Advances in Economics and Econometrics

CHAPTER 1

Empirical Models of Auctions*

Susan Athey and Philip A. Haile

1 INTRODUCTION

Auctions have provided a fruitful area for combining economic theory with econometric analysis in order to understand behavior and inform policy. Early work by Hendricks and Porter (1988) and others made important contributions by testing the empirical implications of auction theory. This work provided convincing evidence of the empirical relevance of private information and confirmed the value of strategic models for understanding firm behavior. However, many important economic questions can be answered only with knowledge of the underlying primitive distributions governing bidder demand and information. Examples include the division of rents in auctions of public resources, whether reserve prices in government auctions are adequate, the effects of mergers on procurement costs, whether changes in auction rules would produce greater revenues, whether bundling of procurement contracts is efficient, the value of seller reputations, the effect of information acquisition costs on bidder participation and profits, whether bidders' private information introduces adverse selection, and whether firms act as if they are risk averse. Many of these questions have important implications well beyond the scope of auctions themselves.

Motivated by a desire to answer these questions, a more recent literature has developed that aims to estimate the primitives of auction models, exploiting restrictions from economic theory as part of the econometric model.[1] Typically, such a "structural" approach incorporates two types of assumptions: (a) economic assumptions, such as behavioral assumptions (e.g. Bayesian Nash equilibrium) and economically motivated restrictions on preferences (e.g., risk

* We thank Liran Einav, Jon Levin, Aviv Nevo and Whitney Newey for helpful comments and Ying Fan and Gustavo Soares for capable research assistance. Research support from the National Science Foundation (grants SES-0112047 and SES-0351500) and the Alfred P. Sloan Foundation is gratefully acknowledged. Any conclusions, findings, or opinions are those of the authors and do not necessarily reflect the views of any funding organization.

[1] A seminal paper in this literature is Paarsch (1992a), which builds on insights in Smiley (1979) and Thiel (1988).

neutrality), and (b) functional form assumptions, imposed either for convenience in estimation or because only a limited set of parameters can be identified. An attractive feature of the recent econometric literature on auctions is that often the second type of assumption can be avoided, both in principle and in practice. In particular, in many cases identification of economic primitives can be obtained without resorting unverifiable parametric assumptions, and nonparametric estimation methods have been developed that perform well in data sets of moderate size. Even when parametric estimation approaches are used in applications, the fact that the literature has provided definitive positive (and sometimes negative) identification results provides important guidance about how to interpret the results. This paper aims to review some of the highlights of this recent literature, focusing on econometric identification and empirical applications.

Fundamental to the structural approach is an interpretation of data through the lens of an economic model. Hence, we begin by defining notation, reviewing the rules of the most prevalent types of auctions, and deriving equilibrium conditions. Next, we discuss three key insights that underlie much of the recent progress in econometrics for auction models. The first is the usefulness of casting the identification problem as one of learning about latent distribution functions based on observation of certain order statistics (e.g., the highest bid or the second-highest valuation). This is a simple observation, but one that has helped to organize the attack on identification of auction models and, in several cases, has led to the discovery of connections between auction models and other familiar models in economics and statistics. The second is the observation that equilibrium can be thought of as a state of mutual best responses. This is again a seemingly trivial observation, but it has enabled economists to obtain surprisingly powerful results by re-casting equilibrium conditions (characterizing a fixed point) in terms of simpler optimality conditions for players facing a distribution (often observable) of equilibrium play by opponents. Finally, we discuss a third fundamental insight: the value of additional variation in the data beyond the realizations of bids. Observable variation in auction characteristics, in the realized value of the object, and in the number of bidders might initially seem to be minor nuisances to be dealt with, perhaps by conditioning or smoothing. In fact, these kinds of variation often can be exploited to aid identification.

Beyond these three central insights, we also discuss some extensions that have proved important for empirical applications. We describe how the econometric approaches can be generalized to account for endogenous participation and unobserved heterogeneity. In addition, we provide a brief discussion of specification tests that can help a researcher evaluate and select among alternative modeling assumptions.

Our discussion of applications begins with Hendricks, Pinkse and Porter's (2003) analysis of oil lease auctions, which exploits the availability of data on the market value of oil (and other minerals) realized ex post from each tract. Combined with data on bids, this enables the authors to quantify the magnitude of the winner's curse in their pure common values model. This work suggests that the subtle inferences required by bidders in common value auctions are economically important, and that they are in fact incorporated in bidding strategies.

