An Estimable Demand System for a Large Auction Market

Matthew Backus[*] — University of Michigan

Gregory Lewis[†] — Harvard University and NBER

[Preliminary — please do not cite]

September 23, 2009

Abstract

Auction mechanisms play an important role in allocating goods. But unlike fixed-price markets, we do not yet have a good demand system in markets with substitutable products, and this limits our ability to make counterfactual predictions. Here we provide a framework for estimating a demand system in large auction markets with numerous persistent buyers, heterogeneous objects and unit demand. We construct a dynamic model of repeated second-price auctions with bidder entry and exit, in which consumers have independent and private valuation vectors over the full set of objects. We prove existence of an equilibrium and characterize equilibrium strategies, and then provide sufficient conditions for the nonparametric identification of the joint distribution of private values. Nonparametric and semiparametric estimation procedures are proposed and tested by Monte Carlo simulation.

[*]Email: backus@umich.edu

[†]Email: glewis@fas.harvard.edu

1 Introduction

Auction markets — and particularly auction platform markets — are responsible for a substantial and growing share of trade in the US economy. E-commerce as a whole has been growing rapidly, with estimated growth of 13.9% in 2006, relative to 6.8% growth rates for offline trade in the same sectors.[1] Of this trade, 93% is accounted for by business-to-business transactions. It is harder to quantify the share held specifically by auction platforms, but in retail eBay alone has revenues of $36 billion from its auctions business in 2007[2]; Google realized $21 billion in revenue from its online advertising platform in 2008[3]; and auction sites such as DoveBid and IronPlanet have sold billions of dollars of used aviation and construction equipment respectively.

Given their importance in the modern economy, one would like to be able to estimate demand in these markets. This would allow us to answer questions of broad economic interest, such as how much welfare has been generated by these platforms; as well as narrower strategic questions, such as how a firm with a fixed inventory should set reserves and time sales to dynamically maximize its revenue. Demand estimation is often also a necessary first step for the evaluation of anti-trust issues, such as the potential impact on the search-keyword advertising market of a merger between Microsoft and Yahoo.

At first glance, auctions data is an extremely rich of information about demand. For any buyer we generally observe all the auctions that they bid in, which provides valuable information about which items they view as close substitutes. For example, a buyer who loses in an auction for a particular Corvette may choose to bid in an auction for a Mustang. This is informative for demand, much in the same way that "second-choice" data is useful in Berry, Levinsohn, and Pakes (2004). Moreover, the choice sets available to buyers vary with high frequency as auctions come and go, which is also useful for identifying substitution patterns.

Yet the strengths of auction market data also pose some difficulties. As Hendricks and Porter (2007) note in their survey article, participants in auction markets are playing a complex dynamic game, where they must continuously adapt to the changing set of available auctions, and try to learn about rival's valuations. Most of the existing tools of structural auction econometrics are focused on independent auctions of homogenous objects, which

[1]Source: US Census Bureau, E-Stats Report 2006. Includes only the manufacturing, merchant wholesale, retail and selected services sectors; excludes eMarketplaces, such as eBay.

[2]Source: eBay Annual Report for 2007

[3]Source: Google Annual Report for 2008

limits their direct applicability to auction demand estimation, where *substitution* across products is important. In this paper, we provide theory, identification and estimation results for estimating demand in auction platform markets with a large number of relatively short-lived buyers. We perform Monte Carlo simulations to show that our estimation approach works well in moderately sized samples.

A key assumption in the analysis will be that the market is *large*, in the sense that there are many bidders and similar objects up for auction. In this case, the true state of the market at any time will be complex and unknown to bidders. Rather than forming beliefs about the unknown true state, they may instead condition their behavior on a simpler publicly observable state vector and on their own private valuation. Similar assumptions underpin the equilibrium notions of Krusell and Smith (1998), Weintraub, Benkard, and Roy (2008) and Fershtman and Pakes (2009). An attractive feature of this approach is that determining what the relevant states are becomes an empirical question.

Then, if bidders have unit demand, we have a simple way to deal with dynamic concerns. Participation has an option value: the expected surplus from future auctions conditional on today's state. This option value is struck when a bidder wins an auction, so she shades her bids accordingly. The challenge of estimating private values is then estimating the long-run option value. We show that this is non-parametrically identified from observations of bids placed when facing different sequences of upcoming auctions.

The paper is related to the empirical literature on dynamic auctions. Jofre-Benet and Pesendorfer (2003) was the first paper to attack estimation in a dynamic auction game, though in a world with a small number of infinitely long-lived bidders. Subsequent to this, a number of papers have looked at dynamics on eBay specifically. Budish (2008) examines the optimality of eBay's market design with respect to the sequencing of sales and information revelation. Zeithammer (2006) developed a model with forward-looking bidders, and showed both theoretically and empirically that bidders shade down current bids in response to the presence of upcoming auctions of similar objects. Ingster (2009) develops a dynamic model of auctions of identical objects, and provides equilibrium characterization and identification results. Sailer (2006) estimates participation costs for bidders facing an infinite sequence of identical auctions. Relative to this literature, our main contribution is the focus on sequential auctions of heterogeneous objects, where bidders have multidimensional persistent private valuations. In short, we're focused on developing a demand system. A different approach has been taken in Adams (2009), who looks at the problem of nonparametric identification when auctions are completely simultaneous. A second related literature is the large literature on demand estimation in durable goods markets (e.g. Berry, Levinsohn, and Pakes (1995)).

The next section describes eBay and motivates the theoretical framework given in Section 3. Section 4 proves non-parametric identification. Section 5 describes our two different estimation approaches, while section 6 gives Monte Carlo simulations for those estimators. Section 7 concludes.

2 Model

Our aim in this section is to create an abstract model of a large auction market, and analyze it. The space of such models is vast, and we narrow in a number of ways. We consider a market in which competing products — such as iPods and Zunes — are sold by second-price sealed bid auctions. These auctions are held in discrete time, with one good auctioned per period over an infinite horizon. Since our focus is on demand, we assume for simplicity that supply is random and exogenous. Bidders are persistent with unit demand, and enter the market with private (possibly correlated) valuations for each of the objects. Winning bidders immediately exit, while losing bidders exit randomly. The market is large in the sense that there are many buyers and many goods auctioned, although the ratio of buyers to goods may be close to one. We aim at characterizing the long-run behavior of this dynamic system.

