Auction Design Goals

CSCI 1440/2440

2024-02-14

We formally define sealed-bid auctions in the framework of mechanism design. This auction format is defined by two rules, an allocation rule and a payment rule. We then state three auction design goals.

1 *Auction Model*

We study auctions, which are mechanisms, with $n \geq 2$ participating agents (bidders). As is usual in a mechanism, a strategy is a function from a type space to a space of possible reports, the latter of 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 bids. Given a profile of types $\mathbf{t} = (t_1, \ldots, t_n)$, a bid profile is denoted by $\mathbf{b} = (b_1, \ldots, b_n) \equiv (s_1(t_1), \ldots, s_n(t_n))$, assuming each agent $i \in [n]$ employs a strategy $s_i : T_i \to B_i$.

Given a bid profile $\mathbf{b}$, the auctioneer computes the auction outcome, which comprises allocations and payments according to two rules. The **allocation rule** describes how the winner(s) of the auction is (are) determined, while the **payment rule** describes what the bidders pay. We write $\mathbf{x}(\mathbf{b}) = (x_1(\mathbf{b}), \ldots, x_n(\mathbf{b})) \in \Delta^n$ to denote the allocation,[1] and $\mathbf{p}(\mathbf{b}) = (p_1(\mathbf{x}(\mathbf{b})), \ldots, p_n(\mathbf{x}(\mathbf{b}))) \in \mathbb{R}_+^n$ to denote the payments; here, $x_i(\mathbf{b}) \in [0,1]$ denotes bidder i's (expected)[2] allocation, and $p_i(\mathbf{x}(\mathbf{b})) \in \mathbb{R}_+$ denotes bidder i's payment.[3] The pair $(\mathbf{x}(\mathbf{b}), \mathbf{p}(\mathbf{b})) \in \Omega$ represents the (expected) outcome of the auction.

Example 1.1. In first- and second-price auctions, it is typical to assume the space of possible bids contains the space of possible types: 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.[4] This is the allocation rule.

For $k \in [n]$, let $b_{(k)} \doteq b_{(k;n)}$ denotes the kth-order statistic, meaning the kth-largest draw among n samples. In particular, $b_{(1)}$ and $b_{(2)}$ denote the highest and second-highest bids, respectively.

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

In the second-price auction, the winner of the auction, i^*, is charged the second-highest bid, $p_{i^*}(b_{i^*}, \mathbf{b}_{-i^*}) = b_{(2)}$, and all other bidders $i \neq i^* \in [n]$ are charged $p_i(b_i, \mathbf{b}_{-i}) = 0$. (Payment rule.)

[1] Δ^n denotes the n-dimensional simplex, meaning all vectors in $[0,1]^n$ whose entries sum to 1. Defining allocations as values in the simplex allows us to naturally incorporate randomization into auctions.

[2] As $x_i(\mathbf{b}) \in (0,1)$ represents the probability bidder i is allocated the good, it is also bidder i's expected allocation, as this probability is the expected value of the corresponding indicator value.

[3] Bidder i's payments depend on her allocation; nevertheless, we abbreviate $\mathbf{p}(\mathbf{x}(\mathbf{b}))$ by $\mathbf{p}(\mathbf{b})$, as the allocation can be reconstructed from the bids.

[4] Ex-ante, it is typical to assume that each highest bidder is allocated with probability $1/|\arg\max_{i \in [n]} b_i|$.

2 *Independent Private Values*

In auctions, an agent's type is her **valuation**, a function from outcomes and (in general all) bidders' types to real values: i.e., $v_i : \Omega \times T \to \mathbb{R}$. We adopt the **independent private values (IPV)** model of valuations. The IPV model makes two simplifying assumptions:

1. Bidders' valuations are *private*, meaning they are not influenced by other bidders' types (i.e., their information).[5]

2. Any one bidder's valuation is *independent* of any other bidder's.

   The first assumption allows us to simplify valuations as follows:

$$v_i(\omega; t_i, \mathbf{t}_{-i}) = v_i(\omega; t_i), \quad \forall i \in [n], \forall \mathbf{t}_{-i} \in T_{-i}.$$

[5] When bidders' valuations are not private, they are called *common*. An example of a good with a common value is a tract of land beneath which there may or may not be oil. Agents may possess private information, from their own probing, about the value of the land, but their values depend on others' information (i.e., types) as well.

