An Empirical Perspective on Auctions

Ken Hendricks, University of Texas
Robert H. Porter, Northwestern University*

June 13, 2006

Abstract

We describe the economics literature on auction markets, with an emphasis on the connection between theory, empirical practice, and public policy, and a discussion of outstanding issues. We describe some basic concepts, to highlight some strengths and weaknesses of the literature, and so indicate where further research may be warranted. We discuss identification and estimation issues, with an emphasis on the connection between theory and empirical practice. We also discuss both structural and reduced form empirical approaches.

Keywords: auctions, bidding, identification, estimation, collusion, bid rigging.

*Forthcoming in the *Handbook of Industrial Organization* (Vol. III), edited by M. Armstrong and R. Porter. We are grateful to Mark Armstrong, Phil Haile, Jakub Kastl, Paul Klemperer and Harry Paarsch for helpful comments.

Contents

1 Introduction

Auctions are an important market institution. Our purpose in this chapter is to review the theoretical and empirical literature on bidding in auction markets, although the discussion of the theory will be limited to instances where the models have a direct bearing on some empirical literature. The emphasis is instead on empirical work.

There are many auction markets in which excellent data are available. There are several different types of auction data sets that have been analyzed in the literature. These include government sales, such as sales of timber and mineral rights, oil and gas rights, treasury bills, and import quotas, privatization sales, spectrum auctions and auctions of SO_2 emission permits. Auctions are important in government procurement, including defense contracts, and contracts for highway construction and repair, and for school milk. In the private sector, auction houses such as Sotheby's and Christies sell art, wine and memorabilia. There are agricultural sales (e.g., eggplants in Marmande, France), estate sales, real estate auctions, and used durable goods sales, for used cars and farm machinery. Burlington Northern has employed auctions to allocate future railcar capacity. The recently created wholesale electricity markets in the United Kingdom, Australia, various regions in the United States, and elsewhere often employ auctions. There are also now many internet auctions, with eBay the most prominent. Finally, auctions are frequently the subject of experimental research, both in the laboratory and in field experiments.

Auction data sets are often better than the typical data set in industrial organization. The auction game is relatively simple, with well-specified rules. The actions of the participants are observed directly, and payoffs can sometimes be inferred.

Why use an auction as a trading mechanism, rather than posting prices, bargaining or contracting? There is usually some uncertainty about the buyers' willingness to pay, and heterogeneity among potential buyers. Also, some degree of product heterogeneity is often present, so that past transactions are not a reliable guide to current market prices. In these circumstances, auctions can be an effective price discovery process.

There are *many* possible auction mechanisms. They can be characterized in terms of (1) a message space (i.e., what information do the bidders send to the seller) and (2) the seller's allocation rule, specifying which, if any, bidders might receive the item, and the probability that they receive it, and payments from (or transfers to) the bidders, as a function of the messages received. For example, consider a seller who has one item to allocate, and where messages are bids. Then one can roughly categorize common auction mechanisms into one of four cells:

	Highest Bid	Second Highest
Open	Dutch	English
Closed	FPSB	Vickrey (SPSB)

In closed auctions, submitted bids are considered simultaneously by the seller. In a first price, sealed bid (FPSB) auction, the item is awarded to the highest bidder, who pays his bid. In a Vickrey or second price sealed bid (SPSB) auction, the highest bidder wins and pays the highest losing bid. In open auctions, bidders call out their bids sequentially. In a Dutch auction, the seller starts the auction at a high price, and lowers the price until a bidder says "stop." That bidder wins the item at the stop price. Thus, only the winning "bid" is submitted. In an English auction, the seller starts the auction at a low price, and the price increases until no bidder is willing to raise it further. As we will describe in more detail in Section 3, there are many ways to run an English auction. Under several variants, the winning bidder acquires the item at the drop out price of the last remaining losing bidder, so that the outcome is like a second price auction. Similarly, the Dutch auction is akin to a FPSB auction. Variations on these mechanisms include secret or publicly announced reserve prices or minimum bids, and entry fees or subsidies.

In this chapter we describe the economics literature on auction markets, with an emphasis on the connection between theory, empirical practice, and public policy, and a discussion of outstanding issues. Auction markets are the subject of a large and distinguished theoretical literature. Klemperer (2004), Krishna (2002) and Milgrom (2004) provide excellent summaries of the theory and applications to public policy issues. There are also older thorough surveys by McAfee and McMillan (1987) and by Wilson (1993). An extensive empirical literature using data from auction markets has emerged recently. Our goal in this chapter is to provide an introduction to the empirical literature. Our intention is to describe some basic concepts, and to highlight some strengths and weaknesses of the literature, and so indicate where further research may be warranted. There are two other contemporaneous reviews of note, which complement our chapter. Athey and Haile (2005) survey various approaches to the estimation of structural models of equilibrium bidding, with an emphasis on nonparametric identification issues. Paarsch and Hong (2006) also discuss structural modeling, with an emphasis on estimation issues, often in a parametric context. We will also discuss identification and estimation issues, but our emphasis will be on the connection between theory and empirical practice. We will also discuss both structural and reduced form empirical approaches. The surveys by Hendricks and Paarsch (1995) and by Laffont (1997) describe some of the early empirical analysis of auction data sets. Finally, Kagel (1995) surveys the considerable literature on laboratory experimental auction markets, and Harrison and List's (2004) survey of field experiments describes some auction market field experiments. We refer to the experimental literature only occasionally, and interested readers should consult the surveys for more detail.

Auctions offer the prospect of a close connection between theory and empirical work. Moreover, much of the theoretical work on auctions has specific positive or normative goals in mind, and so empiricists are not usually required to re-cast the theory before testing theoretical predictions. Given a specific auction mechanism, the positive role of theory is to describe how to bid rationally, which usually involves characterizing the Bayesian Nash equilibrium (BNE) of the bidding game. Given the number of bidders, the joint distribution of their valuations and signals, and some behavioral assumption, the normative role of theory is to characterize optimal or efficient selling mechanisms. These two roles are reflected in the way theoretical work has been employed to guide empirical work and policy advice. Theory helps to shed light on how to interpret patterns in the data, suggests comparative static results that can be tested, and guides optimal mechanism design.

Empirical work also has positive and normative goals. The positive goal is to answer questions such as how agents behave, and whether their valuations are correlated and, if so, the sources of the correlation. Given the auction environment, a bidder's strategy is a mapping from his private information to a bid. Hence, a realization of bidder signals induces a distribution of bids. One can then ask whether an observed bid distribution is consistent with BNE, and test for such properties as independence. With experimental data, the researcher knows what signals were received, but not the preferences of the bidders, and one can compare predicted to actual bids under different assumptions concerning preferences. Thus, the consistency question is well-defined. However, with field data, the researcher does not know what signals were received, and real modeling issues arise in examining the consistency question. The positive analysis can be of more than academic interest, since bid rigging or collusion may be distinguishable from non-cooperative behavior, if the two have different positive implications. One can also ask whether risk aversion is an important feature.

The normative goal of empirical work is to answer questions such as what the revenue maximizing or efficient auction might be. If one knows or one can estimate the relevant features of the auction environment, especially the joint distribution of valuations and signals for potential bidders, and one knows which behavioral model is appropriate, then optimal auction design is feasible. Alternatively, one can test whether the auction design is optimal. For example, McAfee and Vincent (1992) ask whether the reserve price is chosen optimally in the first-price, sealed bid (FPSB) auctions they study.

For want of a better description, there have been two kinds of approaches adopted in the empirical literature, which we term *reduced form* and *structural*. Reduced form analysis tests predictions of the theory to draw inferences about behavior and the bidding environment. The goal of the structural approach is to estimate the data generating process directly, in particular the joint distribution of valuations and signals, usually under the assumption of risk neutrality and BNE. The strategy is to characterize the Bayesian Nash equilibrium of the auction to obtain a functional relationship between signals and bids, and therefore between the distributions of signals and bids. Assuming a BNE equilibrium exists, one can use the relationship between signal and bid distributions to construct a likelihood function for the data. The difficulty, or constraint on applicability, is that equilibrium strategies can be complex, they are often highly non-linear, and there may not be a closed-form representation for the bid strategy, or more precisely, for the inverse bid function. In addition, there are the usual existence and uniqueness issues to worry about. We will survey papers according to this perspective, to give a flavor of what's being done, where progress has been made, and what the outstanding issues are.

This survey is concerned primarily with single object auctions. The data, of course, often consists of many auctions, conducted either sequentially over time or simultaneously. For much of this survey, we follow the literature and treat each auction as an independent event, ignoring any factors, strategic or structural, that link participation and bidding decisions across auctions. Section 2 describes a model of bidding and introduces the notation we employ throughout. Section 3 focuses on structural estimation of second-price auctions. Section 4 discusses structural estimation of first-price auctions. Section 5 looks at tests that distinguish between information environments, and Section 6 describes tests of the theoretical implications of equilibrium bidding. Section 7 reviews empirical work that compares revenues across auction formats, in both single- and multi-unit settings. Section 8 examines work on collusion or bid rigging schemes. We close with a brief discussion of some outstanding empirical issues.

2 Model and Notation

In this section we describe a model of bidding based upon the general symmetric information model introduced by Milgrom and Weber (1982). This model will serve to establish the notation that will be used throughout this chapter. Random variables will be denoted by upper case and realizations by lower case.

Consider a bidding environment in which n potential risk neutral buyers bid to purchase a single item. The buyers are indexed by i. Each bidder i observes a real-valued signal x_i. Here X_i is private information, observed only by bidder i. In addition, there is some random variable V, which may be multidimensional, and which influences the value of the object to the bidders. The values of $(V, X_1, .., X_n)$ are governed by some joint distribution denoted by F. The joint distribution is assumed to be symmetric in the signals and to have a density function f. Let $x_{-i} = (x_1, .., x_{i-1}, x_{i+1}, .., x_n)$. A generic signal is denoted by X with realization x. Bidder i's payoff or valuation is given by $U_i = u_i(V, X_i, X_{-i})$ in the event that bidder i obtains the object being sold. The seller announces a minimum bid, or reserve price, r.

The primitives of the model are the number of bidders n, the distribution function F, and the utility functions $\{u_i\}_{i=1}^n$. These are assumed to be common knowledge among the buyers. The main assumptions of the model are (i) u_i is non-negative, continuous, increasing in each argument and symmetric in the components of X_{-i} and (ii) $(V, X_1, .., X_n)$ are affiliated.[1] These two assumptions imply that the valuations $U_1, ..., U_n$ are affiliated random variables. Let $Y_1, .., Y_{n-1}$ denote the ordering of the largest through the smallest signals from $(X_2, .., X_n)$, the signals of bidder *1*'s rivals. By Theorem 2 of Milgrom and Weber (1982), $(V, X_1, Y_1, .., Y_{n-1})$ are affiliated random variables.

There are several restrictive aspects of the model worth noting. First, the assumption that private information is single dimensional rules out situations in which bidders have private information about the different components of their payoffs. The only empirical work in auctions that we know of that allows for multi-dimensional signals is Cantillon and Pesendorfer (2004). Second, symmetry of F in the signals rules out the possibility that some

[1] The n random variables $X = (X_1, .., X_n)$, with joint density $f(x)$, are affiliated if, for all x and y, $f(x \wedge y)f(x \vee y) \geq f(x)f(y)$, where $\wedge$ denotes the component-wise minimum, and $\vee$ the component-wise maximum. If random variables are affiliated, then they are non-negatively correlated. See Milgrom and Weber (1982).

bidders have more precise signals of common unknown components of the payoffs. We relax this assumption when we discuss auctions with asymmetrically informed bidders. Third, there is an implicit assumption that, in the event that someone else obtains the object, bidder i's valuation is independent of the identity of that agent. Jehiel and Moldovanu (2000) consider the strategic implications of the alternative assumption that the winner's identity matters, as in an auction of an input prior to downstream oligopoly competition, for example. Fourth, the assumption that the number of potential bidders is common knowledge can be relaxed to account for uncertainties about whether rivals are serious bidders. For example, a costly decision to acquire a signal might be private. See Hendricks, Pinkse and Porter (2003), for an example of this sort of model.

A bidding (pure) strategy for bidder i is a correspondence $\beta_i : X_i \longrightarrow \Re_+$. It maps the private signal into a nonnegative real number. As we shall see, empirical work in auctions relies heavily upon the assumption that β_i is a strictly increasing function and hence invertible. Theoretical work in auctions informs us that the conditions under which a Bayesian Nash equilibrium in non-decreasing bid functions exists are affiliation of $(V, X_1, .., X_n)$ and symmetric payoff functions u_i. If either of these conditions are not satisfied, an equilibrium in increasing bid functions may not exist, although we will consider special cases in which such an equilibrium exists. For a discussion, see Athey (2001) or Krishna (2002, Appendix G).

Two special cases of the above model are frequently discussed in the theoretical literature:

1. Private values: $u_i(\nu, x_i, x_{-i}) = x_i$. Bidder i knows his own valuation and is only uncertain about how much others value the item.

2. Pure common values: $u_i(\nu, x_i, x_{-i}) = \nu$ (or one component of ν when it is multi-dimensional). All buyers have the same valuation, which is unknown to them when they bid. If so, then each buyer's signal may be informative. However, each bidder i knows only its own signal and does not observe the other bidders' signals, x_{-i}.

The presence of common components in bidders' payoffs does not imply that valuations are not private. When the common components are known to the bidders, then bidders may be uncertain about their rivals' valuations but they know their own private valuations. When the common components are uncertain then, in terms of our notation, the important strategic issue is whether rival signals are informative about the realization of the common component V, given the realization of their own private signal, x_i. If the distribution of V conditional on x_i is independent of X_{-i}, then the bidding environment is one with private values. In practice, the private valuation can then be defined as the expected value of $u_i(V, X_i, X_{-i})$ given the signal x_i received by bidder i, which in a private values environment is independent of X_{-i}. The location of the signal distribution is arbitrary, and if $E[U_i|X_i = x_i, X_{-i} = x_{-i}] = f(x_i)$ for some monotone increasing function f, signals can be re-normalized to equal $f(x_i)$. In a common values environment, the distribution of V conditional on x_i is not independent of X_{-i} and the realization x_{-i} matters, i.e., $E[U_i|X_i = x_i, X_{-i} = x_{-i}]$ depends on x_{-i}. In summary, if valuations depend on known common components, or on unknown common components about which bidders have no private information, then the bidding environment can be characterized as one of private values (PV). If valuations depend on unknown common components about which bidders have private information, then we will characterize the bidding environment as one of common values (CV), and distinguish the special case in which valuations depend solely on the unknown common component as pure common values (PCV). Within the private and common values bidding environments, we will sometimes distinguish between independent signals (IPV and ICV respectively) and affiliated signals (APV and ACV respectively).[2]

Consider, for example, the bidding environment for offshore oil and gas leases. The argument for the common value case is that the firms are uncertain about common components of the value of the lease being sold, such as the size of any oil or gas deposits under the tract, the prices of oil and gas over the likely production horizon if the lease is productive, and the common costs of exploration and development. The first component is likely to lead to common values, to the extent that firms have private information about the size of the deposits based on the seismic data they obtain, and especially their interpretation of that data. In contrast, firms may not have private information about future prices, and uncertain prices are then not inconsistent with private values.

Alternatively, one might argue that there is little discrepancy in private assessments of

[2] The theoretical literature classifies auctions in terms of the "reduced form" valuation

$$w(x_i, x_{-i}) = E[U_i|X_i = x_i, X_{-i} = x_{-i}]$$

rather than the primitives. Values are private if $w(x_i, x_{-i}) = x_i$ and interdependent otherwise. Common values is a special case of interdependent values.

these common components, and instead that valuations differ because of differences in bidder specific components of valuations. The most likely sources of bidder payoff heterogeneity are the private components of exploration and drilling costs. Bidders are not likely to differ in their valuation of recovered deposits, to the extent that there is a well developed market for oil and gas. Under this alternative view, valuations are best modelled as private, although they may be affiliated because of the common unknown components of payoffs that may be correlated with publicly available information.

An intermediate case, introduced by Wilson (1977), that nests the private and pure common value cases assumes that $u_i(v, x_i, x_{-i}) = u(v, x_i)$. Here bidder i's valuation has a common component and an idiosyncratic component, and the former affects every bidder's payoff in the same way. In terms of oil and gas auctions, the Wilson model allows the bidders' payoffs to depend upon the common, unknown size of the deposit and upon their private drilling costs. Private information, however, must still be real-valued. Consequently, the private signal plays two roles: it is informative of the size of the deposit and of drilling costs. The Wilson model is a restricted class of common values models, but we shall adopt it as the base model since almost all empirical work on auctions can be discussed in the context of this model. Throughout the remainder of this chapter, we will often use common values to refer to the Wilson model when discussing empirical work.

As we shall see, one of the important roles of empirical work on auctions is to determine whether values are private or common. The distinction is fundamental to both the positive goal of characterizing bidding behavior and the normative goal of auction design. In an independent private values environment, the standard logic of competitive markets prevails and more competition raises bids and increases revenues. (Note, however, that this result need not obtain in an APV environment, as shown by Pinkse and Tan (2005).) On the other hand, if the common component is the main determinant of bidder valuations, then more competition does not necessarily lead to higher bids and more revenue. Winning the auction is an informative event, with respect to the value of the item. The information is "bad news," and it has sometimes been called the "winner's curse." A bidder is more likely to bid high and win the item when he overestimates the item's value. In particular, bidding one's ex ante expectation (i.e., not conditioning on the event of winning) is an inferior strategy; one needs to bid more conservatively. As a result, restricting the number of bidders can be a revenue enhancing policy for the seller.

More formally, suppose the common value of the item, v, is unknown to the bidders and independent of the identity of the winner. Given a signal x, bidder *1*'s ex ante estimate of the item is

$$v_1(x) = E[V|X_1 = x] = \int v f_{V|X_1}(v|x)dv$$

where $f_{V|X_1}$ is the posterior conditional density function, derived by Bayes rule. By construction, $E[V_1|v] = v$, where $V_1 = v_1(X_1)$, so that ex ante estimates are conditionally unbiased. Suppose the bidding strategy for each bidder i is monotone increasing in its signal. Then, in a symmetric equilibrium, the winner is the bidder with the highest signal, and hence the highest estimate. But

$$E[\max_i\{V_i\}|v] \geq \max_i\{E[V_i|v]\} = v$$

by Jensen's inequality, since max is a convex function. This is in essence the winner's curse. Winning is bad news (in the sense of Milgrom (1981a)) for bidder *1*, since

$$E[V|X_1 = x, Y_1 < x] < E[V|X_1 = x].$$

Recall that Y_1 is defined to be the maximum signal among bidder *1*'s rivals, and if bidders employ symmetric strategies bidder *1* wins only when his signal is highest, and therefore exceeds Y_1. Thus, independent of strategic considerations, bidder *1* should bid less than his ex ante estimate V_1 to account for the information contained in the event that he has the highest bid.

3 Structural Analysis of Second-Price Auctions

In structural estimation, the empirical researcher posits a theory of equilibrium bidding and estimates the unobserved bidders' utilities and the joint distribution of their signals from bid data. Interest in this exercise is based on the normative objective of finding the optimal selling mechanism, or the positive objective of describing behavior, or distinguishing between alternative behavioral models, such as competition vs. collusion, or risk neutrality vs. risk aversion. This section reviews the work that has been done on English and second-price, sealed bid auctions.

3.1 Theory

It will be useful to first characterize bidder's payoffs and the symmetric Bayesian Nash equilibrium for an English auction. In deriving the equilibrium, we take the perspective of bidder *1*. Define

$$w(x, y) = E[u(V, x)|X_1 = x, Y_1 = y]$$

and let $\underline{x}$ denote the lower bound of the support of X_i. The function $w(x, y)$ denotes the expected payoff of the bidder, given a signal x, in the event that the highest rival signal is y.

There are many ways to run an English auction. In one variant, bidders call out their bids, while in another the auctioneer calls out the bids. Bids may rise continuously or in increments. A bidder's willingness to pay the bid called may be privately communicated to the auctioneer or it may be observed by rival bidders. These institutional details matter since they affect equilibrium behavior. As a result, the different variants of English auctions constitute different data generating mechanisms, and the analyst needs to bear this in mind in interpreting bids and deriving the likelihood function for the data.

Most empirical work is based on a version of the English auction known as a button auction. In this auction, the price rises continuously and bidders stay active as long as they keep their fingers on the button. A bidder wins once all other bidders take their fingers off their buttons, and the price paid is the exit point of the penultimate active bidder. In this context, a strategy for bidder i specifies, at any bid level, whether to stay active or drop out of the bidding. The symmetric equilibrium is characterized by a "no regret" property: a bidder remains active as long as his expected value of the object conditional on winning at the current price exceeds that price, and he drops out otherwise. The critical issue in characterizing the bidder's drop out point is what he observes about rivals' bidding decisions.

One possibility is to assume that each bidder does not observe the prices at which rival bidders drop out. The only inference an active bidder can draw from the fact that the bidding has not ended is that at least one rival is still active. In this case, the equilibrium strategy for bidder *1* is to drop out at

$$\beta(x) = E[u(V, x)|X_1 = x, Y_1 = x] = w(x, x). \tag{1}$$

The reasoning is as follows. Suppose bidder *1*'s rivals bid according to the above strategy and bidding has reached $b = \beta(y)$. Bidder *1*'s payoff conditional upon winning is $w(x, y) - b$, because to win means that everyone else has dropped out and at least one rival has signal y. Affiliation implies that bidder *1*'s payoff is positive if $x > y$ and negative if $x < y$. Hence, he should stay active until the bidding reaches $\beta(x)$ and then drop out. The English auction with unobserved exits is strategically equivalent to a second-price, sealed bid auction. In that case, equation (1) represents bidder *1*'s bid strategy.

Another formulation assumes that active bidders observe the prices at which rivals drop out of the bidding. Furthermore, no bidder who drops out can become active again. In this case, bidder *1*'s strategy at any bid will depend upon how many rivals have previously dropped out and the prices at which they dropped out. Let β_k specify the price at which bidder *1* drops out of the bidding if k rivals have dropped out previously, at the prices $b_1, .., b_k$. The no regret property implies that

$$\beta_k(x) = E[u(V, x)|X_1 = x, Y_1 = ... = Y_{n-k-1} = x, Y_{n-k} = \eta_{k-1}(b_k), .., Y_{n-1} = \eta_0(b_1)]. \quad (2)$$

Here $\eta_k()$ denotes the inverse of the function $\beta_k()$. (To economize on notation, we suppress the dependence of the bid on $b_1, .., b_k$.) The reasoning is similar to the case considered before. At any bid $b = \beta_k(y)$, bidder *1* calculates the expected value of the object conditional on winning and the prices at which k rivals have dropped out in prior rounds. Given those prices, and the inverse bidding rule with $n - k$ remaining rivals, the signals of the bidders who dropped out can be inferred. For example, the first bidder to drop at price b_1 must have had the signal $\eta_0(b_1)$. In this case, however, winning at b means that the $n - k - 1$ rivals who were active in the previous round of bidding must all drop out at b. For this to happen, they must all have the same signal, y. Affiliation implies that bidder *1*'s payoff from staying active is positive if $x > y$ and negative if $x < y$.

