Auction Design Goals

CS 1951k/2951z

2020-02-5

We formally define a single-parameter sealed-bid auction. This auction format is defined by two rules, an allocation rule and a payment rule. Using this formal framework, we state our auction design goals.

1 *Auction Model*

We study an auction with a set of $n \geq 2$ participating agents (bidders). We assume the environment is **single-parameter**, meaning the agents can be characterized by but one parameter. First- and second-price auctions for but one good can easily be modelled this way. But we also often model auctions for multiple goods in this way, such as advertising slots on a web page, because there are tools available for the design of single-parameter auctions that are not straightforward to extend to the multi-parameter case.

We adopt the IPV model, so that each agent $i \in [n]$ has valuation v_i drawn from continuous distribution F_i, with support $T_i = [\underline{v}_i, \overline{v}_i]$, for some $\underline{v}_i, \overline{v}_i \in \mathbb{R}_+$. An agent's valuation is their private information, which, by assumption, can be described by a single parameter.

Each agent i employs a strategy $s_i : T_i \to B_i$, which, as usual, is a function from their space of possible values to their space of possible actions, which, in an auction, are bids. It is assumed that bids are submitted to the auctioneer in sealed envelopes, so that no bidder knows what any other has bid. Given a profile of values $\mathbf{v} = (v_1, \ldots, v_n)$, the vector of bids submitted to the auctioneer is denoted $\mathbf{b} = (b_1, \ldots, b_n) \equiv (s_1(v_1), \ldots, s_n(v_n))$.

Given a bid profile, the auctioneer computes allocations and payments according to the auctions' rules. The **allocation rule** describes how the winner(s) of the auction is determined; the **payment rule** describes what the participants pay. We write $\mathbf{x}(\mathbf{b}) = (x_1(\mathbf{b}), \ldots, x_n(\mathbf{b}))$ to denote the allocation, and $\mathbf{p}(\mathbf{b}) = (p_1(\mathbf{b}), \ldots, p_n(\mathbf{b}))$ to denote payments. In particular, $x_i(\mathbf{b})$ denotes bidder i's allocation,[1] and $p_i(\mathbf{b}) \in \mathbb{R}_+$ denotes bidder i's payment. We call the pair $(\mathbf{x}(\mathbf{b}), \mathbf{p}(\mathbf{b}))$ the **outcome** of the auction.

Finally, we assume a **quasi-linear** model of agent utility,[2] where $u_i(\mathbf{b}; v_i) = v_i x_i(\mathbf{b}) - p_i(\mathbf{b})$. We almost always abbreviate utility as follows: $u_i(\mathbf{b}) \equiv u_i(\mathbf{b}; v_i)$. Further, when the bidding profile is clear from context, we often write: $u_i = v_i x_i - p_i$.

[1] When goods are divisible, bidder i's allocation may be fractional. Alternatively, $x_i(\mathbf{b})$ may represent the probability with which bidder i is allocated an indivisible good.

[2] A **quasi-linear** function is linear in one variable, called the numeraire. A **numeraire** is a basic standard for measuring value (e.g., money).

First- and Second-Price Auctions In the two sealed-bid auctions of primary interest, namely first- and second-price auctions, the space of possible bids contains the space of possible values: i.e., $T_i \subseteq B_i$, for all bidders $i \in N$. Upon receiving bids in sealed envelopes, the auctioneer selects as the winner a bidder with the highest bid, $i^* \in \arg\max_{i \in N} b_i$, allocating to this winner with probability $x_{i^*}(b_{i^*}, b_{-i^*}) = 1$, and breaking ties, as necessary.[3]

Let $b_{(k)}$ denotes the kth-order statistic, meaning the kth-largest draw among n samples. In particular, $b_{(n)}$ and $b_{(n-1)}$ denote the highest and second-highest bids, respectively.

In the first-price auction, the winner of the auction, i^*, is charged their bid (i.e., the highest bid), $p_{i^*}(b_{i^*}, \mathbf{b}_{-i^*}) = b_{(n)}$, and all other bidders $i \neq i^* \in N$ are charged $p_i(b_i, \mathbf{b}_{-i}) = 0$.

In the second-price auction, the winner of the auction, i^*, is charged the second-highest bid, $p_{i^*}(b_{i^*}, \mathbf{b}_{-i^*}) = b_{(n-1)}$, and all other bidders $i \neq i^* \in N$ are charged $p_i(b_i, \mathbf{b}_{-i}) = 0$.

[3] Ex-ante, each bidder in $\arg\max_{i \in N} b_i$ is allocated with probability $\frac{1}{|\arg\max_{i \in N} b_i|}$.

2 Design Goals

There are (at least) three desireable properties of an auction:

Incentive Guarantees The first goal pertains to incentives. In order to compare competing auction designs, we must be able to predict the outcome of an auction, which in turn means, predicting the strategic behavior of the agents in the auction. If we make our auctions simple for agents to reason about, then we may have a better chance of predicting what agents will do. Bidding truthfully is always an option, so one way to make things simple is to ensure that bidding truthfully is an equilibrium strategy. When this property holds, an auction is said to be **incentive compatibile**.