We have chosen this set of assumptions to match some features of the environment on eBay, the world leader in online auctions. In any eBay category (such as digital cameras), there are many different products sold by auction to a large number of anonymous buyers.[4] Although these auctions typically last for many days, and thus overlap — so that at any given point in time there are many auctions occurring simultaneously — they finish at different ending times, in sequence. As Bajari and Hortacsu (2004) and Hendricks and Porter (2007) have noted, this timing, combined with the way the proxy bidding system works, imply that eBay is well approximated as sequence of second-price sealed bid auctions. Given that this is the dominant auction platform design, we thought these a reasonable set of modeling assumptions. Yet our intent is not to model eBay *per se* — and indeed we ignore some important features of the eBay environment — but rather to develop a reasonably motivated and rich abstract model and see what we can learn from the exercise.

[4]eBay hides the identity of the bidders by replacing parts of the username with asterixes.

2.1 Environment

We formalize the above description of the environment in what follows:

Bidders and Payoffs: Bidders have unit demand for a good in the set $\mathcal{A}$, where $|\mathcal{A}| = J$. They have a privately known vector of valuations x for the goods in $\mathcal{A}$, their type. They are risk neutral, and receive a payoff of $x_j - p$ for buying a single good j at price p, and zero otherwise. Bidders are impatient, with a common discount rate δ.

Market: Time is discrete with infinite horizon, $t = 1, 2 \cdots$. In each period t, the following stage game is played. First, a sealed-bid second price auction is held for the current object a_t, in which all bidders participate. Then entry and exit of bidders takes place, and suppliers post new objects. These are described in more detail below.

Auctions: At time t, object $a_t \in \mathcal{A}$ is auctioned. Bidders submit bids $b_t = (b_{1t}, b_{2t} \cdots b_{N_t t})$, where N_t is the number of bidders in the market at t. The winner is the bidder who submits the highest bid, and he pays the second highest bid to the seller.

Entry and Exit: At the end of every period, the winner is assumed to exit with certainty.[5] Losers exogenously exit the market with probability $\rho \in (0, 1)$, receiving a payoff normalized to zero on exit. Simultaneously, E_t new bidders enter, where E_t is random with support on a finite set of integers $\{1, 2, 3 \cdots \overline{E}\}$. The distribution of E_t depends on the total number of participants in the previous period, N_{t-1}. To ensure that N_t doesn't explode, we assume $\exists\ \overline{N}$ such that $\mathbb{P}(E_t > \overline{N} - N_{t-1} | N_{t-1}) = 0$. To ensure the market doesn't collapse nor systematically lose all its high types, we assume that $\mathbb{P}(E_t \geq 2 | N_{t-1} \leq \overline{N} - 2) > 0$. Each entrant draws their valuation vector x identically and independently from a distribution F with strictly positive and continuous density over a compact set $\mathcal{X} = [0, \bar{x}]^J$.

Supply: Supply is essentially the rate at which different products appear on the auction market. With that in mind, we assume that at the end of period t, suppliers list a new object to be auctioned in period $t + m^f + 1$. The object to be auctioned is randomly chosen according to a multinomial distribution over the set of products $\mathcal{A}$.[6]

Information Sets and Bidding Strategies: New entrants are assumed to be able to view

[5] Since they have unit demand, they are indifferent about exiting in any period following a win. But even an $\varepsilon > 0$ participation cost would make exit optimal.

[6] It is easy to extend the model to allow for fluctuating supply by letting the the object be multinomial over $\mathcal{A} \cup \varnothing$, where $\varnothing$ is the event that nothing is listed.

an anonymized history of the game for the last m^h periods.[7] By anonymized, we mean that although the distribution of bids in each period is observable, the identities of the bidders making the bids is not. Letting t_i denote the time of entry of bidder i, we have that bidder i can observe a "window" of past actions and current and upcoming auctions, from $t_i - m^h$ to $t_i + m^f$. The cases $m^h = 0$ and $m^f = 0$ correspond to no observable history and no knowledge of future supply, respectively. So at any time t, the information set of any bidder i consists of their valuation x, the history $h_{it} \in \mathcal{H}^{it}$ and the list of current and upcoming objects $(a_t \cdots a_{t+m})$. We denote the set of all *public* information by Λ. A bid strategy $\beta_i : \mathcal{X} \times \mathcal{H}^{it} \times \mathcal{A}^{m+1} \to \mathbb{R}^+$ is a mapping from any information set to a bid. We will restrict attention to symmetric pure strategies $\beta_i = \beta$.

2.2 Analysis

We analyze this environment in three parts. First, we characterize the long-run properties of the dynamic system for arbitrary symmetric bidding strategies. Second, we define a new notion of equilibrium — a "competitive Markov equilibrium" — and thus characterize the bidding strategies of bidders in this system. Finally, we analyze how these bidding strategies aggregate to show existence of this equilibrium, and then show that this concept agrees with more standard equilibrium definitions as the market becomes "large".

At any point in time t, the state of the system is summarized by who is in the market — a vector of types — and what is to be auctioned now and in the future. Since bidders enter and exit every period over an infinite horizon, if we kept track of specific identities the state-space would grow without bound. Instead, we use an anonymous $\overline{N}$-vector $\mathbf{x}_t$ of types currently in the market, where we 0 as a placeholder when there are fewer than $\overline{N}$ bidders in the market.[8] Type transitions occur first by removing exiting bidders and replacing them with the placeholder, and then adding entrants sequentially, starting from the first open placeholder.[9] The vector of current and upcoming objects is $\mathbf{a_t} = (a_t \cdots a_{t+m})$.

Thus we define the "true state" of the system at t as $\omega_t = (\mathbf{x}_t, \mathbf{a_t}) \in \Omega \equiv \mathcal{X}^{\overline{N}} \times \mathcal{A}^{m+1}$. We use

[7] This mimics eBay, where the history of auctions held in the past 14+ days is public (though anonymized).

[8] The restriction to anonymous states is innocuous given that we assume bidders employ symmetric strategies, and payoffs are also symmetric.