In most auctions, sellers set a reserve price, r. If $r > w(\underline{x}, \underline{x})$, the expected value if all bidders have the lowest possible signal, then bidders with low signals may not participate. Let x^* denote the participation threshold. In the symmetric equilibrium of the button auction described above, Milgrom and Weber (1982) show that participation threshold for each bidder is given by

$$x^*(r) = \inf\{x|E[w(x, Y_1)|X_1 = x, Y_1 < x] \geq r\}.$$

Since affiliation implies that expectations are increasing in x and y, x^* is unique. Bidders with signals less than x^* do not bid, and bidders with signals larger than x^* participate. Therefore, the marginal participant is someone who is just willing to pay r even if he finds out that no one else is willing to pay r. It is straightforward to show that the screening level is increasing in r and, in a common values model, increasing in n, the number of bidders. In the latter case, the participation threshold x^* increases in n to compensate for more rivals having pessimistic signals.

When values are private, the information conveyed by bidders who drop out is irrelevant to each bidder's participation and bidding decisions. Each bidder bids if and only if his value exceeds r, that is $x^*(r) = r$. The equilibrium strategy given in equations (1) and (2) reduces to $\beta(x) = \beta_k(x) = x$. Each bidder stays active until bidding reaches his private value and then drops out. The private values equilibrium is compelling because it is unique if one restricts attention to weakly dominant strategies. It is also an equilibrium in private values auctions with asymmetric bidders.

When values are common, the equilibrium is not unique. Bikhchandani, Haile and Riley (2002) characterizes the set of symmetric separating equilibria in English auctions with common values. They show that there are many ways for the $n-2$ bidders with the lowest valuations to exit the auction. The restriction to symmetric equilibria does pin down the transaction price since the equilibrium to the subgame that begins after the $n-2$ bidders with the lowest valuations drop out is unique. Unfortunately, this subgame has a continuum of asymmetric equilibria.

Consider the following simple example, which Klemperer (1998) calls the wallet game. There are two players, denoted by *1* and *2*. Each player's wallet contains money. Let X_i denote the money in player i's wallet, and assume that X_1 and X_2 are independent random variables uniformly distributed on $[0, \overline{x}]$. Each player knows how much is in his own wallet but does not know the amount in his rival's wallet. They are asked to bid for the sum of the contents of the two wallets in a second-price, sealed bid auction. Klemperer shows that the auction has a continuum of equilibria indexed by $\alpha > 0$ in which

$$\beta_1(x_1) = (1+\alpha)x_1, \qquad \beta_2(x_2) = (1+\frac{1}{\alpha})x_2.$$

Here α can be interpreted as a measure of how aggressively player *1* bids. The symmetric equilibrium, where $\alpha = 1$, is for each player to bid twice the amount in his wallet, as $w(x,x) = 2x$. But, if player *1* bids more aggressively, player *2* will bid less aggressively since he knows that if he wins, the amount in player *1*'s wallet is less than half of the price that he has to pay. Conversely, if player *1* bids less aggressively, then player *2* can afford to bid more aggressively.

The wallet game illustrates the effects of common values. The symmetric equilibrium is no longer in dominant strategies, and it is no longer unique. In fact, there are typically many asymmetric equilibria. The wallet game result is a special case of a more general result established by Milgrom (1981b). He proves that for every continuous, increasing function g, the two player, second-price, common values auction has an equilibrium in which

$$\beta_1(x_1) = w(x_1, g(x_1)),\ \beta_2(x_2) = w(g^{-1}(x_2), x_2).$$

Bikhchandani and Riley (1991) establish similar results for common value English auctions in which bidder exit times are observable.

The symmetric equilibria may not be robust to slight asymmetries in the payoffs. In the wallet game, suppose player *1* values the contents of the wallets by a small amount θ in addition to the total amount of money, whereas player *2* only cares about the money. In that case, Klemperer (1998) shows that the equilibrium consists of player *2* always bidding x_2 and bidder *1* bidding $x_1 + \overline{x}$. Thus, bidder *1* always wins and pays x_2. The intuition is that it is common knowledge that bidder *1* values the object more than bidder *2*, and bidder *1* can use this fact to threaten credibly to always bid more than bidder *2*.

Many oral, ascending price auctions are open outcry auctions in which bidders call out their bids. These are continuous time, dynamic games that are difficult to formalize since bidders can announce a "jump" bid at any time. Bidders can use jump bids to signal information about their valuations (as in Avery (1998)). The idea behind a jump bid is to convince rival bidders to concede the auction before the bidding reaches their values. The strategy can be effective because, if a bidder is certain that his value is less than the value of the jump bidder, then dropping out of the auction is a weakly dominant strategy. As a result, the jump bidder may win the auction at a lower price than he would have had to pay in the equilibrium in which bids increase incrementally. The jump bidder does run the risk of paying more than might have been necessary to win. A jump bid could also be effective if it is costly to continue submitting bids. Then a player might drop out early, if he views his chances of winning as low.

More generally, it is inherently difficult to model continuous time, dynamic games with few limitations on the strategy space. Absent any such limitations, such as those of the button auction, the theoretical literature does not suggest a suitable behavioral model for open outcry auctions.

3.2 Estimation

We now turn to a discussion of estimation issues. The structural econometric exercise we consider consists of finding F, the joint distribution of private signals and the payoff relevant random variable, V, that best rationalizes the bidding data.

Consider first the case of a button auction with symmetric independent private values. The mechanism generating the data consists of the identity function, $\beta(x) = x$, for signals above the reserve price r. Denote the common distribution of bidder values by F_X, which we wish to estimate. As discussed above, the bid levels at which bidders drop out of the auction are often not observed with the exception of course of the bidder with the second-highest valuation. Suppose therefore that the only data available are:

$$\{w_t, n_t, r_t\}_{t=1}^{T} \text{ if } m_t \geq 1,$$

where w_t denotes the winning bid, n_t the number of potential bidders, and r_t the reserve price in auction $t = 1, ..., T$. Here m_t denotes a latent (unobserved) variable, the number of bidders who are active, by which we mean those whose values exceed the reserve price, and therefore submit bids. We assume that we observe data only from auctions in which at least one bid is submitted. The winning bid $w_t = \max\{x_{2:n_t}, r_t\}$, the maximum of the second highest order statistic from a sample of n_t valuations and the reserve price.

Donald and Paarsch (1996) describe how to construct the likelihood function. The researcher needs to take into account three possible outcomes:

1. If $m_t = 0$, $Pr\{m_t = 0\} = F_X(r_t)^{n_t}$.
2. If $m_t = 1$, then $w_t = r_t$ and $Pr\{m_t = 1\} = n_t F_X(r_t)^{n_t - 1}[1 - F_X(r_t)]$.
3. If $m_t > 1$, then $w_t \sim h_t(w) = n_t(n_t - 1)F_X(w)^{n_t - 2}[1 - F_X(w)]f_X(w)$.

The first outcome occurs when no bids are submitted at an auction. We interpret this event as evidence that all of the bidders' valuations are below the reserve price. The second outcome occurs when the winning bid in the auction is the reserve price. We interpret this event as the case where only one bidder values the item more than the reserve price. The third, and final, outcome arises when the winning bid exceeds the reserve price. In this case, $n_t - 2$ bidders have valuations below w_t, the winning bidder's valuation exceeds w_t, and the second highest valuation is equal to w_t. Let

$$D_t = \begin{cases} 1 & \text{if } \; m_t = 1 \\ 0 & \text{if } m_t > 1 \end{cases}$$

Then the likelihood function for the data set, which includes only auctions in which bids are submitted, is:

$$L = \Pi_t [\frac{h_t(w_t)^{1-D_t} \Pr\{m_t = 1\}^{D_t}}{(1 - \Pr\{m_t = 0\})}].$$

Estimation then might proceed by choosing a family for $F_X(\cdot|\theta)$, parameterized by θ, and recovering the distribution function F_X above r using maximum likelihood methods. Note that this method remains valid when the assumptions of risk neutrality and/or symmetry are relaxed since $b_{it} = x_{it}$ remains a dominant strategy. In the asymmetric case, one recovers F_{X_i} above r for each bidder i.

Interest in this exercise is based on the normative objective of selecting the optimal auction design, such as finding the revenue maximizing reserve price. Note that in IPV auctions, the optimal reserve price is independent of n and depends only upon F_X. For example, in the English button auction, expected revenue to the seller who values the item at x_0 is

$$R = x_0 F_X(r)^n + rnF_X(r)^{n-1}[1 - F_X(r)] + \int_r^{\overline{x}} wn(n-1)F_X(w)^{n-2}[1 - F_X(w)]f_X(w)dw.$$

The first term corresponds to events when the item is not sold, the second when the item is sold at the reserve price, and the third when the winning bid exceeds r. Differentiating with respect to r yields the optimal reserve price as the solution to:

$$r = x_0 + [1 - F_X(r)]/f_X(r).$$

Thus the optimal reserve price in an IPV auction depends only on x_0 and F_X, and not the number of bidders. (See, for example, Riley and Samuelson (1981).)

If the bid levels at which losing bidders drop out of the auction are observed, then the likelihood function becomes

$$L = \Pi_{t=1}^T \{[1 - F_X(w_t)](\Pi_{i=2}^{m_t} f_X(b_{it}))F_X^{n_t - m_t}(r_t)\},$$

where b_{it} denotes the drop out point of bidder i for object t, $i = 1, .., n_t$, and bids are ordered so that $b_{1t} \geq b_{2t} ... \geq b_{m_t} \geq r_t$. The highest bid is not observed, in the sense that the winning bidder may have been willing to stay active at a higher price than w_t. If i did not submit a bid, we define $b_{it} \equiv 0$. If more than one bidder is active, then the winning bid is b_{2t}. If only one bidder is active, then the winning bid, and the lower bound on x_{1t}, is the reserve price r_t. If the reserve price is binding, then the likelihood function should be modified as above to account for the sample selection associated with observing auction outcomes only if $m_t \geq 1$.

If the number of potential bidders n_t is not observed, the researcher could assume $n_t = n$ for all t and estimate n as a parameter. Alternatively, one could assume $n = \max_t\{m_t\}$, i.e., the maximum number of active bidders observed in any auction, or a count of all bidders who are ever active. A third approach is to assume that n_t is the realization of a random variable N which has some distribution (e.g., Poisson) and estimate the distributional parameters (here of a truncated Poisson, since $n_t \geq m_t$). Note that in this case, where values are private, bidding behavior is independent of the number of bidders n_t. Hence, one could instead specify the likelihood function as

$$L = \Pi_{t=1}^{T}\{\Pi_{i=1}^{m_t}[f_X(b_{it})/(1 - F_X(r_t))]\}.$$

That is, obeserved bids are draws from the truncated distribution $f_X/(1 - F_X)$ above the reserve price r_t.

Observable heterogeneity in the items being auctioned can be accommodated by conditioning the distribution of values on a vector of covariates, Z_t. In this case, one recovers the family of conditional distributions, $F_X(\cdot|Z_t = z_t)$. For example, one could express

$$x_{it} = \alpha + z_{it}\beta + u_{it}.$$

Then

$$x_{it} \geq b_{it} \Longleftrightarrow u_{it} \geq b_{it} - \alpha - z_{it}\beta.$$

Heterogeneity in the items that is observed by the bidders but not by the researcher is a much bigger problem. The presence of unobserved heterogeneity implies that the signals x_{it} (and therefore the bids) are not independent given t nor is their distribution identical across t. Furthermore, if n_t is not observed and unobserved heterogeneity is present, the researcher has an identification problem: are a few, low bids observed because n_t is low, or because $F_X(\cdot|Z_t = z_t)$ is unfavorable?

When values are private but affiliated, it is still a dominant strategy for bidder i to bid his signal and to participate if x_i exceeds r. From an empirical viewpoint, the researcher needs to treat the auction rather than the individual bid as the unit of observation, and the joint distribution of signals, $F_{<X>}$, rather than F_X as the object of interest. Sample size may become a problem. If all bids are observed, then it is straightforward to specify and maximize a likelihood function to obtain an estimate of $F_{<X>}$. If some of the bids are not observed due to censoring, the specification of the likelihood function is complicated by the need to integrate over some of the components of $F_{<X>}$. With more than two bidders, these integrations are likely to render the maximum likelihood approach infeasible. An alternative approach is simulation methods (McFadden (1989) and Pakes and Pollard (1989)). Note that, because these auctions are dominance solvable, signals do not have to be affiliated.

In many instances, researchers observe *submitted* bids in an open outcry English auction. Any given bidder may submit many bids, or none. Let b_{it} now denote the highest bid submitted by bidder i for object t, $i = 1, ..., n_t$. If bidder i did not submit a bid, let $b_{it} = 0$. But note that bidder i might be active yet not submit a bid. Order the bidders so that $b_{1t} \geq b_{2t} \geq ... \geq b_{m_t} \geq r_t$, where now m_t denotes the number of bidders who submit a bid, as opposed to the number of active bidders. Here the winning bid is submitted by the bidder with the highest valuation, and so $w_t = b_{1t}$. If agents play according to their dominant strategy (i.e., stay active until the high bid surpasses their valuation), then we know that $x_{1t} \geq b_{1t}$, $x_{it} \in [b_{it}, w_t]$ for $i = 2, ..., m_t$, and $w_t \geq x_{it}$ for $i = m_t + 1, ..., n_t$. Both the winning bidder and the losing bidders are willing to pay at least their final bid, whereas the willingness to pay of losing bidders is less than the winning bid. The absence of a bid by bidder i implies that x_{it} is less w_t, but not necessarily less than the reserve price r_t. If we observe n_t, and if we assume that the second highest bidder drops out at the winning bid, or just below, then the likelihood function is as described above for the button version of an English auction. The winning bid is then distributed according to the density function $h_t(w_t)$, as defined at the beginning of this section, when there is more than one active bidder. Maximum likelihood estimation is possible even without observing the losing bids. Losing bid information might be used to obtain a more efficient estimator of F_X, although it is not clear how to do so, as there is not a one-to-one mapping from valuations to bids.

The problem with the likelihood approach in this context is the assumption that the winning bid equals the second highest valuation. If bids rise in discrete steps, or especially if there is jump bidding, with bids increasing faster than the required minimum increment, then the assumption is unlikely to be valid.

The likelihood approach as described here is parametric. Haile and Tamer (2003) propose a nonparametric approach to address the problem of inference when bidders do not indicate whether they are "in" as the ascending auction proceeds, which we describe in the next section. Their approach exploits the information embodied in the losing bids, and they do not require that the winning bid equal the second highest valuation.

We now discuss two empirical studies of bidding in ascending auctions.

Bajari and Hortacsu (2003) study bidding and auction design issues in eBay coin auctions. They treat this environment as a second-price, sealed bid auction. They note that eBay auctions have hard closing times, in that bids are not accepted after a pre-specified time. Moreover, a disproportionate fraction of bids are submitted close to the ending time, a practice known as sniping. Some preliminary empirical analysis indicates that, in their sample, winning bid levels are not correlated with early bidding activity, as measured by a dummy variable indicating whether the losing bids were submitted more than 10 minutes before the closing time (in which case the winning bidder would have had time to respond). They argue that the auctions they study are therefore strategically equivalent to a second price sealed bid auction.

A second potentially controversial assumption is that the coin auctions they study are a pure common value environment. They point to an active resale market, and to some evidence that bids are negatively correlated with the number of people who bid. The latter is likely endogenous, but the correlation persists if they use minimum bids as an instrument. Negative correlation between bids and the number of bidders is not consistent with private values in a second price sealed bid auction, as bidding one's value is a dominant strategy. As Bajari and Hortacsu note, the eBay coin market should probably instead be modeled as a more general common values setting, but to do so would be technically challenging.

Bajari and Hortacsu restrict attention to symmetric equilibria, so that bids are generated according to the strategy described above, whereby $\beta(x) = w(x, x)$, if x is greater than or equal to the screening level $x^*(n, r)$. In a common values environment, the screening level depends on the number of bidders, which we make explicit here. Although this is the unique symmetric equilibrium, as noted above there are a continuum of asymmetric equilibria in second price sealed bid auctions with common values. They also assume that the common value is normally distributed, with mean and variance linear functions of observable covariates, and that signals conditional on the common value are unbiased and normally distributed, with variance proportional to the variance of the common value, v. Given these assumptions, the likelihood function for the bid data is well-defined but difficult to estimate. One complication is that the inverse mapping from bids to signals is no longer the identity function. The likelihood that bidder i submits a bid b_{it} in auction t, conditional on the common value, is given by

$$f_{B|V}(b_{it}|v_t) = f_{X|V}(\eta(b_{it})|v_t)\eta'(b_{it}),$$

where η denotes the inverse bid function, and $f_{X|V}$ is the density of the signals conditional on the common value.[3] Since neither η nor its derivative η' are available in closed form, they must be estimated numerically. A second complication is that, in auctions with binding reserve prices, each bidder's participation decision depends upon the level of competition. As noted above, the screening level is a decreasing function of the number of bidders; the winner's curse causes expected profits to fall. This is not true in private value auctions, where each bidder participates as long as his value exceeds r. The likelihood function for the data is

$$L = \prod_{t=1}^{T} \left\{ \int F_{X|V}(x^*(n_t, r_t)|v)^{n_t - m_t} [\Pi_{i=2}^{m_t} f_{B|V}(b_{it}|v)](1 - F_{X|V}(\eta(w_t)|v))^{1\{m_t \geq 1\}} f_V(v) dv \right\},$$

where $1\{m_t \geq 1\}$ is an indicator variable for at least one active bidder in the auction, and $x^*(n_t, r_t)$ is the screening level, below which the signals of inactive bidders fall. Although eBay observes all bids, the data records only the losing bids and not b_{1t}, the winning bid. As above, if the auction receives only one bid, then b_{1t} is bounded below by the announced reserve price r_t. The likelihood contains no correction for sample selection since auctions that attract no bidders are observed and are assigned a positive probability. Finally, Bajari and Hortacsu argue that in eBay auctions, the number of rival bidders is not known to the bidders. They assume that n_t follows a Poisson distribution, with mean determined endogenously by a zero ex ante expected profit condition. Bidders are assumed to incur an entry cost before learning their signal. The function $w(x, x)$ must be modified to be the expectation of the common value, conditional on the bidder's signal x, where the expectation is now over both the realizations of the rival bidders' signals and the number of rival bidders. The screening level must also be modified to account for this uncertainty.

[3] Affiliation implies that $\beta(x)$ is strictly increasing in x, so the inverse function exists.

Bajari and Hortacsu's analysis of reserve price policy demonstrates the value of structural estimation. One caveat, which is more a complaint about the empirical eBay literature, is that the auctions are treated in isolation, when similar items are often sold at the same time. Thus the optimal reserve price calculation is treated like a monopoly pricing problem, and the bidder participation decision does not account for the outside option of bidding for comparable items.

Hong and Shum (2003) estimate a structural model of an asymmetric, ascending auction in a common value environment, where bidder's preferences depend on an idiosyncratic component and an unknown common component. The basic idea of their paper is to infer bidders' real-valued private signals from the bids at which they drop out, and to build a likelihood function for the observed drop out sequence. One can thus obtain estimates of the parameters of the distribution of bidder valuations. Since the likelihood function involves high dimensional integrals, Hong and Shum adapt the simulated nonlinear least squares method developed by Laffont, Ossard and Vuong (1995) for first-price auctions, which is described in more detail below, to estimate their model.

Hong and Shum tackle a difficult problem using state-of-the-art econometrics. The main weakness is the restricted domain of application. The key step in the analysis infers the bidders' private signals from the bids at which they drop out. From a computational viewpoint, this works only for the class of preferences which admits closed-form solutions for the conditional expectations, since these can be inverted. Hong and Shum use Wilson's (1998) log-additive specification with the common and idiosyncratic components distributed lognormal. From a practical viewpoint, drop out bids are typically not observed, for example because bidders do not have to bid to be active. This is a problem with the data from the FCC auctions that Hong and Shum examine. Hong and Shum assign a dropout price to a given bidder equal to the last submitted bid of the next bidder to drop out on a given local spectrum license.

In a simultaneous English auction, such as the PCS spectrum auction, there may be payoff complementarities across objects, binding budget constraints, or activity rules based on overall participation. There will then be strategic interdependence across objects, and inferring valuations from bids may be difficult. For example, with the FCC activity rules, bidder i may bid more than x_{it} for t in order to retain eligibility to bid on other licenses (a practice known as "parking"). An extreme example of such bidding behavior occurred when Sprint bid on both Oklahoma City C block licenses at once, even though the rules stipulated that they could win at most one of the two licenses.

3.3 Identification

Identification is central to structural estimation, and it informs the interpretation of reduced form estimates of comparative static properties. The models presented in the previous section are identified by functional form assumptions but may not be nonparametrically identified. Nonparametric identification reduces to the following question: given an equilibrium, is there a one-to-one relationship between the joint distribution of bidder values and the joint distribution of bids? One could argue that this standard of identification is too much to ask. After all, economists often make parametric assumptions in estimating demand and supply functions. One difference between these studies and empirical studies of bidding is that the main object of interest in structural models of auctions is often the distribution of the idiosyncratic private component of bidder valuations, and not the deterministic component of bidder valuations. Theory provides some insight on how bidder and auction characteristics should affect bidder valuations, but it offers little guidance on the functional form of the distribution of the idiosyncratic component. Yet the latter is often the primary source of bidder rents and the focus of mechanism design. Consequently, we believe that it is desirable for auction models to be identified nonparametrically. If identification is only by functional form, then a given parametric family of distributions for valuations may not approximate the true unknown distribution, and so specification can play a pivotal role in the analysis.

Athey and Haile (2002) synthesize and extend a fragmented literature on identification of the distribution of bidder values (or information) in a variety of auction settings, including English and second-price auctions. They show when non-parametric identification of valuations is possible, and, if so, what kinds of data are needed. The necessary data depends on whether the value distribution is symmetric or asymmetric, whether values are correlated, and whether the correlating factors are observed. In button auctions with private values, the bid function is the identity function, so bids have a clear interpretation, as losing bids correspond to valuations, and the winning bidder's valuation must not be less than the winning bid. Athey and Haile show, however, that the general private value model is not identified unless *all* bids are observed. This is a problem in English auctions, since the highest valuation is not observed, but known only to be bounded below by the winning bid. However, if bidder valuations are independent and symmetric, and if the winning bid equals the second highest valuation, then only the winning bid and the number of bidders needs to be observed. The winning bid is then the second highest order statistic from a sample of size n whose distribution is uniquely determined by the marginal distribution, F_X. The asymmetric IPV model is also identified if, in addition to the winning bid, the identity of the winner is known. Li, Perrigne and Vuong (2000) show that a restricted class of pure common values auctions is identified if all bids are observed. For example, suppose that signals are independent conditional on the common value, and that signals are additively separable in the common value and the idiosyncratic information. But non-parametric identification fails when there are idiosyncratic value components, i.e., for more general common values settings, or if some bids are not observed.

A second challenge to the structural approach is multiplicity of equilibria. The equilibrium of the second-price, sealed bid auction, and that of the English button auction, is unique if values are private. But, as noted above, in the English button auction with common values, Bikhchandani, Haile and Riley (2002) have shown that there are many ways in which the bidders who don't have one of the two highest values might exit the auction. This multiplicity calls into question the interpretation of the losing bids as the valuations of the losing bidders. Hence, one should not draw strong inferences from the last bid of losing bidders, apart from the highest losing bid. The winning bid is unique in the symmetric equilibrium. But there is a continuum of asymmetric equilibria in which auction outcomes have quite different implications for the distribution of valuations. The latter point also applies to second-price, sealed bid auctions with common values. Thus, in our view, the structural modeling program that identifies the distribution of valuations by inverting the mapping from values or their order statistics to bids is on solid ground when values are private, but it is on shaky ground when values are common. In the latter case, the model has many equilibria, and even after selecting an equilibrium and assuming that bidders play according to this equilibrium in all of the auctions, the model can only be identified by functional form.

