Complementi di Algoritmi e Strutture Dati

First-price and second-price auctions

Docente: *Nicolò Cesa-Bianchi* versione del 21 aprile 2026

This material is partially based on the book: Networks, Crowds, and Markets: Reasoning about a Highly Connected World, by David Easley and Jon Kleinberg. Cambridge University Press, 2010.

The underlying assumption we make when modeling auctions is that each bidder i has an intrinsic value $v_i \in [0,1]$ for the item being auctioned. They are willing to purchase the item for a price up to this value, but not for any higher price.

Ascending-bid auctions: These auctions (also called English auctions) are carried out interactively in real time. The seller gradually raises the price, bidders drop out until finally only one bidder remains, and that bidder wins the object at this final price (i.e., the price at which the last non-winning bidder dropped out). These auctions correspond to non-interactive second-price sealed-bid auctions, where bidders submit simultaneous sealed bids to the seller and the highest bidder wins the object paying the value of the second-highest bid.

Descending-bid auctions: This is also an interactive auction format (sometimes called Dutch auction), in which the seller gradually lowers the price from some high initial value until the first moment when some bidder accepts and pays the current price. These auctions correspond to non-interactive first-price sealed-bid auctions, where bidders submit simultaneous sealed bids to the seller and the highest bidder wins the object paying the value of their bid.

Second-price auctions. If there are n bidders with valuations $v_1, \dots, v_n$ and bids $b_1, \dots, b_n$, the payoff function for bidder j in a second-price auction is

$$f_j(v_j, b_1, \dots, b_n) = \mathbb{I}\left\{ b_j > \max_{i \neq j} b_i \right\} \left(v_j - \max_{i \neq j} b_i \right)$$

and similarly for the other bidders.

We say that bid b_j is dominating for bidder j if

$$f_j(v_j, b_1, \dots, b_n) \ge f_j(v_j, b', b_{-j}) \qquad \text{for all } v_1, b_1, \dots, b_n, b'$$

where (b', b_{-j}) denotes $(b_1, \dots, b_{j-1}, b', b_{j+1}, \dots, b_n)$.

Teorema 1 *In a second-price auction, the bid $b_i = v_i$ is dominating for any bidder i.*

DIMOSTRAZIONE. Consider bidder i with valuation v_i and bid b_i. Consider first $b_i > v_i$. If i is winning with $b_i = v_i$, then increasing the bid does not change the payoff. If i is losing with $b_i = v_i$, then then payoff remains zero unless the new bid goes above the highest bid $\max_{j \neq i} b_j > v_i$. In this case the payoff becomes negative. Hence i should not consider $b_i > v_i$. Now consider $b_i < v_i$. If i is losing with $b_i = v_i$, then decreasing the bid does not change the payoff. If i is winning with $b_i = v_i$, then then payoff remains $v_i - \max_{j \neq i} b_j > 0$ unless the new bid goes below the second-highest bid $\max_{j \neq i} b_j$. In this case the payoff becomes zero. Hence i should not consider $b_i < v_i$. This concludes the proof. $\square$

In first-price auctions, $b_i = v_i$ is not dominant anymore: bidding our valuation makes the payoff always zero irrespective of whether we win or lose the auction. To guarantee an expected positive payoff, bidders should always bid less than their valuation. The mapping from valuations to actual bids is called a shading strategy.

A shading strategy $s : [0,1] \to [0,1]$ is a map from valuations to bids. As we assume that bidders are rational, $s(v) \le v$ holds for all $v \in [0,1]$. Hence $s(0) = 0$. We also assume that s is monotone: $v' > v$ implies $s(v') > s(v)$. In other words, we always increase a bid if the valuation increases.

First-price auctions. If there are two bidders with valuations v_1, v_2 and shading strategies s_1, s_2, the payoff function for bidder 1 in a first-price auction is

$$g_1(v_1, v_2, s_1, s_2) = \mathbb{I}\{s_1(v_1) > s_2(v_2)\}\big(v_1 - s_1(v_1)\big)$$

and similarly for bidder 2.

Assuming the valuations v_1, v_2 are realizations of two random variables V_1, V_2, an **equilibrium** for the two bidders is a pair (s_1, s_2) of strategies such that

$$\begin{aligned} &\mathbb{E}\Big[g_1(v_1, V_2, s_1, s_2) - g_1(v_1, V_2, s', s_2)\Big] \ge 0 \\ &\mathbb{E}\Big[g_2(V_1, v_2, s_1, s_2) - g_2(V_1, v_2, s_1, s')\Big] \ge 0 \end{aligned} \qquad \text{for all } s', v_1, v_2$$

Being an equilibrium is a weaker property for a strategy than being dominant: a strategy is dominant irrespective of what the other players do, whereas a strategy is an equilibrium only if the others also play an equilibrium.