We next discuss the working paper of Haile, Hong, and Shum (2003), who develop and apply tests to discriminate between common values and private values models in first-price auctions. They build on a simple idea: in a common values auction, an increase in the number of competing bidders amplifies the winner's curse. Since the winner's curse is present only in common values auctions, a test for rational responses by bidders to variation in the strength of the winner's curse offers an approach for testing. Equilibrium conditions enable them to isolate responses to the winner's curse, and they show how this idea can be used with several models of endogenous bidder participation. Their preliminary results suggest that common values may not be important, at least for some types of timber contracts.

Next, we discuss Haile and Tamer's (2003) bounds approach to analysis of ascending auctions. Because an actual ascending auction is typically a dynamic game with exceedingly rich strategy and state spaces, the theory of ascending auctions has relied on significant abstractions for tractability. Haile and Tamer (2003), concerned with the potential implications of estimating a misspecified model, propose an approach based on simple intuitive restrictions on equilibrium bidding that hold in a variety of alternative models. They show that these restrictions are sufficient to enable fairly precise inference on bidder demand and on the effects of reserve price policy. Addressing a policy debate regarding reserve prices in timber auctions, they show that actual reserve prices are likely well below the optimal levels, but that raising them would have only a small effect on expected revenues.

In another study of timber auctions, the working paper of Athey, Levin, and Seira (2004) uses variation in auction format (ascending versus first-price auctions) to a) test qualitative predictions of the theory of asymmetric auctions with endogenous participation and b) assess the competitiveness of ascending auctions, widely believed to be more susceptible to collusion. They show that observed bids and participation decisions identify the underlying distributions of bidder valuations and the costs of acquiring the information necessary to participate in an auction. Their preliminary estimates suggest that in several national forests, behavior in ascending auctions is less aggressive than would be consistent with a competitive theory, given a benchmark created using the distributions of valuations estimated from first-price auction data. Although the competitive theory explains part of the revenue gap, an alternative theory such as collusion at ascending auctions is required to rationalize the remainder.

The analysis by Jofre-Bonet and Pesendorfer (2003) of dynamics in procurement auctions provides an elegant generalization of prior approaches for static models. They consider situations in which bidders have capacity constraints, so that winning an auction affects valuations (or costs) in future auctions. Perhaps surprisingly, few additional assumptions are required for identification of the primitives in this kind of model. Their empirical analysis of highway construction auctions reveals significant asymmetries in bidding strategies resulting from asymmetric capacities of bidders at different points in time. They also find a fairly large gap between bids and values, half of which they attribute to bidders' recognition of the option value to losing a contract today: They may use their limited capacity for another contract in the future.

Finally, we discuss the working paper of Hortaçsu (2002), which takes an empirical tack on one of the oldest unresolved questions in the auction literature: whether to sell treasury bills by discriminatory or uniform price auction. The performance of these auctions has a substantial impact on the cost at which governments raise funds. Hortaçsu (2002) extends the econometric approaches of the prior literature to discriminatory multi-unit (share) auctions, building on the theoretical model of Wilson (1979). His preliminary estimates suggest that, for the Turkish treasury auctions he studies, switching to a uniform-price auction would not enhance revenues.

2 ESSENTIAL THEORY

The baseline theoretical framework is a generalization of Milgrom and Weber's (1982) affiliated values model, where a single indivisible good is sold to one of $n \in \left\{\underline{n}, \ldots, \overline{n}\right\}$ risk neutral bidders, with $\overline{n} \geq \underline{n} \geq 2$.[2] We denote random variables in upper case, their realizations in lower case, and vectors in boldface. We let $\mathcal{N} \subset \left\{\underline{n}, \ldots, \overline{n}\right\}$ denote the set of bidders, with N denoting the number of bidders. $\mathcal{N}_{-i}$ will denote the set of competitors faced by bidder i. The utility bidder i would gain by obtaining the good is given by U_i, which we refer to as i's "valuation" and assume to have common support (denoted supp U_i) for all i.