Most second-price auctions are open outcry auctions, and not sealed bid or English button auctions. The free form bidding can generate many possible bid histories for any profile of private signals. The observed bids in these histories are not very informative, even when values are private. Bidders often do not have to indicate whether they are "in" or "out" at every high bid, and even if they do, the bid increments associated with jump bids can be large. Thus, final recorded bids of bidders may be a poor approximation of their values, and even the winning bid may not be a good approximation of the second-highest valuation. As a result, the likelihood function is not well-defined, since the distribution of final bids is not unique given the distribution of values, nor is the distribution of values uniquely defined given the observed distribution of final bids. This is one reason why there have been almost no structural econometric studies of English auctions.

Haile and Tamer (2003) provide a novel solution to the non-uniqueness problem in the IPV model. They exploit two plausible components of equilibrium play. First, the winning bidder may have been willing to bid more than the final bid, as noted above. In addition, none of the other bidders submit a bid more than they were willing to pay. Second, losing bidders were not willing to raise the winning bid by the minimum bid increment. This assumes that there are no costs associated with submitting bids, or with keeping bidding open. The latter feature would not necessarily be satisfied in an equilibrium with jump bidding. The upper bound on the valuations of bidders may also be violated if bidders collude. These two suppositions can be shown to provide bounds on the distribution of valuations, which are assumed to be private.

Let x_i denote the private value of bidder i, and b_i the highest bid submitted by that bidder. If a bidder does not submit a bid, then let b_i be zero. The two behavioral assumptions are then: (A1) $x_i \geq b_i$ for all i and (A2) $x_i \leq w+ \triangle$ for all i except the winning bidder, where $\triangle$ denotes the minimum bid increment. These two assumptions are not sufficient to identify F_X, the distribution of private values, which for convenience are here assumed to be independent and identically distributed. But the two assumptions do put bounds on F_X. The first assumption bounds values below by observed bids, and hence the observed distribution of bids bounds F_X from above. Similarly, the second assumption bounds F_X from below.

The bounds can be derived as follows. Suppose that there are n bidders (as opposed to the number of players who submit bids). Denote the distribution function of the i^{th} highest order statistic from the distribution F_X by $F_{i:n}$. Let $G_{i:n}$ denote the empirical distribution the i^{th} highest bid. Then assumption (A1) implies that $F_{i:n}(x) \leq G_{i:n}(x)$ for all i, n, and x. Similarly, (A2) implies that $F_{2:n}(x) \geq G_{1:n}(x+ \triangle)$ for all n and x. If all losing bidders were not willing to raise the winning bid by the minimum bid increment, it suffices to characterize the highest losing bidder.

These bounds are not the end of the story, since the ordering of bids does not necessarily correspond to the order of valuations. For example, a bidder with an intermediate value might not bid at all, and so register a zero bid, while a lower value bidder might submit a bid early in the auction. Nevertheless, Haile and Tamer show how to pool the bounds across n and i in the case of the upper bound, and across n for the lower bound, in order to derive upper and lower bounds on F_X for all values of x.

If bidders employ the dominant strategy of the button auction, and so the losing bidders' final bids equal their valuation, then it can be shown that the upper and lower bounds coincide, and F_X is uniquely determined. Otherwise, there will be a gap between the bounds on F_X.

An important identifying assumption is that the number of bidders is known. In the Forest Service auctions considered by Haile and Tamer, bidders pre-qualify by submitting a sealed bid, usually at the reserve price, that serves a binding lower bid in the subsequent English auction. Given the number of pre-qualified bidders, Haile and Tamer exploit properties of order statistics to derive bounds on the distribution of valuations. These bounds are surprisingly tight, and they show that they also imply a relatively tight bound on the implied optimal reserve price, i.e., the reserve price that would maximize the seller's revenues. They also provide some Monte Carlo evidence that indicates that their bounds approach can provide much better estimates than a structural model that assumes exit prices reflect valuations, especially when there is jump bidding, or when some bidders are silent (i.e., they never submit a bid, although their valuation exceeds the outstanding high bid at some point before the end of the auction).

In conclusion, the structural maximum likelihood methods that exploit the information embodied in the winning bid, described in the previous sub-section, are valid only if the winning bid can be interpreted as the second highest valuation. The interpretation of losing bids is problematic in many instances. The nonparametric bounds approach of Haile and Tamer indicates how to circumvent both of these potential difficulties in private values environments.

4 Structural Analysis of First Price Auctions

In this section we discuss structural estimation of first price, sealed bid auctions. There are few empirical studies of field data from Dutch auctions. A notable exception is Laffont, Ossard and Vuong (1995), which we discuss below. (See also van den Berg and van der

Klaauw (2000).) We first review the theory of equilibrium bidding, then describe the main estimation approaches, and end with a discussion of identification issues.

4.1 Theory

In a first-price sealed bid auction, each bidder must independently submit a bid to the auctioneer. The high bidder wins and pays his bid, if it exceeds the reserve price. For example, in federal offshore oil and gas auctions, the Department of Interior announces several months in advance that it intends to sell production rights to a set of tracts in a specific geographical area. The announced reserve price for wildcat tracts was \$15 per acre in the 1960s, for example. Firms are invited to submit bids in sealed envelopes at any time prior to the sale date. On the day of the sale, the envelopes are opened, and the values of the bids and the identities of the bidders are announced. The firm or consortium that submits the highest bid on a tract is usually awarded the tract at a price equal to its bid, although the government rejected as inadequate many high bids above the announced reserve price.

Assume that bidder *1*'s rivals in a first-price sealed bid auction are using a common increasing bid function β that has an inverse function η at bids above the reserve price. Then bidder *1*'s profits from bidding b given a signal x are

$$\pi(b,x) = \int_{\underline{x}}^{\eta(b)} [w(x,y) - b] dF_{Y_1|X_1}(y|x).$$

The bidder wins in the event that the highest rival bid is less than b or, equivalently, that the highest rival signal is less than $\eta(b)$. Differentiating with respect to b and imposing the symmetry restriction that bidder *1*'s best reply is $b = \beta(x)$ yields the differential equation

$$[w(x,x) - \beta(x)] f_{Y_1|X_1}(x|x) - \beta'(x) F_{Y_1|X_1}(x|x) = 0. \tag{3}$$

If the auction has a binding reserve price, then bidders who obtain very low signals may not bid. The participation decision in the symmetric equilibrium of the first-price sealed bid auction is the same as in the second-price sealed bid auction. The screening level $x^*(r)$ is the lowest signal at which the expected value of the object conditional on winning is at least the reserve price. We assume that the reserve price is binding so that $x^*(r)$ exists and exceeds $\underline{x}$, the lowest possible signal.

The equilibrium bid function for $x \geq x^*$ is obtained by solving equation (3) subject to the boundary condition $\beta(x^*) = r$. It is given by

$$\beta(x) = rL(x^*|x) + \int_{x^*}^{x} w(s,s)dL(s|x)$$

where

$$L(s|x) = \exp\{-\int_{s}^{x} \frac{f_{Y_1|X_1}(t|t)}{F_{Y_1|X_1}(t|t)}dt\}.$$

We define $\beta(x) = 0$ for $x < x^*$. The symmetric equilibrium exists under fairly weak regularity conditions on F. The assumption of a binding reserve price ensures that there is a unique symmetric equilibrium. See the account in Athey and Haile (2005), for example.

In the special case of affiliated private values, the equilibrium bid function reduces to

$$\beta(x) = x - \frac{\int_r^x F_{Y_1|X_1}(s|x)ds}{F_{Y_1|X_1}(x|x)}$$

which further reduces to

$$\beta(x) = x - \frac{\int_r^x F_X(s)^{n-1}ds}{F_X(x)^{n-1}}$$

in the case of independent private values. The bid falls below the valuation by a mark-down factor. In the case of private values, the mark-down factor, $x - \beta(x)$, is decreasing in the number of bidders, n, and increasing in the dispersion of the value distribution. If there was no dispersion in values, then the equilibrium strategy is to bid one's own value. The equilibrium bid strategy for bidder *1*, in the case of private values, can also be expressed as:

$$\beta(x) = E[\max\{r, Y_1\}|X_1 = x, Y_1 \leq x].$$

The bid equals the expectation of the highest rival signal, where the seller's reserve price is treated like another rival signal, conditional on having the highest signal and therefore being the winning bidder in a symmetric equilibrium. This amount corresponds to the expected payment in a second price private values auction, conditional on being the winning bidder.

In the case of independent private values, Lebrun (1996, 1999) and Maskin and Riley (2000a,b) have extended the existence and uniqueness results to auctions with asymmetric bidders under the assumption that the supports of the bidder distributions are identical and F_{X_i} satisfies certain mild regularity conditions.

4.2 Estimation

The IPV model can be estimated in a variety of ways. Paarsch (1992) derives a maximum likelihood estimator and applies it to procurement auctions for treeplanting contracts in British Columbia. Donald and Paarsch (1993) discuss related econometric issues in more detail. Suppose the data consists of $\{w_t, r_t, n_t\}_{t=1}^T$ for the sample where the number of submitted bids $m_t \geq 1$. Their approach relies on functional form assumptions, as the class of distributions considered guarantee that (β, η) have closed form representations. The winning bid w_t is observed if and only if the highest signal $x_{(1:n_t)} \geq r_t$, otherwise no bids are observed. The probability of the latter event is $F_X(r_t)^{n_t}$. The probability distribution function of w is

$$h_t(w) = n_t F_X(\eta_t(w))^{n_t-1} f_X(\eta_t(w))\eta_t'(w) = \frac{n_t F_X(\eta_t(w))^{n_t}}{(n_t-1)(\eta_t(w)-w)}.$$

Note that the bid and inverse bid functions depend on t since the number of bidders n_t and the reserve price r_t vary across auctions. Note also that the valuation distribution F_X is assumed to be independent of n_t. The bid and inverse bid functions also both depend on the parameters of F_X, which we suppress for notational convenience. Therefore the likelihood function is

$$L = \Pi_t\{h_t(w_t)/[1 - F_X(r_t)^{n_t}]\}.$$

The parameters of F_X are chosen to maximize L subject to $w_t \leq \beta_t(\overline{x})$ for all t, where $\overline{x}$ is the highest possible signal, and subject to the requirement that β, η, and F_X be conformable. Clearly, the main difficulty with this estimation approach lies in computing the inverse η_t, which is typically nonlinear and often does not have a closed form solution. Another technical issue associated with the maximum likelihood method is that the asymptotic distribution of the estimator is non-standard, since the upper bound of the support of the bid distribution depends upon the parameters of interest. Moreover, the likelihood function may be discontinuous at the associated boundary of the parameter space. See Donald and Paarsch (1993, 1996, 2002) for more detail, as well as Chernozhukov and Hong (2004) and Hirano and Porter (2003), who advocate Bayesian estimators.

In his thesis, Bajari (1997) extends the likelihood approach to auctions with asymmetric bidders. His application is to procurement auctions of highway repair contracts in Minnesota. The lowest bid wins, and there is no reserve price, so the data consists of $\{\{b_{it}, z_{it}\}_{i=1}^{n_t}, n_t, z_t\}_{t=1}^T$ where z_{it} is the vector of firm i characteristics in auction t and z_t is the vector of contract t characteristics. The firm specific characteristics are the distances between the locations of firms and contract t, and measures of the firms' committed capacities at the time of auction t. The likelihood of firm i submitting bid b_{it} in auction t is given by

$$f_{B_i}(b_{it}|z_{it}, z_t) = f_{X_i}(\eta_{it}(b_{it})|z_{it}, z_t)\eta'_{it}(b_{it}).$$

In this case, the inverse bid functions and their derivatives are obtained by numerically solving the system of n_t differential equations derived from the firms' first order conditions. Bajari takes a Bayesian approach to estimation, simulating the posterior distribution of the parameters by taking random draws from a prior and evaluating the likelihood function for each draw. The Bayesian approach finesses some of the technical difficulties that arise with maximum likelihood estimators. It is also computationally easier to implement than maximum likelihood and allows the researcher to compare non-nested models such as collusion versus competition in a relatively straightforward way. The prior distribution must be chosen with care, however.

The main drawbacks of the likelihood approach are discussed in Bajari (1998). One difficulty is that the approach is computationally intensive. The need for flexible functional forms means that the inverse bid functions have to be computed numerically. A second difficulty is that the likelihood function typically does not have full support, which means that zero likelihoods often arise in practice. Outliers may be difficult to rationalize, and they can have a disproportionate effect on estimated parameter values.

Laffont, Ossard, and Vuong (1995) develop a simulated non-linear least squares estimator. They exploit the fact that the bid function in the FPSB auction in an IPV environment can be expressed as

$$\beta(x) = E[\max\{r_t, X_{2:n_t}\}|X_{1:n_t} = x].$$

where $X_{k:n_t}$ denotes the k^{th} highest order statistic from a sample of size n_t. The winning bid therefore satisfies

$$w_t = \beta(x_{1:n_t}) = E[\max\{r_t, X_{2:n_t}\}|X_{1:n_t} = x_{1:n_t}].$$

Hence, the expectation of the winning bid, conditional on at least one bid being submitted, is

$$E[W_t|X_{1:n_t} \geq r_t] = E[\max\{r_t, X_{2:n_t}\}|X_{1:n_t} \geq r_t].$$

The expectation involves high dimensional integrals. Instead of computing it, Laffont, Ossard, and Vuong use simulation to approximate the expectation. They choose an "importance function" to weight the simulated samples. They then choose the parameters of F_X to minimize the average squared distance between the observed sample of winning bids, $\{w_t\}$, and the simulated sample mean of $\max\{r_t, x_{2:n_t}\}$, conditional on $x_{1:n_t}$ exceeding the reserve price, after correcting for the estimation error of the simulated sample. Laffont, Ossard, and Vuong apply this method to eggplant sales in Marmande, France. Their method circumvents the need to compute the inverse bid function at each parameter value. As a result, their method allows for a much larger set of distributions than do maximum likelihood methods, although they just consider the lognormal distribution in their application. In their data, as in many applications, the number of bidders n_t is not observed. The authors assume the number is constant, n, treat n as a parameter, and estimate it.

The main restriction in applying their method is that it relies heavily upon the assumptions of the symmetric IPV model. If bidder valuations are not independent draws, or the distributions from which they are drawn are not symmetric, the bid function cannot be expressed as a conditional expectation of a second order statistic. Independence may be a problem in theory, but it may be more of a problem in practice because of unobserved heterogeneity. For example, the bidders may all observe some factor that shifts the location of the distribution of values, where that factor is not observed in the data. In the application, unobserved heterogeneity that is also observed by the seller might result in correlation between the reserve price and the distribution of bidder valuations, in which case the reserve price is not exogenous. Symmetry is also frequently a problem. For example, in the Marmande eggplant market, one buyer was much larger. Laffont, Ossard and Vuong show how to adapt their method to account for this latter issue, under the assumption that the large buyer's valuation can be represented as a draw from F_X^k, where k can be thought of as the number of agents the large buyer represents as an intermediary at the auction. Then the large buyer's valuation is the highest of the k agents he represents, where the agents' valuations are drawn from the same distribution as the other bidders.

A third estimation procedure for private value models has been developed by Elyakime, Laffont, Loisel and Vuong (1994) and by Guerre, Perrigne and Vuong (2000). Their method can be non-parametric, and so not rely upon functional form assumptions. It is computationally easy to implement, nor does it require bidders' values to be independently or identically distributed. Suppose the researcher has data available on all bids submitted,

$$\{\{b_{it}\}_{i=1}^{m_t}, n_t, r_t\}_{t=1}^{T}$$

plus perhaps covariates Z_t for auction t where $m_t \geq 1$. Here n_t again denotes the number of potential bidders, and m_t the number who submit bids. Now assume that $n_t = n$ for all t. In practice, applications of this method often estimate the value distribution separately for each value of n_t. In a symmetric equilibrium, the optimal bid for bidder *1* with signal $x_1 = \eta(b)$ solves

$$(\eta(b) - b) f_{Y_1|X_1}(\eta(b)|\eta(b))\eta'(b) - F_{Y_1|X_1}(\eta(b)|\eta(b)) = 0.$$

Define $M_1 = \beta(Y_1)$ as the maximum bid of bidder *1*'s rivals, and let the conditional distribution of M_1 given bidder *1*'s bid B_1 be denoted by $G_{M_1|B_1}(\cdot|\cdot)$ and its density by $g_{M_1|B_1}(\cdot|\cdot)$. Monotonicity of β and η implies that for any $b \in (r, \beta(\overline{x}))$

$$G_{M_1|B_1}(m|b) = F_{Y_1|X_1}(\eta(m)|\eta(b)).$$

Here $G_{M_1|B_1}(m|b)$ is the probability that the highest bid among bidder *1*'s rivals is less than m conditional upon bidder *1*'s bid of b. The associated density function is given by

$$g_{M_1|B_1}(m|b) = f_{Y_1|X_1}(\eta(m)|\eta(b))\eta'(m).$$

Substituting the above relations into the first order condition for bidder *1* yields

$$\eta(b) = b + \frac{G_{M_1|B_1}(b|b)}{g_{M_1|B_1}(b|b)}.$$

In the special case of IPV, $G_{M_1|B_1} = G^{n-1}$, where G is the marginal bid distribution of individual bidders, $g_{M_1|B_1} = (n-1)gG^{n-2}$, and the inverse bid function for bidder *1* is given by

$$\eta(b) = b + \frac{G(b)}{(n-1)g(b)}.$$

The estimation procedure consists of two steps.

Step 1: Estimate $\widehat{g}_{M_i|B_i}(m|b), \widehat{G}_{M_i|B_i}(m|b)$, either parametrically or non-parametrically, from the observed bids of all bidders. Here bidders are treated symmetrically, and auctions as independent and identical implicitly, except to the extent that one conditions on observable auction characteristics. Form "data,"

$$\widehat{x}_{it} = b_{it} + \frac{\widehat{G}_{M_i|B_i}(b_{it}|b_{it})}{\widehat{g}_{M_i|B_i}(b_{it}|b_{it})}.$$

This yields a sample of pseudo-values $\{\widehat{x}_{it}\}_{i=1}^{n}, t = 1, ..., T$.

Step 2: Given $\{\{b_{it,}\widehat{x}_{it}\}_{i=1}^{n}\}_{t=1}^{T}$, for the subsample with $n_t = n$, estimate $f_X(\widehat{x}_{it})$ and $F_X(\widehat{x}_{it})$. One can also estimate β, parametrically or nonparametrically, as $b_{it} = \beta(\widehat{x}_{it})$. If the value distribution F_X is assumed to be identical across subsamples with different numbers of bidders, the data from these subsamples can be pooled at this stage. In principle, the symmetry assumption (of F_X across n_t) can be tested. (This assumption might equivalently be viewed as an assumption that the variation in the number of bidders is exogenous.) Note that one could not similarly pool estimation of the bid function across subsamples, as the bid function should depend on the number of bidders.

In their application to French timber auctions, Elyakime, Laffont, Loisel and Vuong modified the method to account for the fact that r_t is not announced to the bidders, so that there is a secret reserve price. The seller is treated as another potential bidder, but with an objective function that differs from that of the buyers. This can confound the estimation of F_X, since the bidding behavior of the buyers depends upon the distribution of seller valuations. Also, the researcher must recover two distributions instead of one. The authors simplify matters by assuming that seller always sets the reserve price equal to her valuation, which is optimal in their framework. They also assume that the seller valuation, and hence her behavior, is independent of the buyers' valuations. Their method relies heavily upon the independence assumption.

Since the two-step approach is based on first order conditions, it is easily modified in an IPV model to deal with asymmetric bidders. Bidder *1* with signal x_1 chooses his bid, given that there are n potential bidders, to maximize

$$\pi_1(b, x_1) = (x_1 - b) \prod_{j=2}^{n} F_j(\eta_j(b))$$

His optimal bid solves

$$(\eta_1(b) - b) \sum_{j=2}^{n} \frac{f_{X_j}(\eta_j(b))\eta_j'(b)}{[1 - F_{X_j}(\eta_j(b))]} = 1.$$

Using the same change of variables as above, each bidder's inverse bid function can be expressed as a function of his rivals' bid distributions,

$$\eta_i(b) = b - \frac{1}{\displaystyle\sum_{j \neq i} \frac{g_{B_j}(b)}{[1 - G_{B_j}(b)]}}.$$

In this case, the sample of pseudo-values has to be generated separately for each bidder. The samples can then be used to estimate the bidders' valuation distributions. Applications of this approach include Bajari and Ye (2003), who study collusive bidding in auctions of highway repair contracts in Minnesota, and Bajari, Houghton, and Tadelis (2006), who study auctions of highway repair contracts in California. The latter paper argues that researchers need to be careful in defining contract revenues since, in many cases, subsequent contract changes or renegotiations lead to payment adjustments. The authors assume that ex post adjustments, and ex ante beliefs about these adjustments, are the same for all bidders, and do not depend upon their costs or bids, and hence upon their private information. Given this assumption, contract revenues can be defined as the amount bid plus ex post adjustments, where the latter serves as a proxy for the bidders' expectations, and bidder markups can be estimated accordingly. The authors find that only a fraction of the ex post adjustments are passed through to bids, and attribute the difference from full pass-through to transaction costs.

Krasnokutskaya (2004) and Athey, Levin, and Seira (2004) extend the two-stage approach to allow for auction characteristics that are observed by the bidders but not by the econometrician. The unobserved heterogeneity accounts for the positive correlation in bids within an auction. This is important because failure to account for this correlation leads to an upward bias in bidder markups. Krasnokutskaya takes a semi-parametric approach. In her application to Michigan highway procurement contracts, she assumes that bidder i's costs take the form

$$c_i = x_i v.$$

Here the private signal X_i is independent of X_j, but not necessarily identically distributed, and also independent of the common, and commonly known, component V. The multiplicative structure implies that bidder i's equilibrium bid strategy can be decomposed into a component that is common to all bidders and an idiosyncratic component that is a function of his private signal:

$$\beta(c_i) = v\alpha_i(x_i).$$

Let A_i denote the idiosyncratic component. Krasnokutskaya shows that, under mild regularity conditions, the distribution functions of V and the A_i's are uniquely determined from the joint distribution of two bids. Thus, the probability distribution functions of V and the A_i's can be obtained non-parametrically, based on bid data from auctions with at least two bidders. To generate a sample of pseudo values for bidder i, first draw a random sample of L pseudo bids $\{a_i\}_{i=1}^{L}$ from the estimated distribution of A_i and then use the inverse bid function under the assumption that $v = 1$ to infer the pseudo value $\widehat{x}_i$ associated with each a_i. The sample of pseudo values can then be used to nonparametrically estimate F_{X_i}. Note that, while the distributions of V and the A_i's are identified, one cannot uniquely decompose any individual bid into its two components.

