Chapter 6. Introduction to auctions

Joan Llull

Structural Empirical Methods for Labor Economics
(and Industrial Organization)

IDEA PhD Program

I. Introduction

Auctions are often used as mechanisms for allocation of resources to bidders in an as efficient as possible way. There are many interesting real-world applications that generate interesting data and a great deal of opportunities for empirical investigation, raising interesting issues. As a result, there is an important and growing literature in empirical micro that aims at structurally estimating the parameters of the bidders' valuations. This is of fundamental importance for auction design, because they allow the simulation of efficiency results under different types of auction mechanism. It allows to determine the optimal design, and the parameters of this design (e.g. the reserve prices).

Since the aim of this section is only to provide a brief introduction to the structural estimation of auctions, we focus on two particular types of auctions: first and second price sealed bid auctions for one good, N players, and independent valuations v_i for $i = 1, ..., N$. In the first price sealed bid auctions, each player simultaneously submits a bid, and the object is assigned to the highest bidder, who pays the price she bid. In the second price sealed bid auctions, the highest bidder wins, but pays the bid of the second highest bidder. The literature has studied other types of auctions. For example, English auctions are similar to first price sealed bid auctions except that bidders sequentially call ascending prices, and other players observe bids. In Dutch auctions, the price is reduced until a player accepts the offer, so only the winning bid is ever observed. In Japanese auctions, players exit as the auctioneer raises the price, and the winner pays the price at which the only other remaining bidder exits. While Dutch auctions are strategically equivalent to first price sealed bid auctions, Japanese auctions are not necessarily strategically equivalent to second price sealed bid auctions, because players update their information sets as the auction evolves.

As noted above, we focus on single object first and second price sealed bid auctions. There are N risk-neutral bidders (indexed by $i = 1, ..., N$) and they have valuations v_i, independently drawn from a common distribution $F(\cdot)$. The data consists of the outcomes observed across independent auctions $k = 1, ..., K$ that follow the same paradigm. In this introductory review, we focus on the case where individual i observes her own valuation v_i, but not other players' valuations. The literature also analyzes the case in which the individual, instead, has a signal $x_i \neq v_i$, but does not observe v_i, and also the case in which different players have a (partially) common valuation. Players submit a single bid $b_i \in \mathbb{R}^+$ and do not observe other players' bids. The econometrician only observes bids, either all of them, or only the winning bid or price. The structural estimation of auction models usually rely on the equilibrium bid functions and on distributional assumptions regarding $F(\cdot)$, which often is specified parametrically.

The literature focuses on Perfect Bayesian Equilibria in weakly undominated pure strategies. A bidding strategy is a function that maps valuations into bids. The bidding strategy is the equilibrium solution of a expected utility maximization problem.

II. Equilibrium bid functions

A. Second price sealed bid auctions

In second price sealed bid auctions, it is a weakly dominant strategy for every individual to bid her expected valuation, i.e. $b_i = v_i$. Intuitively, bidding more implies winning some auctions that yield negative expected value but leaves unchanged the expected value of any other auction that would be won, whereas bidding less implies losing some auctions that yield positive expected value but leaves unchanged the expected value of any other auction that she would win.

B. First price sealed bid auctions

In a first price sealed bid auction, best responses are slightly more complicated, because, unlike in the second price counterpart, changing the bid not only affects the probability of winning, but also the price to pay. Let $p(b)$ denote the probability of winning the auction with bid b, defined as:

$$p_i(b) \equiv \Pr(\max\{b_j\}_{j \neq i} \leq b). \tag{1}$$

Then, b_i solves:

$$b_i = \arg\max_b (v_i - b) p_i(b). \tag{2}$$

The resulting b_i is the best response to other player's expected actions. The first order condition yields:

$$(v_i - b_i)p'(b_i) - p(b_i) = 0. \tag{3}$$

Totally differentiating this expression with respect to b and v, we obtain:

$$\frac{db_i}{dv_i} = \frac{-p'(b_i)}{(v_i - b_i)p''(b_i) - 2p'(b_i)} > 0. \tag{4}$$

The last inequality is obtained from observing that the denominator is the second order condition (and, hence, it is negative), and that the winning probability is increasing in the bid. Therefore, if players are in pure strategy equilibrium with an interior solution, then b_i is increasing in v_i. This ensures invertibility, and, hence, identification.

