

NBER WORKING PAPER SERIES

EMPIRICAL MODELS OF AUCTIONS

Susan Athey
Philip A. Haile

Working Paper 12126
<http://www.nber.org/papers/w12126>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2006

We thank Whitney Newey for helpful comments and Ying Fan 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. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2006 by Susan Athey and Philip A. Haile. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Empirical Models of Auctions
Susan Athey and Philip A. Haile
NBER Working Paper No. 12126
March 2006
JEL No. C5, L1, D4

ABSTRACT

Many important economic questions arising in auctions can be answered only with knowledge of the underlying primitive distributions governing bidder demand and information. An active literature has developed aiming to estimate these primitives by exploiting restrictions from economic theory as part of the econometric model used to interpret auction data. We review some highlights of this recent literature, focusing on identification and empirical applications. We describe three insights that underlie much of the recent methodological progress in this area and discuss some of the ways these insights have been extended to richer models allowing more convincing empirical applications. We discuss several recent empirical studies using these methods to address a range of important economic questions.

Susan Athey
Department of Economics
Stanford University
Stanford, CA 94305-6072
and NBER
athey@stanford.edu

Philip A. Haile
Department of Economics
Yale University
37 Hillhouse Avenue
P.O. Box 208264
New Haven, CT 06520
and NBER
philip.haile@yale.edu

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 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

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

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. In particular, 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 modelling 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 $n \in \{\underline{n}, \dots, \bar{n}\}$ risk neutral bidders, with $\bar{n} \geq \underline{n} \geq 2$.² We denote random variables in upper case, their realizations in lower case, and vectors in boldface. We let $\mathcal{N} \subset \{\underline{n}, \dots, \bar{n}\}$ 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 $\text{supp}U_i$) for all i .

Bidder i ’s private information (his “type”) consists of a scalar signal $X_i \in [\underline{x}_i, \bar{x}_i]$. We let $\mathbf{X} = (X_1, \dots, X_n)$ and $\mathbf{X}_{-i} = \mathbf{X} \setminus X_i$. We assume that the random variables $(U_1, \dots, U_n, X_1, \dots, X_n)$ 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

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

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.,

$$X_i = E[U_i | X_i].$$

We will say that the model is *symmetric* if the indices $(1, \dots, n)$ may be permuted without affecting the joint distribution $F_{\mathbf{U}, \mathbf{X}}(U_1, \dots, U_n, X_1, \dots, X_n)$ of bidders’ valuations and signals; otherwise the model is *asymmetric*. The set 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 $t = 1, \dots, T$.⁵ We will then add subscripts t to random variables (e.g., X_{it} , N_t) as needed. Asymptotic arguments will be based on $T \rightarrow \infty$. In practice a stronger assumption will often be required, e.g., that $T_n = \#\{t : N_t = n\} \rightarrow \infty$. In some cases we will imagine for simplicity that these auctions are not only independent but also identically distributed; i.e., that $F_{\mathbf{U}}(U_1, \dots, U_n)$ is the same for every n -bidder auction. In practice this will rarely be the case,

²We discuss an extension to multi-unit auctions in section 4.6 below.

³More formally, random variables $\mathbf{Y} = (Y_1, \dots, Y_n)$ with joint density $f_{\mathbf{Y}}(\cdot)$ are affiliated if for all \mathbf{y} and \mathbf{y}' , $f_{\mathbf{Y}}(\mathbf{y} \vee \mathbf{y}') f_{\mathbf{Y}}(\mathbf{y} \wedge \mathbf{y}') \geq f_{\mathbf{Y}}(\mathbf{y}) f_{\mathbf{Y}}(\mathbf{y}')$, where \vee denotes the component-wise maximum, and \wedge the component-wise minimum.

⁴See, 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.

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

although there are a number 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 $E[U_i|X_1 = x_1, \dots, X_n = x_n] = E[U_i|X_1 = x_1]$ for all x_1, \dots, x_n and all i ; bidders have **common values** if $E[U_i|X_1 = x_1, \dots, X_n = x_n]$ 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 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, $\beta_i(\cdot; \mathcal{N})$, $i = 1, \dots, n$, 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 $\mathbf{B} = \{B_1, \dots, B_n\}$.

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

⁶Affiliation implies that $E[U_i|X_1 = x_1, \dots, X_n = x_n]$ 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.

⁷Common 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.

⁸Some 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.

⁹This 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

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 $\text{supp}U_i = [\underline{u}, \bar{u}]$. (ii) The signals \mathbf{X} are affiliated, with $\text{supp}\mathbf{X} = \times_{i=1}^I \text{supp}X_i$. (iii) $F_{\mathbf{X}}(\cdot)$ has an associated joint density $f_{\mathbf{X}}(\cdot)$ that is strictly positive on the interior of $\text{supp}\mathbf{X}$.

Under Assumption 1, there exists an equilibrium in nondecreasing bidding strategies, and in all models except the asymmetric common value 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 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.,

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

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

$$E[U_i | X_i = x_i, X_j = x_j \forall j \notin \{i \cup \mathcal{E}_i\}, X_k = x_k \forall k \in \mathcal{E}_i], \quad (2.2)$$

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 (1992b), Paarsch (1992a)), 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 $\{Y_1, \dots, Y_n\}$, let $Y^{(k:n)}$ denote the k th 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., $Y^{(n:n)} = \max\{Y_1, \dots, Y_n\}$.

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 $F_Y(\cdot)$ of i.i.d random variables $\{Y_1, \dots, Y_n\}$ 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))

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

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

$$F = \frac{n!}{(n-k)!(k-1)!} \int_0^{\phi} t^{k-1}(1-t)^{n-k} dt \quad \forall s \quad (3.2)$$

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

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

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, \dots, X_n) are independent. This follows from the fact that each bid is a measurable functions 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 2 bids or the second-highest of 22 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 $(U^{(n-1:n)}, U^{(n:n)})$ can be viewed as the two order statistics $(\tilde{U}^{(1:2)}, \tilde{U}^{(2:2)})$ for sample of two i.i.d random variables drawn from the truncated distribution

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

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 $F_{\tilde{U}}(\cdot|u')$ for this sample. Identification of $F_U(\cdot)$ then follows from the fact that

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

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, \dots, u_n) on the interior of the support of $F_{\mathbf{U}}(\cdot)$, with $u_1 < \dots < u_n$. 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 $(u_1, \dots, u_j, \dots, u_n)$ (and each permutation) to a neighborhood of the point $(u_1, \dots, u_j + \epsilon, \dots, u_n)$ (and each permutation). For small ϵ and δ , this change preserves symmetry and produces a valid pdf. Now note that the only order statistic affected in moving from $\tilde{f}_{\mathbf{U}}(\cdot)$ to $f_{\mathbf{U}}(\cdot)$ is $U^{(j:n)}$. Since $U^{(j:n)}$ is unobserved, the distribution of observables is unchanged.¹⁰

¹⁰ Athey and Haile (2002) extend the non-identification results to common values ascending auctions, where we show that the underlying information structure is not identified even under strong assumptions, such as a pure common values model with independent signals.

While this is an important negative result, it may suggest greater pessimism than is warranted for many applications. In many first-price auctions, for example, bids from all n bidders are observable, and we will see in the following section that identification often holds. Furthermore, even when the dimensionality of the bid data is less than n , there are often other observables that can enlarge the dimensionality of the data to match the dimensionality of $F_{\mathbf{U}}(\cdot)$. We discuss this possibility in Section 3.3.

3.2 Equilibrium as Best Responses

Identification outside a second-price sealed-bid or ascending auction presents different challenges, since bids do not directly reveal the underlying private information of bidders. Even the problem of identifying the joint distribution of valuations $F_{\mathbf{U}}(\cdot)$ in a private values auction in which bids are observed from all bidders seems quite challenging at first, since the equilibrium bid function relating the observed bids to the underlying valuations is a function of marginal distributions derived from the joint distribution $F_{\mathbf{U}}(\cdot)$ itself. Smiley (1979) and Paarsch (1992a) proposed early approaches relying on special functional forms. Laffont, Ossard, and Vuong (1995) applied a simulation-based method applicable in symmetric independent private values models.

An important breakthrough, due to Guerre, Perrigne, and Vuong (2000), was the insight that the first-order condition for optimality of a bidder's best response can be rewritten, replacing distributions of primitives with equilibrium bid distributions. Consider a private values auction and let

$$G_{M_i|B_i}(m|b; \mathcal{N}) = \Pr \left(\max_{j \in \mathcal{N} \setminus i} B_j \leq m | B_i = b \right)$$

denote the equilibrium distribution function for the maximum equilibrium bid among a bidder's i 's opponents, conditional on his own equilibrium bid being b . Let $g_n(m|b; \mathcal{N})$ denote the corresponding density. This distribution represents i 's equilibrium beliefs about competing bids. Conditioning on i 's own equilibrium bid is merely a way of conditioning on i 's private information (recall that bids are strictly increasing in types). This conditioning is necessary because bidders' own types may be correlated with those of their opponents, and therefore with the competing bids they face.

Underlying Guerre, Perrigne and Vuong's insight are two simple ideas: (a) equilibrium is achieved when each player best-responds to the equilibrium distribution of opposing bids; and (b) when we assume the data are generated by the model, this equilibrium distribution of opposing bids is observable to the econometrician. For example, in a private values auction, the equilibrium bid of a bidder i with valuation u_i must solve

$$\max_{\tilde{b}} \int_{-\infty}^{\tilde{b}} (u_i - \tilde{b}) g_{M_i|B_i}(m|\beta(u_i; \mathcal{N}); \mathcal{N}) dm \quad (3.3)$$

which has first-order condition

$$u_i = b_i + \frac{G_{M_i|B_i}(b_i|b_i; \mathcal{N})}{g_{M_i|B_i}(b_i|b_i; \mathcal{N})}. \quad (3.4)$$

To interpret this expression note that, rearranging slightly, (3.4) requires that the percentage “markdown” for bidder i , $(u_i - b_i)/b_i$, equal $\frac{G_{M_i|B_i}}{b_i \cdot g_{M_i|B_i}}$, the inverse of the elasticity of the probability of winning with respect to player i 's choice of bid. Interchanging quantities and probabilities, we see that this is precisely the condition characterizing equilibrium pricing in standard oligopsony models, with the probability of winning replacing the residual supply curve. Just as in the oligopsony case, when considering an increase in the price he offers, a bidder here trades off the losses from paying a higher price conditional on winning against an increase in the chance of winning.

If the econometrician observes all bids in the auction as well as bidder identities, everything on the right-hand side of (3.4) is observable, while the left-hand side is the latent valuation associated with the bid b_i . Hence this equation demonstrates the identifiability of the valuations underlying each observed bid. When a bid is observed from each bidder, this immediately implies identification of the joint distribution $F_{\mathbf{U}}(\cdot)$, since

$$F_{\mathbf{U}}(\mathbf{u}) = \Pr \left(b_1 + \frac{G_{M_1|B_1}(b_1|b_1; \mathcal{N})}{g_{M_1|B_1}(b_1|b_1; \mathcal{N})} \leq u_1, \dots, b_n + \frac{G_{M_n|B_n}(b_n|b_n; \mathcal{N})}{g_{M_n|B_n}(b_n|b_n; \mathcal{N})} \leq u_n \right).$$

Note that independence is not required. In fact, the conditions necessary for identification are only the conditions necessary for the theoretical model to have a unique equilibrium characterized by a first-order condition.¹¹

This insight can be combined with those from Section 3.1 when bidders' have independent types. For example in an independent private values environment, the distribution of the winning bid in a first-price sealed-bid or Dutch auction is $G_B^{(n:n)}(\cdot)$. As discussed previously, when bidders' valuations are independent, so are their equilibrium bids. Hence, using the results from Section 3.1, $G_B^{(n:n)}(\cdot)$ is sufficient to uniquely determine the underlying distributions $G_{B_i}(\cdot)$ for each i as long as the set of bidders and identity of the winner are observable to the econometrician (with symmetry, observation of the winning bid and n is sufficient, since $G_B^{(n:n)}(b) = (G_B(b))^n$). Since

$$G_{M_i|B_i}(b|b; \mathcal{N}) = \prod_{j \in \mathcal{N} \setminus i} G_{B_j}(b)$$

under independence, the right-hand side of (3.4) is identified for each bid b_i , implying identification of each $F_{U_i}(\cdot)$ (Guerre, Perrigne, and Vuong (2000), Athey and Haile (2002)).

Of course, the independence assumption will not always be natural. As suggested in the discussion of identification failure in ascending auctions, observation of all n bids will be necessary

¹¹See Athey and Haile (2006) for a detailed discussion of these conditions.

to recover an arbitrary n -dimensional joint distribution $F_{\mathbf{U}}(\cdot)$. However, in a sealed-bid auction, observation of bids from all n bidders is possible and common. Hence, despite the more complicated strategic bidding behavior in a first-price auction, one can often identify richer models than is possible in an ascending auction.