Athey, Levin, and Seira (2004) study timber auctions. They distinguish between two types of bidders, mills and loggers. Mills are thought to have stochastically higher willingness to pay for cutting rights than loggers. Athey, Levin, and Seira adopt a parametric approach. They assume that the distribution of bids for each bidder type conditional on V, the unobserved auction characteristic, is Weibull, and that V has a Gamma distribution. Here V is assumed to be independent of the observable auction characteristics, including the number of bidders. After integrating V out of the likelihood function, the parameters of the distributions of B_i and V are estimated by maximum likelihood. The procedure yields estimates of the conditional distributions of bids and inverse bid functions for loggers and mills. Since v is not observed, it is not possible to use the inverse bid function to infer a bidder's values from his bids. However, one can still recover the conditional distribution function of values for loggers, $F_{X_L|V}$, and for mills, $F_{X_M|V}$, from the identities

$$F_{X_k|V}(x|v) = G_k(\eta_k(b,v)|v), \text{ for } k = L, M.$$

The average bid functions and the unconditional distributions of bidder values are obtained by integrating $\beta(x,v)$ and $F_{X_k|V}$ over v, for $k = L, M$.

4.3 Identification

Laffont and Vuong (1996) discuss identification in first-price sealed bid auctions. In the symmetric common values model, the optimal bid by bidder *1* with signal $x_1 = \eta(b)$ satisfies

$$w(\eta(b), \eta(b)) = b + \frac{G_{M_1|B_1}(b|b)}{g_{M_1|B_1}(b|b)} \equiv \xi(b, G).$$

Recall that in the case of independent private values, this equation reduces to

$$\eta(b) = b + \frac{G(b)}{(n-1)g(b)}.$$

It then follows in the IPV case that F_X is identified from G if η is strictly increasing and $\lim_{b\downarrow r}\eta(b) = r$, i.e., there is not a mass point in the bid distribution at r. The identification condition requires G/g to be well-behaved, which is an equilibrium requirement of the data. Furthermore, a distribution G can be rationalized by F if and only if η satisfies the above properties. A similar argument establishes that affiliated private value models are also identified.

In the common values model,

$$w(\eta(b), \eta(b)) = E[u(x, V)|X = \eta(b), Y = \eta(b)].$$

Thus, $\xi(b, G)$ identifies bidder 1's expected utility conditional on the event that his signal is $\eta(b)$ and the maximum signal of his rivals is also $\eta(b)$. This conditional expectation is invariant to any increasing transformation of the signal. Even in a pure common values model, if X is normalized so that $E[X|V = v] = v$, mean-preserving transformations of $F_{X|V}$ cannot be distinguished in the data. Hence, knowing the value of $w(\eta(b), \eta(b))$ is not sufficient to identify the value of the signal $x = \eta(b)$, and G does not identify F. In fact, as Laffont and Vuong note, by defining a new signal

$$\widetilde{x} = w(x, x)$$

and utility $u(\widetilde{x}) = \widetilde{x}$, one can transform the affiliated common values (ACV) model into an affiliated private value (APV) model. If the reserve price is not binding, then any symmetric ACV model is observationally equivalent to some symmetric APV model. Laffont and Vuong extend this result to asymmetric models.

In summary, Laffont and Vuong (1996) and Athey and Haile (2002) have shown that, while private value models are often identified, common value models are typically not identified, at least not without imposing strong restrictions on the primitives or on the type of data that are available. The lack of identification challenges the usefulness of the structural program for the class of common values models. One alternative, discussed by Hendricks, Pinkse, and Porter (2003) in their study of auctions of offshore oil and gas leases, is to assume a pure common value model and augment bid data with ex post data on lease values. In this case, one can resolve the identification problem by imposing a moment restriction on the joint distribution of (X_{it}, V_t). For example, Wilson (1977) adopts the normalization $E[X_{it}|V_t = v] = v$, where signals are measured so that the mean signal on a tract is equal to the tract's value. Hendricks, Pinkse, and Porter instead assume that

$$E[V_t|X_{it} = x, Z_t = z] = x.$$

The condition states that if a firm obtains a signal x on a tract with observable characteristics z, then the expected value of that tract is equal to the value of the signal. Signals are normalized in terms of ex post value and the bidders' posterior estimates are assumed to be correct on average. Identification of the bid function follows immediately from the above condition and monotonicity of the bid function, since

$$\begin{aligned} x = E[V_t | X_{it} = x, Z_t = z] &= E[V_t | B_{it} = \beta(x, z), Z_t = z] \\ \implies \quad \eta(b, z) &= E[V_t | B_{it} = b, Z_t = z]. \end{aligned} \tag{4}$$

The inverse bid function can be estimated as follows. For every bid level b on a tract with characteristics z, define a neighborhood $B(z)$ of b, and compute the average ex post value of all leases with characteristics z that received a bid in $B(z)$. To implement this idea, the authors employ a kernel estimator of the mean ex post value in the neighborhood of any bid b for tracts with similar characteristics. The inverse bid function can be used to generate the sample of signals $\{\{\widehat{x}_{it}\}_{i=1}^{n_t}\}_{t=1}^{T}$which, together with the sample of ex post values $\{v_t\}_{t=1}^{T}$, can be used to identify F.

Within the class of private value models, there is a second identification problem that remains an open issue. Can the APV model be distinguished from an IPV model with unobserved heterogeneity? In the former model, the correlation in bids is generated by a common factor that is not observed by the bidders; in the latter model, the correlation is generated by a common factor that is observed by the bidders but not by the econometrician. Krasnokutskaya (2004) conjectures that the two models are observationally equivalent.

5 Tests of Private Versus Common Values

The identification results make it important to develop ways of distinguishing between private and common value environments. Another reason is auction design. If the idiosyncratic component of valuations is the main determinant of bids, and values are independent, then under the standard logic of competitive markets more competition implies lower procurement costs. On the other hand, if the common component is the main determinant, more competition does not necessarily lead to lower procurements costs because bidders have to worry about the winner's curse. Restricting the number of bidders can then be a revenue enhancing policy.

Paarsch (1992) attempts to distinguish between a PCV model and an IPV model by estimating the structural parameters of each model and comparing the models using a non-nested hypothesis test. However, he has to adopt restrictive parametric assumptions on the distribution of bidder valuations in order to estimate the structural parameters of the IPV and PCV models. More recently, nonparametric tests have been developed that exploit the fact that the APV and ACV models are not observationally equivalent if reserve prices are binding or if the number of bidders varies across auctions.

In a sealed bid, second-price auction with a binding reserve price, the marginal participant's valuation conditional on winning is

$$r = E[u(V,x)|X_1 = x^*, Y_1 < x^*]$$

but he bids

$$\beta(x^*) = E[u(V,x)|X_1 = x^*, Y_1 = x^*].$$

These two equations yield an interesting, and testable, prediction, first noted by Milgrom and Weber (1982). The bid submitted by a bidder with signal equal to the screening level x^* is equal to the reserve price r if values are private, but $\beta(x^*)$ is strictly larger than r if values are affiliated. In the latter case, the lower bound of the support of the distribution of submitted bids in the second-price auction exceeds r. Moreover, the lower bound is increasing in the number of bidders and in any variable that is affiliated with V. Similarly, in an English auction with common values, the support of the distribution of prices should exhibit a gap above the reserve price r. When only one bidder is active, the price is r. When two or more bidders are willing to bid, the price is strictly above r.

Hendricks, Pinkse, and Porter (2003) show that an analogous result holds in first-price auctions. In a private values model, the function ξ defined in Section 4.3 must satisfy the boundary condition, $\lim_{b\downarrow r}\xi(b, G(b)) = r$, which implies

$$\lim_{b\downarrow r}\frac{G_{M_{it}|B_{it}}(b|b)}{g_{M_{it}|B_{it}}(b|b)} \longrightarrow 0.$$

By contrast, in a common values environment, ξ is discontinuous at r, which implies that

$$\lim_{b\downarrow r}\frac{G_{M_{it}|B_{it}}(b|b)}{g_{M_{it}|B_{it}}(b|b)} \longrightarrow c > 0$$

for some constant c. Therefore, it is possible in principle to distinguish between the two models by examining the behavior of $G_{M_{it}|B_{it}}$ near the reserve price. In practice, reserve prices are often set so that the number of bids near the reserve are too few to implement the test with confidence.

Haile, Hong, and Shum (2004) develop a test that exploits exogenous variation in the number of bidders. Let $\widehat{G}_n$ denote an estimate of $G_{M_i|B_i}$ for auctions with n bidders and define

$$\widehat{x}_{it} = \xi(b_{it}, \widehat{G}_n(b_{it}))$$

as the pseudo-value corresponding to bidder i's bid in auction t. Under the hypothesis of private values, this is an estimate of bidder i's valuation of item t but, under the hypothesis of common values, it is an estimate of the latent conditional expectation $w(\eta(b_{it}), \eta(b_{it}))$. This expectation is decreasing in n if the common component is important, since winning against more bidders is worse news. Let $\widehat{F}_n$ denote the estimate of the empirical distribution of pseudo-values in auctions with n bidders. Haile, Hong, and Shum propose testing whether $\widehat{F}_n$ is invariant to n, as implied by the private value hypothesis, against the alternative that $\widehat{F}_n$ is strictly increasing in n, as implied by the common values hypothesis. A similar test can be applied to bids in SPSB auctions and to losing bids in English button auctions. Note that a maintained hypothesis is that the variation in the number of bidders is exogenous, and hence independent of the distribution of values. Haile, Hong and Shum propose an instrumental variables procedure for cases where the number of bidders is endogenous.

Some early studies test private versus common values in a FPSB format by testing for monotonicity of the bidding strategy β in the number of bidders n. In an independent private values environment, β is strictly increasing in n, as competition causes bidders to bid more aggressively. Recall that the optimal strategy in a FPSB auction with IPV and a minimum bid r can be expressed as

$$\beta(x) = E[\max\{r, X_{2:n}\}|X_{1:n} = x].$$

This expectation is increasing in n, as it is the expectation of a monotone function of the highest order statistic from a sample of size $n-1$. But what does theory predict about the behavior of β with respect to n in a common value environment?

Wilson (1993) discusses several examples in which equilibrium bid strategies may be decreasing or non-monotone in the number of potential bidders. Consider the pure common value example in which F_V and $F_{X|V}$ are lognormal distributions and the reserve price is not binding. Wilson (1993) then shows that if F_V, the prior distribution of the common value, is diffuse,

$$\beta(x_i) = k(n)x_i,$$

so that the equilibrium bid is proportional to the signal, where the factor of proportionality depends on the number of bidders and the informativeness of their signals. Signals are more informative the lower the variance of $F_{X|V}$, and bidding is more aggressive the lower this variance. Wilson provides some examples in which $k(n)$ is increasing for small n, and then decreasing after $n = 2$ or 4. When the number of bidders is small, increasing competition leads to more aggressive bids. But this effect diminishes as the number of bidders increases, and eventually winner's curse considerations dominate the competitive effect. There are no general results, however.

It is important not to focus on the winning bid in the above exercises, since the expected winning bid in a CV environment, $E[\beta(X_{1:n})]$, is monotone increasing in n for many examples. (There are exceptions, however. Bulow and Klemperer (2002) provide one counter-example.) One should use the vector of all bids or a summary statistic, such as the average bid. Several studies have examined the sale of oil and gas leases. The basic idea of the test is to regress bids or average bid on the number of bids submitted n, and n^2, or some other function of n, and Z, a vector of observable characteristics (e.g., area and date of sale, tract acreage, water depth, etc.). An IPV model often matches the data better than CV, in the sense that β is increasing in n over the observed range of n (typically, from 1 to 18 bids).

Hong and Shum (2002) study the effects of winner's curse considerations on bidding in procurement auctions for highway repair contracts and quantify its importance. They use the Wilson (1998) log-additive model of preferences, which has the PV and PCV models as special cases. Identification is achieved by adopting parametric distributional assumptions. Given the functional form assumptions, they follow Laffont, Ossard, and Vuong (1995) and use simulation methods to generate a bid distribution from the equilibrium relationship between bids and signals. But, instead of trying to match the mean of the bid distribution, Hong and Shum exploit the monotonicity of the bid function and focus on matching the quantiles of the equilibrium bid distribution and the actual bid distribution. Simulating the quantiles of the equilibrium bid distribution is substantially easier than simulating its moments. They then show that procurement costs rise as competition intensifies, consistent with a CV model but not with private values.

The main limitation of tests that exploit the variation in the number of bidders is that, in most data sets, the number of bidders is endogenous. The problem is that the number of bidders is often correlated with unobserved characteristics of the item. For example, in the Hong and Shum study on highway contracts, the basic empirical regularity observed in the data is that bids are higher on contracts with more bidders. If the contracts are identical, then this fact can be attributed to the winner's curse, which counsels less aggressive bidding as the number of competitors increase. However, contracts are not identical, and the larger, more valuable contracts are likely to attract more bidders. The absence of good covariates like engineering estimates to control for contract heterogeneity means that their results may be due to the fact that the number of bidders is an imperfect proxy for the contract size. A similar problem arises with bidding for oil and gas auctions. Even if tract characteristics include an ex post proxy for V, ex ante bidder errors in expectation will be positively correlated with the number of bids submitted. If bidders are optimistic, more bid, and they bid aggressively. One solution is to employ a proxy or instrument for the number of potential bidders, but an instrumental variable may not be available. Many factors that are correlated with the participation decisions of firms will also be correlated with their valuations, and hence their bid levels. In their study of timber auctions, Haile, Hong and Shum adopt an instrumental variable strategy, exploiting variation in the number of nearby mills.

A second problem is that the comparative statics is with respect to a known number of *potential* bidders, not the number of submitted bids. If there is a binding reserve price, or if the support of V falls below zero, not all of the potential bidders will necessarily bid, and then there is a endogeneity problem. In the studies that test the monotonicity of the bid function in the number of bidders, the bias is towards the private values model, to the extent that the number of bids is positively correlated with bid levels. For example, in the oil and gas auctions, a potential bidder is more likely to submit a bid if tract is more likely to contain oil. Haile, Hong, and Shum propose a clever test that exploits the fact that bidder's probability of participation does not depend upon the number of bidders in a private value environment but decreases with the number of rivals in a common value environment.

There is also a serious issue in many applications about whether the number of potential rivals is known by the bidders. In the OCS auctions, a bidder is classified as a potential bidder on a tract if it conducts detailed seismic studies. But bidders do not necessarily conduct such studies on every available tract. Furthermore, the decision to conduct a detailed analysis is essentially private. Thus, a firm may not know which tracts have been searched by its rivals. In this case, if the probability of search is higher on tracts that are perceived to be more likely to contain oil, the number of bids is likely to be positively correlated with bid levels even after controlling for tract values. The reason is the same as that given above.

Finally, the result that in a private values setting bids are increasing in the number of bidders relies on the assumption that private values are independently distributed. Pinkse and Tan (2005) show that if private valuations are instead affiliated, then equilibrium bids in a first price sealed bid auction can be decreasing in the number of bidders. Hence, bidding strategies that are decreasing in the number of bidders are not necessarily inconsistent with private values.

6 Tests of the Theory

Structural models force a commitment to the theoretical model. Parametric methods entail joint hypotheses regarding preferences, behavioral assumptions, and the functional form of the value distribution. Nonparametric methods relax the last assumption, but retain the preference and behavioral assumptions. The behavioral assumption of noncooperative bidding is crucial. If some bidders are colluding, some bids may be phony. Also, the effective number of bidders may be more than actual number submitting bids, due to coalitions. The existence of bidding coalitions may violate the symmetry assumption, since coalitions may have different objectives induced by their formation (when individual rationality or incentive compatibility constraints bind), or by the aggregation of their preferences or their information.

The theory of equilibrium bidding under the assumption of affiliated signals does deliver two tests that do not depend upon whether values are private or common. In first-price auctions, monotonicity of the equilibrium bid function implies that

$$\xi(b, \widehat{G}) = b + \frac{\widehat{G}_{M_i|B_i}(b|b)}{\widehat{g}_{M_i|B_i}(b|b)}$$

is increasing in b, and the boundary condition on the equilibrium bid function implies that

$$\lim_{b \downarrow r} \xi(b, \widehat{G}) \geq r.$$

These are consistency checks, akin to overidentifying restrictions. The monotonicity test has low power, however, as it can be satisfied by non-equilibrium strategies. For example, in a private values model, consider the strategy of bidding one's value, $\beta(x) = x$ if $x \geq r$. This strategy is monotone, and it satisfies the boundary condition $\eta(r) = r$. No such test is available in second-price auctions, where bids are simply interpreted as conditional expectations.

The second test is a variant of the test proposed by Haile, Hong, and Shum for distinguishing between private and common value models. In first price auctions, equilibrium bidding implies that the empirical distribution of pseudo-values, $\widehat{F}_n$, is non-decreasing in n. If this is not the case, then either the data are not consistent with equilibrium bidding or valuations are not common. A similar test applies to bids in second-price auctions and to losing bids in English auctions. Note that this test requires the number of bidders to be exogenous.

The lack of ex post data on bidder values makes it difficult to test the theory in private value auctions using field data, absent some alternative hypothesis. Bajari and Hortacsu (2005) use bid data from private value auction experiments to test the theory. Experimental data has two advantages over field data. First, the variation in the number of bidders is exogenous. Second, the realizations of the private values of the bidders, and the value distribution itself, are known. The main disadvantages are that bidders may not have sufficient experience, and the stakes are low. The authors evaluate the performance of several models of bidding by comparing the estimated valuations under the various models to the actual values. For example, in the case of risk neutral bidders with IPV, the estimated values are generated according to $\widehat{x} = \xi(b, \widehat{G})$. They find that the model that performs best is Bayesian Nash equilibrium with risk averse bidders.

6.1 Pure Common Value Auctions

The theory of bidding in pure common value auctions can be tested when data on ex post values are available. Much of the early empirical interest in auctions with common values focused on the possibility that bidders may bid naively and be afflicted by the winner's curse. A (rather trivial) test of rationality, under which bidders anticipate the information revealed by winning, looks at individual rationality constraints. In private value first price sealed bid laboratory experiments, where valuations are observed by the researcher, individual rationality simply means not bidding more than the value. (Violations of this form of rationality are especially distressing. If the stakes are low, subjects may place too much value on winning for its own sake.) In field studies of common value auctions, the test often reduces to checking whether realized profits are positive.

Capen, Clapp and Campbell (1971) is perhaps most influential empirical auction paper addressing this issue. The authors compute ex post returns on offshore oil and gas auctions in the early years of these sales. They claim that bidders suffered from the winner's curse since, by their calculations, ex post returns were negative, i.e., the internal rate of return was less than that of Treasury bills. Their finding suggests that bidders violated the basic tenets of rational bidding. However, there are several reasons to doubt this conclusion. First, their measure of ex post returns is based on incomplete production histories. The wells they study were productive for many more years. Mead and his coauthors (1980) compute, based on longer well histories, real internal rates of return of approximately 7 percent. Hendricks, Porter and Boudreau (1987) also find positive returns for most firms on wildcat tracts. In addition, negative ex post returns in small samples are not necessarily indicative of irrationality, if some unpredictable adverse common payoff shock occurred.

In a series of papers, we (with a number of different coauthors) have used data on bids, drilling costs, and oil production from federal sales of oil and gas leases on the Outer Continental Shelf (OCS) to study the impact of the winner's curse on bidding behavior and to re-examine the rationality issue. Oil and gas leases are classified into two categories. Wildcat tracts are located in previously unexplored areas. Prior to a wildcat auction, firms are allowed to conduct seismic studies, but they are not permitted to drill any exploratory wells. The seismic studies provide noisy, but roughly equally informative signals about the amount of oil and gas on a lease. We argue that wildcat auctions are likely to satisfy the symmetry assumption on the signal distribution. Drainage leases are adjacent to wildcat tracts where oil and gas deposits have been discovered previously. Firms that own adjacent tracts possess drilling information that makes them better informed about the value of the drainage tract than other firms, who are likely to have access only to seismic information. We argue that these auctions can be modeled by assuming one bidder has a private, informative signal and all other bidders have no private information.

For auctions of both types of leases, we ask the following question: do bidders behave as predicted by game theoretic models? In the case of drainage auctions, we study this question by examining whether the predictions of the theory are consistent with the data.

In auctions with asymmetric information, knowing which bidders have better information allows us to generate testable predictions of equilibrium bidding based on bidders' identities. In the case of wildcat auctions, we address this question by testing whether the first-order conditions for optimality hold.

6.1.1 The Asymmetric Case

In the data set on drainage lease bidding, we can distinguish between firms owning adjacent wildcat leases (neighbors) and others (non-neighbors). This observable asymmetry between potential bidders appears to matter. Compared to wildcat auctions, drainage auctions have less entry (i.e., fewer bids submitted per tract), yet higher ex post returns. Drainage leases are more likely to be explored, and much more likely to be productive if explored. Although drainage tracts receive higher bids, their higher productivity translates into higher profits. Some numbers for the sample period 1954-79 are as follows (Porter, 1995):

	Wildcat	**Drainage**
Number of tracts	2,510 (2,255 sold)	295 (237 sold)
Number of bids	3.52	2.45
Mean high bid	\$5.5 million	\$7.4 million
Number of tracts drilled	1,757 (78%)	211 (89%)
Number of productive tracts	881 (50%)	147 (70%)
Discounted revenues if productive	\$19.5 million	\$17.3 million

What is the explanation for lower entry and higher average returns on drainage leases? Two alternatives come to mind. One is asymmetries in information: neighbor firms are better informed about the value of the drainage tracts than non-neighbors. An alternative explanation is asymmetries in payoffs: neighbors have lower costs due to economies of scope. Either alternative can explain why entry does not drive ex post returns down to wildcat levels. Entry of non-neighbors might drive their own returns down to zero, but asymmetries protect the rents of neighbors, especially if they coordinate their bids.

Consider an environment in which one bidder, which we will refer to as I for informed, has a private, informative signal, X, concerning the unknown common value, V, and another bidder, called U for uninformed, has access only to public information. We will refer to I as the neighbor firm (if there is more than one, we assume that they collude) and U as a non-neighbor firm. In the OCS drainage auctions, more than one non-neighbor firm might potentially bid, but an important implication of equilibrium bidding is that the distributions of the neighbor bid and the maximum of the non-neighbor bids do not depend upon the number of uninformed bidders. (This result follows from the assumption that non-neighbor firms do not have any private information.) Hence we can assume without loss of generality that there is only one non-neighbor firm. Assume that r is the announced minimum bid, or reserve price, and that $E[V] > r$. If $E[V] < r$, then non-neighbor firms would never submit a bid.

The non-neighbor firm faces an acute case of the winner's curse in trying to win the tract from the better informed neighbor firm at a profitable price. Should it bid and, if so, how should it bid? Engelbrecht-Wiggans, Milgrom, and Weber (1983) address this question. They show that U has to participate in equilibrium. If U does not participate, then I would bid the reserve price whenever $E[V|X = x]$ exceeds r. But, in that case, U could bid slightly more than r, win the auction for certain, and obtain a positive payoff of $E[V] - r$. Second, U has to bid randomly. If U used a pure strategy and bid $b > r$, then I's best reply is to bid slightly more than b if $E[V|X = x]$ exceeds b. But then U would suffer an extreme form of the winner's curse, as it would win only if $E[V|X = x]$ is less than b, i.e., the more informed bidder expects V to be less than U's bid b. Finally, U earns zero expected profit from its mixed strategy. In other words, there is a positive probability that U does not bid.