Another natural incentive criterion of auctions (actually mechanisms, more generally) is called **individual rationality**. This property ensures that no participant can be made worse off by participating in the auction. In other words, all bidders' utilities are guaranteed to be non-negative. Like incentive compatibility, individual rationality can hold either ex-ante or ex-post. In the former case, it is sometimes called **Bayesian individual rationality**.

Economic Performance Guarantees The second goal is that the auction format achieve some objective. A popular choice is welfare maximization.[4] **Welfare** is defined as the total expected utility of all participants, including the auctioneer, which assuming incentive

[4] When welfare is maximized, an economy is said to be **efficient**.

compatibility so that $\mathbf{b} = \mathbf{v}$, can be written as follows:

$$\mathbb{E}_{\mathbf{v}}\left[\sum_{i \in N} u_i(\mathbf{v}) + \sum_{i \in N} p_i(\mathbf{v})\right] = \mathbb{E}_{\mathbf{v}}\left[\sum_{i \in N} (v_i x_i(\mathbf{v}) - p_i(\mathbf{v})) + \sum_{i \in N} p_i(\mathbf{v})\right]$$
$$= \mathbb{E}_{\mathbf{v}}\left[\sum_{i \in N} v_i x_i(\mathbf{v})\right]$$

Revenue maximization is a popular alternative (Think Amazon, Facebook, Google, etc.). **Revenue** is defined as the total expected payments, only:

$$\mathbb{E}_{\mathbf{v}}\left[\sum_{i \in N} p_i(\mathbf{v})\right]$$

Computational Performance Guarantees In addition to aiming for economic efficiency, we also aim to design auctions that are computationally efficient, meaning they run in polynomial time and space.

The allocation algorithm for first- and second-price auctions—namely, allocate to the highest bidders—satisfies this requirement: it is $O(n)$ in time and space in the worst case. (There may be an n-way tie, and a tie-breaking rule may randomly select among all the tied bidders.) Similarly, calculating payments in these auctions is $O(n)$ (or $O(1)$, if you store the results of solving allocation problem).

3 *A Solution via Mathematical Programming*

A k**-good** auction is one in which k copies of a homogeneous good are on offer. Thus, it is a natural generalization of the usual first- and second-price auction setting.

Solving for the welfare- (or revenue-) maximizing k-good auction can be viewed as a constrained optimization problem. Let $F = \prod_{i \in [n]} F_i$. Total expected welfare is then:

$$\mathop{\mathbb{E}}_{\mathbf{v} \sim F}\left[\sum_{i \in N} v_i x(v_i, \mathbf{v}_{-i})\right],$$

while total expected revenue is:

$$\mathop{\mathbb{E}}_{\mathbf{v} \sim F}\left[\sum_{i \in N} p_i(v_i, \mathbf{v}_{-i})\right].$$

In both cases, the decision variables are the allocation rule $\mathbf{x}$ and payment rule $\mathbf{p}$.

The constraints are as follows:[5]

1. **Incentive compatibility**. Truthful bidding maximizes utility.[6]

$$u_i(v_i, \mathbf{v}_{-i}) \geq u_i(t_i, \mathbf{v}_{-i}), \quad \forall i \in N, \forall v_i, t_i \in T_i, \forall \mathbf{v}_{-i} \in T_{-i};$$

[5] Because of incentive compatibility, we *assume* everyone bids their true value: i.e., $b_i = v_i, \forall i \in N$.

[6] Here, we are assuming $B_i = T_i, \forall i \in N$.

2. **Individual rationality**. Utility is non-negative (assuming truthful bidding, which is ensured by the IC constraints):

$$u_i(v_i, \mathbf{v}_{-i}) \geq 0, \quad \forall i \in N, \forall \mathbf{v} \in T;$$

3. **Allocation constraints**. The allocation variables vary with the auction set up. In a k-good auction, they are $0/1$ variables in principle, but as there could be ties, we represent them as probabilities.

   The probability of winning must be in $[0, 1]$:

$$0 \leq x_i(v_i, \mathbf{v}_{-i}) \leq 1, \quad \forall i \in N, \forall \mathbf{v} \in T.$$

4. **Ex-post feasibility**. Goods are not overallocated:

$$\sum_{i \in N} x_i(v_i, \mathbf{v}_{-i}) \leq 1, \quad \forall \mathbf{v} \in T;$$

One of these objective functions together with these constraints comprise a mathematical program that can be used to solve for an optimal k-good auction. The good news is, the objective function and the constraints are linear. The bad news is, there are an exponential number of constraints (assuming we discretize the value space). But do not despair. Roger Myerson won a Nobel prize in part for his elegant solution to this auction design problem. Stay tuned!