Even when all bids are observable, however, there are limits to what can be identified. In a common values auction, for example, if we let

$$v_i(x, y; \mathcal{N}) = E \left[U_i | X_i = x, \max_{j \in \mathcal{N} \setminus i} \beta_j(X_j; \mathcal{N}) = \beta_i(y; \mathcal{N}) \right] \quad (3.5)$$

a bidder's optimization problem (the analog of (3.3)) can be written

$$\max_b \int_{-\infty}^{\tilde{b}} \left(v_i(x_i, \beta_i^{-1}(m); \mathcal{N}) - \tilde{b} \right) g_{M_i|B_i}(m|\beta_i(x_i); \mathcal{N}) dm$$

giving the first-order condition

$$v_i(x_i, x_i; \mathcal{N}) = b_i + \frac{G_{M_i|B_i}(b_i|b_i; \mathcal{N})}{g_{M_i|B_i}(b_i|b_i; \mathcal{N})}. \quad (3.6)$$

This looks encouraging, since the right-hand-side again consists only of observables. However the left-hand side is not a primitive in a common values auction. By its definition in (3.5), the left-hand side depends on equilibrium bidding strategies. Furthermore, the primitive of interest for a common values auction is typically the joint distribution $F_{\mathbf{U}, \mathbf{X}}(\cdot)$. For example, this is the distribution needed to predict outcomes under alternative selling procedures, or to assess the division of surplus. A simple counting exercise again suggests the problem here: while the arguments above do imply that one can identify the joint distribution of the n random variables $(v_1(X_1, X_1; \mathcal{N}), \dots, v_n(X_n, X_n; \mathcal{N}))$ for any given \mathcal{N} , the distribution of interest $F_{\mathbf{U}, \mathbf{X}}(\cdot; \mathcal{N})$ governs $2n$ random variables for each \mathcal{N} . It is easy to confirm that this intuition is correct: without additional restrictions, observation of all bids is insufficient to identify $F_{\mathbf{U}, \mathbf{X}}(\cdot; \mathcal{N})$. A simple proof is to observe that whatever the true $F_{\mathbf{U}, \mathbf{X}}(\cdot; \mathcal{N})$, the model will be observationally equivalent to a private values model in which $U_i = X_i = v_i(X_i, X_i; \mathcal{N})$.

While this is an important negative result, in many applications one observes more than just bids. As we discuss in the following section, such additional observables can help to overcome the limitations of bid data alone.

3.3 The Value of Data beyond Bids

3.3.1 Bidder Covariates

In Section 3.1 we saw that if the only observables are bids, identification fails in an ascending auction when the assumption of independent types is dropped. This problem can be overcome if

there are bidder-specific covariates with sufficient variation. Indeed, observation of such variation and the transaction price alone can suffice.

Let W_i denote a scalar covariate (extension to the non-scalar case is straightforward) affecting bidder i 's private value for the good—for example, his location relative to a construction site.¹² Each bidder i has a valuation given by

$$U_i = g_i(W_i) + A_i$$

where each $g_i(\cdot)$ is a function unknown to the econometrician, and the private stochastic components (A_1, \dots, A_n) are drawn from an arbitrary joint distribution $F_{\mathbf{A}}(\cdot)$. Assume that (A_1, \dots, A_n) are independent of $\mathbf{W} = (\mathbf{W}_1, \dots, \mathbf{W}_n)$.

To see the role of the bidder-specific covariates, suppose for the moment that each $g_i(\cdot)$ were known and that we could somehow observe the order statistic $U^{(n:n)}$ (the intended exit price of the winning bidder), but no other valuations/bids. Given \mathbf{w} , $U^{(n:n)}$ has cumulative distribution

$$\begin{aligned} F_U^{(n:n)}(u|\mathbf{w}) &= \Pr(U^{(n:n)} \leq u|\mathbf{w}) \\ &= F_{\mathbf{U}}(u, \dots, u|\mathbf{w}) \\ &= \Pr(g_i(w_i) + A_i \leq u \forall i) \\ &= F_{\mathbf{A}}(u - g_1(\mathbf{w}_1), \dots, u - g_n(\mathbf{w}_n)). \end{aligned}$$

Observing this probability reveals a great deal when there is sufficient independent variation in $(g_1(\mathbf{w}_1), \dots, g_n(\mathbf{w}_n))$; indeed, such variation could “trace out” the value of joint distribution $F_{\mathbf{A}}(\cdot)$ at every possible value of its arguments. So what about the assumption that $g_i(\cdot)$ is known? This is not necessary with sufficient variation in covariates: at sufficiently low values of $g_j(w_j) \forall j \neq i$, bidder i will have the largest valuation with probability arbitrarily close to one, so that variation in \mathbf{w}_i and the point of evaluation u would trace out the function $g_i(\cdot)$.¹³

This is not to suggest that in practice we should hope to have a large number of observations at extreme realizations of covariate values. Rather, the point is that variation in covariates can enable even limited bid data to reveal significant information about the underlying structure — enough to construct finite sample estimates with valid asymptotic justifications.

Of course in practice one cannot observe $U^{(n:n)}$ in an ascending auction, but given the analysis above one might hope that another order statistic would also suffice. Athey and Haile (2006) have shown that this is the case. They make the following assumptions which, in particular, spell out sufficient conditions for the “sufficient variation in covariates” condition referred to above:

¹²Extension to non-scalar covariates is straightforward.

¹³This argument is essentially that made by Heckman and Honoré (1989) and Heckman and Honoré (1990) for the nonparametric identification of competing risks models and the Roy model.

- (i) $U_i = g_i(\mathbf{W}_i) + A_i$, $i = 1, \dots, n$.
- (ii) \mathbf{A} has support \mathbb{R}^n and a continuously differentiable density.
- (iii) A_i and \mathbf{W}_j are independent for all i, j .
- (iv) $\text{supp}(g_1(\mathbf{W}_1), \dots, g_n(\mathbf{W}_n)) = \mathbb{R}^n$.
- (v) For all i , $g_i(\cdot)$ is continuously differentiable and satisfies $\lim_{\mathbf{w}_i \rightarrow (\infty, \dots, \infty)} g_i(\mathbf{w}_i) = \infty$ and $\lim_{\mathbf{w}_i \rightarrow (-\infty, \dots, -\infty)} g_i(\mathbf{w}_i) = -\infty$.

Under these assumptions, $F_{\mathbf{A}}(\cdot)$ and each $g_i(\cdot)$, $i = 1, \dots, n$, are identified up to a location normalization from observation of $U^{(j:n)}$ and \mathbf{W} , for any single value of $j \in \{1, \dots, n\}$. Hence, for example, observation of the winning bid alone in an ascending auction is sufficient to identify and arbitrary joint distribution $F_{\mathbf{U}}(\cdot)$ under these assumptions (see Athey and Haile (2006) for the proof).

3.3.2 Ex Post Values

Another source of variation in the data that can be particularly useful in a pure common values setting is the ex post realization of the good's value.¹⁴ For example, data on ex post values have been collected for U.S. Forest Service timber auctions (Athey and Levin (2001)), auctions of real estate (McAfee, Quan, and Vincent (2002)), and U.S. offshore oil lease auctions (Hendricks, Pinkse, and Porter (2003)). In a pure common values auction, $U_i = U_0 \forall i$, so the joint distribution $F_{\mathbf{U}, \mathbf{X}}(\cdot)$ (which we can rewrite $F_{U_0, \mathbf{X}}(\cdot)$) governs $n + 1$ random variables. Observation of bids from all n bidders as well as the realization of U_0 , therefore, at least offers hope of identification.

Hendricks, Pinkse, and Porter (2003) have provided conditions under which this is the case. Consider a symmetric pure common values model in which, conditional on the realization of U_0 , bidders' signals are independent draws from a distribution $F_X(\cdot|U_0)$.¹⁵ Fixing the number of bidders, one can impose the normalization

$$E[U_i | X_i = \max_{j \neq i} X_j = x, n] = x \tag{3.7}$$

without loss of generality. In a first-price sealed-bid auction, the first-order condition (3.6) then gives

$$E \left[U_0 | X_i = \max_{j \neq i} X_j = x_i, n \right] = x_i = b_i + \frac{G_{M_i|B_i}(b_i|b_i; n)}{g_{M_i|B_i}(b_i|b_i; n)}. \tag{3.8}$$

¹⁴Related possibilities include noisy ex ante estimates (e.g., Smiley (1979)) or noisy ex post estimates (e.g., Yin (2003)).

¹⁵In the literature this special case of a pure common values model is referred to as the symmetric "mineral rights model."

As discussed in Section 3.2, observation of all bids identifies the right-hand side of (3.8) and, therefore, the realizations of X_i associated with each bid. Observation of U_0 for each auction then immediately enables identification of the joint distribution $F_{U_0, \mathbf{X}}(\cdot)$.

The use of data on ex post values can be extended to environments in which not all bids are observed, exploiting the independence of signals conditional on U_0 . In the case of an ascending auction, however, bidders' ability to observe their opponents' bids presents a complication. In particular, there is no normalization like (3.7) that equates signals to bids for all bidders, since bidders update their beliefs about opponents' signals as the auction proceeds. In a two-bidder auction, this problem does not arise. In fact, with $n = 2$ there is no multiplicity of equilibria in weakly undominated strategies (Bikhchandani, Haile, and Riley (2002)) and for $i = 1, 2$

$$\begin{aligned} b_i &= \beta_i(x_i; 2) = E[U_0 | X_1 = X_2 = x_i, N = 2] \\ &= E[U_0 | \beta_1(X_1; 2) = \beta_2(X_2; 2) = \beta_i(x_i; 2), N = 2] \\ &\equiv \zeta(b_i, b_i, 2). \end{aligned}$$

Under the normalization (3.7), $B^{(1:2)} = X^{(1:2)}$ and, conditional on U_0 , $X^{(1:2)}$ is an order statistic from a sample of independent draws from $F_{X|U_0}(\cdot)$. Exploiting equation (3.1), identification is then obtained for a symmetric two-bidder auction when the transaction price (lowest bid) and *ex post* value are observed. This approach, based on the lowest observed bid, can be extended to the case with $n > 2$, although relying on the assumption $B^{(1:n)} = U^{(1:n)}$ may be particularly questionable when $n > 2$ (see the discussion in Section 4.3 below).¹⁶

3.4 Important Extensions

The insights discussed in the preceding sections have been applied and extended in many different ways. We will discuss some of these in the context of individual applications below. Before doing so, however, we highlight several general topics that have received increasing attention in recent years—topics that are also particularly relevant for developing convincing empirical applications of the basic methods.

3.4.1 Endogenous Participation

Bidders will not participate in an auction unless they expect doing so to be profitable. In the baseline model, bidding is always profitable; however, a number of factors such as reserve prices, participation fees, or costly information acquisition can cause some potential bidders to stay out of

¹⁶See Athey and Haile (2006) for a variation.

an auction. The distinction between actual bidders (those who place a bid) and potential bidders (those who could in principle bid) has been emphasized throughout the literature on structural econometrics of auctions (e.g., Paarsch (1992b)). Indeed, a common goal of work in this literature has been to better inform mechanism design choices accounting for the endogeneity of participation (e.g., Paarsch (1997), Haile and Tamer (2003), Athey, Levin, and Seira (2004)).

Recently, more attention has been given to the variety of possible models of endogenous participation. Bidders may decide to participate based on exogenous randomization (mixed strategies) as in Li and Zheng (2005) or Athey, Levin, and Seira (2004); based on preliminary private signals of the object's value (e.g., Hendricks, Pinkse, and Porter (2003)); based on their idiosyncratic costs of acquiring information; or based on public information about the object that is observable or unobservable to the econometrician (e.g., Haile, Hong, and Shum (2003)). Each of these possibilities has implications that must be accounted for in interpreting bids. In addition, one lesson of the work thus far is that incorporating equilibrium participation decisions can often aid identification and estimation of the bidding model (see, e.g., Section 4.2 below). Furthermore, endogenous participation can have important implications for mechanism design choices, such as the auction format (Athey, Levin, and Seira (2004)). More work is needed in this area, for example to develop approaches for evaluating alternative assumptions on bidder participation.