Consider next I's behavior and let G_U denote U's mixed strategy. The payoff to I from a bid of $b \geq r$ is given by

$$\pi_I(b, x) = (E[V|X = x] - b)G_U(b).$$

This payoff depends on the private signal X only through the expectation of the value of the tract. Define $H = E[V|X]$ as the sufficient statistic for I's information and let F_H denote its distribution. (Note that $E[H] = E[V]$, by the law of iterated expectations.) Then I's payoff can be written as

$$\pi_I(b, h) = (h - b)G_U(b),$$

and a pure bid strategy is a function $\beta_I(h)$. Let η_I denote its inverse.

As far as U is concerned, $\beta_I(H)$ is a random variable, which depends on the realization of the H, with distribution function $G_I(b) = F_H(\eta_I(b))$. Consequently, U's payoff from bidding $b \geq r$ is

$$\pi_U(b) = E[H - b|H < \eta_I(b)]F_H(\eta_I(b)).$$

Since I has to bid in such a way that U earns zero profits at any bid above r, its equilibrium bid strategy is given by

$$\beta_I(h) = \begin{cases} 0 & \text{if } \ h < r \\ r & \text{if } \ h \in (r, \widehat{h}) \\ \quad E[H|H < h] \ = h - \frac{\int^h F_H(s)ds}{F_H(h)} & \text{if } h > \widehat{h} \end{cases}$$

where $\widehat{h}$ satisfies $E[H|H < h] = r$. Clearly β is monotone increasing in h for $h \geq \widehat{h}$. Therefore, $G_I(b) = F_H(\eta_I(b))$ for bids above the reserve price. Furthermore, as $h \longrightarrow \infty$, $\beta(h) \longrightarrow E[V]$. Therefore, the positive support of G_I is $(r, E[V])$. The remaining probability is concentrated at 0 (no bid) and at the reserve price r. The probability mass at these two bids are $F_H(r)$ and $F_H(\widehat{h}) - F_H(r)$, respectively.

We turn next to characterizing U's mixed strategy. It has to bid so that I wants to bid according to β. Since $\beta(h)$ maximizes

$$\pi_I(b, h) = (h - b)G_U(b) \ \text{ s.t. } \ b \geq r,$$

$\beta(h)$ must solve

$$(h - b)g_U(b) - G_U(b) = 0 \ \text{ for all } b > r.$$

Substituting $\beta_I(h)$ for b, this first order condition can be expressed as

$$\frac{g_U(\beta_I(h))\beta_I'(h)dh}{G_U(\beta_I(h))} = \frac{f_H(h)dh}{F_H(h)}$$

Note that this equation holds for $h \geq \widehat{h}$. Integrating both sides of this equation from $\widehat{h}$ to h and applying the boundary condition $\beta_I(\widehat{h}) = r$, yields

$$\frac{G_U(\beta_I(h))}{G_U(r)} = \frac{F_H(h)}{F_H(\widehat{h})} \ \text{ for any } h \geq \widehat{h}.$$

Since $\lim_{h \longrightarrow \infty} G_U(\beta(h)) = 1$, it follows that $G_U(r) = F_H(\widehat{h})$ and hence $G_U(\beta_I(h)) = F_H(h)$. Equivalently, $G_U(b) = F_H(\eta_I(b))$ for $b > r$. Continuity of expected profits implies that both bidders cannot have a mass point at r. Hence, the remaining probability of G_U is concentrated at 0. If there are n uninformed bidders, then G_U is the distribution of the maximum uninformed bid.

The equilibrium generates a number of testable predictions for drainage auctions. First, non-neighbor firms should have a lower aggregate participation rate than the neighbor firms.

In the data, the participation rates for these two types of firms are 57% and 87% respectively. Second, neighbor firms should have a higher win rate than non-neighbor firms. In the data, neighbor firms account for 57% of the high bids. Third, non-neighbor firms should earn zero profits on average. The data are consistent with this prediction, at least prior to 1973, assuming a 5% real discount rate. Fourth, non-neighbor firms should earn negative expected profits on leases where neighbor firm(s) did not bid, since no neighbor bid is bad news. This pattern is found in the data, as the non-neighbor firms incurred significant (both statistically and economically) losses when no neighbor firm bid. Finally, the distribution of neighbor bids should have a mass point at r and should coincide with the distribution of the maximum non-neighbor bid above r. This last prediction is not consistent with the data. The two distributions coincide at higher bid levels but non-neighbor firms submitted relatively few low bids. (There were few low non-neighbor bids of any sort, and not just few low maximum non-neighbor bids.) Hendricks and Porter (1988) provide more detail.

We conclude that the simple theory of asymmetric information and common values can account for several features of the data, but not all. The simple theory ignores one important institutional feature, namely, that the government reserves the right to reject the high bid. In 1959-79, the government rejected 58 of 295 high bids on drainage leases, approximately 20%. Rejection usually occurred when only one bid was submitted and the value of the bid was low.

We model the rejection policy as an unknown reserve price R with lower bound r, the announced reserve price, and upper bound $\bar{r}$. The high bid is rejected if $b < R$. The secret reserve price complicates the analysis considerably. The payoff of I is

$$\pi_I(b,x) = (E[V|X = x, b \geq R] - b)\Pr\{b \geq R|X = x\}G_U(b).$$

In this case, H is no longer a sufficient statistic for optimal bidding behavior. The private signal X may contain information that is relevant to the distributions of both V and R. Hendricks, Porter and Wilson (1994) assume that (V, X, R) are affiliated. This assumption implies that (i) $E[V|X, R]$ is non-decreasing in X and R, and (ii) the distribution $F_{R|X}$ satisfies the monotone likelihood ratio property with respect to realizations of X. These two properties are used to establish that $\beta_I(X)$ is non-decreasing. This in turn implies that (V, R, β_I) are affiliated and permits a characterization of the equilibrium bid distributions.

The main effects of the secret reserve price are captured in the special case where R is independent of (V, X), in which case H is a sufficient statistic for X. Let G_R denote the distribution of R. Then I's payoff from a bid of b is now

$$\pi_I(b,h) = (h-b)G_U(b)G_R(b).$$

Define

$$\beta_0(h) = \arg\max(h-b)G_R(b)$$

as the optimal strategy of I absent competition from uninformed bidders. The optimal bid function with U present is $\max\{\beta_0(h), \beta_I(h)\}$. In some cases, $\beta_0(h) > \beta_I(h)$ when $h < \widetilde{h}$ for some $\widetilde{h} > r$. Then more aggressive bidding by I on the interval $[r, \widetilde{h})$ means that $\pi_U(b) < 0$ for $b \in (r, \widetilde{b})$ where $\widetilde{b} = \beta(\widetilde{h})$.

The predictions for the bid distributions have to be modified as follows: (i) G_I does not have a mass point if $G_R(r)$ is sufficiently close to 0; (ii) $G_U(b) \leq G_I(b)$ with strict inequality for $b \in (r, \overline{r})$ (stochastic dominance) and (iii) $G_U = G_I$ for $b > \overline{r}$. In the data, the bid distributions essentially coincide above \$4 million. The probability of rejection at this bid level is very low. Stochastic dominance is evident, as depicted in Figure 1 of Hendricks, Porter and Wilson (1994), and significant. We conclude that the asymmetric information common value model with a secret reserve price can account for the main features of the data.

What about the assumption that there is only one neighbor? In the sample, the average number of neighboring wildcat tracts is 3.78. Yet non-neighbor bidding is independent of the number of neighbor firms. If neighbor firms are symmetrically informed, and if they bid competitively, then non-neighbor firms with no private information should not bid. At a minimum, one would expect less aggressive non-neighbor bidding as the number of neighbors increases, as the winner's curse is more severe. One could view participation by non-neighbor firms as evidence of coordination on the part of neighbor firms. Further, the high neighbor bid is independent of number of neighbors, and on tracts where more than one neighbor bids, neighbor profits are much higher. This evidence suggests coordination by the neighbor firms, with occasional phony bids to guard against rejection. Alternatively, the neighbor firms may themselves be asymmetric in some way. Note that it is relatively easy for neighbor firms to coordinate, as they could allocate production as a means of making side payments.

What about the alternative model in which there are cost asymmetries rather than information asymmetries? Several facts argue against this model. First, to avoid common pool problems, production on adjacent tracts is often unitized, with revenues allocated in proportion to area owned above the common pool. Second, non-neighbor firm profits are approximately zero. On tracts won by non-neighbors when neighbors do not bid, profits are negative because tracts are less likely to contain oil or gas deposits. This is inconsistent with unobserved cost differences being the predominant source of asymmetry.

The tests described above are reduced form. There is a role for structural estimation in this environment, in designing the optimal selling mechanism. But knowing only that information is asymmetric, and who has superior information, can be sufficient for optimal mechanism design. Hendricks, Porter and Tan (1993) discuss this issue.

6.1.2 The Symmetric Case

For wildcat tracts, Hendricks, Pinkse and Porter (2003) implement several tests of bidder rationality by comparing bids with ex post outcomes. One basic test is that actual rents, $v_t - w_t$, measured as the difference between discounted net revenues and the winning bid, should on average be positive. A second related test is that firms should expect to earn positive rents conditional on submitting the winning bid. To compute this expectation, we estimate the function

$$\phi(b) = E[V_t | B_{it} = b, M_{it} < b],$$

and evaluate it at b_{it}. The profit margin that firm i expects to earn upon winning tract t is $\phi(b_{it}) - b_{it}$. Averaging across firms and tracts, this profit margin should be positive. We find that actual rents and expected rents conditional on winning are significantly positive, and there is no evidence of adverse selection associated with winning. The magnitude of the actual rents are approximately equal to total entry costs, so it appears that on average firms earn zero profits.

Did bidders bid less than their expected tract value? To address this question, we estimate the function

$$\varsigma(b) = E[V_t | B_{it} = b]$$

and test the restriction that $\varsigma(b) > b$ at all bid levels. We also test the stronger restriction that firms should always bid less than the expected value of the tract conditional on winning, $\phi(b) > b$. We find that both $\varsigma(b) - b$ and $\phi(b) - b$ are significantly positive throughout the relevant range of bids, and furthermore that $\varsigma(b)$ is significantly larger than $\phi(b)$ at every bid level. The difference between these two functions is a measure of the informativeness of winning. This difference is greater on tracts with more potential bidders.

Our final test is a test of equilibrium bidding. Recall that, in a symmetric equilibrium, bids satisfy the first-order condition

$$E[V|X_i = \eta(b), Y_i = \eta(b)] = \xi(b, G).$$

We estimate the conditional expectation function,

$$\zeta(b) = E[V|B_i = b, M_i = b].$$

Since β is monotone, η is monotone as well, and hence these two equations imply that

$$\zeta(b) = \xi(b, G)$$

at any $b > r$. Equality of these two functions is an empirically testable implication of equilibrium bidding in wildcat auctions. We find that Bayesian Nash equilibrium cannot be rejected when competition is high (i.e., more than six potential bidders, by our measure), but equilibrium behavior is rejected when competition is low (six bidders or less). In the latter case, the data suggest that firms may have overbid.

In summary, we conclude that the winner's curse is evident in the data, but also that bidders are aware of its presence and bid accordingly, for the most part.

7 Revenues and Auction Design

The early theoretical work on auctions by Vickrey (1962) and Ortega-Reichert (1968) studies the four standard single object auction formats, namely, Dutch, English, FPSB, and SPSB, under the assumption that bidders are risk neutral, their valuations are private, and their signals are independent and identically distributed. The authors obtain a striking and surprising result: the expected revenues from all four formats are identical. In each case, the expected revenue to the seller in equilibrium is simply the expected value of $\max\{r, X_{2:n}\}$, where $X_{k:n}$ denotes the k^{th} highest order statistic among the random variables $\{X_1, ..., X_n\}$. (The reserve price r can be interpreted as the bid of the seller, and as the second highest estimate in the event that no rival submits a bid. The revenue comparison assumes that r is constant across auction formats.) Myerson (1981) and Riley and Samuelson (1981) subsequently demonstrated that the revenue equivalence result holds under bidder risk neutrality in any auction format in which the bidder with the highest valuation wins and the bidder with the lowest possible valuation, $\underline{x}$, has zero expected payoff.

A number of empirical studies have examined the prediction of revenue equivalence. The prediction has also received the attention of the experimental literature. One strand of the empirical literature compares bid data from different auction formats. Although there are not many instances where different auction formats are run in parallel for similar auction environments, one example is U.S. Forest Service timber auctions (Hansen (1986)). Hansen's econometric approach to test revenue equivalence is to regress winning bid on sales characteristics, Z_t, and an indicator variable

$$D_t = \begin{cases} 1 & \text{if auction t is FPSB} \\ 0 & \text{if auction t is English} \end{cases}.$$

Hansen finds that the coefficient on D_t is significantly positive. Revenues in the FPSB auctions are about 10% higher than in English auctions. But, as Hansen argues, the Forest Service does not select auction type randomly. If one corrects for sample selection, where the Forest Service region is employed as an instrument for auction format, the coefficient on D_t is approximately zero. Hansen concludes that one cannot reject revenue equivalence. Athey, Levin and Seira (2004) find that, in more recent periods, revenues are higher in the FPSB auctions, even correcting for sample selection.

There are other implications of equilibrium bidding that could be tested. Higher moments of the distribution of seller revenues are not the same across auction formats. In particular, Matthews (1980) shows the FPSB winning bid distribution has lower variance than revenue in the SPSB and English auctions. Therefore, one might examine higher order moments of the distribution of prices, or properties of the entire distribution of prices. For example, Haile's (2001) method of moments estimator for the distribution of valuations in an English auction (with an active resale market) employs information on both the first and second moments of the bid distribution. But we know of no instance in which comparative statics properties across auction formats concerning higher order moments have been examined.

Rejection of revenue equivalence then leads one to ask why it occurs. What feature of equilibrium bidding in the symmetric, risk neutral IPV model is being rejected? One possibility is risk aversion (Holt (1980)). Bidding one's valuation remains a dominant strategy in an English auction, but expected prices are higher in a FPSB, as risk averse bidders increase their bids to increase the probability of winning the auction. In a private values environment, bidders know their valuation net of their bid in the event they win, but are uncertain whether they will win the auction given their bid. Risk aversion therefore leads to more aggressive bidding. (In contrast, risk aversion may lead to less aggressive bidding in common value auctions, due to the winner's curse.) Haile (2001) argues that the presence of a resale market, and the consequent incentives to signal via bids in the original auction to affect subsequent resale negotiations, can account for FPSB revenue superiority.

Athey, Levin, and Seira relax the assumption of symmetric bidders. They assume that a mill's willingness to pay for cutting rights is stochastically greater than a logger's willingness to pay since the latter has to sell its logs to a mill. Under this assumption, revenue equivalence does not hold and the auction format matters. In open, ascending price auctions, the bidder with the highest valuation wins but, in first-price, sealed bid auctions, this need not be the outcome. Mills shade their bids further below their valuations than do loggers in a FPSB auction, because of their favorable distribution of values. As a result, loggers are more likely to participate, and win, in a sealed bid auction than in an open auction, and sealed bid auctions may yield greater revenue. The main difficulty the authors face in testing these theoretical predictions is, again, that the choice of auction format is not random. The authors deal with this problem by using a non-parametric matching estimator, and assuming that the choice of auction format is a function of observable sale characteristics. The latter assumption is plausible since the data provide a rich set of sale characteristics. After controlling for sale characteristics, they find that the participation and win rates of loggers are higher in sealed bid auctions than in open auctions. In addition, winning bids in sealed bid auctions are significantly higher in Montana and Idaho, but similar to winning bids in open auctions in California.

Athey, Levin, and Seira also explore the possibility that the FPSB revenue superiority in the Northern states may have been due to mills bidding less competitively in the open auctions. As we argue below, collusion can be easier to maintain in an English auction. To test this hypothesis, the authors use structural estimates of entry costs and the value distributions of loggers and mills from sealed bid auctions to predict average sale prices in the oral auctions under the assumption that mills behave competitively, and under the assumption that they bid collusively. For each tract, the mean parameter of the Poisson distribution determining the number of loggers is calibrated so that loggers are indifferent between entering and staying out. Thus, in contrast to other analyses of collusion in auctions which treat the number of bidders as fixed, collusion by mills affects the entry behavior of loggers. As a result, the impact of collusion on sale prices is not as large as it might have been had mills colluded in secret. Collusion is nevertheless profitable because loggers' valuations are stochastically lower than the mills' valuations. The main finding is that the average price in the Northern oral auctions lies below the average predicted price obtained under competition, and above that obtained under collusion. In California, average prices are not significantly different from the average predicted price obtained under competition.

The Athey, Levin, and Seira paper illustrates the usefulness of structural estimation for building counterfactuals. Shneyerov (2005) pursues a similar strategy in studying the effects of auction format on revenues in municipal bond auctions. The auctions are first-price sealed bid, and Shneyerov wants to predict the revenue changes from switching from the FPSB format to either second-price sealed bid or an English button auction. Milgrom and Weber (1982) show that, in affiliated common value environments, the expected revenues of first-price auctions are lower than the expected revenues of sealed bid, second-price auctions, which in turn are lower than the expected revenues of English button auctions. The idea is that the SPSB format reveals rival information by making the winner's payment equal the highest rival bid. In an English button auction, rival signals are revealed by their exit decisions, further allaying the winner's curse. Before constructing the counterfactual, Shneyerov estimates the empirical distributions of the random variables, $\xi(B, \widehat{G}_n)$, and tests the null hypothesis of private values against the alternative of affiliated values using one of the tests proposed by Haile, Hong, and Shum (2004). He rejects the null, so that revenue gains are possible. The difficulty that he faces in constructing the counterfactual is that, in auctions with common values, the distribution of latent bidder valuations is not identified. Shneyerov observes, however, that $\xi(b, \widehat{G}_n)$ is an estimate of the latent conditional expectation, $w(\eta(b), \eta(b))$, which is the amount that a bidder with signal $x = \eta(b)$ would bid in a second-price auction with n bidders. Hence, the pseudo-values generated by the two-step approach of Elyakime, Laffont, Loisel, and Vuong (1994) and Guerre, Perrigne, and Vuong (2000) can be interpreted as counterfactual bids in a second-price auction, and used to identify expected revenue from this auction format. Shneyerov finds that expected revenues from the second-price auction are approximately 20 percent higher than in the first-price auctions.

The second main result of Myerson (1981) and Riley and Samuelson (1981) is that the revenue maximizing mechanism in an IPV environment can be implemented by any of the standard auction formats with an appropriately specified reserve price. As we discussed in Section 3.1, the optimal reserve price for a seller who values the item at x_0 is the solution to:

$$r = x_0 + [1 - F_X(r)]/f_X(r).$$

The optimal reserve price in an IPV auction depends only on x_0 and F_X, and not the number of bidders. Given an estimate of F_X, the seller can choose r to satisfy the above equation. Paarsch (1997) and Haile and Tamer (2003) estimate the optimal reserve price in an IPV environment.

McAfee and Vincent (1992) and McAfee, Quan, and Vincent (2002) study the problem of setting optimal reserve prices in common value auctions when the seller does not have an estimate of F. This case is certainly relevant since, as noted earlier, F cannot be identified in these environments from bid data alone. McAfee and Vincent (1992) consider pure common value, first-price auctions and develop a distribution-free test statistic that the seller can compute from data on bids and ex post values to determine whether the reserve price is too low or too high. The test statistic allows for endogenous entry and stochastic participation. They apply the test to the data on federal offshore oil and gas auctions. They find that the federal government's revenue-maximizing reserve price was too low. In their 2002 paper, the authors extend the test to common value environments and to a broader class of auctions, including first price, second price and oral auctions. The extension does not require data on ex post values. They apply the test to real estate auctions.

7.1 Multi-Unit Auctions

The revenue equivalence theorem has been tested in sequential auctions under the assumption of unit demands. That is, each bidder's valuation refers to their valuation of the first unit they acquire, and the marginal value of any additional unit is assumed to be zero. Suppose that only the winning bid is revealed after each round. Let w_l denote the winning bid in the l^{th} auction. It can be shown that in an IPV environment with risk neutral and symmetric bidders who have unit demands, the prices in a k-item, sequential first price auction should form a martingale, that is $E[w_l|w_{l-1}] = w_{l-1}$ for $l = 2, ..., k$, and furthermore, that $E[w_l] = E[X_{k+1:n}]$ for all l. This claim is also true of a SPSB sequence. In an English auction, the price sequence depends on what information is revealed to the bidders.

Ashenfelter (1989) tests this arbitrage condition for wine auctions. Identical cases of wine were sold sequentially via English auctions by an auction house. He finds evidence of a "declining price anomaly." Prices systematically and significantly decline over the sale sequence, although the magnitude of the decline is not large, a few percent. This pattern appears to violate the martingale prediction, as well as a more general notion of the implications of no arbitrage opportunities. Other studies that have examined this prediction include Ashenfelter and Genesove (1992) and Beggs and Graddy (1997), who respcetively find that condominium and art prices also decline. In contrast, Donald, Paarsch and Robert (2006) find evidence of rising prices in auctions of Siberian timber export permits. Ashenfelter and Graddy (2003) provide a recent survey.

Is there a resolution to the anomaly? The martingale result relies on risk neutrality, the IPV assumption and the assumption of unit demands. A number of papers have shown that the conclusions are not robust. McAfee and Vincent (1993) study the effects of risk aversion. They show that in order to obtain declining prices, preferences must exhibit non-decreasing absolute risk aversion, which is not thought to be plausible. Milgrom and Weber (2000) show that affiliation implies that prices should (weakly) increase, as winner's curse concerns are mitigated as information is revealed. Finally, one might also question whether buyers only want one unit, and if not whether their valuations are independent across units. In some auctions with declining prices, initial winners have the option to buy more than one unit at their bid price, and the price decline may reflect the value of this option, as noted by Black and de Meza (1992).

The literature on the declining price anomaly illustrates the interaction between theory and empirical work in auctions. Here the theory generates a simple prediction, which is rejected by the data, inducing further theoretical work on the robustness of the predictions. It also illustrates an important theme of empirical work on auctions: tests of reduced form predictions are most compelling when the predictions are robust to the various maintained hypotheses.

Beggs and Graddy (1997) examine sequential sales of paintings, in part to see whether the order within the sequence matters. They find that sale rates and sale prices relative to pre-sale catalog estimates rise with bid order for the objects sold at the beginning of the day. Thereafter, sale rates and prices relative to estimates fall, similar to the declining price anomaly. Perhaps not surprisingly, there is a tendency to place higher value objects earlier in the sale, although not right at the beginning. Note that there is an implicit assumption in the econometric model that looking at prices relative to pre-sale estimates is sufficient to control for heterogeneity across dissimilar paintings. In sequential sales of identical objects, the variance of prices will evolve over time. An open issue is whether higher order moments also vary with the pre-sale estimate, i.e. whether there is heteroskedasticity, as might result if the dispersion of valuations relative to the pre-sale estimate varies with the value of the object.