[9] For example, suppose we have $\overline{N} = 3$, and there are two bidders with (unidimensional valuations) 1 and 2 respectively. Then we have $\mathbf{x_t} = (1, 2, 0)$; and if at the end of the period bidder 1 exits and two new bidders with values 3 and 4 respectively enter, we will have $\mathbf{x}_{t+1} = (3, 2, 4)$.

the term "true state" to distinguish it from a later definition of state. Given a fixed bidding strategy β, the true state ω_t evolves as a Markov Process on the compact set Ω. Denote by $\mathcal{F}$ the Borel σ-field over Ω. The assumptions made thus far imply that this process has an invariant measure.

Lemma 1 (Ergodic Distribution of True States). *For any bidding strategy β, there is a unique invariant measure μ_β on the measurable space $(\Omega, \mathcal{F})$, with $\mu_\beta(A) > 0 \; \forall A \in \mathcal{F}$.*

This lemma basically says that regardless of the strategies played, all the states in Ω are recurrent. The rough argument for this is that all the states "communicate" — you can get from one to the other. To see this, given $A, B \in \mathcal{F}$, to get from a state in A to one in B, you need the set of upcoming auctions to change; those bidders in A but not B to exit; and those bidders in B but not A to enter. But since supply and entry are completely random, and exit is almost random with *only* the winner determined by the strategies, it is not hard to show that this can happen in a finite number of steps with positive probability.

Now consider how a bidder should decide to bid. Fixing the strategies of his opponents in turn fixes the transition function of ω_t. Then all the bidder needs to know about the opposing bids today and in the future, as well as the objects to be auctioned, is summarized in a single payoff-relevant vector, ω_t. In principle, since ω_t is not observable, the bidder should solve a filtering problem in every period to determine his best guess of the distribution of ω_t given his private history h_{it}. However, this is both technically problematic and seemingly unrealistic. On the technical side, the (vanishingly small) possibility of infinitely lived bidders implies that private histories may never recur, and that therefore the econometrician simply cannot make inferences about the relationship between bids and valuations, even with infinite data.

But more importantly, expecting this level of inference from bidders seems unrealistic. Instead, we suppose that bidders assume that now and in the future, they are to compete with a random draw from a population of types over current and upcoming objects, where they may condition their beliefs about this population on some simple and publicly observable "state" variables. These beliefs will be correct in equilibrium (i.e. bidders will have rational expectations). One specification would be that the state consists solely of the current object, $s_t = a_t$, and bidders assume that they face a draw from the stationary distribution of types, and that they will bid according to $\beta(x, s_t)$. More complex specifications could allow bidders to be forward-looking (e.g. $s_t = \mathbf{a_t}$), adjusting their bids in response to future supply; and backward-looking, where the bidders make inferences about the population they face based on recent outcomes. Importantly, bidders will take the state transitions as given, even though in fact their actions influence future states.

This kind of assumption was introduced in the macroeconomics literature by Krusell and Smith (1998) as a behavioral assumption, arguing that agents will make inferences based on simple functionals of all the information available in the environment, at least when doing so loses little information. We also provide a "large-market" justification of the assumption in the discussion below. As we argue there, this vastly simplifies the problem while still approximating closely our intuition about the working of these markets.

Formally, we define a *coarsening* T of Λ as a measurable stationary function T from $(\Lambda, \mathcal{I})$ to $(S, \mathcal{P}(S))$, where $\mathcal{I}$ is the Borel σ-field over the public information Λ, S is a finite set and $\mathcal{P}(S)$ is the associated power set. We call elements $s \in S$ the "states". T must have the property that it partitions different current objects into different states, so that minimally the agents condition on the object currently under auction. T creates states with the properties we desire: they are stationary functions of Λ, which includes only a finite public history, and therefore can be conditioned on by everyone, and may recur. Moreover, there are a finite number of these states, which avoids some measure-theoretic complexities here and is important for the identification arguments that follow later in the paper. Then we can give a formal definition of the equilibrium concept:

Definition 1 (**Competitive Markov Equilibrium**). *A competitive Markov equilibrium (CME) with respect to a coarsening T consists of:*

(i) Correct beliefs about the conditional ergodic distribution of the true state ω given s; and the transition matrix Q with elements $Q_{ij} = \mathbb{P}(s' = j | s = i)$ for s' the state tomorrow.

(ii) Symmetric Markovian strategies $\beta(x, s)$ that maximize expected payoffs given beliefs.

Let's unpack this a bit. The CME is extremely similar to a Bayes-Nash equilibrium, requiring that strategies are optimal given beliefs; and that beliefs are correct in the sense of Bayes rule. The key differences are that here bidders condition only on the state s in forming beliefs, even when this is coarser than the public information available to them; and that they do not account for how their actions may influence the state transitions Q. This is the is the "competitive" part of the name, as it corresponds to the case in perfect competition where agents do not recognize that their joint production decisions determine the price.

Here, bidders behave as though they were small, and do not endogenize the impact of their own actions on the future states.

It turns out that under these assumptions, the equilibrium bidding strategy $\beta(x,s)$ has a intuitive and simple form. Fix a CME. By Lemma 1, whatever the bidding strategies, there is an ergodic distribution of the true state, and hence a conditional distribution of types and upcoming objects $\mu_\beta|s$ for any $s \in S$.[10] Moreover, equilibrium play implies a transition rule across states, summarized in a transition matrix Q. So in sum, bidders have rational expectations about the distribution of the true state ω_t both now and in the future, based on the current state s_t and their beliefs about its evolution.

This allows us to define an (ex-post) value function for a type x in state s:

$$v(x,s) = \max_b G_1(b|s)\left(x_a - E[B^1|B^1 < b]\right) + \delta(1 - G_1(b|s))(1-\rho)\sum_{j=1}^{S} v(x,j)Q_{sj} \qquad (1)$$

where $G_1(\cdot|s)$ is the distribution of the highest opposing bid today given the state; x_a is the bidder's valuation of the object currently under auction, and Q_{sj} is the equilibrium probability of a transition from state s to state j where (in an abuse of notation) we allow S to also denote the cardinality of S. The first term in the value function is the probability of winning — the probability that the highest opposing bid is lower — times the surplus conditional on winning, equal to current valuation less expected payment.