3.4.2 Unobserved Heterogeneity

Unobserved heterogeneity is a concern in almost all empirical work. In auctions, one may often expect that there are characteristics of the object for sale that are observed by all bidders but not by the econometrician. Such unobservables can create significant challenges. In an ascending auction, unobserved heterogeneity generally leads to correlation in bidders' valuations, often causing identification to fail (see Section 3.1). In a first-price sealed-bid auction, suppose bidders have private values given by

$$U_i = V_0 + A_i \tag{3.9}$$

where V_0 is an auction-specific characteristic observed by all bidders and A_i is a private idiosyncratic factor affecting only bidder i 's valuation. In this environment, bidders will use their knowledge of V_0 when constructing their beliefs about the competing bids they face. In particular, equilibrium bids will solve the first-order condition

$$u_i = b_i + \frac{G_{M_i|B_i}(b_i|b_i; v_0)}{g_{M_i|B_i}(b_i|b_i; v_0)} \tag{3.10}$$

where $G_{M_i|B_i}(b_i|b_i; v_0) = \Pr(\max_{j \neq i} B_j \leq b_i | B_i = b_i, V_0 = v_0)$. If the econometrician does not observe V_0 , it is not clear how the the right-hand side of (3.10) could be consistently estimated.

Although this is a challenging problem, several fruitful ideas have recently been explored. One approach builds on insights from the econometrics literatures on measurement error with repeated measurements and duration models with multiple spells. These literatures consider multiple observations for each of many units, with observations within each unit reflecting both a common (unobserved) shock as well as idiosyncratic shocks. Using deconvolution methods (e.g., Kotlarski (1966)), these literatures have shown how to separate the unit-specific factors (V_0 in the auction case) from the observation-specific factors (the A_i here). Krasnokutskaya (2004), for example, shows how to combine these techniques with the insights from Section 3.2 to obtain identification and consistent estimators for first-price sealed-bid auctions. The model of Athey, Levin, and Seira (2004), discussed in Section 4.4 below, builds on these ideas in an application to timber auctions.

Another possible approach is to utilize additional observables that control for the unobserved heterogeneity through another endogenous outcome. For example, suppose bidder participation is monotonic in the unobservable V_0 , conditional on some set of observables, \mathbf{Z} ; i.e.,

$$N = \phi(\mathbf{Z}, V_0)$$

with ϕ strictly increasing in V_0 . In this case, conditioning on (\mathbf{Z}, N) indirectly fixes the realization of V_0 , making it possible to estimate the right-hand side of (3.10) (Campo, Perrigne, and Vuong (2003), Haile, Hong, and Shum (2003)). This can be thought of as using a set of observable outcomes as “control functions” for the unobservable.¹⁷

Each of these approaches requires assumptions on unobservables: independence and a separable structure like (3.9) in the case of the deconvolution approach; strict monotonicity of a participation equation (or some other relation) in the control function approach. Each set of assumptions has strengths and limitations. Given the importance of unobserved heterogeneity in practice, we suspect that other useful methods will be explored in the future.

3.4.3 Specification Testing

As mentioned in the introduction, a “structural” approach to empirical work typically 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 neutrality), and (b) functional form assumptions, imposed either for convenience in estimation or because only a limited set of parameters can be identified. As we have seen, for some kinds of auction models, the second type of assumption is not essential. However, not all models are nonparametrically identified, and typically even those that are make assumptions like independence or separability

¹⁷Indeed, an alternative approach is to condition directly on an estimated residual from the participation equation, rather than on the observables (\mathbf{Z}, N) .

on unobservables. Because the conclusions one reaches in empirical work can be sensitive to such assumptions, it is desirable to test them when possible.

A natural approach to specification testing is to examine overidentifying restrictions. For example, in a model in which the winning bid alone is sufficient to identify the distribution of bidder valuations, observability of other bids (which would also identify this distribution on their own) provide overidentifying restrictions that can be tested. In other cases, it is possible to nest one model (e.g., symmetric private values) within another (asymmetric private values), leading to a testable restriction. Perhaps surprisingly, even without overidentifying restrictions (or even identifying restrictions), tests are sometimes possible. We will see an example in Section 4.2 below. Athey and Haile (2002) and Athey and Haile (2006) discuss a variety of testable restrictions that can be used to evaluate or decide between alternative specifications. More work is needed in this area, for example, to develop appropriate statistical tests.

4 Applications

In the following sections we illustrate some of the ways the ideas above have been applied and extended to address important economic questions.

4.1 Assessing the Winner’s Curse in Common Value First-Price Auctions

Since the early work of Capen, Clapp, and Campbell (1971), considerable attention has been given to U.S. Interior Department auctions of rights to drill for offshore oil and gas. One topic of particular interest is the “winner’s curse”—the fact that, in a common values auction, winning bidders tend to be those who have overestimated the good’s value. Some of the most influential early empirical work on auctions (e.g., Hendricks and Porter (1988)) explored implications of rational bidding in the presence of the winner’s curse and provided compelling evidence of the empirical relevance of asymmetric information and strategic behavior. Recent developments in methods for structural empirical work on auctions have opened up opportunities for additional testing. This might be surprising, since sometimes empirical work is described as falling in one of two categories: structural estimation, or model testing. However, these are not mutually exclusive. Structure from economic theory often provides restrictions that enable one to test hypotheses that could not otherwise be considered. In auctions, for example, maintaining one set of restrictions (e.g., Bayesian Nash equilibrium bidding) often provides overidentifying restrictions that can be used to test other economic hypotheses that could not be examined without this structure.¹⁸

¹⁸Athey and Haile (2006) discusses a wide range of testable restrictions in auction models. Other examples from the literature include Haile (2001), Campo, Guerre, Perrigne, and Vuong (2002), and Haile, Hong, and Shum (2003),

Hendricks, Pinkse, and Porter (2003) take this approach to develop several tests of rational responses to the winner’s curse in offshore “wildcat” sales in Texas and Louisiana. In these auctions, held between 1954 and 1970, leases of tracts in previously unexplored areas of the outer continental shelf were offered for sale. The auctions are conducted in first-price sealed-bid format, and all bids are recorded.

The authors make use of the fact that the actual volume of oil and other minerals extracted from each tract is metered to determine royalties and, therefore, observable ex post. If idiosyncratic determinants of costs are small, then a pure common values model may be appropriate, and the product of the realized volumes and market prices (less a measure of extraction costs) provides a measure of the ex post value, U_0 .

Hendricks, Pinkse, and Porter (2003) apply the mineral rights model to these auctions (see Section 3.3.2 above)¹⁹ and develop a sequence of tests of increasingly demanding implications of rational strategic bidding.²⁰

Positive Average Rents With rational bidding, winning bidders should make positive profits on average. That is, up to sampling variation,

$$\frac{1}{T} \sum_t \left[u_{0t} - b^{(n_t:n_t)} \right]$$

should be positive. Hendricks, Pinkse, and Porter (2003) find average rents of around \$3.7 million. Taking this further, not only should the actual winners obtain positive expected rents, but every observed bid should be consistent with a rational expectation of making positive profits in the event that the bid wins. This is a more subtle restriction, since conditioning on the (usually counterfactual) event that a bid wins the auction means conditioning on the implied (in equilibrium) information about opponents’ signals, and therefore about U_0 .

Let

$$\omega(b) \equiv E \left[U_0 | B_{it} = b, \max_{j \neq i} B_{jt} \leq b; n \right]$$

be the expected value of the object for bidder i conditional on his winning the auction with bid b . This expectation is identified directly from the data on ex post values and bids. With symmetric

discussed in section 4.2 below.

¹⁹They allow for two classes of bidders: “large” and “fringe,” and for endogenous participation. Their analysis focuses on the large bidders, although bids from fringe bidders are included when they construct distributions of opposing bids. The model of participation offers a formal “purification” of a mixed strategy participation equilibrium. We abstract from both of these features for simplicity.

²⁰Hendricks, Pinkse, and Porter (2003) use an estimate of the number of potential bidders (the number of firms ever to bid in a neighborhood of the tract in question) as a control for unobserved heterogeneity of tracts and perform the analysis below conditional on “high” and “low” values of this variable.

strictly increasing equilibrium bidding strategies $\beta(\cdot; n)$, this is equivalent to

$$E \left[U_0 | X_{it} = \beta^{-1}(b; n), \max_{j \neq i} X_{jt} \leq \beta^{-1}(b; n); n \right]$$

Hendricks, Pinkse, and Porter (2003) propose testing the restriction

$$E_B [\omega(B) - B] > 0$$

where B represents the bid of a generic bidder.

Because each U_{0t} and all bids are observable, various parametric or nonparametric regression techniques could be used to estimate $\omega(b)$. Hendricks, Pinkse, and Porter (2003) employ local-linear regression (see, e.g., Loader (1999)), yielding estimates $\hat{\omega}(b_{it})$ for each observed bid b_{it} . They then examine whether

$$\frac{1}{T} \sum_t \sum_{i=1}^{n_t} [\hat{\omega}(b_{it}) - b_{it}] > 0.$$

They find an average expected margin of around \$3.7 million—essentially the same as the average rents to winning bidders. This compares to winning bids averaging \$3 to \$8 million. Using a block bootstrap to construct standard errors,²¹ the hypothesis of zero average margins can be rejected in favor of positive margins.

Positive Expected Rents for All Bids. The tests of rationality above can be sharpened by asking whether inequalities hold for all bids, not just on average. For example,

$$E[U_0 | B_i = b] - b$$

should be positive for each observed bid b . Likewise, if bidders recognize that winning the auction is informative about others' signals, $\omega(b) - b$ should be positive for all observed bids b , not just on average. Again using local linear regression to estimate the conditional expectations, Hendricks, Pinkse, and Porter (2003) find support for both restrictions.

Equilibrium Bidding The most demanding test uses the structural model, combining several the ideas discussed in Sections 3.2, 3.3, and 3.4.3. Let

$$\zeta(b, m, n) = E \left[U_0 | B_i = b, \max_{j \neq i} B = m, n \right]. \quad (4.1)$$

²¹Blocking by auction is necessary since bids in this model will be correlated within auction.

With strictly increasing equilibrium bidding strategies

$$\begin{aligned}
E \left[U_0 | X_i = \max_{j \neq i} X_j = x_i, n \right] &= E \left[U_0 | \beta(X_i; n) = \max_{j \neq i} \beta(X_j; n) = \beta(x_i; n), n \right] \\
&= E \left[U_0 | B_i = \max_{j \neq i} B_j = b_i, n \right] \\
&= \zeta(b_i, b_i, n).
\end{aligned}$$

The first-order condition (3.8) can then be written

$$\zeta(b_i, b_i, n) = b_i + \frac{G_{M|B}(b_i|b_i; n)}{g_{M|B}(b_i|b_i; n)} \equiv \xi(b_i, n). \quad (4.2)$$

Because the joint distribution of $(U_0, B_1, \dots, B_n, N)$ is observable, $\zeta(b_i, b_i, n)$ is identified directly through equation (4.1). Indeed, $\zeta(b_i, b_i, n)$ is just a conditional expectation of the observable U_0 given that the observable bids satisfy $B_i = \max_{j \neq i} B_j = b_i$. Since $\xi(b_i, n)$ is also identified from the bidding data (without the ex post values) under the assumption of equilibrium bidding (recall Section 3.2), the overidentifying restriction $\zeta(b_i, b_i, n) = \xi(b_i, n)$ can be tested.

The right-hand side of (4.2) can be estimated using the kernel methods now standard in the literature on first-price auctions (see, e.g., Guerre, Perrigne, and Vuong (2000), Li, Perrigne, and Vuong (2000), and Li, Perrigne, and Vuong (2002)). In particular, let

$$\hat{G}_{M,B}(b, b; n) = \frac{1}{nT_n h_G} \sum_{t=1}^{T_n} \sum_{i=1}^n K \left(\frac{b - b_{it}}{h_G} \right) \mathbf{1} \{ m_{it} < b, n_t = n \} \quad (4.3)$$

$$\hat{g}_{M,B}(b, b; n) = \frac{1}{nT_n h_g^2} \sum_{t=1}^{T_n} \sum_{i=1}^n \mathbf{1} \{ n_t = n \} K \left(\frac{b - b_{it}}{h_g}, \frac{b - m_{it}}{h_g} \right) \quad (4.4)$$

where M_{it} denotes $\max_{j \neq i} B_{jt}$, $K(\cdot)$ is a kernel, and h_G and h_g are appropriately chosen bandwidth sequences. Under standard conditions, $\frac{\hat{G}_{M,B}(b, b; n)}{\hat{g}_{M,B}(b, b; n)}$ is a consistent estimator of $\frac{G_{M|B}(b|b; n)}{g_{M|B}(b|b; n)}$, so that

$$\hat{\xi}_{it} \equiv b_{it} + \frac{\hat{G}_{M,B}(b_{it}, b_{it}; n)}{\hat{g}_{M,B}(b_{it}, b_{it}; n)}$$

is a consistent estimate of $\xi(b_i, n)$.²² Using local linear regression to estimate the left-hand side of (4.2) and the bootstrap to construct critical values, Hendricks, Pinkse, and Porter (2003) fail to reject the null hypothesis of equality, providing support for the equilibrium bidding hypothesis.