Moreover, since each bidder's valuation depends only her own type, we can fold a bidder's type information into her valuation, $v_i(\omega; t_i, \mathbf{t}_{-i}) = v_i(\omega)$, for all $i \in [n]$, and then draw valuations, rather than types, from distributions.

The second assumption decorrelates bidders' valuations (formerly, types) from one another: when a bidder learns her own valuation/type, she does not learn anything about the other bidders' valuations/types. The mathematical implication of this assumption is that we can draw bidder i's valuation v_i from a distribution F_i, which is independent of the distributions F_j of all other bidders $j \neq i \in [n]$: i.e., $v_i \sim F_i$, for all $i \in [n]$.

The IPV model is applicable in auction settings where goods do not have any resale value, or where reselling is infeasible. If goods cannot be resold, then how other bidders value them is irrelevant.

3 *Quasilinear Utilities*

Recall that an agent i's utility in a mechanism is a function of both the outcome ω and the various agents' types $\mathbf{t}$: i.e., $u_i : \Omega \times T \to \mathbb{R}$. In the case of auctions, assuming IPV, these types are valuations, which ascribe values to outcomes: i.e., $v_i : \Omega \to \mathbb{R}$.

The utility model that we adopt for auctions is called **quasilinear**. A quasilinear utility function is a potentially nonlinear function that is nonetheless linear in the quantity of a select good, called the **numeraire**. A numeraire is a standard for measuring value, which in the modern economy, translates to money.

In auctions, the outcomes are allocations and payments. By assuming quasilinear utility functions, we separate these two components of an auction outcome into a potentially nonlinear value of the outcome less the corresponding payment: $u_i(\omega) = v_i(\omega) - p_i(\omega)$. For this addition (technically, subtraction) to make sense, valuations must be expressed in terms of money. Indeed, quasilinear utility gains and losses are measured in terms of money.

This utility model, in which a valuation v_i associates a value with all outcomes, is very general. It can be used to model auctions for multiple goods, where outcomes are assignments of **bundles**, or subsets of goods, to bidders. In such auctions, the bidders' valuations can be very complex, as they ascribe value to a combinatorial number of outcomes. Correspondingly, they require complex (i.e., multiparameter) representations, such as a table of values, one per outcome, potentially exponentially-many of them, or a neural network.

A simpler representation—which is certainly relevant to single-good auctions, and is also sometimes used in auctions for multiple goods—is one in which a bidder i's valuation is a single number, that roughly corresponds to her value for winning some part of the auction. In a so-called **single-parameter** auction, we typically express utility as $u_i(\mathbf{b}; v_i) = v_i - p_i(\mathbf{b})$. Overloading notation, expected utility is given by $u_i(\mathbf{b}; v_i) = x_i(\mathbf{b})(v_i - p_i(\mathbf{b}))$, or simply $u_i = x_i(v_i - p_i)$, when the bid profile is clear from context.

4 *Design Goals*

There are three primary goals that guide auction design: incentive, economic performance, and computational performance guarantees. For the most part, these goals are relevant in all mechanisms, not only auctions; however, specific economic performance goals, such as maximizing revenue, may not be relevant in all mechanisms.

Agent Strategy Guarantees: Ensuring Simplicity The first goal pertains to agent strategies. 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 being able to predict what agents will do. Bidding truthfully is always an option, so one way to make things simple[6] is to ensure that bidding truthfully is an equilibrium strategy. When this property holds, an auction is said to be **incentive compatible**:

[6] Arguably, as simple as possible.

Definition 4.1. Dominant-strategy incentive compatibility (DSIC) means truthful bidding maximizes utility.[7]

[7] We assume $B_i = T_i, \forall i \in [n]$.

$$u_i(v_i, \mathbf{v}_{-i}) \geq u_i(t_i, \mathbf{v}_{-i}), \quad \forall i \in [n], \forall v_i, t_i \in T_i, \forall \mathbf{v}_{-i} \in T_{-i}.$$

Bayesian incentive compatibility (BIC) means truthful bidding maximizes *expected* utility.

$$\mathop{\mathbb{E}}_{\mathbf{v}_{-i} \sim F_{\mathbf{v}_{-i}|v_i}} \left[u_i(v_i, \mathbf{v}_{-i}) \right] \geq \mathop{\mathbb{E}}_{\mathbf{v}_{-i} \sim F_{\mathbf{v}_{-i}|v_i}} \left[u_i(t_i, \mathbf{v}_{-i}) \right], \quad \forall i \in [n], \forall v_i, t_i \in T_i.$$