The revenue equivalence result also extends to simultaneous auctions of multiple identical objects when bidders have unit demands. Milgrom and Weber (2000) prove the following result in auctions where k units are auctioned simultaneously in an IPV environment with risk neutral bidders. Suppose the auction rules and equilibrium are such that the k highest types always win and that the lowest type has zero expected payoff. Then the expected payment of type x of bidder i at the Bayesian Nash equilibrium of the discriminatory auction (where the winning bidders pay their bid) is the same as in Vickrey's auction, where the winning bidders pay the highest rejected bid. The expected payment of that bidder is the expected valuation of the marginal losing bidder (including the seller, as represented by the reserve price), conditional on being among the set of winning bidders:

$$E[\max\{X_{k+1:n}, r\} | x \geq \max\{X_{k+1:n}, r\}]$$

The Vickrey auction is also known as a uniform price auction. Another uniform price auction stipulates that the k highest bidders win and pay the lowest accepted bid, that is the lowest bid submitted by one of the k winning bidders. Note that, in the symmetric equilibria, the outcomes are the same and they are always efficient, because the items are allocated to the bidders who value them the most.

Treasury bills are an important example of a simultaneous item auction that has been run under both discriminatory and uniform pricing rules in a number of countries, such as the United States, Sweden, Switzerland, Mexico, Turkey and Korea. However, the unit demand assumption in these auctions typically fails. Bidders submit bid schedules, listing their willingness to pay for a varying number of items. Ausubel and Cramton (2002) provide examples under which either the uniform or the discriminatory format may yield higher revenues. Thus, the question as to which format yields more revenue is an empirical issue. Bikhchandani and Huang (1993) survey empirical studies that have addressed this issue using reduced form methods.

Several authors have recently employed structural estimation methods to recover the bidders' willingness to pay schedules. Hortacsu (2002) examines bidding in Turkish treasury auctions when a discriminatory format was employed. He assumes that values are private, and that buyers' private information is a real number (i.e., not multi-dimensional). An increase in the private signal leads to a monotonic increase in the bidder's marginal valuation schedule. Bidder's demand functions are assumed to be monotonic functions of their private signals. Thus, if a bidder's rivals adopt a given bid schedule, the bidder's best response can be computed for a given realization of his private information. As in the single unit case, one can define an inverse bid function, where now

$$v_i(\varphi_i(b,x),x) = b + \frac{G_i(b,\varphi_i(b,x))}{\partial G_i(b,\varphi_i(b,x))/\partial b}.$$

Here $\varphi_i(b,x)$ denotes the bid function specifying a quantity demanded by bidder i at price b, given a private signal x, and $v_i(q,x)$ is the marginal valuation of acquiring q units given the private signal. $G_i(b,\varphi_i(b,x))$ denotes the probability that the lowest winning bid is less than or equal to b, given that bidder i adopts the strategy $\varphi_i(b,x)$. This equation generalizes the expression for the inverse bid function in a single unit private value auction. Hortacsu proposes a method based on resampling to estimate $G_i(b,\varphi_i(b,x))$ and its derivative for each bidder, and thus recover a pseudo-sample of marginal valuations. Hortacsu shows that his method yields consistent estimates. After recovering the distribution of bidder valuations, he computes an upper bound on the counterfactual seller revenues that would be obtained in a uniform price auction, under the assumption that bidders would not bid more than their marginal values. This bound is less than the observed revenue in the discriminatory auction, suggesting that a uniform price format would not result in higher revenues. Such counterfactual comparisons are feasible only with structural estimates. Other structural estimation papers include Castellanos and Oviedo (2004), who study Mexican treasury auctions, and Fevrier, Preget, and Visser (2004), who study French treasury auctions.

Several empirical studies of treasury bill and electricity markets, in which bidders submit bid schedules rather than bid prices, assume that the schedules are continuous functions. In practice, bidders often submit step functions with only a few steps, in some instances with many fewer steps than they would be permitted to use. Kastl (2005) argues that in some contexts the restriction to continuous bid schedules can mask important strategic issues, and that revenue comparisons based on continuous schedules may not be valid. In particular, revenues in uniform price auctions can exceed the bound derived by Hortacsu. His study of uniform price Czech treasury auctions indicates that the distinction between continuous and discrete schedules can matter empirically.

Kastl's empirical work exploits the fact that, although bid functions are not continuous, the price and quantity associated with any bid point are continuous variables. He characterizes the necessary conditions the price and quantity choices must satisfy, and his empirical work exploits the first order condition for quantities. Wolak (2003, 2005) proposes an empirical method that accounts for both the price and quantity first order conditions associated with each bid point, and he applies the method to determine the value of forward contracts in the Australian national electricity market.

8 Collusion

There is evidence of collusion in many auction markets. Examples include highway construction contracts, school milk delivery, timber sales, and spectrum auctions. Auction markets may be especially vulnerable to collusion. An auction is an effective price discovery mechanism under competition, and so can be attractive to sellers who are uncertain about the price they should receive. But the sellers may then be unable to determine when they are the victim of a bid rigging scheme. Collusion among bidders, if successful, can benefit the participants at the expense of their suppliers. The social losses usually outweigh the benefits. For example, if a bidding ring lowers the prices they offer to pay for products or services, relative to competitive levels, then sellers suffer a loss, and trade will be less likely to occur. The outcome may also be inefficient, if bidders outside the ring are more likely to win. Further, potential sellers will be less inclined to offer their products or services in the future. There will be social welfare losses to the extent that current and potential future gains from trade are foregone.

In this section, we describe various collusive schemes, the factors that facilitate or inhibit collusion among bidders in auction markets, as well as circumstances where detection is possible. Collusion in this instance means explicit, as opposed to tacit, cooperation, involving direct communication and perhaps side payments. These actions are usually surreptitious, either because they are illegal under antitrust laws, or because they are most effective if they are kept secret from the intended victims. We describe some results from the theoretical literature, and discuss a number of recent empirical studies. Two surveys that touch on some of the same issues, although not just for auction markets, are by Harrington (2005) and Porter (2005).

8.1 Collusive Mechanisms

Collusion can take many different forms in auctions. The form often depends on the auction rules and characteristics of the environment (Hendricks and Porter, 1989). But all bidding rings face a set of typical cartel problems: private information about the gains from trade, conflicting objectives, internal enforcement, detection by authorities or seller, and entry. An important feature of many collusive agreements, and a determinant of their success, is the need to reconcile disparate interests. Interests may differ for a number of reasons. For example, firms may have adopted technologies of differing vintages for historical reasons, they may serve non-overlapping and heterogeneous customer bases, or their payoffs may be subject to imperfectly correlated shocks. Side payments can solve some internal problems, but they may not be legal. Also, the contractual terms associated with side payments may not be enforceable. An unhappy conspirator whose loyalty cannot be purchased is more likely to report the collusion to antitrust authorities.

Theoretical studies have focussed primarily on private information as a source of difficulties for a ring. In this literature, the ring's objective in an auction is viewed as a problem of mechanism design. The mechanism design approach is natural for legal cartels, which can write binding contracts to enforce side payments or allocative agreements, and to a lesser extent for illegal conspiracies where members communicate and may make side payments. Since most illegal rings operate in many auctions over time, they may be able to use repeated game strategies to enforce the mechanism played in an individual auction (Athey and Bagwell (2001), Athey, Bagwell and Sanchirico (2004)). In a mechanism design framework, allocations and transfers are determined by internal messages. The collective goal is to win the auction when it is optimal to do so, and in those cases to allocate the good and divide the spoils among the participants. The ring mechanism must be incentive compatible and individually rational. That is, ring members must have an incentive to reveal their private information, and they must also participate voluntarily, so that there is no incentive to defect. The specification of these constraints, and the extent to which they are binding, will depend upon the auction rules and on the characteristics of the environment.

8.1.1 Private Values

In a private values environment, the cartel has to overcome an adverse selection problem: the participants do not know how much fellow ring members are willing to pay for the item being auctioned, and each member wants to exploit this private information to argue for a bigger share of the spoils. Graham and Marshall (1987) and Mailath and Zemsky (1991) study collusion in second-price, IPV auctions. They show that ex post efficient collusion by any subset of bidders is possible. In a symmetric environment, the ring can easily implement the incentive compatible, ex post efficient cartel mechanism with a pre-sale knockout auction. In the knockout auction, the ring members bid for the right to be the ring's representative bidder in the seller's auction. The conspirator who bids the highest amount wins this right, and the winner pays an amount to the other ring members based on the bids submitted. Since the knockout auction has a symmetric equilibrium with monotone strategies, the winner is the member who has the highest valuation. The participation constraints are evaluated at the interim stage, after the members have obtained their values but before they have decided to participate in the ring. The payoff to a bidder who decides not to participate in a ring of size m is given by her equilibrium payoff in the seller's auction when she bids against a ring of size $m-1$ and $n-m$ non-ring bidders. This payoff is easily computed in a second-price, private value auction because participants have a dominant strategy to bid their value. Hence, a noncolluding bidder faces the same high bid whether her rivals form a ring (of any size) or not, assuming the ring selects the member with the highest valuation. The knockout auction satisfies the interim participation constraints, so bidders prefer the cartel mechanism to bidding non-cooperatively in the seller's auction.

McAfee and McMillan (1992)) study collusion in first-price IPV auctions. They show that ex post efficient collusion is possible in these auctions if the ring includes all bidders. The all-inclusive ring can implement the incentive compatible, ex post efficient cartel mechanism using a first-price knockout auction in which the winning bidder pays his bid to the losing bidders, who share equally in this bid. Hendricks, Porter, and Tan (2006) extend this result to affiliated private value environments. McAfee and McMillan assume that the bidders have to make their participation decisions before they obtain their private information so that the relevant participation constraints are ex ante. They show that the bidders' expected cartel payoffs exceed the equilibrium payoff they would receive if everyone bid individually and noncooperatively in the first price auction, and so everyone will want to join the ring. An alternative specification of the participation constraint is the equilibrium payoff in a first-price sealed bid auction in which one bidder bids against a ring of size $n-1$. These payoffs are difficult to compute since the auction involves asymmetric bidders. McAfee and McMillan study a simple model in which each buyer's private value is an independent Bernoulli random variable. They show that the noncolluding bidder is better off ex ante than ring members. Thus, the all-inclusive ring is not stable, which may explain why most rings in first-price auctions are partial.

An alternative to a pre-sale knockout auction is a post-sale knockout, such as the one used by a bidding ring involving rare book dealers in England in 1919. After one large estate sale, the ring held a series of knockout auctions. Successively smaller subsets of the dealers conspired to deprive the seller, and then their fellow conspirators, of some of the gains. The book dealers differed according to expertise and scale of operation, and the larger and more experienced dealers stayed longer in the knockout process. The participants in the various knockout auctions shared the increases in bids over prices in the previous round. The original seller received less than 20 percent of the final settlement prices. (Note that this figure probably overstates the damage to the seller, as the ring settlement mechanism could induce conspirators to bid more than their valuations in the knockout auction.) Why did the larger ring members conspire with the smaller members? If they had not, the larger dealers would have had to outbid the smaller dealers at the original auction, and it would have been cheaper to share some of the collusive gains with them. But it is also in their interest to share only enough to buy the loyalty of the smaller dealers, and not the full difference between the original purchase price and what the larger dealers were willing to pay. A sequence of knockouts would limit the amount shared, assuming that smaller dealers were not aware that they would be excluded from later knockouts. (Porter (1992) provides a brief account.)

An imperfect solution that involves no side payments is the territorial division of bidding privileges, by region, by point in time, by incumbency, or even by pure randomization (via the submission of many identical bids). A scheme that assigns customers or territories to the participants would grant individual firms wide latitude within their own territories. Comanor and Schankerman (1976) provide evidence that rotating bid arrangements can be stable. Here firms take turns submitting serious bids for the ring. The serious bid may be optimized against the seller's acceptance rule, or against other bidders who are not coconspirators. Other ring members may submit complementary bids to create the appearance of competition. McAfee and McMillan (1992) show that it may be optimal for a weak all-inclusive cartel (that is, one that cannot make side payments) to submit many identical bids at the reserve price, and so rely on the auctioneer to randomly select among them. In the 1950s, General Electric and Westinghouse assigned low bid privileges for electrical equipment contracts based on a phases-of-the-moon system (Smith, 1961), a practice that is unlikely to reflect differences in values.

Pesendorfer (2000) argues that a weak conspiracy that cannot make side payments may be forced to maintain relatively constant market shares, despite some losses from not allocating bidding rights to the low cost firm, in order to maintain internal discipline. He shows that, if there are many items being sold, the ring can achieve approximate internal efficiency via a ranking mechanism. That is, members rank items, and bidding privileges for individual items are assigned based on the submitted rankings. The ring does not achieve full internal efficiency, as minimal market shares must be guaranteed to ensure that participation constraints are satisfied. He compares Florida and Texas bid rigging schemes for providing school milk, and shows that market shares were less stable in Florida, where dairies used side payments.

8.1.2 Common Values

In a pure common value auction environment, a ring's internal allocation problem is simpler, because all members value the item identically. As a result, cartels in these auctions can adopt equal division rules, in which all members share equally in the spoils. Given this sharing rule, cartel members have no incentive to misrepresent their information. They share a common goal, which is to bid only when the expected value of the item conditional on their pooled information exceeds the reserve price. The equal division sharing rule can be implemented without transfer payments by having every member submit the same bid (e.g., the reserve price) and letting the seller randomly select the winner, or by rotating bidding privileges over several auctions. However, participation constraints can be a problem.

Hendricks, Porter and Tan (2006) study all-inclusive cartel formation in a first-price auction with common values. The ring is assumed to form after the bidders have obtained their private information so the relevant participation constraints are interim. In the standard mechanism, bidders compare the equilibrium payoff (conditional on their private information) they would receive if everyone bid individually and noncooperatively in the first price auction to their expected cartel payoff. We show that the equal division sharing rule does not generally satisfy these participation constraints, nor does the knockout tournament. Indeed, no incentive compatible, ex post efficient cartel mechanism may satisfy the participation constraints. The reason is that bidders have to bid cautiously in the first-price auction due to the winner's curse. As a result, a bidder with favorable information when commonly available signals are pessimistic is often able to earn a higher expected payoff from competitive bidding in the seller's auction. He will have to pay a somewhat higher price paid to the seller, but the surplus is not shared with other ring members, who are not likely to pose a competitive threat.

We also consider a second model of the participation constraint. Cramton and Palfrey (1995) have argued that the participation choices of bidders are informative about their types. If one or more bidders chooses not to participate, bidders should revise their beliefs accordingly before bidding in the first-price auction. The revision in beliefs will affect bidding behavior and hence payoffs, and should be anticipated by the bidders when they make their participation decisions. Cramton and Palfrey refer to this issue as the information leakage problem. They formalize the problem by assuming that bidders first play a veto game in which they simultaneously vote for or against the proposed cartel mechanism. If the mechanism is unanimously ratified, then the all-inclusive ring forms and the mechanism is implemented. If at least one bidder votes against the mechanism, then the all-inclusive ring does not form, bidders revise their beliefs, and they bid individually and noncooperatively in the first-price auction. The cartel mechanism is ratifiable if it is not possible to construct a veto set such that the cartel payoff of every type in the set is less than their equilibrium payoff in the first-price sealed bid auction. The main difficulty with applying this solution concept in our context is computing the equilibrium payoffs to the veto set. If some bidders vote for and others vote against the cartel, then the seller's auction will involve asymmetric bidders. We show, however, that in the case of pure common values, one does not need to compute the equilibrium to prove that the cartel mechanism is ratifiable. The intuition behind this result is that the lowest type in the veto set makes zero profits in the auction. Since this type makes positive profits in the cartel mechanism, it is not possible to construct a veto set that makes everyone in the set better off. Hence, the all-inclusive cartel should always form. This conclusion generalizes to other voting models of joint venture formation.

The motivating example for our study is joint bidding in U.S. federal auctions of wildcat oil and gas leases in the Outer Continental Shelf off the coasts of Louisiana and Texas during the period 1954 to 1970. Joint bidding ventures for all firms were legal during this period. The potential gains from joint bidding appear to be substantial. The stakes are large, and the risks significant. By pooling geological data and expertise in interpreting the data, firms could reduce the risk of buying dry leases and, by pooling financial resources, they can bid for more leases and diversify away more of the tract-specific uncertainties. Despite these gains, solo bidding was the dominant form of bidding for the most active participants. Joint bids involving pairs of the most active firms represented less than 15% of all their bids, even though joint bidding agreements were legal. Furthermore, if these firms bid jointly, they almost always did so in pairs, and not in all-inclusive partnerships.

The standard model of participation constraints (i.e., passive beliefs) seems to be a better approximation to the firms' joint venture decisions in OCS auctions than the learning model. The reason is that the joint venture cover blocks of tracts, typically 25 to 50 tracts, and not individual tracts. If a bidder refuses to join the ring, then the other bidders may infer that the non-participating bidder has obtained favorable information about one or more tracts in the area, but they do not know which tracts and, since most are not worth bidding for, the inference is likely to have little impact on beliefs about individual tracts. The learning model of participation constraints suggests that the situation would have been quite different if the joint bidding decisions were taken on a tract by tract basis. In that case, refusal to bid jointly on a tract would cause beliefs about that tract, and therefore bidding behavior, to change. Information leakage would have forced the firms to participate. But, given passive beliefs, our results provide an explanation for the surprising low incidence of joint bidding, particularly on marginal tracts. Consistent with this explanation, we document a positive correlation between the incidence of joint bidding and the value of tracts. This correlation probably reflects the incentive for firms to find financial partners on tracts where the high bid is likely to be large. However, it may also reflect the fact that potential bidders are more likely to know their competition (i.e., who intends to bid) on high value tracts.

Solo bidding does not imply the absence of collusion. In testimony before Congress in the mid 1970s, Darius Gaskins of the Department of Interior argued that the collusive effects of joint ventures should not be measured solely in terms of tracts receiving joint bids. Joint bidding negotiations could allow partners to coordinate their solo bids. We find evidence of bid coordination by bidders who bid jointly in a sale. In particular, bidders are unlikely to submit competing solo bids if they have submitted a joint bid in another region in the same sale.

Finally, rings in common value environments may have to worry about a moral hazard problem, since each member has an incentive to free ride on the costly information gathering activities of other members. These difficulties may also explain the low incidence of joint bidding in the OCS auctions.

8.2 Enforcement

To succeed, a ring has to keep its members from deviating, by ensuring that it is in each of the conspirator's self-interest to adhere to the agreement. Robinson (1985) points out that enforcement is easier in SPSB or English private values (PV) auctions. Suppose that the ring is ex post efficient, and therefore designates the member with the highest valuation to be its representative in the seller's auction. This agent's dominant strategy is to bid his valuation, independent of the competition he faces. Then the other ring members cannot gain from deviating and outbidding the designated bidder. The success of the ring depends only on how many potential bidders refrain from bidding, thereby lowering the expected price paid by the serious bidder when he wins. Similarly, in an English auction, the serious bidder only needs to outbid other submitted bids. There is an internal enforcement problem only if the serious bidder does not have the highest valuation among the ring members.

In FPSB PV auctions, the ring bidder optimally bids below his valuation. His bid is decreasing in the size of the ring, as more potential competition is eliminated. If the serious bid is low enough, another ring member might profitably deviate because the serious bid is below his valuation. A ring may then form only when the participants know that they will be competing against each other in subsequent auctions. In that case, the ring can credibly threaten to punish non-compliance in current play by expelling the deviator from the ring, or by dissolving the ring.

Marshall and Marx (2004) argue that collusion might even be possible in a one-shot FPSB auction, if side payments are feasible and if shill bids are employed to create the proper incentives within the FPSB auction. Shill bids, which are also known as complementary or phony bids, are submitted by ring members other than the designated bidder. By design, a shill bid is less than the serious ring bid. The idea is that the designated ring representative bids more than would be optimal against outsiders, to ensure that other ring members do not defect. The purpose of a shill bid, just below the high ring bid, is to dissuade the designated bidder from bidding lower. Thus complementary bids may be employed by a ring to enforce internal discipline, and not just to create the appearance of competition.

In a multiple-unit simultaneous ascending bid format, such as the mechanism employed by the US Federal Communications Commission (FCC) to sell spectrum for PCS (personal communications services), punishments can be wide-ranging. Defections in the bidding for one object can induce responses elsewhere. In the FCC DEF block spectrum auction, a territorial division was achieved within the bidding process itself (Cramton and Schwartz, 2000). The FCC employed a simultaneous ascending bid procedure, in which bidding was kept open on all licenses throughout the auction, and firms with enough eligibility could switch between licenses. Some bidders used trailing digits on their bids to communicate their future intentions. For example, one response to a new bidder in one's territory was to outbid that bidder on at least one other license where it held the standing high bid. The response bid's last three digits would be the identifying code of the original market, and the intended message was the offer to not compete on this license if the rival stays out of your territory. No overt communication is involved, unless the parties need to resolve how to interpret bid signals, and a territorial allocation can be achieved at relatively low bid prices. Gertner (1995) describes how these bidding strategies can be self-enforcing, and result in low prices. The auction rules could be amended to prevent this sort of signaling, for example by requiring new bids to be a fixed amount or fraction higher than the current high bid. There could also be a fixed ending time to the auction, in which case it would not be possible to retaliate after that time.

In a multi-unit uniform price auction, price is determined by a market clearing condition, where available supply equals demand. A version of this mechanism was employed by the England and Wales electricity auction market. (Wolfram (1998, 1999) provides an account.) Bidders can implicitly make it costly for rivals to steal market share by bidding low prices for infra-marginal supplies. A generating unit is infra-marginal if it is likely to be called on to supply power, but unlikely to be decisive in determining the market price. In multi-unit uniform price auctions, the gains from deviation can be limited by pricing infra-marginal units low, via "hockey stick" bidding, with low infra-marginal bids and high marginal bids.

An auction designer can combat this sort of bidding coordination by employing a different rationing rule (Kremer and Nyborg (2004)) or a downward sloping demand curve (LiCalzi and Pavan (2005), McAdams (2005)). In this instance, a discriminatory auction, in which each supplying unit is paid the amount of its bid, might also induce more competitive bidding.

The ring also has to be able to control entry to prevent outsiders from capturing the benefits of collusion. Asymmetries in information or payoffs can act as a barrier to entry. For example, in auctions of drainage leases, neighbors with informational advantages over non-neighbors have obvious gains from coordination, and they would not have to worry that entry of non-neighbors would dissipate all of the gains. In general, firms with a favorable distribution of values, or those with better information, have an incentive to collude, and they will not want to share the gains with disadvantaged firms that do not pose a significant competitive threat.

8.3 Detection

Collusive schemes are often illegal, and a problem faced by antitrust authorities (such as the U.S. Department of Justice) is to detect their presence. Most cartels encounter operational problems. It is the manner in which a conspiracy deals with these problems that often facilitates the detection of the scheme. In some instances, one can do more than just look for direct evidence of the exertion of market power, such as high and persistent profits. In this subsection, we describe several papers that propose a variety of methods that are designed to distinguish between collusive and competitive bidding.

While statistical evidence that bidders seem to be colluding is not sufficient to establish guilt in a criminal price fixing case, such evidence can be used to guide an investigation. For example, suppose that examination of the data suggests the presence of a bid rigging scheme. The firms might be alerted that they are suspected of colluding, and informed that corporate leniency programs offer much reduced sentences and fines to those who confess first, thereby creating incentives for a "race to the courthouse." Econometric evidence may also be decisive in civil cases, where the burden of proof is less onerous.

Alternatively, antitrust authorities may pursue policies that inhibit successful collusion, by altering characteristics of the economic environment, such as pursuing an active merger policy. Further, the potential suppliers of the ring can alter the rules of the auctions they employ in response to the presence of collusion. For example, the seller can raise the minimum bid, adopt a secret reserve price, or not publicize bids to make it harder for the ring to detect cheating and maintain discipline.