²²Hendricks, Pinkse, and Porter (2003) do not condition on the number of bidders n but on whether the auction is one with a large or small number of potential bidder, treating this as the information available to bidders.

How Large is the Winner’s Curse? Hendricks, Pinkse, and Porter (2003) also suggest ways of illustrating the magnitude of the winner’s curse based on the difference

$$E[U_0|X_i = x_i, n] - E\left[U_0|X_i = x_i, \max_{j \neq i} X_j \leq x_i, n\right]. \quad (4.5)$$

Under the pure common values assumption, this represents the adjustment a bidder must make to account for the information implied by winning the auction.²³ Exploiting the monotonicity of equilibrium bidding, this difference is equal to

$$E[U_0|B_i = b_i, n] - E\left[U_0|B_i = b_i, \max_{j \neq i} B_j \leq b_i, n\right] \quad (4.6)$$

and each of the expectations can be estimated using a variety of regression techniques. Using local linear regression, Hendricks, Pinkse, and Porter (2003) find that this difference is positive and larger on tracts with more potential bidders. On tracts with few potential bidders, the estimated difference is about \$2 million for a tract receiving a bid of \$1 million. For tracts with more potential bidders, the difference in (4.6) of \$5 million (at a \$1 million bid) is considerably larger. Hence, the magnitude of the winner’s curse is significant. Moreover, it is larger when there is greater anticipated competition, as must be the case if the maintained assumption of symmetric pure common values is correct.

4.2 Discriminating between Private and Common Values

The distinction between private and common values models is fundamental and was, in fact, the motivation behind Paarsch (1992a)’s early influential work on structural estimation of auction models. Discriminating between the two classes of models is subtle, however. Laffont and Vuong (1996) argued that they are empirically indistinguishable without a priori parametric assumptions. When data beyond bids are available, however, this is not always correct. For example, Hendricks, Pinkse, and Porter (2003) suggested a testing approach for auctions with binding reserve prices.²⁴

Haile, Hong, and Shum (2003) have proposed another approach based on the fact that the winner’s curse is present only in common values auctions and becomes more severe as the level of competition increases. In particular, while the expectation

$$v(x, x; n) = E[U_i|X_i = \max_{j \neq i} X_j = x; n] \quad (4.7)$$

²³As pointed out by Hendricks, Pinkse, and Porter (2003), positive values for this difference (or even monotonicity of this difference in n) do not provide evidence against a private values assumption, since the interpretation of the ex post value data presumes pure common values. See Athey and Haile (2006) for additional discussion. Hendricks, Pinkse, and Porter (2003) do suggest a testing approach that could be used when there is a binding reserve price. We discuss this in greater detail in Athey and Haile (2006). We discuss a different approach to discriminating between private and common values models in section 4.2 below.

²⁴See Hendricks and Porter (forthcoming) or Athey and Haile (2006) for details.

is invariant to n in a private values model, it strictly decreases in n in any common values environment.²⁵ This observation provides a testing approach that can be applied with or without binding reserve prices, and which can allow for a range of models of endogenous participation.

The basic idea is simple: estimate the right-hand side of (3.6) and compare estimates of $v(x, x; n)$ at different values of n .²⁶ Let

$$H_n(\check{v}) = \Pr(v(X_i, X_i; n) \leq \check{v}).$$

Under the null hypothesis of private values, $H_n(\check{v})$ should be invariant to n , while under the common values alternative,

$$H_n(\check{v}) < H_{n+1}(\check{v}) \quad \forall \check{v}, n.$$

Tests of equal distributions against a one-sided alternative of stochastic dominance can then provide tests of the private values null against the common values alternative.²⁷ This testing problem nonstandard, however, due to the fact that one cannot observe realizations of the random variable $v(X, X; n)$, but only estimates obtained through observed bids and the first-order condition (3.6). Haile, Hong, and Shum (2003) explore a variety of testing approaches. A consistent test demonstrating good finite-sample power and coverage properties in Monte Carlo simulations is a modified Kolmogorov-Smirnov test, using subsampling to construct critical values.

Haile, Hong, and Shum (2003) apply this approach to two types of auctions held by the U.S. Forest Service (USFS). In all auctions, a contract to harvest all specified timber on a tract of land is sold to the high bidder by first-price sealed-bid auction. Before each auction, the Forest Service conducts a “cruise” of the tract in order to publish detailed pre-auction estimates of timber volumes by species and other determinants of the contract’s value. Under one type of contract, known as a lumpsum sale, the winning bidder simply pays its bid. Consequently, bidders frequently conduct their own cruises of the tract, providing a likely source of private information. While such information might be of a private or common values nature, intuition might suggest that these auctions might have a significant common values component. With the second type of contract, payments are not made until the timber is harvested, and the amount paid is based on the actual volume of timber harvested, not the ex ante estimates. Bids in these “scaled” sales are real unit

²⁵The idea of using variation in n to detect the winner’s curse precedes the first precise statements, given in early drafts of Athey and Haile (2002) and Haile, Hong, and Shum (2003). However, the observation was often misinterpreted as providing a way to test for common values based directly on the relation between bids and n (see, e.g., Brannman, Klein, and Weiss (1987), Paarsch (1992a), Paarsch (1992b), Laffont (1997)). Pinkse and Tan (2005) showed that the relation of bids to n is indeterminate in both private and common values models.

²⁶The approach described here uses all bids. Athey and Haile (2002) discuss a variation that could be applied when certain subsets of bids are observed, in either first-price or ascending auctions.

²⁷As discussed in Athey and Haile (2006), there are several other hypotheses of interest in auction models that take the form of equality of estimated distributions.

prices that will be applied to the actual volumes. So although bidders may face considerable uncertainty over the volume of timber on the tract, this common uncertainty may have limited importance to the value of the tract. Consequently, bidders are less likely to conduct a private cruise for a scaled sale. Several prior empirical studies used these facts to motivate an assumption of private values.²⁸

Several practical issues must be overcome to implement the testing approach in this application. One is that tracts are heterogeneous, as evidenced by the variation in tract characteristics published in the Forest Service cruise report. These characteristics could be conditioned on using standard nonparametric methods, but they are sufficiently numerous that this approach is impractical here. Haile, Hong, and Shum (2003) propose a semi-parametric alternative based on the observation that if valuations take an additively separable form

$$u_{it} = \gamma(\mathbf{z}_t) + \epsilon_{it}$$

then equilibrium bids have the same additively separable form

$$b_{it} = \gamma(\mathbf{z}_t) + \tilde{b}_{it}.$$

Here \tilde{b}_{it} represents the equilibrium bid that bidder i *would have made* in auction t if \mathbf{z}_t were such that $\gamma(\mathbf{z}_t) = 0$. Such a value of \mathbf{z}_t need not actually exist and, of course, bids could be recentered at any other value of \mathbf{z}_t . This additive separability implies that the effects of covariates on bids can be controlled for using a first-stage regression of bids on covariates. In particular, subtracting the estimate $\hat{\gamma}(\mathbf{z}_t)$ from each b_{it} then gives a sample of “homogenized” bids. Haile, Hong, and Shum (2003) use linear regression for this step.

A second issue is the likelihood that the number of bidders is endogenous. Several models of endogenous participation exist in the literature, and it is fairly straightforward to extend the approach to allow for binding reserve prices, entry fees, and costly signal acquisition (see Haile, Hong, and Shum (2003) for details). A more difficult case is that in which participation is partially determined by unobservable factors that also affect the distribution of bidders’ signals and valuations. Such a situation creates two problems. First, if the number of bidders is correlated with valuations, the expectation $v(x, x; n)$ would vary with n even in a pure private values model, or might fail to vary with n in a common values model. This is a familiar problem of an endogenous “treatment.” The second problem is the problem discussed in Section 3.4.2: unobserved heterogeneity threatens the identification of the distribution of $v(X, X; n)$.

²⁸See e.g., Baldwin, Marshall, and Richard (1997), Haile (2001), Haile and Tamer (2003), and Athey, Levin, and Seira (2004). Baldwin (1995), Athey and Levin (2001), and Haile (2001) have provided evidence that a common value element is introduced to scaled sales through misestimation of timber volumes by the Forest Service, or by resale opportunities.

Concerned that participation may be affected by unobserved heterogeneity,²⁹ Haile, Hong, and Shum (2003) propose an instrumental variables approach that combines an exclusion restriction with a monotonicity assumption: that, conditional on the observables, participation is strictly increasing in the unobservable.³⁰ Their approach requires an instrument, w , for the number of bidders, and they use counts of nearby sawmills and logging firms, following Haile (2001).

With this structure, bidders' expectations of the value of winning the auction can be written

$$v(x, x; n, w) = E[U_i | X_i = \max_{j \neq i} X_j = x; n, w].$$

These expectations can be estimated using the first-order condition

$$v(x_i, x_i; n, w) = b_i + \frac{G_{M|B}(b_i, b_i; n, w)}{g_{M|B}(b_i, b_i; n, w)}$$

using the methods described above but conditioning on both the number of bidders and the value of the instrument w . The resulting estimates $\hat{v}(x, x; n, w)$ are then pooled over the realizations of n and compared across values of w , enabling a comparison of auctions with different levels of participation due to exogenous factors. Their preliminary results suggest that common values may not be empirically important for some types of timber auctions.

4.3 Bidder Demand and Reserve Price Policy for Ascending Auctions

Reserve price policy has been one source of tension for the USFS timber sales program. Critics of Forest Service policies have suggested that reserve prices are systematically too low, resulting in a loss of revenues and a subsidy to the timber industry. However, the USFS has a range of mandated objectives, including forest management goals and provision of a steady supply of timber, which lead to a desire to ensure that most tracts it offers for sale are actually harvested. The Forest Service has sometimes argued that historical data are useless for guiding the trade-offs between these objectives and revenues:

“Studies indicate it is nearly impossible to use sale records to determine if marginal sales made in the past would have been purchased under a different [reserve price] structure.”
(U.S. Forest Service (1995))

Economic theory tells us that there is something to this view: the effects of reserve prices in a competitive environment are subtle and unlikely to be measured correctly with a descriptive

²⁹Evidence that unobserved heterogeneity affects participation in other USFS timber auctions is given in Haile (2001).

³⁰Recall Section 3.4.2. See Haile, Hong, and Shum (2003) for details. This approach is related to prior insights in Olley and Pakes (1996) and Campo, Perrigne, and Vuong (2003).

empirical analysis. However, this is exactly the kind of policy question that structural empirical work on auctions has the potential to inform. Haile and Tamer (2003) explore this possibility using data from auctions in Washington and Oregon.

One immediate concern in studying ascending auctions with a structural model is the suitability of the standard “button auction” model, discussed in Section 2. Because bidders typically call out bids rather than indicating their participation continuously, researchers may have quite noisy measures of the prices at which each bidder actually drops out of an auction—the “bids” in the button auction model. For example, even a bidder with a relatively high valuation may appear to exit at a low price if active bidding by his opponents pushes the price beyond his willingness to pay while he stands by. Further concerns arise from the use of minimum bid increments and the prevalence of “jump bidding.” Hence, in many applications there may be no observables that correspond to the equilibrium “bids” (exit prices) in the standard model.

A natural response to this problem would be to explore an alternative model better suited for interpreting observed bids. Capturing the free-form dynamics of a typical ascending auction, however, presents serious challenges. Haile and Tamer (2003) explore an alternative. They consider a symmetric independent private values environment and rely on only two behavioral assumptions. First, no bidder bids more than his valuation. Second, no bidder lets an opponent win at a price he would be willing to beat. More precisely, if we let $\Delta \geq 0$ denote the minimum bid increment,

Assumption HT1. $b_i \leq u_i \quad \forall i.$

Assumption HT2. $b^{n-1:n} \leq b^{n:n} + \Delta.$

The first assumption seems uncontroversial. The second reflects the defining characteristic of an ascending auction: the ability to observe and respond to the bids of competitors. These assumptions can be interpreted as axioms or as necessary conditions for equilibrium in a variety of models. While these assumptions permit bidding as in the dominant strategy equilibrium of the standard model, they do not require it: bids need not be equal to valuations or even monotonic in valuations, and the price need not equal the second highest valuation.