III. Identification

A. Second price sealed bid auctions

Let $F(v)$ denote the distribution of valuations. In a second price sealed bid auction, the distribution of valuations is trivially identified if all bids are observed because, as noted before, all players bid their valuation.

The case in which only the winning price is observed (in each of K auctions in which the same equilibrium is played) is a bit more convoluted. If only the winning price is observed, which equals to the second highest bid, then the probability distribution of the second highest valuation, denoted by $F_{N-1,N}(v)$ is identified trivially from the distribution of paid prices. Let $f_{N-1,N}(v)$ denote the corresponding density. Given symmetry and independence, this density equals:

$$f_{N-1,N}(v) = N(N-1)F(v)^{N-2}(1 - F(v))f(v). \tag{5}$$

Intuitively, the $N(N-1)$ comes from the combinatorial possibilities in which v is the second highest valuation, $F(v)^{N-2}$ is the probability that $N-2$ valuations are lower than v, and $1 - F(v)$ is the probability that one of them is higher. Given a boundary condition $F_{N-1,N}(\underline{v}) = F(\underline{v}) = 0$, and noting that $f(v) > 0$ above the boundary condition, the identification of $F(v)$ comes from solving the differential equation above.

B. First price sealed bid auctions

In first price sealed bid auctions, identification comes from the first order condition above. If all bids are observed, then $p(b)$ is trivially identified. Hence, v_i is identified from (3):

$$v_i = b_i + \frac{p(b_i)}{p'(b_i)} = b_i + \frac{G(b_i)}{(N-1)g(b_i)}, \tag{6}$$

where $G(\cdot)$ and $g(\cdot)$ are respectively the cdf and pdf of observed bids, and the latter equality is obtained from noting that $p(b) = G(b)^{N-1}$. Therefore, the probability distribution of $(v_1, ..., v_N)$ is identified off the bidding distribution $G(b)$.

If only the winning bid is recorded, the distribution of winning bids, $H(b)$ is identified from the outcomes observed in the different auctions. Since the winning bid is defined as the highest one, $H(b)$ is just the probability that all the bids in all the auctions are less than or equal to b, such that:

$$H(b) = \Pr(b_i^k \leq b \forall i = 1, ..., N) = G(b)^N. \tag{7}$$

Therefore:

$$G(b) = H(b)^{\frac{1}{N}}, \tag{8}$$

and:

$$g(b) = \frac{1}{N} H(b)^{\frac{1}{N}-1} h(b), \tag{9}$$

where $h(b)$ is the density of winning bids. Replacing these two in (6), this shows that the bidding distribution is identified off the winning bid distribution.

IV. Estimation

Estimation strategies range from minimum distance to maximum likelihood. In minimum distance estimation, once the distribution of bids is derived, up to parameter values, these parameters are estimated comparing sample moments of the observed bids against theoretical moments of the $G(\cdot)$ distribution for each parameter value. Sometimes, this moments are trivial functions of the parameters, and standard minimum distance methods are easy to implement. Other times, it is too costly to derive the theoretical moments from the distribution, and we proceed with simulated method of moments, using Monte Carlo approaches. In particular, for a given set of parameters, valuations are drawn for all players. This valuations correspond to bids, given the equilibrium bidding strategies of each player. Keeping the seed fixed, iterate over parameters to minimize the distance between simulated and data moments. Maximum likelihood approaches are also feasible, but have the complication that the upper bound of the support often depends on parameter values, which lead to estimates that, while consistent, are not asymptotically normal.