Bidder i's private information (his "type") consists of a scalar signal $X_i \in \left[\underline{x}_i, \bar{x}_i\right]$. We let $\mathbf{X} = (X_1, \ldots, X_n)$ and $\mathbf{X}_{-i} = \mathbf{X} \backslash X_i$. We assume that the random variables $(U_1, \ldots, U_n, X_1, \ldots, X_n)$ are affiliated, i.e., that higher realizations of one variable make higher realizations of the others more likely.[3] Signals are further assumed to be informative in the sense that the expectation

$$E\left[U_i | X_i = x_i, \mathbf{X}_{-i} = \mathbf{x}_{-i}\right]$$

is strictly increasing in x_i for all realizations $\mathbf{x}_{-i}$ of i's opponents' signals. Since signals play a purely informational role, it is without loss of generality to impose a normalization, e.g.,

$$X_i = E\left[U_i | X_i\right].$$

We will say that the model is *symmetric* if the indices $(1, \ldots n)$ may be permuted without affecting the joint distribution $F_{\mathbf{U},\mathbf{X}}(U_1, \ldots, U_n, X_1, \ldots, X_n)$ of bidders' valuations and signals; otherwise the model is *asymmetric*. The set

[2] We discuss an extension to multi-unit auctions in section 4.6 below.
[3] More formally, random variables $\mathbf{Y} = (Y_1, \ldots, Y_n)$ with joint density $f_{\mathbf{Y}}(\cdot)$ are affiliated if for all $\mathbf{y}$ and $\mathbf{y}'$, $f_{\mathbf{Y}}\left(\mathbf{y} \vee \mathbf{y}'\right) f_{\mathbf{Y}}\left(\mathbf{y} \wedge \mathbf{y}'\right) \geq f_{\mathbf{Y}}(\mathbf{y})\, f_{\mathbf{Y}}\left(\mathbf{y}'\right)$, where $\vee$ denotes the component-wise maximum, and $\wedge$ the component-wise minimum.

of bidders and the joint distribution $F_{\mathbf{U},\mathbf{X}}(\cdot;\mathcal{N})$ are assumed to be common knowledge among bidders.[4]

When we come to discuss estimation, we will generally assume a sequence of independent auctions indexed by $t = 1, \ldots, T$.[5] We will then add subscripts t to random variables (e.g., X_{it}, N_t) as needed. Asymptotic arguments will be based on $T \to \infty$. In practice a stronger assumption will often be required, e.g., that $T_n = \#\{t : N_t = n\} \to \infty$. In some cases we will imagine for simplicity that these auctions are not only independent but also identically distributed, i.e., that $F_{\mathbf{U}}(U_1, \ldots, U_n)$ is the same for every n-bidder auction. In practice this will rarely be the case, although there are a number of approaches available to account for observable (and to some degree, unobservable) differences across auctions t.

Within this general framework we will make a distinction between *private values* and *common values* auctions. The distinction concerns the nature of bidders' private information. In a private values auction, a bidder's private information concerns only factors idiosyncratic to that bidder; in a common values auction, each bidder's private information concerns factors that affect all bidders' valuations. More precisely,

Definition 2.1 *Bidders have* ***private values*** *if* $E[U_i|X_1 = x_1, \ldots, X_n = x_n] = E[U_i|X_1 = x_1]$ *for all* $x_1, \ldots, x_n$ *and all* i*; bidders have* ***common values*** *if* $E[U_i|X_1 = x_1, \ldots, X_n = x_n]$ *strictly increases in* x_j *for all* i, j, *and* x_j.[6]

Common values models apply whenever information about valuations is dispersed among bidders.[7] They include the special case of *pure common values,* where the value of the good is the same (but unknown) for all bidders.[8] Note that the distinction between private and common values is separate from the question of whether bidders' information is correlated. Bidders may have highly correlated private values, or could have pure common values but independent signals. In addition, the distinction is separate from the question of whether bidders' valuations are affected by shared factors. For example, even in a private values model, bidder valuations might all be affected by characteristics of the good for sale that are known to all bidders, or be subject to future macroeconomic shocks, about which bidders have identical priors. In either

4 See, e.g., Hendricks, Pinkse, and Porter (2003), Athey and Haile (2006), Song (2004), and Li and Zheng (2005) for applications relaxing the assumption that $\mathcal{N}$ is known by bidders.