Now let $\widetilde{v}(x,s) = \delta(1-\rho)\sum_{j=1}^{S} v(x,j)Q_{sj}$ denote the discounted ex-ante value function (i.e. before exit and state transitions are determined). Fixing the opponent bids and the associated state transitions, we can take an FOC in the bidder problem to get the optimal strategies:

Lemma 2 (**Equilibrium Strategies**). *In equilibrium, bidders bid their valuation less their continuation value:*

$$\beta(x,s) = x_t - \widetilde{v}(x,s) \qquad (2)$$

To get some intuition for this result, think about this process as a single auction, where the winner gets the object, and the losers are awarded a prize with value equal to the continuation value. Re-normalizing the prizes, it's like a standard second-price auction where the winner gets the object less continuation value, and losers get nothing. Then of course the weakly

[10] This is immediate because S is finite.

dominant strategy is to bid the value of the prize, which is just the value of the object less the continuation value. We are now in a position to state an existence result:

Theorem 1 (**Existence**). *For any coarsening T, there exists a CME in symmetric strategies. In the particular case where there is only one product (i.e. $|J| = 1$), the CME is unique.*

The proof is non-trivial. Proving existence requires overcoming the standard problem in auction settings that payoffs are not continuous in actions; as well as additional difficulties created by the dynamics. Our approach is to use the characterization of the bidding strategies in (2) to show the existence of a fixed point of the operator $\Gamma(\beta) = x_t - \widetilde{v}_\beta(x, s)$, where we are now explicit in noting the dependence of the value function on the strategies β. When there is only one product — and hence a unidimensional type space — any two monotone bidding functions produce the same winners and hence the same state transitions. The implication is that the invariant measure over types does not depend on the bidding function, and this helps us to show that the operator Γ is a contraction mapping, thus yielding a unique fixed point via the Banach fixed point theorem.

However, with a multidimensional type space things are more messy, as different strategies may produce different invariant measures. We show that if two strategies are "close", the invariant measures they generate are also "close" in the sense of weak convergence. We also show that if the bidding strategies are smooth with bounded derivatives, the distribution of highest bids has no atoms, therefore avoiding payoff discontinuities. Leveraging these two results, we show continuity of the Γ operator for those strategies, and then we can use the Schauder fixed-point theorem on a carefully defined function space to get the result.

Finally, to conclude this section, we wish to argue that this concept is a sensible large-market approximation to a more standard concept, like a Bayes-Nash Equilibrium (BNE). So consider a modification of this model in which instead of one auction being held in each period, we instead had n auctions of the same good, and n times as many entrants (i.e. the buyer/seller ratio was held fixed). Bidders are randomly assigned to each of the n auctions, as is common in the theory literature.[11] In the limit as $n \to \infty$, the CME assumption that bidders cannot affect the state variable becomes exact. Since history is anonymous, for any coarsening T the actions of a single bidder become irrelevant. Similarly, the distribution

[11] See for example the classic paper of Wolinsky (1988), or more recently Satterthwaite and Shneyerov (2007).

of types in every period is exactly the stationary distribution — by the usual abuse of the law of large numbers — and so there is no reason to learn from the past. Indeed, the state variable, however defined, will not vary with the public history.

The limiting CME will be an equilibrium of the anonymous discounted sequential game, in the language of Jovanovic and Rosenthal (1988). If you believe, as we do, that the anonymous game is a reasonable approximation to behavior in this market, then the CME is attractive. It allows additional flexibility in the strategies by letting bidders condition on features of the recent history. On the other hand, if bidders pay attention to the specific actions of other bidders — as would be the case in a small or concentrated market — it is less sensible.

The key take home of this section is that in equilibrium, bidders shade their bids down from their values, where the extent of shading depends on their continuation value in the current state. Notice immediately that this is starkly different from the "usual" model of second-price sealed bid auctions, where bids may be interpreted as valuations. Indeed, valuations are strictly higher than bids, implying that nonparametric estimates of the value distribution obtained by treating auctions as independent will be systematically biased upward. In the next section we develop nonparametric identification results for large auction markets.

3 Nonparametric Identification

We give a nonparametric identification result in the spirit of Athey and Haile (2002). We think this is useful because it makes explicit the assumptions that are needed to identify the primitives of the dynamic game, as well as providing some guidance as to a sensible estimation strategy.

Theorem 2 (**Identification**). *If δ is known, and the econometrician observes the state, all bids and bidder identities generated by play under a fixed equilibrium, the valuation distribution F is non-parametrically identified. Moreover, the private valuation of any bidder observed bidding in every state $s \in S$ is identified.*

The idea of the proof is as follows. Call an individual bid vector "complete" if it includes a bid for every state $s \in S$. Not all bid vectors will be complete, due to exit, and the set of complete observations is a selected sample. Now suppose we could identify the valuations of those bidders with complete observations; and could also determine the probability that their observation was complete. Then by re-weighting the valuation density for complete bidders by the inverse of that probability, we would get the underlying valuation density.[12]

The key is showing that for complete observations there is an inversion from bids to valuations. Now it turns out that we can express their continuation values as the solution to a linear system of equations based on (1). The linear system has a unique solution, and then from the vector of value functions we can deduce the valuations by re-arranging (2).[13] This is easiest to see in a simple example.

Example of Identification: There are two goods, so $\mathcal{A} = \{a_1, a_2\}$. The exit probability ρ is constant across states. Supply is binomial and independent of state, with q the probability of good 1. Bidders condition on the product identity in the current and next auction, implying four states: $1 = \{1, 1\}$, $2 = \{1, 2\}$, $3 = \{2, 1\}$ and $4 = \{2, 2\}$. Let the bids for a given bidder be $b_1 \cdots b_4$. Then the value function is:

$$v(x, i) = G_1(b_i|i)\,(x_i - \mathbb{E}[B_1|B_1 < b_i, i]) + \delta(1-\rho)(1 - G_1(b_i|i)) \sum_{j=1}^{4} Q_{ij}\; v(x, j)$$

where Q is the transition matrix between states. Substituting out x_i using the bidding function and rearranging gives:

$$v(x, i) - \delta(1-\rho) \sum_{j=1}^{4} T_{ij}\; v(x, j) = G_1(b_i|i)\,(b_i - \mathbb{E}[B_1|B_1 < b_i, i]) \tag{3}$$

Let $V = [v(x, 1), v(x, 2, v(x, 3), v(x, 4)]^T$, and let U be given by:

$$U \equiv \begin{bmatrix} G_1(b_1|1)\,(b_1 - \mathbb{E}[B_1|B_1 < b_1, 1]) \\ G_1(b_2|2)\,(b_2 - \mathbb{E}[B_1|B_1 < b_2, 2]) \\ G_1(b_3|3)\,(b_3 - \mathbb{E}[B_1|B_1 < b_3, 3]) \\ G_1(b_4|4)\,(b_4 - \mathbb{E}[B_1|B_1 < b_4, 4]) \end{bmatrix}$$

i.e. the expected difference between bid and payment in any one period. Then (3) can be represented as:

$$(I - \delta(1-\rho)Q)V = U$$

[12]It is necessary that the probability of a complete observation is non-zero for all types; this is true on the interior of X.