While relying only on these two behavioral assumptions has some appeal in terms of the robustness of the estimates to the details of the true data generating process, this flexibility comes at a price. In particular, under these assumptions, typically only bounds on $F_U(\cdot)$ are identified. Identification of these bounds follows from an application of (3.1). In an n -bidder auction, Assumption HT1 implies $b^{(i:n)} \leq u^{(i:n)}$ for all i . Hence,

$$G_B^{(i:n)}(u) \geq F_U^{(i:n)}(u) \quad \forall i, n, u. \tag{4.8}$$

Applying the monotonic transformation $\phi(\cdot; i, n)$ defined in (3.2) to both sides gives

$$\phi\left(G_B^{(i:n)}(u); i, n\right) \geq F_U(u) \quad \forall i, n, u.$$

For each u , this actually provides as many upper bounds on $F_U(u)$ as there are combinations of i and n . The most informative bound (i.e., the smallest upper bound) is obtained by taking the minimum at each value of u :

$$F_U^+(u) = \min_{i,n} \phi\left(G_B^{(i:n)}(u); i, n\right). \quad (4.9)$$

To obtain a lower bound, let $G_\Delta^{(n:n)}(\cdot) = \Pr(B^{(n:n)} + \Delta \leq b)$ and observe that Assumption HT2 implies

$$G_\Delta^{(n:n)}(u) \leq F_U^{(n-1:n)}(u) \quad \forall n, u.$$

Applying $\phi(\cdot; n-1, n)$ to both sides gives

$$\phi\left(G_\Delta^{(n:n)}(u); n-1, n\right) \leq F_U(u) \quad \forall n, u.$$

As before, the most informative bound can be constructed by taking the pointwise maximum:

$$F_U^-(u) = \max_n \phi\left(G_\Delta^{(n:n)}(u); n-1, n\right). \quad (4.10)$$

Haile and Tamer (2003) point out that estimates obtained using the standard ascending auction model to interpret bids need not lie within these bounds, even asymptotically. For example, if there is no minimum bid increment and the econometrician assumes only that the transaction price equals $u^{(n-1:n)}$, the natural nonparametric estimator $\hat{F}_U(\cdot)$ satisfies $\text{plim} \hat{F}_U(u) \leq F_U^-(u) \forall u$. On the other hand, if $B^{(n:n)} = B^{(n-1:n)}$ (as implied by the button auction model), the bounds $F_U^-(\cdot)$ and $F_U^+(\cdot)$ collapse to the true distribution $F_U(\cdot)$, providing point identification.

To estimate the bounds, one can substitute the empirical distributions

$$\hat{G}_B^{(i:n)}(b) = \frac{1}{T_n} \sum_{t=1}^T \mathbf{1}\{n_t = n, b^{(i:n_t)} \leq b\}$$

and

$$\hat{G}_\Delta^{(n:n)}(b) = \frac{1}{T_n} \sum_{t=1}^T \mathbf{1}\{n_t = n, b^{(n_t:n_t)} + \Delta_t \leq b\}$$

for the corresponding population CDFs in (4.9) and (4.10).³¹ It is straightforward to show consistency of these estimators, and Haile and Tamer (2003) suggest smooth approximations of the min

³¹A variation on this estimation approach is applicable when one assumes the full structure of the button auction model. See Haile and Tamer (2002, 2003) or Athey and Haile (2006).




Figure 4.1:

and max functions in (4.9) and (4.10) to reduce finite sample bias. They also demonstrate that the bootstrap can be used for inference.

Note that while we have suppressed the presence of auction heterogeneity in the discussion above, the entire discussion can be repeated conditioning on any vector of auction characteristics. Haile and Tamer (2003) condition on auction-specific covariates \mathbf{Z} , obtained from pre-sale assessments published by the Forest Service, using standard kernel smoothing methods. This leads to construction of bounds on the conditional distribution function $F_U(u; \mathbf{z})$ for all (u, \mathbf{z}) . They obtain fairly precise bounds on the effects of auction covariates on valuations and on the distribution of the idiosyncratic components of bidders' private values. Figure 4.1 shows the estimated bounds and 95 percent confidence bands for $F_U(\cdot; \mathbf{z})$ at the mean value of \mathbf{Z} .

Using bounds on $F_U(\cdot)$ to explore alternative reserve prices is straightforward. Assumptions HT1 and HT2 are sufficient to ensure that an ascending auction is revenue equivalent (Myerson (1981)) to a standard auction with the same reserve price. Since revenues are monotonic in $F_U(\cdot)$, one can simulate auction revenues using a standard (e.g., second-price sealed-bid) auction and valuations drawn from the bounds $F_U^-(\cdot)$ and $F_U^+(\cdot)$.

A more subtle problem is placing bounds on the profit-maximizing reserve price, defined by the equation³²

$$r^* = c_0 + \frac{1 - F_U(r^*)}{f_U(r^*)} \quad (4.11)$$

where c_0 is the seller's valuation (or marginal cost) of the good. Since nondegenerate bounds on $F_U(\cdot)$ place no restriction on its derivative $f_U(\cdot)$ at any given point, it is not obvious how to proceed. Note, however, that r^* is also the solution to the (seemingly unrelated) problem

$$\max_r \pi(r)$$

where

$$\pi(r) = (r - c_0)(1 - F_U(r)).$$

Since $F_U(r)$ must lie between $F^-(r)$ and $F^+(r)$, $\pi(r)$ must lie between


$$\pi_1(r) = (r - c_0)(1 - F_U^+(r)) \quad \text{and} \quad \pi_2(r) = (r - c_0)(1 - F_U^-(r)).$$

Under the additional assumption that $\pi(r)$ is strictly quasi-concave in r (ensuring a unique solution to (4.11)) the functions $\pi_1(\cdot)$ and $\pi_2(\cdot)$ bounding $\pi(\cdot)$ can be used to place bounds on r^* . The idea is easily seen in Figure 4.2. Since $\pi(\cdot)$ must lie between $\pi_1(\cdot)$ and $\pi_2(\cdot)$, it must reach a peak of at least π_1^* . Such a peak cannot be reached outside the interval $[r^-, r^+]$, although prices arbitrarily close to either of these endpoints could be the true optimum r^* . Haile and Tamer (2003) provide a formal arguments and propose consistent estimators of the resulting bounds on r^* .

Haile and Tamer (2003) apply this approach to USFS timber auctions, using a range of values for c_0 based on USFS estimates. While the estimated bounds on the optimal reserve price are quite wide, they are tight enough to determine that the actual reserve prices used by the Forest Service are likely well below the optimal levels. The important question, however, concerns not the difference between actual and optimal reserve prices, but the differences in revenues and probability of a sale under alternative policies. Their results indicate that, except when the actual reserve price is below c_0 , the revenue gains from raising the reserve prices would actually be quite small. This suggests that getting the reserve price "right" may be relatively unimportant in these auctions, at least as long as it meets the minimal requirement of exceeding the opportunity cost c_0 .

³²See, e.g., Riley and Samuelson (1981) for a derivation. Haile and Tamer (2003) show that r^* remains optimal in their incomplete model if Assumptions 1 and 2 are interpreted as a partial characterization of equilibrium behavior in the true but unspecified auction mechanism.

Figure 4.2



4.4 Comparing Auction Formats: Entry, Revenue, and Competitiveness

The relative performance of ascending and first-price auctions is a central question in auction design. The Revenue Equivalence Theorem (Vickrey (1961)) implies that if bidders are risk-neutral, have independent and identically distributed values, and bid competitively, the two auction formats yield the same winner, the same expected revenue, and even the same bidder participation. However, if these assumptions are relaxed, auction choice becomes relevant, and the relative performance of alternative mechanisms depends both on details of the market and the objective (e.g., revenues or efficiency).

Athey, Levin, and Seira (2004) (ALS) study the impact of auction format using data from U.S. Forest Service timber program, which has historically used both ascending and sealed-bid auctions, sometimes randomizing. ALS maintain the independent private values assumption but focus on two potentially important departures from Vickrey's assumptions: bidder asymmetry and collusion.

Motivated by institutional features of the timber industry, ALS distinguish between two types of bidders, “mills” (firms with manufacturing capability) and “loggers.” They highlight differences in participation patterns across auction formats that arise in the presence of bidder asymmetry, and they further explore the common hypothesis that collusion may be more likely in ascending auctions.

ALS analyze two distinct regions of the Forest Service during the 1980s. Auction format was determined in different ways in the two regions; however, in each region, ALS argue that conditional on the observable features of the tracts, auction format should be uncorrelated with unobserved sale characteristics. In the Northern region (comprised of Idaho and Montana), during the period of study the selection of auction format was explicitly randomized for a subset of the tracts in the sample, where assignment to the randomization pool was based on observable characteristics in most cases. In California, the volume of the sale was the primary determinant of auction format. ALS compare outcomes across auction format in each region separately, and they control flexibly for the factors that influence the selection of auction format in the empirical analysis.

For the Northern region, ALS establish that controlling for auction characteristics, relative to ascending auctions, sealed bid auctions have the following features: (1) 10% more logger participants, (2) .4% more mill participants, (3) the fraction of sales won by loggers is 3.4% higher, and (4) 14% higher revenue.³³ The findings in California are similar, with the notable exception that revenue is almost identical in the first-price and ascending auctions. The large revenue difference in the Northern region is somewhat surprising, and it raises the question of whether a competitive model of equilibrium bidding can rationalize the gap.

ALS develop a model with several key features to reflect these facts. First, mills and loggers draw their valuations from different distributions (asymmetric independent private values). Second, a firm must choose whether to pay a cost to learn its valuation for each tract to be auctioned. Bidders make this decision based on expected profits, which differ by bidder type and auction format. Third, two models of conduct are considered: a baseline model in which bidders behave competitively, and a collusive alternative.

The timeline of the model is as follows: First, the auctioneer announces the auction format. Next, each firm i chooses whether to pay cost K to learn its valuation. Firms cannot bid unless they pay this cost. Third, bidder i draws private value U_i , where if i is a logger, $U_i \sim F_{U_L}(\cdot)$, and otherwise $U_i \sim F_{U_M}(\cdot)$. Fourth, participating bidders learn who else has made the information acquisition investment before bidding. Fifth, each bidder who acquired the information chooses

³³In first-price auctions, the Forest Service records the identities of all bidders in each auction. In ascending auctions, all bidders who wish to participate must register in advance and place a deposit, and the Forest Service also records the identities of these bidders.

a bid. For the purposes of the theory, ALS assume that $F_{U_M}(\cdot)$ dominates $F_{U_L}(\cdot)$ according to a conditional stochastic dominance order; this assumption is confirmed (not imposed) in the structural estimation.

Now consider the information acquisition equilibrium in this model, which may be in mixed or pure strategies. As in standard entry models in industrial organization, there will typically be multiple equilibria in the information acquisition stage. To narrow the range of possibilities, ALS assume that all bidders of the same “type” (that is, mills or loggers) use the same entry strategies. They refer to such equilibria as *type-symmetric*. ALS provide a sufficient condition under which there is a unique type-symmetric equilibrium, and in this equilibrium, either loggers never enter or mills enter with probability one. They verify that this condition is satisfied given the parameter values consistent with their bidding data. For simplicity, we assume for the rest of this discussion that the condition holds, and that parameters are such that loggers enter with positive probability (so that mills must enter with probability one).

Relative to ascending auctions, first-price auctions lead to inefficient allocations at the bidding stage. As established by Maskin and Riley (2000), when the set of participants is fixed, “strong” bidders (mills) bid less aggressively, because when loggers and mills bid against one another, the competition for any logger is tougher than that for any mill. Anticipating this, in first-price auctions (relative to ascending auctions), (i) mill entry is the same, but the winning bidder is less likely to be a mill, while (ii) loggers enter (weakly) more often and the winning bidder is more likely to be a logger. In addition, (iii) the ascending auction equilibrium is socially efficient (subject to the type-symmetry constraint), while first-price auctions yield inefficient entry and bidding. Finally, the model does not generate unambiguous predictions about revenue. Since both participation and bidding vary with the auction format, a revenue comparison depends on all the primitives of the model (value distributions and entry costs). A goal of ALS is to estimate these primitives in order to assess the revenue gain (if any) from sealed bidding, as well as the efficiency distortion that sealed bidding induces in both entry and bidding.

Turning to the empirical work, ALS show that conditional on observable characteristics of an auction, bids are positively correlated within a first-price auction. To account for this correlation, ALS select a model of independent private values with unobserved heterogeneity. As described above, this model is not identified in data from ascending auctions; thus, ALS focus their structural estimation on first-price auctions and use their estimates of the underlying parameters of the model (value distributions and information acquisition costs K) to make out-of-sample predictions about outcomes at ascending auctions, predictions that can be compared to actual outcomes.³⁴

³⁴In addition to the issue of non-identification of the ascending auction model, there is another advantage to the approach of first estimating primitives from data on first-price auctions and second using these primitives to predict

Due to limitations in sample size, ALS use a parametric model to fit the bid distributions, where the bids are drawn from a Weibull distribution with Gamma-distributed heterogeneity. In this model, the hazard rate of the bid distributions is a multiplicative function of the unobserved heterogeneity. This allows the implied difference between equilibrium bids and values to vary across auctions.³⁵ Once the bid distributions have been estimated, they follow the approaches of Guerre, Perrigne, and Vuong (2000) and Krasnokutskaya (2004) to construct the value distributions implied by the estimated bid distributions. No parametric assumptions are imposed in this second step.

