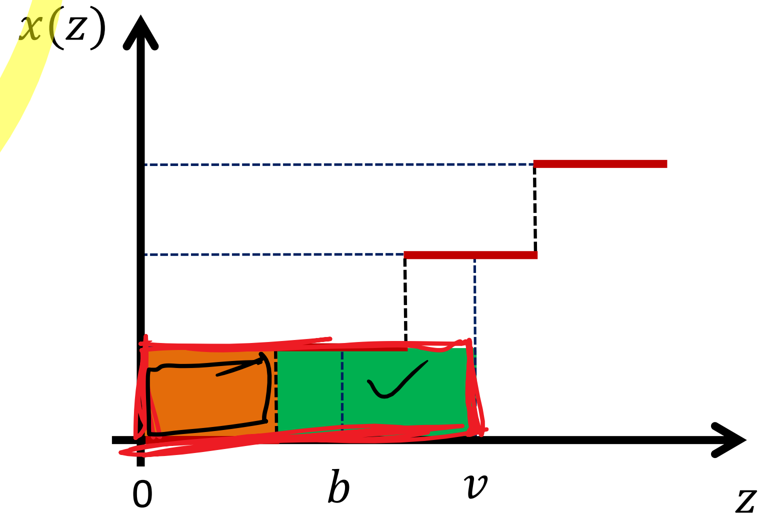

# Sponsored search auctions

$$p(b) = \sum_{y \in [0,b]} y \cdot (\text{jump of } x \text{ at } y)$$

- $y$ enumerates the break points: the bids that are smaller than $b$
  - In other words, $y$ enumerates the slots from worst to best
- jump of $x$ at $y$: the difference in CTR between two consecutive slots
- The total payment of the $i$-th highest bidder is:

$$\boxed{p_i(b_i, \boldsymbol{b}_{-i}) = \sum_{j=i}^{k} b_{j+1}(a_j - a_{j+1})}$$

sum of
orange
areas

# Summary

- **Auctions:** allocation rule + payment rule
- An allocation rule is **implementable** is there exists a payment rule, so that together they define a truthful auction
- An allocation rule is **monotone**, if larger bids give more stuff
- **Single-item auctions:** first-price is not truthful, second-price is truthful and maximizes the social welfare (sells to the bidder with the highest value)
- **Sponsored search auctions:** generalized second-price auction is not truthful
- **Myerson's Lemma:** a characterization of truthful mechanisms in single-parameter environments
- Using Myerson's Lemma we can design a truthful sponsored search auction