[13]A similar argument was made in a different context by Pesendorfer and Schmidt-Dengler (2003).

where ρ, Q and U can be estimated directly from the data. Then standard results imply $(I - \delta(1-\rho)Q)$ is invertible and therefore the existence of a unique solution for V. Then finally, to identify the valuation for good i, we look at the bid in states when good i was auctioned, and add the ex-ante continuation value to the bid:

$$x_1 = b_1 + \delta(1-\rho)\left(qV(1) + (1-q)V(2)\right) = b_2 + \delta(1-\rho)\left(qV(3) + (1-q)V(4)\right)$$
$$x_2 = b_3 + \delta(1-\rho)\left(qV(1) + (1-q)V(2)\right) = b_4 + \delta(1-\rho)\left(qV(3) + (1-q)V(4)\right)$$

Notice that the valuations are in fact *overidentified* in this case, where $|J| < |S|$. This provides a potential test of the theory.

4 Estimation Strategy

Suppose that the econometrician has a good dataset, observing all bids placed in each auction, the state variable and the identity of all bidders. How should estimation proceed? We propose two different approaches. The first approach is nonparametric, following the logic of the identification section by inverting from observed bids to valuations. We look directly at the individual-level micro-data, treating the record of all bids placed by a given bidder as an observation. The individual-level data may differ in its dimensions: for some bidders, we may only see a single bid, while for others we may see many bids. The structural model implies that at most we should see S distinct bids by any one bidder, a different bid for every state. In the language of the section above, these S length bid vectors are "complete observations".

For complete observations, we can invert from the bid vector to a valuation vector via the first order condition provided we have estimates of the transition matrix and the distribution of opposing bids. This is very much like the approach of Guerre, Perrigne, and Vuong (2000). One important difference is that the set of complete observations is a selected sample of the bidders — bidders with high valuations are more likely to win and exit quickly, and therefore less likely to be observed bidding in every state. For this reason, it is necessary to re-weight the density of the estimated valuations in order to get an estimate of the type density.

The nonparametric approach is very clean and makes no parametric assumptions, but requires a fair number of complete observations. This may be impractical in markets with many states and high turnover in participants. Many bidders on eBay, for example, participate in only one or two auctions before either winning or giving up. We therefore also outline a semiparametric estimation approach based on simulated method of moments, as is used elsewhere for demand estimation in industrial organization and marketing. In that approach, we impose a parametric structure on the distribution of types, and then choose parameters to match moments implied by the structural model with those observed in the data. In both cases, we proceed in two steps. The first step is identical for both approaches; the second varies depending on the method.

4.1 First Step: Transitions, exits and payments

Let $i = 1 \cdots I$ index bidders, and let $t = 1 \cdots T$ index auctions. In the first step, we non-parametrically estimate the probability of winning with a bid of b in state s, $G_1(b|s)$; the expected payment conditional on winning, $\mathbb{E}[B_1|B_1 < b, s]$; the Markov transition matrix Q; the invariant measure over states π; and the probability of exit conditional on losing, $\rho = [\rho_1, \rho_2 \cdots \rho_S]$. This first step can be summarized as estimating elements of the per period payoffs and the transition probabilities, and is similar to that of both Bajari, Benkard, and Levin (2007) and Pakes, Ostrovsky, and Berry (2007) in their papers on dynamic games estimation.

All of these are conditional moments, and provided the conditioning variable is discrete — as the state variable is — we can consistently estimate the conditional moment from the relevant empirical analogue. So for example, to estimate an element of the transition matrix Q_{ij}, we have:

$$\widehat{Q_{ij}} = \frac{\sum_{t=1}^{T} 1(s_{t-1} = i) 1(s_t = j)}{\sum_{t=1}^{T} 1(s_{t-1} = i)}$$

where $t = 1 \cdots T$ indexes auctions and $1(\cdot)$ is an indicator function. The only "difficult" object to estimate is $\mathbb{E}[B_1|B_1 < b, s]$ because for fixed s the conditioning variable b is continuous. This can be done state-by-state using any nonparametric approach, such as kernel density or sieve estimation.

4.2 Second Step: Nonparametric Approach

The key to the non-parametric approach is to treat the data as a sequence of (short) time series, one for each bidder. Drop bidders not observed bidding in every state, so you have a dataset of of S-dimensional bid vectors $i = 1 \cdots I_c$ for I_c the total number of complete observations.

Now for each observation i, use the first-stage estimates to construct a vector $\widehat{u}_i = [\widehat{u}_{i1}, \widehat{u}_{i2} \cdots \widehat{u}_{iS}]$, where $\widehat{u}_{is} = b_{is} - \widehat{\mathbb{E}}[B_1|B_1 < b, s]$. Then, as we show in the proof of Theorem 2, the value function for bidder i, $\widehat{v}_i = [\widehat{v}(x_i, 1), \widehat{v}(x_i, 2) \cdots \widehat{v}(x_i, S)]$, is the solution to the linear system $\widehat{v}_i = (I - \delta(1-\rho)\widehat{Q})^{-1}\widehat{u}_i$. Moreover, we have from (2) that given a J-length sub-vector $\widetilde{b}_i$ of b_i consisting of bids on different objects, an associated sub-vector $\widetilde{\rho}$ of ρ, and an associated $J \times S$ submatrix $\widetilde{Q}$ of Q consisting of the transitions associated with the bids in $\widetilde{b}_i$, we get $x_i = \widetilde{b}_i + \delta(1-\widetilde{\rho})\widetilde{Q}v_i$. Substituting in our estimates on the right hand side of this expression, we get an estimate $\widehat{x_i}$ of the valuation of each bidder.