Now consider the identification of information acquisition (entry) costs.³⁶ Let $\tilde{\mathcal{N}}$ be a random variable whose realization is \mathcal{N} , the number of participants. In any entry model that generates variation in $\tilde{\mathcal{N}}$ that is independent of valuations, our basic identification results for first-price auctions imply that a bidder's ex ante gross expected profit $\Pi_i(\mathcal{N})$ from entering the auction is identified. In particular,

$$\Pi_i(\mathcal{N}) = E_{U_i} [(U_i - \beta_i(U_i; \mathcal{N})) G_{M_i|B_i}(\beta_i(U_i; \mathcal{N}) | \beta_i(U_i; \mathcal{N}); \mathcal{N})]$$

with the right-hand side determined by the observed bid distribution and the first-order conditions for equilibrium bidding. Identification of $\Pi_i(\mathcal{N})$ requires no assumptions about the nature of the signal acquisition equilibrium (or equilibrium selection). Estimates of $\Pi_i(\mathcal{N})$ can then be used to calculate all equilibria of an entry game for given entry costs.

ALS also observe that for any firms that are indifferent about acquiring a signal (as in a mixed strategy equilibrium), the expected profit from entry must be zero. Thus, entry costs are identified by $\Pi_i(\mathcal{N})$ and the distribution of $\tilde{\mathcal{N}}$, which is directly observable. In particular, for any firm i that is indifferent about acquiring a signal, entry costs must be equal to

$$K = \bar{\Pi}_i = \sum_{\mathcal{N}:i \in \mathcal{N}} \Pr(\tilde{\mathcal{N}} = \mathcal{N} | i \in \tilde{\mathcal{N}}) \Pi_i(\mathcal{N}).$$

ascending auction outcomes. The empirical anomaly ALS wish to explain is that revenue is “too high” in the first-price auction relative to the ascending auction in the Northern region. By avoiding the use of a conduct assumption in ascending auctions for the estimation, it is possible to consider several alternative conduct assumptions for open auctions, including a model of collusion. In contrast, the conduct assumption of competition is imposed in the first-price auction. A competitive model of first-price auctions leads to implied valuations that are lower than with a collusive model, since in a collusive model there is a large gap between bids and valuations. In turn, it will generate a lower out-of-sample prediction about revenue in ascending auctions. This gives the best chance for the competitive model to rationalize the revenue gap.

³⁵ALS assess the goodness of fit of this model using a conditional Kolmogorov test due to Andrews (1997), showing that they cannot reject the Gamma-Weibull model against nonparametric alternatives.

³⁶Li (2003) considers parametric estimation of value distributions in a symmetric entry model. Hendricks, Pinkse and Porter (2003) have also considered a variation on this model in a common values setting in which bidders choose whether to invest in a signal based on noisier (in a precise sense) preliminary estimates of their valuations. These papers do not provide an analysis of identification and estimation of entry costs, or of entry equilibria.

Thus, in contrast to much of the empirical industrial organization literature on entry,³⁷ which draws inferences solely from entry decisions, the level of entry costs can be inferred. Hence it is possible to conduct counterfactual simulations about changes in these costs on the competitiveness of markets and bidder rents.

With estimates of the value distributions and of K (both of which are allowed to vary with auction covariates \mathbf{Z}), it is possible to predict outcomes in an ascending auction. In particular, for each ascending auction tract, ALS calculate the information acquisition equilibrium (which generates a distribution over information acquisition behavior) and the expected revenue from the ascending auction. They then compare the actual outcomes in the ascending auction to those predicted by the model.

For California sales, ALS find that the competitive model can rationalize the bidding data: they cannot reject the hypothesis that the bids in ascending auctions are equal to those predicted using the primitives estimated from the first-price auction data. This is a striking finding, since the predictions are out-of-sample in two ways: first, the predictions are for tracts that were not used in the original estimation, and second, the predictions are for a different game. The actual revenue per unit of volume is \$119, while the structural model predicts \$115 with a standard error of \$9.

For Northern sales, the structural model predicts that if behavior in ascending auctions is competitive, sealed bid auctions raise 8.4% more per unit of volume than ascending auctions. In contrast, the actual difference is 12%.³⁸ ALS can reject the hypothesis that the competitive model fits the ascending auction data. Thus motivated, they turn to consider an alternative, namely that ascending auctions are less competitive.³⁹ Bidder collusion has been a long-standing concern in timber auctions; the prevailing view is that ascending auctions are more prone to collusion because bidders are face-to-face and can respond immediately to opponents' behavior.⁴⁰ ALS extend the theory to allow for collusion by mills at ascending auctions. For simplicity, ALS consider a model where in ascending auctions, mills collude perfectly in bidding after independently gathering

³⁷See, e.g., Bresnahan and Reiss (1987) or Berry (1992).

³⁸This is the gap without adjusting for covariates; the model includes the effect of covariates for the ascending auctions when it makes predictions.

³⁹ALS argue that several factors that seem plausible in the context of timber auctions, but are omitted from their model, such as common values and bidder risk-aversion, are not good candidates to rationalize the findings.

⁴⁰There is a small but growing empirical literature on collusion at auctions. A variety of approaches have been suggested to assess whether bidding data are consistent with models of competition or collusion. Examples include Porter and Zona (1993), Porter and Zona (1999), Bajari and Ye (2003), Pesendorfer (2000); see Bajari and Summers (2002) for a survey. Baldwin, Marshall, and Richard (1997) also analyze collusion in U.S. Forest Service timber auctions using data from open auctions; they argue that collusion provides a better fit than competition. Some of these approaches require prior knowledge about the existence and structure of a cartel, while others interpret departures from symmetric bidding behavior as evidence of collusion. The method of ALS differs in that they use behavior under one set of auction rules as a benchmark from which to evaluate the competitiveness of behavior under an alternative set of rules.

information. A convenient feature of this model of collusion is that it does not affect predictions about entry and allocation; the only effect of collusion on outcomes is to increase the predicted revenue difference between auction formats.

ALS find that this model of collusion can easily depress ascending auction prices below the observed levels. The actual ascending auction data can be rationalized with a model where mills collude in 25% of the ascending auctions.

Turning to the welfare differences between ascending and first-price auctions, ALS find that for a fixed set of participants (that is, ignoring the predicted differences in participation between ascending and first-price auctions), the calibrated model predicts relatively small discrepancies between sealed bid auctions and competitive ascending auctions. Sealed bid auctions raise more revenue and distort the allocation away from efficiency and in favor of loggers, but the effects are small (less than 1%). The differences are larger when accounting for equilibrium entry behavior: ALS predict that sealed bidding increases revenue by roughly 5% relative to a competitive ascending auction, at minimal cost to social surplus. Strikingly, even a mild degree of collusion by the mills at ascending auctions—the behavioral assumption most consistent with the observed outcomes in the Northern forests—results in much larger revenue differences (on the order of 20%). This suggests that susceptibility to collusion is an important consideration in the choice of auction format.

4.5 Dynamics in Procurement Auctions

Virtually all of the structural empirical literature has examined models of auctions as isolated games—one-time interactions between the bidders. This is obviously false in most applications. Moreover, significant intertemporal considerations seem likely to arise in many auction markets where the distribution of valuations changes as a function of observable auction outcomes.⁴¹ One way this could arise is through learning-by-doing, where the winning firm expects to have higher valuations or lower costs in the future. In other settings, capacity constraints or diseconomies of scale may be important, causing a winning bidder to expect to have a lower valuation (or higher cost) in the future. In both cases, rational bidders will consider the effect of winning an auction on their continuation values when choosing participation and bidding strategies.

Jofre-Bonet and Pesendorfer (2003) consider this type of dynamic model in what, to our knowledge, is the first structural model of dynamic auctions with long-lived bidders and valuations that change as a function of auction outcomes. Despite the complexity of the dynamic problem, Jofre-Bonet and Pesendorfer (2003) establish that if the firms' discount factors are known to the

⁴¹There are other potentially interesting sources of dynamics. Even in a stationary environment, dynamic considerations arise if firms engage in collusion. In addition, bidders' valuation distributions may change over time in a way that is private information to each bidder. This can create dynamic links in bidder strategies.

econometrician, the distributions of bidder valuations are identified under very general conditions. The authors cleverly show how to combine the insights of Guerre, Perrigne, and Vuong (2000), outlined in Section 3.2, above with those developed by Hotz and Miller (2003) for single-agent dynamic discrete choice models.

A simplified version⁴² of Jofre-Bonet and Pesendorfer (2003)’s model can be described as follows. Firms participate in a series of auctions taking place over an infinite horizon. They discount the future at rate δ , assumed to be known to the econometrician. In each period t , an item is sold by first-price auction to one of n bidders. For simplicity we ignore reserve prices and assume that all objects to be auctioned have the same observable characteristics. Capacities directly affect the distribution of bidder valuations. Conditional on capacities, bidder valuations are independent across bidders and over time. Let $c_{i,t}$ be bidder i ’s publicly observable capacity in period t , and let $F_U(\cdot|c_{i,t})$ be the conditional distribution of bidder i ’s valuation in period t .

The transition function for bidder capacities can be described as follows. Let k be the identity of the winning bidder in period t and let \mathbf{c}_t be the vector of bidder capacities in period t . Then

$$c_{i,t+1} = \omega_i(\mathbf{c}_t, k).$$

The authors focus on Markov-perfect equilibrium with the vector of all bidders’ capacities as the state variable. Bidders’ equilibrium strategies depart from those in a static equilibrium because bidders anticipate that the identity of today’s winner will affect future valuations of all bidders through changes in capacities. Hence, the distribution of outcomes in future auctions is affected by today’s outcome too. Since bidders are symmetric except for differences induced by different capacities, Jofre-Bonet and Pesendorfer (2003) consider exchangeable strategies, which can be written $\beta(u_{i,t}, \mathbf{c}_t)$.⁴³

The formal analysis begins by using dynamic programming to represent bidder payoffs. Suppress \mathcal{N} in the notation. For a given vector of bidder capacities \mathbf{c} , let $G_{M_i}(\cdot|\mathbf{c})$ be the equilibrium distribution of the maximum opponent bid for bidder i . Let $\omega(\mathbf{c}, k) = (\omega_1(\mathbf{c}, k), \dots, \omega_n(\mathbf{c}, k))$. When opponent strategies are held fixed, the interim expected discounted sum of future profits for

⁴²Their paper studies procurement auctions. In order to maintain consistency with the rest of this paper, we state the problem where bidders are buyers. We also make a further simplification, studying only long-lived bidders; they also consider “fringe bidders” who bid rarely and are myopic.