5 We discuss relaxation of the independence across auctions in section 4.5.

6 Affiliation implies that $E[U_i|X_1 = x_1, \ldots, X_n = x_n]$ is increasing in x_j for all i, j, and x_j. For simplicity, our definition of common values rules out cases where strict monotonicity holds for some realizations of types but not others.

7 Common values models include all environments in which a winner's curse arises – i.e., where winning an auction reveals to the winner new information about his own valuation for the object.

8 Some authors (e.g., Krishna (2002)) use the term "interdependent values" to refer to the class of models we call "common values," motivated in part by inconsistencies in the literature in the use of the latter term.

case, because bidders have no private information about the shared factor, a private values model still applies.

We follow the literature and restrict attention to (perfect) Bayesian Nash equilibria in weakly undominated pure strategies, $\beta_i(\cdot;\mathcal{N})$, $i = 1, \ldots, n$, mapping each bidder's signal (and, implicitly, any public information) into a bid. In symmetric models we further restrict attention to symmetric equilibria, where $\beta_i(\cdot) = \beta(\cdot)\ \forall i$. We will denote a bidder i's equilibrium bid by B_i, with $\mathbf{B} = \{B_1, \ldots, B_n\}$.

In a first-price sealed-bid auction, bids are submitted simultaneously, and the good is awarded to the high bidder at a price equal to his bid (as long as this exceeds any reserve price, r).[9] For first-price auctions we make the following additional assumptions:

Assumption 1. *(First-Price Auction Assumptions) (i) For all i, U_i has compact, convex support denoted* $\operatorname{supp} U_i = [\underline{u}, \bar{u}]$. *(ii) The signals $\mathbf{X}$ are affiliated, with* $\operatorname{supp}\mathbf{X} = \times_{i=1}^{n} \operatorname{supp} X_i$. *(iii) $F_{\mathbf{X}}(\cdot)$ has an associated joint density $f_{\mathbf{X}}(\cdot)$ that is strictly positive on the interior of* $\operatorname{supp}\mathbf{X}$.

Under Assumption 1, there exists an equilibrium in nondecreasing bidding strategies, and in all models except the asymmetric common values model (which we will not discuss here), existence of an equilibrium in strictly increasing strategies has been established (see Athey and Haile (2006) for a more detailed discussion). We will restrict attention to equilibria in strictly increasing strategies and will derive the first-order conditions characterizing equilibrium bidding in Section 3.2 below. An important feature of equilibrium in first-price auctions is that bidders "shade" their bids by bidding less than their valuations; thus, a key step in developing econometric approaches to first-price auctions is estimation of the equilibrium bid functions that relate the observable bids to the latent primitives.

A second prevalent auction format is the oral ascending bid, or "English" auction. Ascending auctions are typically modeled following Milgrom and Weber (1982). In their model (sometimes referred to as a "clock auction" or "button auction" model) the price rises continuously and exogenously while bidders raise their hands or depress a button to indicate their willingness to buy at the current price. As the auction proceeds, bidders exit by lowering their hands or releasing their buttons. Exits are observable and irreversible, and the auction ends when only one bidder remains. This bidder wins the auction and pays a price equal to that at which the auction stopped, i.e., at his final opponent's exit price. Bids are synonymous with exits, so the auction ends at the second highest bid.

A bidding strategy in this model specifies a price at which to exit, conditional on one's own signal and on any information revealed by previous exits. If bidders

[9] This auction game is strategically equivalent to a Dutch (descending) auction. An important difference for empirical work, however, is the fact that only one bid could be observed, since only one bid (the winner's) is ever made.

use strategies that are strictly increasing in their signals, the price at which a bidder exits reveals his signal to the others. This matters in a common values auction, since the observed exit prices cause remaining bidders to update their beliefs about their own valuations. The prices at which bidders plan to exit thus change as the auction proceeds. In a private values auction there is no such updating, and each bidder has a weakly dominant strategy to bid up to his valuation, i.e.,

$$\beta_i(x_i; \mathcal{N}) = E[U_i \mid X_i = x_i] = x_i \equiv u_i. \tag{2.1}$$

With common values there are multiple equilibria, even with the restriction to strictly increasing, weakly undominated strategies; however, in any such equilibrium if i is one of the last two bidders to exit, his exit price b_i is equal to

$$E[U_i \mid X_i = x_i,\ X_j = x_i\ \forall j \notin \{i \cup \mathcal{E}_i\},\ X_k = x_k\, \forall k \in \mathcal{E}_i], \tag{2.2}$$

where $\mathcal{E}_i$ denotes the set of bidders who exit before i (Bikhchandani, Haile, and Riley (2002)).