The set of bidders with complete observations is a selected sample, and we need to correct for this. In the appendix, we show how given previously estimated objects, we can derive an expression for $P(A, x)$, the probability that a type x is observed bidding in the set of states in A, for $A \subseteq S$. The probability that a given type generates a complete observation is $P(S, x)$. We correct for the selection bias by assigning a weight equal to $1/\widehat{P}(S, \widehat{x_i})$ to each $\widehat{x}_i$, and then use weighted kernel density estimation to back out the type density $f(x)$.

We omit an analysis of the asymptotic properties of this estimator, because both it takes us into the realm of non-parametrics with dependent data and because we suspect that the semiparametric approach outlined below is more likely to be used in practice. We do show, however, that it performs extremely well in small samples in our Monte Carlo experiments, below.

4.3 Second Step: Semiparametric Approach:

The semiparametric approach proceeds in the opposite direction. Instead of inverting bids to valuations, we take draws from a parametrized type distribution, simulate bids, and match the moments of the simulated bid distribution with those observed in the data.

The first step is working out how to simulate bids for a given type. Our idea is to solve for the optimal biding function in this environment. To do this, look at the value function representation in (1). Using our first-stage estimates of $G_1(b|s)$, $\mathbb{E}[B_1|B_1 < b, s]$, $\rho(s)$ and Q, for any type x you can determine the value to following any strategy $\beta(x, s)$. Better yet, it is possible to solve for the optimal bidding strategy by value function iteration on the Bellman equation. This is computationally cheap since the iteration process obeys a contraction mapping.

Additionally, one can solve for the probability that these bids appear in any subset of states A.[14] This is important because we know certain types will be overrepresented in certain states: the selection problem. Given these tools, and a parametric model for the type distribution determined by a finite dimensional parameter θ, one can simulate the bid distribution for any state. Then any minimum distance estimator that chooses θ to minimize the distance between thoughtfully chosen functionals of the observed and actual bid distributions will converge to the true θ_0, provided it is identified.

One such estimation approach would be GMM. In that case we treat the data as a time series of multidimensional observations $t = 1 \cdots T$, and match sample moments with simulated moments. For example, we could construct a moment by looking at the mean bid in auctions held in state 1. Another moment would be the covariance between successive bids by the same bidder across states. With this data, our "covariates" are bidder identity and the state, so it makes sense to construct moments that either condition on or interact with these variables. For this approach, the asymptotic theory of Hansen (1982) applies, showing that provided the true parameter is identified and the environment is strictly stationary, the GMM estimation approach will recover the truth asymptotically. Note that Lemma 1 proves the stationarity of the environment, so that is satisfied. But we need also appeal to the results of Pakes and Pollard (1989) to argue that the simulation error has no impact on the asymptotic distribution of the estimator as the number of simulations grows large; and to either Andrews (1994) or Ai and Chen (2003) to argue that the non-parametric first stage does not stop us from getting $\sqrt{N}$ consistency.

From a computational point of view, there are ways to speed up the estimation. One important bottleneck is that for every new parameter update and sample of bidders, we must solve a dynamic optimization problem for a large sample of simulated types. We can speed this up using an importance sampling approach. Instead of drawing the types anew on each iteration and solving out for their bidding strategies, instead choose a set of R types initially to uniformly span some plausible region of X and compute their optimal bids. The choice of initial region is up to the researcher: one suggestion might be to regress prices on states, and then take the region of types spanned by the coefficient estimate plus four standard deviations on either side. This need only be done once. Then, to compute the simulated

[14] See the appendix for a formal derivation of $P(A, x_i)$.

moments, we weight the types according to their relative likelihood under θ and compute the simulated moments as weighted sums.

4.4 Characteristic Space Approach

At the end of the day, we are trying to estimate the distribution of valuations over different products. As in the more general demand literature, this can be overly demanding of the data if the product space is large. Even after imposing a multivariate normal parametric structure, for example, we need to estimate a variance covariance matrix with $J(J+1)/2$ parameters. Given this, we may want to project valuations down onto product characteristics (e.g. McFadden (1974)).

To do this, we assume that valuations depend on the characteristics of the goods z_t, as well as on an unobserved heterogeneity term and an idiosyncratic component:

$$x_{it} = z_t \beta_i + \xi_t + \epsilon_{it} \tag{4}$$

where we index individuals by i and auctions by t as before. The ξ_i term captures auction specific heterogeneity — something observable to all the bidders, but not to the econometrician. The ϵ_{it} terms capture idiosyncratic individual preferences for the particular object being auctioned, while the β_i terms are the individual's type. This is pretty much a standard random coefficients demand specification. Bidders differ in their tastes for the characteristics, and we would like to recover the distribution of the random coefficients. The characteristic space X has dimension $|X|$, where we assume $|X| \leq |J| \leq |S|$.

An implication of the characteristic-based approach is that the payoff-relevant state variables may change. For example, if consumers of digital cameras are assumed to have preferences only over the camera resolution, then if two different cameras may have the same mega-pixels, states in which those different cameras are up for auction may be pooled.

Since presumably this is a case where there are few complete observations, we take the parametric approach. The first stage is exactly as before: estimate transitions, exit and entry and per period payoffs just as before. In the second stage, the parametric specification must be of a (joint) distribution for the β_i (although, as often done elsewhere, the β's could be assumed independent), as well as a distribution for the unobserved heterogeneity ξ_t and idiosyncratic errors ϵ_{it}. Given this specification, one can simulate types and errors, which imply valuations x_{it} from (4), and through the first-stage estimates, bids. This is in the spirit of the simulated GMM approach in Berry, Levinsohn, and Pakes (1995).

The idiosyncratic terms do necessitate adding additional moments, however. In order to identify the ϵ_{it} terms, one should include moments based on the variance of bids by the same bidder in the same state (since only ϵ_{it}) varies in this case. To get the ξ_t terms, one could look at the variance of bids within auctions.