⁴³Jofre-Bonet and Pesendorfer (2003) show that an equilibrium exists within the parametric framework they use for estimation, and they sketch a proof for the general case. Uniqueness is an open question.

bidder i is given by

$$W_i(u_i, \mathbf{c}) = \max_{b_i} \left\{ (u_i - b_i)G_{M_i}(b_i|\mathbf{c}) + \delta \sum_{j=1}^n \Pr(j \text{ wins} | b_i, \mathbf{c}) \int_{u'_i} W_i(u'_i, \omega(\mathbf{c}, j)) f_{U_i}(u'_i | \omega_i(\mathbf{c}, j)) du'_i \right\}$$

where, given current capacities, the second term sums over the possible identities of the winner to form an expectation of the continuation value to player i . The *ex ante* value function is then written

$$V_i(\mathbf{c}) = \int W_i(u_i, \mathbf{c}) f_{U_i}(u_i|\mathbf{c}) du_i$$

or, substituting,

$$V_i(\mathbf{c}) = \int \left\{ \max_{b_i} \left\{ (u_i - b_i)G_{M_i}(b_i|\mathbf{c}) + \delta V_i(\omega(\mathbf{c}, i)) + \delta \sum_{j \neq i} \Pr(j \text{ wins} | b_i, \mathbf{c}) [V_i(\omega(\mathbf{c}, j)) - V_i(\omega(\mathbf{c}, i))] \right\} \right\} f_{U_i}(u_i|\mathbf{c}) du_i.$$

The second step of the analysis entails solving for the *ex ante* value functions in terms of observables, which requires a novel extension of the two-step indirect approach proposed by Guerre, Perrigne, and Vuong (2000). We begin by stating a bidder optimization problem in a given auction, letting $\bar{b}_j(\mathbf{c})$ and $\underline{b}(\mathbf{c})$ denote the largest and smallest equilibrium bid when capacities are \mathbf{c} , respectively. Bidder i solves:

$$\max_{b_i} \left\{ (u_i - b_i)G_{M_i}(b_i|\mathbf{c}) + \delta V_i(\omega(\mathbf{c}, i)) + \delta \sum_{j \neq i} \left(\int_{b_i}^{\bar{b}_j(\mathbf{c})} \prod_{k \neq i, j} G_{B_k}(b_j|\mathbf{c}) g_{B_j}(b_j|\mathbf{c}) db_j \right) [V_i(\omega(\mathbf{c}, j)) - V_i(\omega(\mathbf{c}, i))] \right\}$$

which has first-order condition

$$u_i = b_i + \frac{G_{M_i}(b_i|\mathbf{c})}{g_{M_i}(b_i|\mathbf{c})} + \delta \sum_{j \neq i} \frac{G_{M_i}(b_i|\mathbf{c})}{g_{M_i}(b_i|\mathbf{c})} \frac{g_{B_j}(b_i|\mathbf{c})}{G_{B_j}(b_i|\mathbf{c})} (V_i(\omega(\mathbf{c}, j)) - V_i(\omega(\mathbf{c}, i))). \quad (4.12)$$

If the inverse bidding strategy in(4.12) is substituted into the *ex ante* value function, a change of variables yields

$$V_i(\mathbf{c}) = \int_{\underline{b}(\mathbf{c})}^{\bar{b}(\mathbf{c})} \frac{G_{M_i}(b_i|\mathbf{c})}{g_{M_i}(b_i|\mathbf{c})} G_{M_i}(b_i|\mathbf{c}) dG_{B_i}(b_i|\mathbf{c}) + \delta \sum_{j \neq i} V_i(\omega(\mathbf{c}, j)) \left\{ \int_{b_i}^{\bar{b}_j(\mathbf{c})} \prod_{k \neq i, j} G_{B_k}(b_j|\mathbf{c}) g_{B_j}(b_j|\mathbf{c}) db_j + \int_{\underline{b}(\mathbf{c})}^{\bar{b}(\mathbf{c})} \frac{G_{M_i}(b_i|\mathbf{c})}{g_{M_i}(b_i|\mathbf{c})} \frac{g_{B_j}(b_i|\mathbf{c})}{G_{B_j}(b_i|\mathbf{c})} G_{M_i}(b_i|\mathbf{c}) dG_{B_i}(b_i|\mathbf{c}) \right\}.$$

This expresses each $V_i(\mathbf{c})$ as a linear function of $V_i(\cdot)$ evaluated at other capacity vectors. This linear function has coefficients that depend on the observable bid distributions, so that the equations can be solved explicitly and value functions can be expressed as functions of observables.

Given *ex ante* value functions, the first-order condition (4.12) can be used to establish identification of the distributions $F_U(u_i|\mathbf{c})$. Since we have assumed that the discount factor δ is known, $F_{U_i}(\cdot|\mathbf{c})$ is identified from the distributions of observed bids (conditional on capacities).

The application of these techniques in Jofre-Bonet and Pesendorfer (2003) considers California highway construction contracts. They use a parametric estimation approach for parsimony, since there are a number of covariates to be included. They solve for the value functions by discretizing the set of possible capacities. Then, calculating the value functions entails solving a system of linear equations. They use a quadratic approximation to the value function in order to further simplify the estimation.

Jofre-Bonet and Pesendorfer (2003) assess the importance of private information, capacity constraints, and the inefficiencies that arise due to the asymmetries induced by capacity differences among bidders under the assumption of forward-looking equilibrium behavior. As usual, it is impossible to *test* whether bidders are forward looking when the discount factor δ is assumed known. However, their results suggest that there is significant asymmetry in bidding behavior introduced by variation in capacities over time. Bidders with greater excess capacity, and thus higher valuations on average, bid less aggressively. Asymmetric bidding leads to inefficient allocations, since the highest-valuation bidder may not win the auction. In addition, their estimates imply that the average markup (that is, the average gap between bid and value) is equal to 40%, half of which they attribute to bidder recognition of the option value to losing a contract today. The interpretation is that bidders recognize that they may use their limited capacity by winning another contract in the future.

4.6 Mechanism Choice for Treasury bill Auctions

Most of the empirical literature on auctions focuses on the case of single-unit auctions. Recently, however, auctions of multiple units of identical goods (“multi-unit auctions”) have gained increasing attention, in part because of the role they play in important public and private activities.⁴⁴ For example, multi-unit auctions have been used in restructured electricity markets to allocate electric power generation to different plants (see, e.g., Wolfram (1998), Borenstein, Bushnell, and Wolak (2002), or Wolak (2003)). Multi-unit auctions have long been used to allocate Treasury bills in

⁴⁴See Cantillon and Pesendorfer (2003) for a recent analysis of auctions of contracts that are not perfect substitutes, but may be either substitutes or complements.

the U.S. and elsewhere. Indeed, at least since Friedman (1960), economists debated the optimal design of Treasury bill auctions. This question is potentially relevant to the design of markets for other types of securities as well. A number of complexities arise in the multi-unit auction case, including nonlinearities in cost functions and incentives to exercise market power by withholding demand. Thus, even when the choice is limited to standard *discriminatory* (“pay your bid”) or *uniform-price* auctions, the revenue maximizing (or cost minimizing) design is typically ambiguous, depending on the primitives of the problem.

Here, we focus on the case of treasury auctions. In these auctions, a large number of identical securities is sold in a mechanism in which each bidder i offers a downward sloping schedule of price-quantity combinations (b_{ij}, q_j) , where b_{ij} is the price he is willing to pay for his q_j th unit. This is typically referred to as i ’s “demand curve” even though this is a strategically chosen schedule that need not correspond to the usual notion of a demand curve in economics. The discriminatory format is the most common auction form in practice, although recently the U.S. adopted uniform-price auctions after conducting an experiment to evaluate alternative formats. In the discriminatory auction, each bidder who offers more than the market clearing bid for a unit receives that unit at the price he offered. Thus, different prices are paid for different units of the same security, even for the same bidder. In a uniform-price auction, the market clearing price is paid on all units sold. In addition to U.S. treasury bill auctions, electricity auctions are typically uniform price, and some firms have used uniform price auctions in initial public offerings.⁴⁵

Of course, the auction format has an important effect on bidder strategies. Bidding one’s true marginal valuation for each unit is not an equilibrium in either auction. In a discriminatory auction this is obvious, since such “truthful” bidding would lead to zero surplus to any bidder. In a uniform-price auction, bidders also have an incentive to shade bids below marginal valuations, since a bidder’s own bid on a marginal unit may set the price for all inframarginal units.⁴⁶ The revenue ranking of the two mechanisms depends on the true distributions of bidder valuations (Ausubel and Cramton (2002)).

A working paper by Hortaçsu (2002) proposes to explore this question empirically by estimating the underlying distributions. He does this by developing a structural model of a discriminatory auction based on the “share auction” analysis of Wilson (1979).⁴⁷ Suppose that a fixed quantity,

⁴⁵In the finance literature, these are often referred to as “Dutch auctions,” conflicting with economists’ use of this term for descending price single-unit auctions.

⁴⁶Identification and estimation in uniform-price auctions has been explored recently by Wolak (2003) and Kastl (2005).

⁴⁷Unfortunately, the theory of multi-unit auctions is not as complete as the theory of first-price auctions and ascending auctions. Swinkels, Jackson, Simon, and Zame (2002) and Jackson and Swinkels (forthcoming) establish existence of equilibrium in mixed strategies, but existence of pure strategy Nash equilibria in monotone strategies has been established for only a limited class of models, and there are examples with multiple equilibria (e.g. Back

Q , of securities is to be offered. Each bidder i observes a signal X_i (possibly multi-dimensional) that determines his marginal valuations for all possible units of the good.⁴⁸ In particular his marginal valuation for a y th unit of the good is given by $v_i(y; x_i)$ for all $y \leq Q$.

Each bidder i 's strategy specifies, for each possible value of his signal x_i and each potential price b , a demand function $q_i(b) = \varphi_i(b; x_i)$ giving the quantity demanded at price b . Hence, "actions" in this model are functions, and a strategy maps the realization of a signal into a demand function. Given a set of demand functions $q_1(\cdot), \dots, q_n(\cdot)$, the market clearing price p^c is set by the seller to equate supply and demand, i.e.,

$$Q = \sum_i q_i(p^c).$$

Viewed differently, the market clearing price is that at which bidder i 's demand curve intersects his *residual supply curve*

$$Q_{R_i}(p) = Q - \sum_{j \neq i} q_j(p).$$

Of course, the residual supply curve is stochastic and determined by the equilibrium bidding strategies of his competitors, as well as the realizations of their signals. Let

$$G_i(b, y) = \Pr \left(y \leq Q - \sum_{j \neq i} \varphi_j(b; X_j) \right) \quad (4.13)$$

so that $G_i(b, y)$ is the probability that, given equilibrium bidding by i 's opponents, the market clearing price falls below b if i himself demands quantity y at price b .

When his type is x_i , bidder i 's optimal strategy solves the problem

$$\max_{q_i(\cdot)} \int_0^\infty \left(\int_0^{q_i(p^c)} (v_i(y, x_i) - q_i^{-1}(y)) dy \right) \frac{\partial G_i(p^c, q_i(p^c))}{\partial p^c} dp^c$$

The optimal bidding strategy can then be characterized by the Euler-Lagrange necessary condition

$$v_i(\varphi(b; x_i), x_i) = b + \frac{G_i(b, \varphi(b; x_i))}{\frac{\partial}{\partial b} G_i(b, \varphi(b; x_i))}.$$

This is similar to the first-order condition (3.4) used by Guerre, Perrigne, and Vuong (2000) to establish identification of the first-price sealed-bid auction, which can also be thought of as a

and Zender (1993)). Thus, most existing econometric approaches to these auctions require an explicit assumption about existence of equilibrium. In some cases, empirically verifiable properties of equilibrium bid distributions will guarantee that the assumptions of the econometric model are consistent with equilibrium behavior.

⁴⁸Thus, this is a private values model. For parametric structural approaches to common values models, see the recent studies by Février, Préget, and Visser (2002) and Armantier and Sbaï (2003). Common values models may be appropriate for many securities auctions, although this is ultimately an empirical question—one for which testing approaches have not been developed. Hortaçsu (2002) motivates his assumption of private values using institutional details of the Turkish treasury bill auctions he studies.

single-unit discriminatory auction. Because the demand functions $q_j(b) = \varphi_i(b; x_j)$ are observable to the econometrician, $G_i(b, y)$ is identified from equation (4.13). For a given quantity y demanded at price b by bidder i , we can then express the marginal valuation of bidder i at quantity y as

$$v_i(y, x_i) = b + \frac{G_i(b, y)}{\frac{\partial}{\partial b} G_i(b, y)}.$$

Thus, the distributions of each $v_i(y, X_i)$ are identified; in particular, if for each quantity y we define B_i^y to be a (observable) random variable $\varphi_i^{-1}(y; X_i)$, $\Pr(v_i(y, X_i) \leq \tilde{v})$

$$\Pr(v_i(y, X_i) \leq \tilde{v}) = \Pr\left(B_i^y + \frac{G_i(B_i^y, y)}{\frac{\partial}{\partial b} G_i(b, y)|_{b=B_i^y}} \leq \tilde{v}\right) \quad \forall \tilde{v}.$$

The distribution of the marginal valuation functions can then be used to do counterfactual experiments and evaluate policies.

Hortaçsu (2002) describes several different approaches to estimation. He also proposes a method for placing an upper bound on the revenue that could be obtained with a uniform price auction when one knows the underlying distributions of marginal valuations. This approach circumvents the difficult problem of solving for the equilibrium of the uniform price auction by exploiting the fact that a bidder would not submit a bid function offering more than her marginal valuation for each unit. The revenue that would be obtained if bidders bid their marginal valuations for each unit in a uniform auction can then serve as an upper bound on the equilibrium revenue. Preliminary evidence shows that in the Turkish treasury bill auctions Hortaçsu (2002) studies, this upper bound on the uniform-price auction revenue is still lower than that from the actual discriminatory auction.

A number of authors have considered variants on the model of Hortaçsu (2002) in order to account for differences between theory and practice. In many applications, the demand functions bidders submit are restricted to be step function. Recently, Wolak (2004), McAdams (2005), and Kastl (2005) have explored empirical models explicitly accounting for this discreteness.