In a private values ascending auction, the Milgrom-Weber model predicts a trivial relation between a bidder's valuation and his bid. Even in this case, however, identification can present challenges, due to the fact that the auction ends before the winner bids. While bids directly reveal valuations in this model, they do not reveal all of them. Furthermore, in many applications one may not be comfortable imposing the structure of the Milgrom-Weber model. In many ascending auctions, prices are called out by bidders rather than by the auctioneer, and bidders are free to make a bid at any point, regardless of their activity (or lack thereof) earlier in the auction. This raises doubts about the interpretation of bids (e.g., the highest price offered by each bidder) as representing each bidder's maximum willingness to pay. In Section 4.3 we will show that progress can still be made in some cases using a relaxation of Milgrom and Weber's model.

3 FOUNDATIONS OF IDENTIFICATION

3.1 Bids as Order Statistics

A simple but important insight, made early in the literature (e.g., Paarsch (1992a), Paarsch (1992b)), is that bid data can usefully be thought of in terms of order statistics. In particular, many identification problems involve the recovery of the latent distribution of a set of random variables from the distribution of a limited set of observable order statistics. An order statistic of particular interest is the transaction price (winning bid). This bid is the most commonly available datum, and it is the only bid one could observe in a Dutch auction. Thus, an important question is whether (or when) the joint distribution of bidder valuations can be recovered from the distribution of the winning bid alone.

We introduce some additional notation in order to discuss order statistics more formally. Given any set of random variables $\{Y_1, \ldots Y_n\}$, let $Y^{(k:n)}$ denote the kth order statistic, with $F_Y^{(k:n)}(\cdot)$ denoting the corresponding marginal CDF.

We follow the convention of indexing order statistics lowest to highest so that, e.g., $Y^{(n:n)} = \max\{Y_1, \ldots Y_n\}$.

Order statistics are particularly informative in the case of independent random variables. Independence reduces the dimensionality of the primitive joint distribution of interest. For example, the joint distribution $F_{\mathbf{Y}}(\cdot)$ of i.i.d random variables $\{Y_1, \ldots Y_n\}$ is the product of identical marginal distributions $F_Y(\cdot)$. This suggests that the distribution of a single statistic might be sufficient to uncover $F_{\mathbf{Y}}(\cdot)$. This is obviously correct in the case that one observes the maximum, $Y^{(n:n)}$, since $F_Y(y) = \left(F_Y^{(n:n)}(y)\right)^{1/n}$. In fact, it is well known that the distribution of any single order statistic from an i.i.d. sample of size n from an arbitrary distribution $F_Y(\cdot)$ has the distribution (see, e.g., Arnold, Balakrishnan, and Nagaraja (1992))

$$F_Y^{(k:n)}(s) = \frac{n!}{(n-k)!(k-1)!} \int_0^{F_Y(s)} t^{k-1}(1-t)^{n-k}\, dt \quad \forall s. \tag{3.1}$$

It is easy to verify that the right-hand side is strictly increasing in $F_Y(s)$. Hence, for any k and n, we can define a function $\phi(F; k, n) : [0, 1] \to [0, 1]$ implicitly by the equation

$$F = \frac{n!}{(n-k)!(k-1)!} \int_0^{\phi} t^{k-1}(1-t)^{n-k}\, dt. \tag{3.2}$$

Then $F_Y(y) = \phi\left(F_Y^{(k:n)}(y); k, n\right)$ for all y; i.e., knowledge of the distribution of a single order statistic uniquely determines the underlying parent distribution.

Athey and Haile (2002) point out that this observation is immediately useful for the standard model of the ascending auction in the symmetric independent private values setting, where each bidder's valuation is an independent draw from a CDF $F_U(\cdot)$. The equilibrium transaction price is equal to the second highest valuation, $u^{(n-1:n)}$. Since $F_U^{(n-1:n)}(u)$ uniquely determines $F_U(u)$ for all u (by (3.1)), $F_U(\cdot)$ is identified, even if one observes just the transaction price and the number of bidders. This identification result immediately extends to cases in which valuations are affected by auction-specific observables, which we denote $\mathbf{Z}$. In that case, (3.1) implies that the underlying parent distribution $F_U(\cdot|\mathbf{z})$ is uniquely determined by $F_U^{(n-1:n)}(\cdot|\mathbf{z})$ for all $\mathbf{z}$.