Absent the direct evidence of a conspirator, a conspiracy may be difficult to detect. The constancy of market shares or geographical specialization, while consistent with a collusive assignment, are not in and of themselves evidence of collusion. There is a tendency to view bid rotation or incumbency bidding advantages as evidence of presence of collusion. However, these bidding patterns can be consistent with non-cooperative bidding. For example, bid rotation can be a competitive outcome in auctions of highway construction contracts where bidders' cost functions exhibit decreasing returns to scale. Firms with idle capacity are more likely to win the contract, but having won the contract, are less likely to bid or less likely to win another until some existing contracts are completed (Porter and Zona, 1993).

Similarly, incumbency patterns can reflect unobserved asymmetries among bidders. Those who won in the past may have done so because of location or other advantages that persist through time. Incumbents may have lower costs due to experience, or they may have an advantage with buyers who are reluctant to switch suppliers (Porter and Zona, 1999). An empirical challenge is to develop tests that can discriminate between collusive and non-cooperative explanations for rotation or incumbency patterns.

Most early empirical studies identified behavior that is difficult to reconcile with a non-cooperative bidding. An extreme example involves the submission of several identical bids. Mund (1960) and Comanor and Schankerman (1976) describe several instances of identical bids "independently" submitted in government procurement auctions. In 1955, five companies submitted identical sealed bids of $108,222.58 for an order of 5,640 one hundred capsule bottles of antibiotic tetracycline (Scherer and Ross, 1990, p. 267). The submission of identical bids is virtually a zero probability event in a Bayesian Nash equilibrium if there is any dispersion in information or valuations across bidders.

Baldwin, Marshall and Richard (1997) study US timber English auctions to see whether bidding was competitive or collusive, using structural methods. In an English auction, collusion can reduce the bid necessary to win. Under competition, the winning bid in a button IPV auction equals the second highest valuation. If the ring members are a subset of the potential bidders, the winning bid is affected only if the two highest valuation bidders are conspirators. Then the winning bid is no longer the second order statistic of the valuations. If the highest non-ring valuation is the k^{th} highest valuation overall, the winning bid will be the k^{th} highest order statistic for $k > 1$, as the serious ring bidder outbids the highest non-conspirator. If the highest valuation bidder is not a ring member, $k = 1$, and the winning bid is the second highest valuation. Note that there is an implicit assumption of efficient collusion, where the cartel designates the member with highest value.

Baldwin, Marshall and Richard do not have firm specific information, and they assume that the valuation distribution is symmetric. Functional form then plays an important role in distinguishing between competition and collusion. In the former case, the winning bid is the second order statistic from the valuation distribution, whereas in the latter case it is a mixture of the second and higher order statistics. One cannot distinguish between competition and collusion without some assumption concerning the functional form of the distribution of values. However, their method could be useful when there are observable bidder asymmetries.

The extension of the method to first price sealed bid auctions would be complicated, as bidding value is no longer a dominant strategy.

If the seller knows that a knockout auction has preceded the sale, it should set a higher reserve price. If the ring is not all-inclusive, it may also want to keep its presence unknown to other bidders, say because new potential bidders may enter the market in response to profitable opportunities. Therefore, it is in the interest of the ring to keep its meeting secret. Ring members may create the appearance of competition in order to avoid detection. Bidding rings may submit phony, or complementary, bids that are designed merely to be lower than the serious bid submitted by the ring. Then only the highest bid from the ring is serious. According to Preston McAfee, one conspiracy was investigated by the U.S. Department of Justice after a bidder submitted an envelope containing his own bid plus his notes from a pre-auction meeting. But phony bids, unlike serious bids, may not be related to the likely profits of the bidder in the event that it wins. Porter and Zona (1993) describe a bidding ring involving highway-paving jobs on Long Island in New York. A subset of the firms participated in pre-auction meetings in order to assign low bidding privileges for specific procurement contracts. The conspirators often submitted complementary bids above the low bid. How does this behavior manifest itself? The fact that the ring was not all-inclusive, plus the presence of phony bids, aids detection.

Porter and Zona split data in two dimensions, distinguishing cartel from non-cartel bidders (those absent from meetings), and the low bid from a given set of bidders from higher bids from that set. Serious bids should depend on the costs of doing the job. Porter and Zona do not have good contract-specific information, and so they look at the rank distribution of bids within the two groups of firms, as opposed to bid levels. They test whether the identity of the low bidder is determined by same process as the ranking of higher bids. Bids by non-cartel members pass this test, whereas those from the ring do not.

The order of the bids submitted by non-conspirators was related to observable cost measures, such as firm capacity and the utilization rate of that capacity. The lowest non-conspirator bid was most likely to be submitted by the firm with the lowest cost. The lowest conspirator bid was determined by a similar process. In contrast, the order of the higher bids submitted by ring members was not correlated with the same cost measures.

Note that a ring could pass this test, were members aware that the ranking of bids might be examined. For example, suppose the ring is efficient and assigns low bidding rights to the low cost firm in the pre-auction meeting. If the serious bidder intends to bid some multiple of its costs in the auction, then the other conspirators could submit phony bids that were the same multiple of their costs.

Note also that the bid rank test exploits information concerning firms' participation in pre-auction meetings. However, it is not necessary to have this information. The rank test could be performed for any partition of the firms. Bajari and Ye (2003) conduct a similar exercise for several combinations of the larger bidders in Minnesota highway construction auctions, in order to determine who might be colluding.

In addition to creating the appearance of competition, complementary bids may also be intended to manipulate the expectations of the buyer. Feinstein, Block and Nold (1985) note that many agencies estimate the cost of projects on the basis of bidding on similar past projects. Multiple phony bids close to a relatively high bid may lead an unaware buyer to think that costs are higher than they are, and therefore not suspect collusion. Their analysis of data from North Carolina highway construction auctions suggests that contractors were indeed manipulating the information received by the buying agency.

If there is entry, and the entrants are not party to the collusive agreement, then the non-inclusive nature of the cartel may lead to evidence of its existence. As noted above, Porter and Zona distinguish complementary bids by a ring from non-winning bids submitted by a competitive fringe. Hendricks and Porter (1988) describe another example in our study of drainage auctions. Recall that an oil or gas lease is said to be a drainage lease if there has been prior exploration in the area. In that instance, the firms with prior drilling experience will have an informational advantage over firms that have access only to seismic data. In the offshore oil and gas drainage auctions, the identities of the firms owning the mineral rights on neighboring tracts ("neighbors") are known, and their numbers limited by the number of tracts previously sold and explored. Neighbors can gain from coordination, and they do not have to worry that the entry of non-neighbors will dissipate all of the gains. We find that neighbors earn high profits, whereas non-neighbors approximately break even. Despite relatively high overall returns, there is less entry (i.e., fewer bids are submitted per tract) than on wildcat leases, where bidders share similar information sources. The lower entry rates on drainage leases are consistent with asymmetries of information acting as an entry barrier.

If neighbors bid non-cooperatively in the drainage auctions, then there should not be entry by non-neighboring firms, if the latter do not have access to private drilling information. Yet there is entry by non-neighbors. Further, non-neighbors' bids are independent of the number of neighboring firms, rather than a decreasing function as winner's curse considerations would dictate. In addition, there are often multiple bids from the neighbors on a single drainage tract, yet their returns are an increasing function of the number of their bids submitted. Finally, the highest neighbor bid is independent of the number of neighbors, and their average bid level is a decreasing function of this number. This latter fact is consistent with the neighbors submitting only one serious bid, and the probability of submitting complementary bids being an increasing function of the number of neighboring leases in order to create the appearance of competition.

Porter and Zona (1999) provide evidence that the bidding behavior of some Ohio dairies for school milk contracts in the 1980s was more consistent with collusion than with competition. Bidder payoff asymmetries, based on plant locations, matter, as processed milk is costly to ship. Milk's value is low relative to its weight, and therefore competition is localized. Many dairies nevertheless bid for both local school district contracts and more distant school districts. That is, they submit bids relatively near their plants and they also submit bids well beyond their local territories. Porter and Zona's analysis of bidding levels shows that the more distant bids of the three Cincinnati dairies tended to be lower than their local bids, even after controlling for various covariates. In contrast, other dairies' bids are an increasing function of the distance from the school district to the firm's nearest plant. These features of bidding are consistent with a territorial allocation of nearby school districts by dairies with plants in the Cincinnati area to restrict competition, and relatively competitive bidding at more distant locations, which were perhaps outside the area of territorial allocation. If bidding for local districts had been competitive, local bids should have been lower than distant bids, because shipping costs were lower and because the Cincinnati area has many potential local suppliers. The effect of collusion is to relax the constraint on bids by removing marginal rivals. The effect is an "inverted price umbrella," consistent with a local monopoly constrained only by more distant rivals.

8.4 Collusion By Sellers

We have discussed collusion among bidders in auctions. There are three other forms of auction market manipulation that may also blunt the incentives to engage in trade, and hence compromise the social value of auctions as a market institution. First, the recent investigation of Sotheby's and Christie's indicates that collusion among bidders is not the only potential antitrust concern in auction markets. To the extent that the market for providing auction services is concentrated, there may be collusion among auctioneers to raise service fees to potential sellers or buyers.

Second, recent events indicate that eBay auction rules can be manipulated by sellers who unilaterally submit phony or shill bids on their own items, in an attempt to obtain higher prices. Also, groups of sellers can inflate their eBay seller rankings by giving each other glowing reviews for service and product quality. These activities might best be characterized as fraudulent, rather than raising antitrust concerns. But they also reduce faith in the institution.

A third form of manipulation is the corruption of the seller by one or more bidders, or corruption of the auctioneer if he is an agent for the seller. For example, in a sealed bid auction a bidder could bribe the auctioneer to reveal the other bids. Armed with this knowledge, the bidder would not have to bid more than necessary to win. Burguet and Perry (2002) provide an analysis of this situation. Concerns about potential corruption of the bidding process are an important component of auction design, such as public procurement rules that limit the discretion of the auctioneer.

9 Further Research Issues

In this section, we conclude with a brief discussion of some outstanding issues.

9.1 Scoring Rules

One extension of the standard static model of single-unit auctions considers auctions with scoring rules. In some cases, the seller (or the buyer in procurement auctions) is interested in soliciting and evaluating bids on more than one dimension. For example, the state of Louisiana sells its onshore oil and gas leases using an auction in which bidders submit a pair of numbers, a bonus and a royalty rate. The state's scoring rule is secret, but presumably it will select a bid with a high bonus bid if it believes that the likelihood of finding oil is low, and a bid with a high royalty rate if it believes that the likelihood of finding oil is high. In motivating their theoretical study of scoring auctions, Asker and Cantillon (2004) observe that 38 states in the United States evaluate bids for highway construction contracts on the basis of their costs as well as time to completion, weighted by a road user cost. In this case, the procurement authorities announce and commit to a specific scoring rule. The multi-dimensionality of the bid space is often associated with multi-dimensionality of the bidder type space, which greatly complicates the theoretical and empirical analysis.

Athey and Levin (2001) study scale auctions of timber, which employ scoring rules. In these auctions, a bidder submits a price for each species of tree, and her total bid is computed by multiplying these unit prices by the quantities announced by the Forest Service (FS). The tract is awarded to the bidder who is willing to pay the highest total bid, but her payment is based on the realized volumes of each species. Therefore, if she believes that actual relative volumes will differ from those announced by the FS, she has an incentive to skew her bid, allocating more of her total bid to species that she thinks the FS has overestimated. A risk neutral bidder would allocate all of her total bid to overestimated species, unless severely skewed bids are likely to be rejected. The decision is non-trivial if the bidder is risk averse, which seems relevant since the motivation for using scale auctions is to reduce the risk borne by the winners. Athey and Levin characterize equilibrium bidding when a tract has two species. The restriction to two species implies that a bidder's information can be summarized by her estimate of the proportion of one of the two species (ignoring uncertainty about the total volume of wood), so bidder types are one-dimensional. Given certain regularity conditions, the authors show that the skew (i.e., the difference in unit prices of the two species) and the total bid are increasing functions of the bidder's estimate. This theoretical result forms the basis of several predictions that are testable given the availability of the actual species volumes cut. The empirical results confirm the importance of equilibrium analysis.

9.2 Entry and Dynamics

To date, researchers have used auctions to study the strategic effects of private information and the relevance of Bayesian Nash equilibrium. We believe that auctions can provide an important testing ground for studying two other important issues in industrial organization: entry and dynamics. Auctions are held repeatedly, and firms have to make frequent entry decisions. Auctions also provide a rich variety of settings (e.g., sealed bid versus oral, private versus common values) for studying entry decisions and outcomes. Recently, Bajari, Benkard, and Levin (2005), Pakes, Ostrovsky, and Berry (2005), Pesendorfer and Schmidt-Dengler (2004), and Aguirregabiria and Mira (2004) have developed estimators for dynamic games based on the approach taken by Hotz and Miller (1993). These papers make estimation of dynamic oligopoly models much more feasible, although the assumptions under which the estimators are valid can be quite restrictive. In particular, the dynamic game should consist of repeated stage games played in a stationary environment in which the unobservable shocks are independent across players and time. Private value auctions such as highway procurement auctions can come close to satisfying these assumptions.

Most of the empirical literature on auctions assumes that bidders treat auctions as static games, choosing their bids in each auction to maximize profits in that auction. This assumption is plausible in environments where there is no learning, and the time interval between auctions is sufficiently long that the outcome in the current auction is unlikely to influence the state of play in subsequent auctions. The state of play in an auction consists of the set of bidders and their value distributions. However, in some auctions, such as highway procurement contracts, the time between auctions can be quite short, often measured in weeks or even days. In these cases, the presence of capacity constraints or decreasing returns to scale introduces a dynamic element to the game. The bidder who wins the current auction may not have sufficient capacity to bid in the next auction, or it will have costs that are stochastically higher. As a result, the losers in the current auction are more likely to win the next auction. Bidders should anticipate the impact of the current auction outcome on the state of competition in future auctions and bid accordingly. The equilibrium mapping of valuations into bids in the dynamic game is not the same as in the static game, and estimating the distribution of costs under the assumption that bidders behave myopically may yield incorrect results. In particular, bidders may bid less aggressively if they anticipate that winning the current auction will put them at a competitive disadvantage in future auctions.

Jofre-Bonet and Pesendorfer (2003) estimate such a dynamic model using data from highway procurement auctions in California. The average duration between auctions was about 3 days. The state variables for each bidder are its backlog, measured as the dollar amount of work remaining from previous contract awards, the firm's size, measured by the number of plants in the region, and distance from the nearest plant to the project. Bidders' costs and contract characteristics are assumed to be independently distributed conditional on the state variables. In any given auction, bidders who are larger, those with closer plants, and those with less backlog have stochastically lower costs. The state of play changes with bidding outcomes and the location of contracts. Bidders who have an advantage in one auction may be at a disadvantage in another auction. The bidders are restricted to play symmetric Markovian strategies. Symmetry implies that the equilibrium is invariant to permutations of the bidder identities across states of play. The primitives of the model are the conditional cost distributions and the discount rate. Extending the first order approach of Laffont and Vuong, the authors show that the inverse bid function in the dynamic game is equal to the bid plus a markdown factor that consists of two terms. The first is the usual term that accounts for the level of competition in the current auction. The second term accounts for the incremental impact of the outcome of the current auction on future payoffs. The latter term is weighted by the discount rate. Hence, forward looking bidders in the dynamic game bid less aggressively than myopic bidders.

One problem associated with estimating the dynamic auction model is that the primitives of model are not nonparametrically identified. Jofre-Bonet and Pesendorfer fix the discount rate in order to identify the conditional distribution of costs. This is an important issue, since rational behavior has to be assumed and cannot be tested. In particular, the data cannot distinguish between myopic bidders and forward-looking bidders, although the inferred costs and conditional distribution of costs will depend upon the assumed value of the discount rate. A potential source of identification is variation in the time between auctions. For example, in some inclement regions, highway procurement contracts are auctioned prior to the construction season, and have to be completed before the onset of winter. Hence, bidders in the last auction of the bidding season effectively can bid myopically, and ignore the impact of that auction outcome on future play.

Much of the empirical literature studies auctions under the assumption that the number of bidders is fixed. However, the number of bidders may be determined as part of the equilibrium to the auction game. In practice, participating in an auction can be costly. For example, bidders in timber auctions may have to cruise the stand to determine the type and volumes of available timber; bidders in oil and gas auctions have to conduct and analyze geological surveys, and bidders in highway procurement auctions have to develop cost estimates. When these costs are nontrivial and sunk at the bidding stage, the auction has to provide bidders with the prospect of sufficient rents that they will be able cover their entry costs, at least in expectation. The question then arises: does the auction attract too many or too few bidders? This issue is especially important when bidders are asymmetric since, in this case, the Revenue Equivalence Theorem does not hold and auction design matters. For example, the Athey, Levin and Seira (2004) study of oral and sealed bid timber auctions distinguishes between two kinds of bidders, loggers and mills. They argue that a mill's willingness to pay for cutting rights is stochastically greater than that of loggers, since the latter has to sell its logs to a mill. In open, ascending price auctions, the bidder with the highest valuation wins but, in first-price, sealed bid auctions, this need not be the outcome. As a result, loggers are more likely to participate, and win, in a sealed bid auction than in an open auction.

Standard market models use entry decisions by firms to draw inferences about post-entry behavior. In auctions, post-entry behavior is observable and can help identify entry costs and behavior. For example, Athey, Levin, and Seira examine the effect of potential collusion by mills on entry rates of loggers; Athey and Levin (2005) and Krasnokutskaya and Seim (2005) examine the effects on entry of policies that exclude or subsidize certain classes of bidders. A fundamental problem that arises in estimating entry models is the multiplicity of equilibria. If entry is modeled as a simultaneous move, complete information game, then the entry game typically has many asymmetric pure strategy equilibria in which some firms are certain to enter and others stay out. There are also mixed strategy equilibria in which participation is random. The multiplicity of equilibria implies that the mapping from the unobservables to outcomes is not unique, and so the likelihood function is not well-defined. The literature offers a number of ways of modifying the model to generates a well-defined likelihood function for the data: Bresnahan and Reiss (1990, 1991) and Berry (1992) restrict the payoffs of the players so that the number of entrants in the set of pure strategy equilibria is unique, and define the likelihood function in terms of this event; Seim (2005) "purifies" the mixed strategy equilibrium as an (often unique) Bayesian Nash equilibrium by assuming that players have different and private entry costs; Bajari, Hong, and Ryan (2004) append a set of equilibrium selection rules to the model and estimate the joint probability distribution over outcomes and selection rules.[4] But, in auctions, the specific rules governing entry and

[4]Ciliberto and Tamer (2004) develop an estimation approach that allows for multiple equilibria but may sacrifice point identification. They use the theoretical entry model to define upper and lower bounds on the probabilities of the outcomes, and then estimate the parameters of the model by minimizing the distance

bidding decisions in auctions may operate as a selection rule. Bidders may find it easier to coordinate on an asymmetric equilibrium in oral auctions where the bidders have to register and show up to bid than in sealed bid auctions, where the bidders mail or phone in their bids. In the latter case, the mixed strategy equilibrium may be a better description of entry behavior. For example, Athey, Levin, and Seira assume that loggers play a mixed entry strategy, since the data are more consistent with an equilibrium in which too many or too few bidders enter ex post. Similarly, Li (2005) considers structural estimation of an IPV model with a binding reserve price, and entry determined by a symmetric mixed strategy equilibrium.

The structure of eBay auctions, and the availability of data from these auctions, gives researchers an opportunity to study participation and bidding decisions in a dynamic context. Sellers in an eBay auction choose the duration of the auction, but their choices are restricted to 1, 3, 5, 7, or 10 days. The discreteness of the duration menu, combined with the random arrival of sellers, means that the auctions for a product like TI-83 calculators or the Compaq Ipaq H3850 personal digital assistant can be ordered sequentially by their closing times. In fact, eBay often uses this ordering in presenting the choice set of auctions to potential bidders. A number of studies (e.g., Roth and Ockenfels (2002), Bajari and Hortacsu (2003)) have shown that most bidders "snipe," that is, they wait until the closing time of an auction to bid. The winning bidders usually then exit, since most bidders have unit demands. Losing bidders either exit (e.g., buy the item at the posted price in a buy-it-now auction or at a retail store) or bid again in a subsequent auction. If they bid again, and depending to some extent on the length of the time interval between closing times, losing bidders sometimes choose to participate in the same auction but more frequently participate in different auctions. In other words, most losing bidders do not participate in the next-to-close auction. The randomness in the participation decisions may be due in part to unobserved product heterogeneity, but more likely it reflects randomness in participation costs. A losing bidder has to wait until the next closing time to submit a bid, and the opportunity cost of participating in an auction will vary over time. The bidder may decide that he is better off coming back later in the day to submit a bid. Thus, to a first approximation, eBay auctions for a given product are a sequence of second-price, sealed bid auctions in which bidders arrive randomly and participate randomly.

between the empirical frequencies and these bounds.

As in the case of the static models, the first step of the research program consists of characterizing the Markov perfect equilibrium (or equilibria) to the dynamic game. This task is difficult since, in contrast to the highway auctions, the items in eBay auctions are close substitutes. If the intervals between auction closing times is short, losing bidders can learn about their rivals' valuations from their bids. The possibility for learning gives bidders a strategic incentive to bid in ways that make it difficult for rivals to draw inferences. The second step is to estimate the primitives of the model (the bidders' arrival and exit rates, the distribution of valuations, and the distribution of participation costs) treating the participation and bid data as the outcome of equilibrium play. The problems inherent in extending the structural program to dynamic auctions like eBay may prove intractable. However, the normative and positive goals of the research program are worthy of the effort.

We end this survey by noting that the theoretical and empirical literature surveyed in this chapter has emphasized the role of private information as the main source of rents in auctions, and the key determinant of strategic play. However, in many auctions, bidder asymmetries may be a more important source of market power. For example, Moshkin, Porter and Zona (2005) argue that dairies' bids to supply milk products to school districts in the Cincinnati area are largely determined by the distances between the diaries' plants and the schools, as well as other factors that are known by their rivals. They model the auction as a Bertrand pricing game in which costs are common knowledge and, in equilibrium, the lowest cost supplier wins the market by bidding the cost of the marginal rival. Collusion among bidders will affect the identity of the marginal rival. (They also consider a more general setting where fixed entry costs are private information, and where there may be some uncertainty about who will be awarded the contract.) Cost differences among suppliers is the traditional explanation for market power. Since bidder asymmetries in private value auctions can be accommodated using the first-order approach developed by Laffont and Vuong, it is possible to examine the relative importance of private information and bidder asymmetries as sources of rent.

References

[1] Aguirregabiria, V. and P. Mira (2004) "Sequential Estimation of Dynamic Discrete Games," mimeo, Boston University.

[2] Asker, J. and E. Cantillon (2004) "Properties of Scoring Auctions," mimeo, Harvard University.

[3] Ashenfelter, O. (1989) "How Auctions Work for Wine and Art," *Journal of Economic Perspectives*, 3: 23-36.

[4] Ashenfelter, O. and D. Genesove (1992) "Testing for Price Anomalies in Real Estate Auctions," *American Economic Review Papers and Proceedings*, 82: 501-5.

[5] Ashenfelter, O. and K. Graddy (2003) "Auctions and the Price of Art," *Journal of Economic Literature*, 41: 763-87.