5 Conclusion

We have discussed some of the key insights behind the recent methodological advances in the econometric analysis of auction data. In addition, we presented six extensions and empirical applications illustrating the range of economic questions now being addressed using these methods. We believe that the empirical analysis of auctions will remain a fruitful area for some time to come. The popularity of online auctions has generated a wealth of new data, and it has suggested new extensions to the theory and econometric methods to account for the institutional features of that environment (e.g. Song (2003), Song (2004)). Policy debates about natural resource auctions continue to arise

(e.g. oil, timber), and the subtleties of procurement auctions have led to interesting theoretical extensions of standard models and new questions that can be asked of bidding data (e.g., Athey and Levin (2001), Bajari, Houghton, and Tadelis (2004), Asker and Cantillon (2004), Asker and Cantillon (2005)). The choice of auction design for treasury bills has changed in recent years, partly in response to academic arguments, although the issue is far from settled. Recent applications of multi-unit auctions to allocate goods (e.g., radio spectrum) that may be substitutes or complements have inspired new theory, experimental work, and empirical work (e.g., Cantillon and Pesendorfer (2003)).

References

- ANDREWS, D. (1997): “A Conditional Kolmogorov Test,” *Econometrica*, 65, 1097–1128.
- ARNOLD, B., N. BALAKRISHNAN, AND H. NAGARAJA (1992): *A First Course in Order Statistics*. Wiley & Sons, New York.
- ASKER, J., AND E. CANTILLON (2004): “Properties of Scoring Auctions,” Working Paper (New York University).
- (2005): “Optimal Procurement when both Price and Quality Matter,” Working Paper (New York University).
- ATHEY, S., AND P. A. HAILE (2002): “Identification of Standard Auction Models,” *Econometrica*, 70, 2107–2140.
- (2006): “Nonparametric Approaches to Auctions,” in *Handbook of Econometrics*, ed. by J. J. Heckman, and E. Leamer, vol. 6. Elsevier.
- ATHEY, S., AND J. LEVIN (2001): “Information and Competition in U.S. Forest Service Timber Auctions,” *Journal of Political Economy*, 109, 375–417.
- ATHEY, S., J. LEVIN, AND E. SEIRA (2004): “Comparing Open and Sealed Bid Auctions: Theory and Evidence from Timber Auctions,” Working Paper (Stanford University).
- AUSUBEL, L. M., AND P. CRAMTON (2002): “Demand Reduction and Inefficiency in Multi-Unit Auctions,” Working Paper 96–07 (University of Maryland).
- BACK, K., AND J. F. ZENDER (1993): “Auctions of Divisible Goods: On the Rationale for the U.S. Treasury Experiment,” *Review of Financial Studies*, 6, 733–764.

- BAJARI, P., S. HOUGHTON, AND S. TADELIS (2004): “Bidding for Incomplete Contracts,” Working Paper (Duke University).
- BAJARI, P., AND G. SUMMERS (2002): “Detecting Collusion in Procurement Auctions,” *Antitrust Law Journal*, 70, 143–170.
- BAJARI, P., AND L. YE (2003): “Deciding Between Competition and Collusion,” *Review of Economics and Statistics*, 85, 971–989.
- BALDWIN, L. H. (1995): “Risk Aversion in Forest Service Timber Auctions,” Working Paper (RAND Corporation).
- BALDWIN, L. H., R. C. MARSHALL, AND J.-F. RICHARD (1997): “Bidder Collusion in U.S. Forest Service Timber Sales,” *Journal of Political Economy*, 105, 657–699.
- BIKHCHANDANI, S., P. A. HAILE, AND J. G. RILEY (2002): “Symmetric Separating Equilibria in English Auctions,” *Games and Economic Behavior*, 38, 19–27.
- BORENSTEIN, S., J. B. BUSHNELL, AND F. A. WOLAK (2002): “Measuring Market Inefficiencies in California’s Restructured Wholesale Electricity Market,” *American Economic Review*, 92, 1376–1405.
- BRANNMAN, L., J. D. KLEIN, AND L. W. WEISS (1987): “The Price Effects of Increased Competition in Auction Markets,” *Review of Economics and Statistics*, 69, 24–32.
- CAMPO, S., E. GUERRE, I. M. PERRIGNE, AND Q. VUONG (2002): “Semiparametric Estimation of First-Price Auctions with Risk Averse Bidders,” working paper (Pennsylvania State University).
- CAMPO, S., I. M. PERRIGNE, AND Q. VUONG (2003): “Asymmetry in First-Price Auctions with Affiliated Private Values,” *Journal of Applied Econometrics*, 18, 197–207.
- CANTILLON, E., AND M. PESENDORFER (2003): “Combination Bidding in Multi-Unit Auctions,” Working Paper (London School of Economics and Political Science).
- CAPEN, E. C., R. V. CLAPP, AND W. M. CAMPBELL (1971): “Competitive Bidding in High-Risk Situations,” *Journal of Petroleum Technology*, 23, 641–653.
- FRIEDMAN, M. (1960): *A Program for Monetary Stability*. Fordham University Press, New York.
- GUERRE, E., I. M. PERRIGNE, AND Q. VUONG (2000): “Optimal Nonparametric Estimation of First-Price Auctions,” *Econometrica*, 68, 525–574.

- HAILE, P. A. (2001): “Auctions with Resale Markets: An Application to U.S. Forest Service Timber Sales,” *American Economic Review*, 91, 399–427.
- HAILE, P. A., H. HONG, AND M. SHUM (2003): “Nonparametric Tests for Common Values in First-Price Sealed-Bid Auctions,” NBER Working Paper 10105.
- HAILE, P. A., AND E. T. TAMER (2003): “Inference with an Incomplete Model of English Auctions,” *Journal of Political Economy*, 111, 1–52.
- HECKMAN, J. J., AND B. E. HONORÉ (1989): “The Identifiability of the Competing Risks Model,” *Biometrika*, 76, 325–330.
- (1990): “The Empirical Content of the Roy Model,” *Econometrica*, 58, 1121–1149.
- HENDRICKS, K., J. PINKSE, AND R. H. PORTER (2003): “Empirical Implications of Equilibrium Bidding in First-Price, Symmetric, Common Value Auctions,” *Review of Economic Studies*, 70, 115–145.
- HENDRICKS, K., AND R. H. PORTER (1988): “An Empirical Study of an Auction with Asymmetric Information,” *American Economic Review*, 78, 865–883.
- (forthcoming): “An Empirical Perspective on Auctions,” in *Handbook of Industrial Organization*, ed. by M. Armstrong, and R. Porter, vol. 3. Elsevier.
- HORTAÇSU, A. (2002): “Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market,” Working Paper (University of Chicago).
- HOTZ, V. J., AND R. MILLER (2003): “Conditional Choice Probabilities and the Estimation of Dynamic Models,” *Review of Economic Studies*, 60, 497–531.
- JACKSON, M. O., AND J. M. SWINKELS (forthcoming): “Existence of Equilibria in Single and Double Private Value Auctions,” *Econometrica*.
- JOFRE-BONET, M., AND M. PESENDORFER (2003): “Estimation of a Dynamic Auction Game,” *Econometrica*, 71, 1443–1489.
- KASTL, J. (2005): “Discrete Bids and Empirical Inference in Divisible Good Auctions,” working paper (Northwestern University).
- KOTLARSKI, I. (1966): “On Some Characterization of Probability Distributions in Hilbert Spaces,” *Annali di Matematica Pura ed Applicata*, 74, 129–134.

- KRASNOKUTSKAYA, E. (2004): “Auction Models with Unobserved Heterogeneity: Application to the Michigan Highway Procurement Auctions,” Working Paper (University of Pennsylvania).
- KRISHNA, V. (2002): *Auction Theory*. Academic Press, San Diego.
- LAFFONT, J.-J. (1997): “Game Theory and Empirical Economics: The Case of Auction Data,” *European Economic Review*, 41, 1–35.
- LAFFONT, J.-J., H. OSSARD, AND Q. VUONG (1995): “Econometrics of First-Price Auctions,” *Econometrica*, 63, 953–980.
- LAFFONT, J.-J., AND Q. VUONG (1996): “Structural Analysis of Auction Data,” *American Economic Review, Papers and Proceedings*, 86, 414–420.
- LI, T., I. PERRIGNE, AND Q. VUONG (2000): “Conditionally Independent Private Information in OCS Wildcat Auctions,” *Journal of Econometrics*, 98, 129–161.
- (2002): “Structural Estimation of the Affiliated Private Value Auction Model,” *RAND Journal of Economics*, 33, 171–193.
- LI, T., AND X. ZHENG (2005): “Procurement Auctions with Entry and an Uncertain Number of Actual Bidders: Theory, Structural Inference, and an Application,” Working Paper (Indiana University).
- LOADER, C. (1999): *Local Regression and Likelihood*. Springer, New York.
- MASKIN, E. S., AND J. G. RILEY (2000): “Asymmetric Auctions,” *Review of Economic Studies*, 67, 413–438.
- MCADAMS, D. (2005): “Identification and Testable Restrictions in Private Value Multi-Unit Auctions,” Working Paper (MIT).
- MCAFEE, R. P., D. C. QUAN, AND D. R. VINCENT (2002): “Minimum Acceptable Bids, with Application to Real Estate Auctions,” *Journal of Industrial Economics*, 50, 391–416.
- MEILIJSON, I. (1981): “Estimation of the Lifetime Distribution of the Parts from the Autopsy Statistics of the Machine,” *Journal of Applied Probability*, 18, 829–838.
- MILGROM, P. R., AND R. J. WEBER (1982): “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 50, 1089–1122.

- MYERSON, R. B. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- OLLEY, G. S., AND A. PAKES (1996): “The Dynamics of Productivity in the Telecommunications Equipment Industry,” *Econometrica*, 64, 1263–1297.
- PAARSCH, H. J. (1992a): “Deciding between the Common and Private Value Paradigms in Empirical Models of Auctions,” *Journal of Econometrics*, 51, 191–215.
- (1992b): “Empirical Models of Auctions and an Application to British Columbian Timber Sales,” *Research Report 9212 (University of Western Ontario)*.
- (1997): “Deriving an Estimate of the Optimal Reserve Price: An Application to British Columbian Timber Sales,” *Journal of Econometrics*, 78, 333–57.
- PESENDORFER, M. (2000): “A Study of Collusion in First-Price Auctions,” *Review of Economic Studies*, 67, 381–411.
- PINKSE, J., AND G. TAN (2005): “The Affiliation Effect in First-Price Auctions,” *Econometrica*, 73, 263–277.
- PORTER, R. H., AND J. D. ZONA (1993): “Detection of Bid Rigging in Procurement Auctions,” *Journal of Political Economy*, 101, 518–538.
- (1999): “Ohio School Milk Markets: An Analysis of Bidding,” *RAND Journal of Economics*, 30, 263–288.
- RILEY, J. G., AND W. F. SAMUELSON (1981): “Optimal Auctions,” *American Economic Review*, 71, 381–392.
- SMILEY, A. K. (1979): *Competitive Bidding Under Uncertainty: The Case of Offshore Oil*. Ballinger, Cambridge.
- SONG, U. (2003): “Nonparametric Estimation of an eBay Auction Model with an Unknown Number of Bidders,” Working Paper (University of Wisconsin).
- (2004): “Nonparametric Identification and Estimation of a First-Price Auction Model with an Uncertain Number of Bidders,” Working Paper (University of Wisconsin).
- SWINKELS, J., M. JACKSON, L. SIMON, AND W. ZAME (2002): “Equilibrium, Communication, and Endogenous Sharing Rules in Discontinuous Games of Incomplete Information,” *Econometrica*, 70, 1711–1740.

- THIEL, S. E. (1988): "Some Evidence on the Winner's Curse," *American Economic Review*, 78, 884-895.
- U.S. FOREST SERVICE, . (1995): *Timber Program Issues*. U.S. Forest Service, Department of Agriculture, Washington.
- VICKREY, W. (1961): "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16, 8-37.
- WILSON, R. B. (1979): "Auctions of Shares," *Quarterly Journal of Economics*, 93, 675-689.
- WOLAK, F. A. (2003): "Identification and Estimation of Cost Functions Using Observed Bid Data: An Application to Electricity Markets," in *Advances in Economics and Econometrics - Theory and Applications, Eighth World Congress*, ed. by M. Dewatripont, L. P. Hansen, and S. J. Turnovsky, vol. 2, pp. 115-149. Cambridge University Press, Cambridge, Econometric Society Monographs.
- (2004): "Quantifying the Supply-Side Benefits from Forward Contracting in Wholesale Electricity Markets," Working Paper (Stanford University).
- WOLFRAM, C. D. (1998): "Strategic Bidding in a Multi-Unit Auction, An Empirical Analysis of Bids to Supply Electricity in England and Wales," *RAND Journal of Economics*, 29, 703-725.
- YIN, P.-L. (2003): "Information Dispersion and Auction Prices," Working Paper (Harvard Business School).