5 Monte Carlo

We perform Monte Carlo exercises to test both our nonparametric and parametric estimation approaches in small samples. In the simple case there are two goods, and bidders condition their bids only on the identity of the good and their private information, so there are two corresponding public states. In each period, each of these goods is equally likely to be listed. The number of entrants k is always 3 each period, but exit is random with losing bidders exiting with probability $\rho = 0.25$ and winning bidders exiting with certainty. The discount rate δ is set to 0.99. We consider alternative parameterizations for k and ρ. Bidders' private valuations are distributed bivariate normal with mean μ and covariance matrix Σ, which we also allow to vary in different Monte Carlo experiments.

Data is generated by first solving for the bidding function via policy iteration – though we have only been able to prove this converges in general for $|J| = 1$ (see Theorem 1), it has converged in all of our experiments to date. Then for each Monte Carlo iteration we simulate a dataset of 500 auctions, after discounting an initial 10,000 auctions as a "burn in". This amounts to, in expectation, 250 auctions per product, which seems like a moderate amount of data, especially given the volume of transactions in many online auction markets.

We run both estimation routines assuming that the econometrician knows the entry process, the discount rate δ and that the transitions between states are random rather than Markov. In the common first stage, we estimate the probability of exit ρ, and the probability of good one being listed, q. Then in the nonparametric second stage approach, we estimate the marginals of the type density using a Gaussian kernel density estimation approach with automatic bandwidth choice by cross-validation. In the parametric estimation, we (correctly) specify a multivariate normal distribution for the types, and take a minimum distance (MD) approach by matching the mean and variance of bids for each state in which a bidder may be

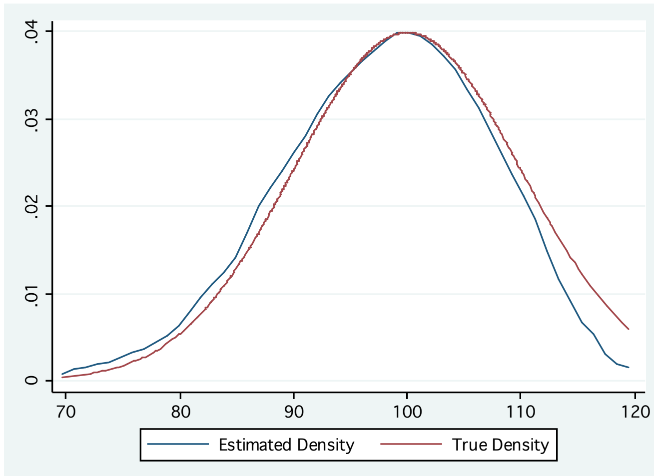

Figure 1: **Monte Carlo Simulation** The figure shows the true and estimated marginal density of valuations for product 1 for a randomly chosen Monte Carlo simulation of 500 auctions with specification E.

observed, as well as the covariance between bids in each state for bidders observed in both.

Our results for the simple case are presented in Table 1 in the appendix for a variety of parameterizations. Experiment A is a baseline case with $k = 3$, $\rho = 1/4$, and symmetric μ and Σ. Experiment B slowed the rate of entry and random exit, allowing for fewer, but longer-lived bidders. Experiments C and D add positive and negative correlation between valuations, respectively, and finally experiment E adds positive correlation as well as significant asymmetry in the means and variance terms.

Estimates from the structural models proposed in this paper appear in panel 1 for each experiment. Panel 2 presents, for comparison, estimates from a naive approach that treats each observed bid as a draw from the distribution of private valuations.

The estimates of p and ρ are straightforward and precise. Estimates of μ and Σ come from the parametric approach, and are found to be accurate and encouragingly precise for all parameterizations of the model. The naive approach gets both the means and the variances wrong in a statistically significant way. This bias derives from three sources that we have dealt with in our model: repeat bidding, selection by type, and bid shading according to the option value of losing.

In the far-right hand columns appear estimates of the mean integrated squared error[15] (MISE), which is a measure of goodness-of-fit for our nonparametric approach. We report

[15] For a candidate distribution $f_n(x)$ and a true distribution $f(x)$ this is calculated according to $MISE = \int (f_n(x) - f(x))^2 dF(x)$

MISE for the marginal distributions of types in both dimensions. In comparing these numbers, note that MISE is not a normalized measure, and is therefore not comparable across specifications. We can, however, compare the MISE generated by the structural approach with the MISE generated by the naive approach, and we see substantial improvement, with the structural model often doing a full order of magnitude better. For a graphical intuition of the fit achieved, Figure 1 depicts the true and the estimated marginal distribution of types for good 1 from experiment E of the simple model.

We also run a parallel set of Monte Carlo experiments for a forward-looking model, which is based on the identification example presented earlier in the paper. The simple model is augmented by bidders' awareness of the good being auctioned in the next period, which means we now have four public states. Results are presented in Table 2. While the more complicated structure of the model has a cost in precision, both estimators continue to dramatically outperform the naive approach under all specifications. Note that the cost of precision is much sharper for the nonparametric approach– this is because, with four states instead of two, the sample of bidders for which we observe bids in every state is much smaller.

6 Conclusion

We have developed a demand system for a large auction market. By developing a tractable theoretical framework that seems suitable for large auction markets, we have been able to provide a structural model in which the distribution of bidders is stationary and their valuations are identified non-parametrically from bids.

From the Monte Carlo exercises, we learn that accounting for dynamics and allowing for multiple products turns out to be important in a number of ways. First, surplus measures would be systematically biased downward without accounting for dynamics, because bidders shade their bid below their valuations to account for surplus value. Second, the dynamic process generates selection on types, and without a model it would be impossible to correct for this. It is easy to oversample low types and again get biased estimates of valuations.

Hopefully understanding these basic empirical problems will be helpful for future research. We hope also that this is a first step towards satisfying models for other markets: those where bidders have market power, or non-unit demand. We would certainly like to take this model to data and see how it performs.

References

Adams, C. P. (2009): "Identification of the Joint Value Distribution from Auction Data," Working paper, FTC.

Ai, C., and X. Chen (2003): "Efficient Estimation of Models with Conditional Moment Restrictions Containing Unknown Functions," *Econometrica*, 71(6), 1795–1843.

Andrews, D. W. K. (1994): "Asymptotics for Semiparametric Econometric Models Via Stochastic Equicontinuity," *Econometrica*, 62(1), 43–72.

Athey, S., and P. A. Haile (2002): "Identification of Standard Auction Models," *Econometrica*, 70(6), 2107–2140.