Teorema 2 *If the valuations V_1, V_2 for the two bidders are independently drawn from the uniform distribution over the $[0,1]$ interval, then (s, s) with $s : v \mapsto v/2$ is an equilibrium for the bidders in a first-price auction.*

Dimostrazione. Given the realized valuation v_1, the expected payoff for player 1 is

$$\begin{aligned} \mathbb{E}\big[g_1(v_1, V_2, s, s)\big] &= \mathbb{P}\big(s(v_1) > s(V_2)\big)\big(v_1 - s(v_1)\big) \\ &= \mathbb{P}\big(v_1 > V_2\big)\big(v_1 - s(v_1)\big) \\ &= v_1\big(v_1 - s(v_1)\big) \end{aligned}$$

where we used the monotonicity assumption: $s(v) > s(v')$ if and only if $v > v'$, and our assumption on the uniform distribution for V_1, V_2. The equilibrium condition for player 1 then states that

$$v_1\big(v_1 - s(v_1)\big) \ge v_1\big(v_1 - s'(v_1)\big)$$

Since bidder 2 is never going to bid more than $s(1)$, we can assume that s' satisfies the additional condition $s'(v) \in [0, s(1)]$ for all $v \in [0,1]$. Indeed, bidding higher than $s(1)$ would only reduce the payoff of bidder 1, without increasing the probability of winning. This implies that we can find $v \in [0,1]$ such that $s'(v_1) = s(v)$. Hence, the equilibrium condition becomes

$$\mathbb{P}\big(s(v_1) > s(V_2)\big)\big(v_1 - s(v_1)\big) = v_1\big(v_1 - s(v_1)\big) \ge v\big(v_1 - s(v)\big) = \mathbb{P}\big(s'(v_1) > s(V_2)\big)\big(v_1 - s'(v_1)\big)$$

where we used that $v = \mathbb{P}\big(v > V_2\big) = \mathbb{P}\big(s(v) > s(V_2)\big) = \mathbb{P}\big(s'(v_1) > s(V_2)\big)$. Substituting $s(v) = v/2$ we get

$$\frac{v_1^2}{2} \geq vv_1 - \frac{v^2}{2}$$

Multiplying by 2 both sides and rearranging we obtain $v_1^2 + v^2 - 2vv_1 \geq 0$ which is always true. This concludes the proof. □

This analysis generalizes to the case of $n \geq 2$ bidders. In this case, the equilibrium strategy takes the form

$$s : v \mapsto \left(\frac{n-1}{n}\right) v$$

We now compare the revenue of the seller in the first-price and second-price auctions. Since the seller's revenue is based on the values of the highest and second-highest bids, which in turn depend on the highest and second-highest valuations, we need to know the expectations of these quantities. Computing these expectations is complicated, but the form of the answer is very simple.

Lemma 3 *Suppose n numbers are drawn independently from the uniform distribution on the interval $[0,1]$ and then sorted from smallest to largest. The expected value of the number in the k-th position on this list is $\frac{k}{n+1}$.*

Now, if the seller runs a second-price auction, and the bidders follow their dominant strategies and bid truthfully, the seller's expected revenue will be the expectation of the second-highest value. Since this will be the value in position $n-1$ in the sorted order of the n random values from smallest to largest, the expected value is $\frac{n-1}{n+1}$. On the other hand, if the seller runs a first-price auction, then in equilibrium we expect the winning bidder to submit a bid $b^* = \frac{n-1}{n} v^*$, where v^* is the corresponding valuation. Note that v^* is the realization of V^* with expected value $\frac{n}{n+1}$, since it is the largest of n numbers drawn independently from the unit interval, and so the seller's expected revenue is $\frac{n-1}{n} \frac{n}{n+1} = \frac{n-1}{n+1}$. So the two auctions provide exactly the same expected revenue to the seller.

Revenue Equivalence. As far as seller revenue is concerned, this calculation is in a sense the tip of the iceberg: it is a reflection of a much broader and deeper principle known in the auction literature as revenue equivalence. Roughly speaking, revenue equivalence asserts that a seller's revenue will be the same across a broad class of auctions and arbitrary independent distributions of bidder values, when bidders follow equilibrium strategies.

Reserve Prices. In our discussion of how a seller should choose an auction format, we have implicitly assumed that the seller must sell the object. Let's briefly consider how the seller's expected revenue changes if he has the option of holding onto the item and choosing not to sell it. To be able to reason about the seller's payoff in the event that this happens, let's assume that the seller values the item at $u \geq 0$, which is thus the payoff he gets from keeping the item rather than selling it. It's clear that if $u > 0$, then the seller should not use a simple first-price or second-price auction. In either case, the winning bid might be less than u, and the seller would not want to sell the object.

Instead, it is better for the seller to announce a reserve price of r before running the auction. With a reserve price, the item is sold to the highest bidder if the highest bid is above r; otherwise, the item is not sold. In a first-price auction with a reserve price, the winning bidder (if there is one) still pays her bid. In a second-price auction with a reserve price, the winning bidder (if there is one) pays the maximum of the second-place bid and the reserve price r.

It is in fact useful for the seller to declare a reserve price even if his value for the item is $u = 0$. Let's consider how to reason about the optimal value for the reserve price in the case of a second-price auction. First, it is not hard to go back over the argument that truthful bidding is a dominant strategy in second-price auctions and check that it still holds in the presence of a reserve price. Essentially, it is as if the seller were another "simulated" bidder who always bids r; and since truthful bidding is optimal regardless of how other bidders behave, the presence of this additional simulated bidder has no effect on how any of the real bidders should behave.

Now, what value should the seller choose for the reserve price? If the item is worth u to the seller, then clearly he should set $r \geq u$. But in fact the reserve price that maximizes the seller's expected revenue is strictly greater than u. Indeed, taking $u = 0$ for simplicity, the seller's revenue is

$$B_2 \times \mathbb{I}\{B_2 \geq r\} + r \times \mathbb{I}\{B_2 < r \leq B_1\} + 0 \times \mathbb{I}\{B_1 < r\}$$

where we use B_1 and B_2 to denote the largest and second-largest bid. This shows that a good choice of r is as close as possible to (but not larger than) the highest bid B_1.