Another natural incentive criterion is called **individual rationality**. This also property is sometimes called **voluntary participation**, as it ensures that no participant can be made worse off by participating in the mechanism. In other words, all agents' 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 called **Bayesian individual rationality**:

Definition 4.2. Individual rationality (IR) means utility is non-negative (assuming truthful bidding, which is ensured by the IC constraints):

$$u_i(v_i, \mathbf{v}_{-i}) \geq 0, \quad \forall i \in [n], \forall \mathbf{v} \in T.$$

Bayesian individual rationality (BIR) means utility is non-negative (assuming truthful bidding, which is ensured by the IC constraints):

$$\mathop{\mathbb{E}}_{\mathbf{v}_{-i} \sim F_{\mathbf{v}_{-i}|v_i}} \left[u_i(v_i, \mathbf{v}_{-i}) \right] \geq 0, \quad \forall i \in [n], \forall v_i \in T_i.$$

Economic Performance Guarantees The second goal is that the auction format implement desirable economic properties. One such property is **fairness**. While a precise definition of fairness continues to elude us, it is worth noting that simplicity ensures a kind of fairness, as it levels playing field, because more sophisticated bidders—those who have more time or money to devote to strategizing—cannot benefit as much from doing so as they might otherwise.

A second desirable economic property is **efficiency**, which usually means optimizing societal welfare. There are numerous measures of societal welfare, ranging from **egalitarian**[8] to **utilitarian**.[9]

Utilitarian welfare is defined as the total expected utility of all participants, including the auctioneer, which assuming incentive compatibility (i.e., $\mathbf{b} = \mathbf{v}$) can be written as follows:

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

In the context of auctions specifically, not mechanisms in general, revenue maximization is also commonly adopted as an objective.[10]

[8] John Locke, John Rawls

[9] Jeremy Bentham, John Stuart Mill

[10] 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 striving for economic efficiency, we also aim to design *computationally* efficient auctions, 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.[11] Calculating payments based on the allocation is then $O(1)$.

[11] There may be an n-way tie, and a tie-breaking rule may randomly select among all the tied bidders.

More generally, auctions for multiple goods that require a search for an optimal assignment of goods to bidders among combinatorially-many possibilities, cannot be run in polynomial time or space.

A 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. We optimize over direct mechanisms only.

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:

1. **Dominant strategy incentive compatibility (DSIC)**. Truthful bidding maximizes utility.[12]

   $$u_i(v_i, \mathbf{v}_{-i}) \geq u_i(t_i, \mathbf{v}_{-i}), \quad \forall i \in [n], \forall v_i, t_i \in T_i, \forall \mathbf{v}_{-i} \in T_{-i}.$$

   (Or **Bayesian incentive compatibility (BIC)**. Truthful bidding maximizes *expected* utility.)

   $$\mathop{\mathbb{E}}_{\mathbf{t}_{-i}\sim F_{\mathbf{t}_{-i}|t_i}}[u_i(v_i, \mathbf{v}_{-i})] \geq \mathop{\mathbb{E}}_{\mathbf{t}_{-i}\sim F_{\mathbf{t}_{-i}|t_i}}[u_i(t_i, \mathbf{v}_{-i})], \quad \forall i \in [n], \forall v_i, t_i \in T_i.$$

[12] We also assume $B_i = T_i, \forall i \in [n]$.

2. **Individual rationality (IR)**. Utility is non-negative (assuming truthful bidding, which is ensured by the IC constraints):[13]

$$u_i(v_i, \mathbf{v}_{-i}) \geq 0, \quad \forall i \in [n], \forall \mathbf{v} \in T.$$

(Or **Bayesian individual rationality (BIR)**. Utility is non-negative (assuming truthful bidding, which again is ensured by the IC constraints):)

$$\mathop{\mathbb{E}}_{\mathbf{t}_{-i} \sim F_{\mathbf{t}_{-i}|t_i}} [u_i(v_i, \mathbf{v}_{-i})] \geq 0, \quad \forall i \in [n], \forall v_i \in T_i.$$

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,[14] assuming we discretize the value space, or an infinite number of them, if not! But do not despair. Roger Myerson won a Nobel prize in part for his elegant solution to this auction design problem.[15] Stay tuned!

[13] Because our goal is to ensure incentive compatibility, we *assume* everyone bids their true value: i.e., $b_i = v_i, \forall i \in [n]$.

[14] exponential in the number of bidders

[15] In single-parameter environments, only, unsurprisingly.