Independence is the key assumption. If the symmetry assumption is dropped but independence is maintained, $F_{\mathbf{U}}(\cdot)$ is again the product of n marginal distributions $F_{U_i}(\cdot)$, and one can show that when all U_i have the same support, observation of

$$\Pr\left(U^{n-1:n} \leq u, i \text{ is winner}; \mathcal{N}\right)$$

for each $i \in \mathcal{N}$ is sufficient to identify each $F_{U_i}(\cdot)$ in the standard ascending auction model (Athey and Haile (2002)). In an asymmetric model, identification requires having some information about which bidders' actions are observed; here, the identity of the winner is sufficient. Athey and Haile (2006) sketch the formal argument, which is based on results for an isomorphic model studied by Meilijson (1981).

In a first-price auction, the observations here regarding distributions of order statistics are not enough by themselves to demonstrate identification, since bids do not directly reveal bidders' private information. However a hint at their value can be seen by noting that the joint distribution of *bids* is identified from observation of a single order statistic of the bids when bidders' signals $(X_1, \ldots, X_n)$ are independent. This follows from the fact that each bid is a measurable function of the latent signal, which implies that bids are independent. We discuss this further below.

Note that these results require observation of n. This is easy to understand: In interpreting the second-highest bid (for example) it is essential to know whether this is the second highest of two bids or the second-highest of twenty-two bids! However, observation of an additional order statistic can eliminate this requirement (Song (2003)). Consider a symmetric independent private values ascending auction and suppose, for example, that in addition to the winning bid ($B^{(n:n)} = U^{(n-1:n)}$) the next highest bid (equal to $U^{(n-2:n)}$ in equilibrium) is also observed. The number of bidders n, however, is not known. Observe that, given $U^{(n-2:n)} = u'$, the pair $\left(U^{(n-1:n)}, U^{(n:n)}\right)$ can be viewed as the two order statistics $\left(\tilde{U}^{(1:2)}, \tilde{U}^{(2:2)}\right)$ for sample of two i.i.d random variables drawn from the truncated distribution

$$F_{\tilde{U}}\left(\cdot|u'\right) = \frac{F_U\left(\cdot\right) - F_U\left(u'\right)}{1 - F_U\left(u'\right)}.$$

Although $\tilde{U}^{(2:2)}$ is not observed, equation (3.1) implies that observation of the transaction price $\tilde{U}^{(1:2)}$ alone is sufficient to identify the parent distribution $F_{\tilde{U}}\left(\cdot|u'\right)$ for this sample. Identification of $F_U(\cdot)$ then follows from the fact that

$$\lim_{u' \downarrow \inf \operatorname{supp} U^{(n-2:n)}} F_{\tilde{U}}\left(\cdot|u'\right) = F_U\left(\cdot\right).$$

Note that as long as the distribution $F_U(\cdot)$ does not vary with n, this argument does require that n be fixed or have a particular stochastic structure.

When the independence assumption is dropped, Athey and Haile (2002) show that identification fails (even with symmetric private values) when one observes only a subset of bidders' valuations and the set of bidders, $\mathcal{N}$. This is particularly important in an ascending auction, where the winning bidder's valuation cannot be observed. Intuitively, without independence, the joint distribution of interest is n-dimensional, so data of lower dimension will not be adequate. To see this more precisely (following Athey and Haile (2002)), consider a symmetric n-bidder environment and suppose all order statistics of bidders' valuations are observed except $U^{(j:n)}$ for some j. Take a point $(u_1, u_2, \ldots, u_n)$ on the interior of the support of $F_{\mathbf{U}}(\cdot)$, with $u_1 < \cdots < u_n$. Define a joint density function $\tilde{f}_{\mathbf{U}}(\cdot)$ by shifting mass δ in the true density $f_{\mathbf{U}}(\cdot)$ from a neighborhood of $\left(u_1, \ldots, u_j, \ldots, u_n\right)$ (and each permutation) to a neighborhood of the