[6] Ashenfelter, O. and K. Graddy (2004) "Anatomy of the Rise and Fall of a Price-Fixing Conspiracy: Auctions at Sotheby's and Christie's," mimeo, Princeton University.

[7] Athey, S. (2001) "Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information," *Econometrica*, 69: 861-89.

[8] Athey, S. and Bagwell, K. (2001) "Optimal Collusion with Private Information." *RAND Journal of Economics*, 32: 428-65.

[9] Athey, S., K. Bagwell and Sanchirico, C. (2004) "Collusion and Price Rigidity." *Review of Economic Studies*, 71: 317-49.

[10] Athey, S. and P. Haile (2002) "Identification of Standard Auction Models," *Econometrica*, 70: 2107-40.

[11] Athey, S. and P. Haile (2005) "Nonparametric Approaches to Auctions," *Handbook of Econometrics*, Vol. 6, forthcoming.

[12] Athey, S. and J. Levin (2001) "Information and Competition in U.S. Forest Service Timber Auctions," *Journal of Political Economy*, 109: 375-417.

[13] Athey, S. and J. Levin (2005) "Set-Asides and Subsidies in Auctions," mimeo, Stanford University.

[14] Athey, S., J. Levin and E. Seira (2004) "Comparing Open and Sealed Bid Auctions: Theory and Evidence from Timber Auctions," mimeo, Stanford University.

[15] Ausubel, L. and P. Cramton (2002) "Demand Reduction and Inefficiency in Multi-Unit Auction," mimeo, University of Maryland.

[16] Avery, C. (1998) "Strategic Jump Bidding in English Auctions," *Review of Economic Studies*, 65: 185-210.

[17] Bajari, P. (1997) "The First Price Auction with Asymmetric Bidders: Theory and Applications," unpublished Ph.D. thesis, University of Minnesota.

[18] Bajari, P. (1998) "Econometrics of Sealed-Bid Auctions," *Proceedings of the Business and Economic Statistics Section of the American Statistical Association*, 41-9.

[19] Bajari, P., L. Benkard and J. Levin (2005) "Estimating Dynamic Models of Imperfect Competition," mimeo, Stanford University.

[20] Bajari, P., S. Houghton and S. Tadelis (2006) "Bidding for Incomplete Contracts: An Empirical Analysis," mimeo, University of Michigan.

[21] Bajari, P., H. Hong and S. Ryan (2004) "Identification and Estimation of Discrete Games of Complete Information," mimeo, Duke University.

[22] Bajari, P. and A. Hortacsu (2003) "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions," *RAND Journal of Economics*, 34: 329-55.

[23] Bajari, P. and A. Hortacsu (2004) "Economic Insights from Internet Auctions," *Journal of Economic Literature*, 42: 457-86.

[24] Bajari, P. and A. Hortacsu (2005) "Are Structural Estimates of Auction Models Reasonable? Evidence from Experimental Data," *Journal of Political Economy*, 113: 703-41.

[25] Bajari, P. and L. Ye (2003) "Deciding Between Competition and Collusion," *Review of Economics and Statistics*, 85: 971-89.

[26] Baldwin, L., R. Marshall and J.F. Richard (1997) "Bidder Collusion at Forest Service Timber Sales," *Journal of Political Economy,* 105: 657-99.

[27] Beggs, A. and K. Graddy (1997) "Declining Values and the Afternoon Effect: Evidence from Art Auctions," *RAND Journal of Economics*, 28: 544-65.

[28] Berry, S. (1992) "Estimation of a Model of Entry in the Airline Industry," *Econometrica*, 60: 889-917.

[29] Bikhchandani, S., P. Haile and J. Riley (2002) "Symmetric Separating Equilibria in English Auctions," *Games and Economic Behavior*, 38: 19-27.

[30] Bikhchandani, S. and C. Huang (1993) "The Economics of Treasury Securities Markets," *Journal of Economic Perspectives,* 7: 117-34.

[31] Bikhchandani, S. and J. Riley (1991) "Equilibria in Open Common Value Auctions," *Journal of Economic Theory*, 53: 101-30.

[32] Black, J. and D. de Meza (1992) "Systematic Price Differences Between Successive Auctions are No Anomaly," *Journal of Economics & Management Strategy*, 1: 607-28.

[33] Bresnahan, T. and P. Reiss (1990) "Entry in Monopoly Markets," *Review of Economic Studies*, 57: 531-53.

[34] Bresnahan, T. and P. Reiss (1991) "Empirical Models of Discrete Games," *Journal of Econometrics*, 48: 57-81.

[35] Bulow, J. and P. Klemperer (2002) "Prices and the Winner's Curse," *RAND Journal of Economics*, 33: 1-21.

[36] Burguet, R. and M. Perry (2002) "Bribery and Favoritism by Auctioneers in Sealed Bid Auctions," mimeo, Rutgers University.

[37] Cantillon, E. and M. Pesendorfer (2004) "Combination Bidding in Multi-Unit Auctions," mimeo, London School of Economics.

[38] Capen, E., R. Clapp, and W. Campbell (1971) "Competitive Bidding in High-Risk Situations," *Journal of Petroleum Technology*, 23: 641-53.

[39] Castellanos, S. and M. Oviedo (2004) "Optimal Bidding in the Mexican Treasury Securities Primary Auctions: Results from a Structural Econometric Approach," mimeo, Yale University.

[40] Chernozhukov, V. and H. Hong (2004) "Likelihood Estimation and Inference in a Class of Nonregular Econometric Models," *Econometrica*, 72: 1445-80.

[41] Ciliberto, F. and E. Tamer (2004) "Market Structure and Multiple Equilibria in Airline Markets," mimeo, Northwestern University.

[42] Comanor, W. and M. Schankerman (1976) "Identical Bids and Cartel Behavior," *Bell Journal of Economics*, 7: 281-6.

[43] Cramton, P. (1995) "Money Out of Thin Air: The Nationwide Narrowband PCS Auction," *Journal of Economics & Management Strategy*, 4: 267-343.

[44] Cramton, P. and T. Palfrey (1995) "Ratifiable Mechanisms: Learning from Disagreement," *Games and Economic Behavior*, 10: 255-83.

[45] Cramton, P. and J. Schwartz (2000) "Collusive Bidding: Lessons from the FCC Spectrum Auctions," *Journal of Regulatory Economics*, 17: 229-52.

[46] Das, S. and R. Sundaram (1997) "Auction Theory: A Summary with Applications to Treasury Markets," NBER Working Paper 5873.

[47] Donald, S. and H. Paarsch (1993) "Piecewise Pseudo-Maximum Likelihood Estimation in Empirical Models of Auctions," *International Economic Review*, 34: 121-48.

[48] Donald, S. and H. Paarsch (1996) "Identification, Estimation, and Testing in Parametric Empirical Models of Auctions within the Independent Private Values Paradigm," *Econometric Theory,* 12: 517-67.

[49] Donald, S. and H. Paarsch (2002) "Superconsistent Estimation and Inference in Structural Econometric Models Using Extreme Order Statistics," *Journal of Econometrics,* 109: 305-40.

[50] Donald, S., H. Paarsch and J. Robert (2006) "An Empirical Model of the Multi-Unit, Sequential Clock Auction," *Journal of Applied Econometrics,* forthcoming.

[51] Elyakime, B., J.J. Laffont, P. Loisel and Q. Vuong (1994) "First-Price, Sealed-Bid Auctions with Secret Reservation Prices," *Annales d'Economie et de Statistique,* 34: 115-41.

[52] Engelbrecht-Wiggans, R., P. Milgrom, and R. Weber (1983) "Competitive Bidding and Proprietary Information," *Journal of Mathematical Economics*, 11: 161-9.

[53] Feinstein, J., M. Block and F. Nold (1985) "Asymmetric Information and Collusive Behavior in Auction Markets," *American Economic Review*, 75: 441-60.

[54] Fevrier, P., R. Preget, and M. Visser (2004) "Econometrics of Share Auctions," mimeo, University of Chicago.

[55] Gertner, R. (1995) "Revenue and Efficiency Differences Between Sequential and Simultaneous Multiple-Unit Auctions with Limited Competition," mimeo, University of Chicago.

[56] Graham, D. and R. Marshall (1987) "Collusive Bidder Behavior at Single-Object Second-Price and English Auctions," *Journal of Political Economy*, 95: 1217-39.

[57] Guerre, E., I. Perrigne and Q. Vuong (2000) "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica*, 68: 525-74.

[58] Haile, P. (2001) "Auctions with Resale Markets: An Application to U.S. Forest Service Timber Sales," *American Economic Review*, 91: 399-427.

[59] Haile, P., H. Hong and M. Shum (2004) "Nonparametric Tests for Common Values in First-Price Sealed Bid Auctions," mimeo, Yale University.

[60] Haile, P. and E. Tamer (2003) "Inference in an Incomplete Model of English Auctions," *Journal of Political Economy*, 111: 1-51.

[61] Hansen, R. (1986) "Sealed-Bid Versus Open Auctions: The Evidence," *Economic Inquiry*, 24: 125-42.

[62] Harrington, J. (2005) "Detecting Cartels," mimeo, Johns Hopkins University.

[63] Harrison, G. and J. List (2004) "Field Experiments," *Journal of Economic Literature*, 62: 1009-55.

[64] Hendricks, K. and H. Paarsch (1995) "A Survey of Recent Empirical Work Concerning Auctions," *Canadian Journal of Economics*, 28: 315-38.

[65] Hendricks, K., J. Pinkse and R. Porter (2003) "Empirical Implications of Equilibrium Bidding in First-Price, Symmetric, Common Value Auctions," *Review of Economic Studies*, 70: 115-45.

[66] Hendricks, K. and R. Porter (1988) "An Empirical Study of An Auction with Asymmetric Information," *American Economic Review*, 78: 865-83.

[67] Hendricks, K. and R. Porter (1989) "Collusion in Auctions," *Annales d'Economie et de Statistique*, 16-17: 217-30.

[68] Hendricks, K., R. Porter and B. Boudreau (1987) "Information, Returns and Bidding Behavior in OCS Auctions: 1954-1969," *Journal of Industrial Economics*, 35: 517-42.

[69] Hendricks, K., R. Porter and G. Tan (1993) "Optimal Selling Strategies for oil and Gas Leases with an Informed Buyer," *American Economic Review Papers and Proceedings*, 83: 234-9.

[70] Hendricks, K., R. Porter and G. Tan (2006) "Bidding Rings and the Winner's Curse," mimeo, Northwestern University.

[71] Hendricks, K., R. Porter and C. Wilson (1994) "Auctions For Oil and Gas Leases with an Informed Bidder and a Random Reservation Price," *Econometrica*, 62: 1415-44.

[72] Hirano, K. and J. Porter (2003) "Asymptotic Efficiency in Parametric Structural Models with Parameter-Dependent Support," *Econometrica*, 71: 1307-38.

[73] Holt, C. (1980) "Competitive Bidding for Contracts under Alternative Auction Procedures," *Journal of Political Economy*, 88: 433-45.

[74] Hong, H. and M. Shum (2002) "Increasing Competition and the Winner's Curse: Evidence from Procurement," *Review of Economic Studies*, 69: 871-98.

[75] Hong, H. and M. Shum (2003) "Econometric Models of Ascending Auctions," *Journal of Econometrics*, 112: 327-58.

[76] Hortacsu, A. (2002) "Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market," mimeo, University of Chicago.

[77] Hotz, J. and R. Miller (1993) "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, 60: 497-531.

[78] Jehiel, P. and B. Moldovanu (2003) "Auctions with Downstream Interaction Among Buyers," *RAND Journal of Economics*, 31: 768-91.

[79] Jofre-Bonet, M. and M. Pesendorfer (2003) "Estimation of a Dynamic Auction Game," *Econometrica*, 71: 1443-89.

[80] Kagel, J. (1995) "Auctions: A Survey of Experimental Research," in: J. Kagel and A. Roth, eds., *Handbook of Experimental Economics*. Princeton, N.J.: Princeton University Press, 501-85.

[81] Kastl, J. (2005) "Discrete Bids and Empirical Inference in Divisible Goods Auctions," mimeo, Northwestern University.

[82] Klemperer, P. (1998) "Auctions with Almost Common Values: The "Wallet Game" and its Applications," *European Economic Review*, 42: 757-69.

[83] Klemperer, P. (2004) *Auctions: Theory and Practice*. Princeton, NJ: Princeton University Press.

[84] Krasnokutskaya, E. (2004) "Identification and Estimation in Highway Procurement Auctions Under Unobserved Auction Heterogeneity," mimeo, University of Pennsylvania.

[85] Krasnokutskaya, E. and K. Seim (2005) "Bid Preference Programs and Participation in Highway Procurement Programs," mimeo, University of Pennsylvania.

[86] Kremer, I. and K. Nyborg (2004) "Divisible-Good Auctions: The Role of Allocation Rules," *RAND Journal of Economics*, 35: 147-59.

[87] Krishna, V. (2002) *Auction Theory*. San Diego: Academic Press.

[88] Laffont, J.J. (1997) "Game Theory and Empirical Economics: The Case of Auction Data," *European Economic Review*, 4: 1-35.

[89] Laffont, J.J., H. Ossard and Q. Vuong (1995) "Econometrics of First-Price Auctions," *Econometrica*, 63: 953-80.

[90] Laffont, J.J. and Q. Vuong (1996) "Structural Analysis of Auction Data," *American Economic Review Papers and Proceedings*, 86: 414-20.

[91] Lebrun, B. (1996) "Existence of an Equilibrium in First Price Auctions," *Economic Theory*, 7: 421-43.

[92] Lebrun, B. (1999) "First Price Auctions in the Asymmetric N Bidder Case," *International Economic Review*, 40: 125-42.

[93] Li, T. (2005) "Econometrics of First-Price Auctions with Entry and Binding Reservation Prices," *Journal of Econometrics*, 126: 173-200.

[94] Li, T., I. Perrigne and Q. Vuong (2000) "Conditionally Independent Private Information in OCS Wildcat Auctions," *Journal of Econometrics*, 98: 129-61.

[95] Li, T., I. Perrigne and Q. Vuong (2002) "Structural Estimation of the Affiliated Private Value Auction Model," *RAND Journal of Economics*, 33: 171-93.

[96] LiCalzi, M. and A. Pavan (2005) "Tilting the Supply Schedule to Enhance Competition in Uniform Price Auctions," *European Economic Review*, 49: 227-50.

[97] Mailath, G. and P. Zemsky (1991) "Collusion in Second Price Auctions with Heterogeneous Bidders," *Games and Economic Behavior*, 4: 467-86.

[98] Marshall, R. and L. Marx (2004) "Bidder Collusion," mimeo, Duke University.

[99] Maskin, E. and J. Riley (2000a) "Asymmetric Auctions," *Review of Economic Studies*, 67: 413-38.

[100] Maskin, E. and J. Riley (2000b) "Equilibrium in Sealed High Bid Auctions," *Review of Economic Studies*, 67: 439-54.

[101] Matthews, S. (1980) "Risk Aversion and the Efficiency of First and Second Price Auctions," mimeo, Northwestern University.

[102] Matthews, S. (1995) "A Technical Primer on Auction Theory I: Independent Private Values," mimeo, Northwestern University.

[103] McAdams, D. (2005) "Adjustable Supply in Uniform-Price Auctions: The Value of Non-Commitment," mimeo, Massachusetts Institute of Technology.

[104] McAfee, P. and J. McMillan (1987) "Auctions and Bidding," *Journal of Economic Literature*, 25: 699-738.

[105] McAfee, P. and J. McMillan (1992) "Bidding Rings," *American Economic Review*, 82: 579-99.

[106] McAfee, P. and J. McMillan (1996) "Analyzing the Airwaves Auction," *Journal of Economic Perspectives*, 10: 159-75.

[107] McAfee, P., D. Quan and D. Vincent (2002) "How To Set Minimum Acceptable Bids, with an Application to Real Estate Auctions," *Journal of Industrial Economics,* 50: 391-416.

[108] McAfee, P. and D. Vincent (1992) "Updating the Reserve Price in Common-Value Auctions," *American Economic Review Papers and Proceedings*, 82: 512-8.

[109] McAfee, P. and D. Vincent (1993) "The Declining Price Anomaly," *Journal of Economic Theory*, 60: 191-212.

[110] McFadden, D. (1989) "A Method of Simulated Moments for Estimation of Multinomial Discrete Response Models Without Numerical Integration," *Econometrica*, 57: 995-1026.

[111] McMillan, J. (1994) "Selling Spectrum Rights," *Journal of Economic Perspectives*, 8: 145-62.

[112] Meade, W., P. Sorenson, R. Jones and A. Moseidjord (1980) *Competition and Performance in OCS Oil and Gas Lease Sales and Development, 1954-1969*, Final Report to US Geological Survey, Reston, VA.

[113] Milgrom, P. (1981a) "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 12: 380-91.

[114] Milgrom, P. (1981b) "Rational Expectations, Information Acquisition, and Competitive Bidding," *Econometrica*, 49: 921-43.

[115] Milgrom, P. (1985) "The Economics of Competitive Bidding: A Selective Survey," in: L. Hurwicz, D. Schmeidler and H. Sonnenschein, eds., *Social Goals and Social Organization*. Cambridge: Cambridge University Press.

[116] Milgrom, P. (1987) "Auction Theory," in: T. Bewley, ed., *Advances in Economic Theory: Fifth World Congress*. Cambridge: Cambridge University Press.

[117] Milgrom, P. (1989) "Auctions and Bidding: A Primer," *Journal of Economic Perspectives*, 3: 3-22.

[118] Milgrom, P. (2004) *Putting Auction Theory to Work*. Cambridge: Cambridge University Press.

[119] Milgrom, P. and R. Weber (1982) "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50: 1089-122.

[120] Milgrom, P. and R. Weber (2000) "A Theory of Auctions and Competitive Bidding, II" in: P. Klemperer, ed., *The Economic Theory of Auctions*. Cheltenham, U.K.: Edward Elgar.

[121] Moshkin, N., R. Porter and D. Zona (2005) "Bid Rigging in School Milk Auctions: Evidence from Ohio," mimeo, Northwestern University.

[122] Mund, V. (1960) "Identical Bid Prices," *Journal of Political Economy*, 68: 150-69.

[123] Myerson, R. (1981) "Optimal Auction Design," *Mathematics of Operations Research*, 6: 58-73.

[124] Ortega Reichert, A. (1968) *Models for Competitive Bidding under Uncertainty*. Ph.D. Dissertation, Stanford University.

[125] Paarsch, H. (1992) "Deciding between the Common and Private Value Paradigms in Empirical Models of Auctions," *Journal of Econometrics*, 51: 191-215.

[126] Paarsch, H. (1997) "Deriving an Estimate of the Optimal Reserve Price: An Application to British Columbia Timber Sales," *Journal of Econometrics*, 78: 333-57.

[127] Paarsch, H. and H. Hong (2006) *An Introduction to the Structural Econometrics of Auction Data*. Cambridge, MA: MIT Press.

[128] Pakes, A., M. Ostrovsky and S. Berry (2005) "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples)," mimeo, Harvard University.

[129] Pakes, A. and D. Pollard (1989) "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*, 57: 1027-57.

[130] Perrigne, I. and Q. Vuong (1999) "Structural Econometrics of First-Price Auctions: A Survey and Methods," *Canadian Journal of Agricultural Economics*, 47: 203-23.

[131] Pesendorfer, M. (2000) "A Study of Collusion in First-Price Auctions," *Review of Economic Studies*, 67: 381-411.

[132] Pesendorfer, M. and P. Schmidt-Dengler (2004) "Least Squares Estimators for Dynamic Games," mimeo, London School of Economics.

[133] Pinkse, J. and G. Tan (2005) "The Affiliation Effect in First-Price Auctions," *Econometrica*, 73: 263-77.

[134] Porter, R. (1992) "Review of *Anatomy of an Auction: Rare Books at Ruxley Lodge, 1919* by Arthur Freeman and Janet Ing Freeman," *Journal of Political Economy*, 100: 433-6.

[135] Porter, R. (1995) "The Role of Information in U.S. Offshore Oil and Gas Lease Auctions," *Econometrica*, 63: 1-27.

[136] Porter, R. (2005) "Detecting Collusion," *Review of Industrial Organization*, 26: 147-67.

[137] Porter, R. and D. Zona (1993) "Detection of Bid Rigging in Procurement Auctions," *Journal of Political Economy*, 101: 518-38.

[138] Porter, R. and D. Zona (1999) "Ohio School Milk Markets: An Analysis of Bidding," *RAND Journal of Economics*, 30: 263-88.

[139] Riley, J. (1989) "Expected Revenues from Open and Sealed Bid Auctions," *Journal of Economic Perspectives*, 3: 41-50.

[140] Riley, J. and W. Samuelson (1981) "Optimal Auctions," *American Economic Review*, 71: 381-92.

[141] Robinson, M. (1985) "Collusion and the Choice of Auction," *RAND Journal of Economics*, 16: 141-5.

[142] Roth, A. and A. Ockenfels (2002) "Last Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet," *American Economic Review*, 92: 1093-103.

[143] Scherer, F. M. and D. Ross (1990) *Industrial Market Structure and Economic Performance*, 3rd ed., Boston, MA: Houghton Miflin.

[144] Seim, K. (2005) "An Empirical Model of Firm Entry with Endogenous Product-Type Choices," mimeo, Stanford University.

[145] Shneyerov, A. (2005) "An Empirical Study of Auction Revenue Rankings: The Case of Municipal Bonds," mimeo, University of British Columbia.

[146] Smiley, A. (1979) *Competitive Bidding Under Uncertainty: The Case of Offshore Oil*. Cambridge, Mass.: Ballinger.

[147] Smith, R. (1961) "The Incredible Electrical Conspiracy," *Fortune*, 63: April, 132-80, and May, 161-224.

[148] Van den Berg, G. and B. van der Klaauw (2000) "Structural Analysis of Dutch Flower Auctions," mimeo, Free University Amsterdam.

[149] Vickrey, W. (1961) "Counterspeculation, Auctions and Competitive Sealed Tenders," *Journal of Finance*, 16: 8-37.

[150] Vickrey, W. (1962) "Auctions and Bidding Games," in *Recent Advances in Game Theory*, Princeton Conference Series, 29, Princeton, NJ: Princeton University Press, 15-27.

[151] Wilson, R. (1977) "A Bidding Model of Perfect Competition," *Review of Economic Studies*, 44: 511-18.

[152] Wilson, R. (1993) "Strategic Analysis of Auctions," in R. Aumann and S. Hart, eds., *Handbook of Game Theory*, Amsterdam: North-Holland.

[153] Wilson, R. (1998) "Sequential Equilibria of Asymmetric Ascending Auctions," *Economic Theory*, 12: 433-40.

[154] Wolak, F. (2003) "Identification and Estimation of Cost Functions Using Observed Bid Data: An Application to Electricity," in M Dewatripont, L. Hansen and S. Turnovsky, eds., *Advances in Econometrics: Theory and Applications, Eighth World Congress, Vol. 2*, Cambridge: Cambridge University Press.

[155] Wolak, F. (2005) "Quantifying the Supply Side Benefits from Forward Contracting in Wholesale Electricity Markets," mimeo, Stanford University.

[156] Wolfram, C. (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: 703-25.

[157] Wolfram, C. (1999) "Measuring Duopoly Power in the British Electricity Spot Market," *American Economic Review*, 89: 805-26.