Bajari, P., L. Benkard, and J. Levin (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75(5), 1331–1370.

Bajari, P., and A. Hortacsu (2004): "Economic Insights from Internet Auctions," *Journal of Economic Literature*, 42(2), 457–486.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63(4), 841–890.

——— (2004): "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market," *The Journal of Political Economy*, pp. 68—105.

Budish, E. (2008): "Sequencing and Information Revelation in Auctions for Imperfect Substitutes: Understanding eBay's Market Design," Working paper, Harvard University.

Fershtman, C., and A. Pakes (2009): "Finite State Dynamic Games with Asymmetric Information: A Framework for Applied Work," Working Paper, Harvard University.

Guerre, E., I. Perrigne, and Q. Vuong (2000): "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica*, 68(3), 525–574.

Hansen, L. P. (1982): "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1029—1054.

Hendricks, K., and R. Porter (2007): "A Survey of Empirical Work in Auctions," in *Handbook of Industrial Organization*, ed. by M. Armstrong, and R. Porter, vol. III. Vickers.

Ingster, A. (2009): "A Structural Study of a Dynamic Auction Market," Working paper, NYU Stern.

Isaac, R. (1963): "A General Version of Doeblin's Condition," *Ann. Math. Statist.*, 34(2), 668–671.

Jofre-Benet, M., and M. Pesendorfer (2003): "Estimation of a Dynamic Auction Game," *Econometrica*, 71(5), 1443–1489.

Jovanovic, B., and R. W. Rosenthal (1988): "Anonymous Sequential Games," *Journal of Mathematical Economics*, 17(1), 77—87.

Karr, A. F. (1975): "Weak Convergence of a Sequnce of Markov Chains," *Probability Theory and Related Fields*, 33(1).

Krusell, P., and A. A. Smith (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106(5), 867—896.

McFadden, D. (1974): "Conditional logit analysis of qualitative choice behavior," in *Frontiers in Econometrics*, ed. by P. Zarembka, pp. 105—142. Academic Press.

Pakes, A., M. Ostrovsky, and S. Berry (2007): "Simple Estimators for the Parameters of Discrete Dynamic Games, with Entry/Exit Examples," *Rand Journal of Economics*, 38(2), 373—399.

Pakes, A., and D. Pollard (1989): "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*, 57(5), 1027–1057.

Pesendorfer, M., and P. Schmidt-Dengler (2003): "Identification and Estimation of Dynamic Games," Working Paper, LSE.

Rust, J. (1994): "Structural Estimation of Markov Decision Processes," in *Handbook of Econometrics*, ed. by R. Engle, and D.L.McFadden, vol. IV. Elsevier Science.

Sailer, K. (2006): "Searching the eBay Marketplace," Working Paper, University of Munich.

Satterthwaite, M., and A. Shneyerov (2007): "Dynamic Matching, Two-Sided Incomplete Informationa and Participation Costs: Existence and Convergence to Perfect Competition," *Econometrica*, 75(1), 155–200.

Weintraub, G., C. L. Benkard, and B. V. Roy (2008): "Markov Perfect Industry Dynamics with Many Firms," *Econometrica*, 76(6), 1375–1411.

Wolinsky, A. (1988): "Dynamic Markets with Competitive Bidding," *The Review of Economic Studies*, 55(1), 71—84.

Zeithammer, R. (2006): "Forward-looking bidding in online auctions," *Journal of Marketing Research*, 43(3), 462–476.

Appendix

Proof of Lemma 1: Fix a bidding strategy β. From Isaac (1963), it will suffice to show that a generalization of the Doeblin condition holds: there exists a probability measure m such that $\forall \delta > 0$, $\forall A \in \mathcal{F},\ m(A) > \delta \Rightarrow m(l(A)) > 0$, where $l[A] = \{\omega : lim\ inf_k P^k(\omega, A) > 0\}$ and $P^k(\omega, A)$ is the probability of reaching A from ω in k steps. In turn, it suffices to show that there exists $\bar{k}$ such that $P^{\bar{k}}(\omega, A) > 0\ \forall \omega \in \Omega, \forall A \in \mathcal{F}$; since if this holds, then $lim\ inf_k P^k(\omega, A) > 0$ is bounded below by $inf_{\omega \in \Omega, A \in \mathcal{F}} P^{\bar{k}}(\omega, A) > 0$, implying $l[A] = \Omega$ and $m(l[A]) = 1$ for any probability measure m.

So fix any $\omega \in \Omega$ and $A \in \mathcal{F}$. Consider the following sequence of transitions. First, all the bidders simultaneously exit and exactly two bidders in A enter. In each successive period, no-one but the winner exits (where the winner is determined by β); exactly two bidders in A enter; and upcoming supply is added according to A. Each of these transitions has positive probability since we have assumed $P(E_t \geq 2 | N_{t-1} < \bar{N} - 2$, the density of types and goods is positive everywhere and $0 < \rho < 1$. It is clear that this process will reach A in a finite number of periods $\bar{k}$ for any ω and A, as required. □

Proof of Lemma 2: The expected payoff function $\pi(b, s)$ has the form:

$$\pi(b, s) = \int_0^b (x_a - B^1) f(B^1|s) dB^1 + \left(\int_b^\infty f(B^1|s) dB^1 \right) (1 - \rho) \delta \sum_{j=1}^{S} v(x, j) Q_{sj}$$

Writing $\widetilde{v}(x, s)$ for $(1-\rho)\delta \sum_{j=1}^{S} v(x, j) Q_{sj}$, and applying Leibnitz's rule in taking an FOC in b, we get $(x_a - b) f(B^1|s) - f(B^1|s)\widetilde{v}(x, s) = 0$, which has unique solution $b = x_a - \widetilde{v}(x, s)$. □

Proof of Theorem 1: Note: for this proof, we use $v(x, s)$ instead of $\widetilde{v}(x, s)$ for the ex-ante value function. Before the main proof, we state a useful lemma:

Lemma 3. *Let $v_\beta(x, s)$ be the ex-post value function when agents use continuous strategies $\beta(x, s) = x_a v_\beta(x, s)$. If $\frac{\partial v_\beta(x,s)}{\partial x_j} < \delta$ for all (x, s), then the stationary distribution of highest bids has continuous density in every state.*

Proof: Towards a contradiction, suppose that the stationary distribution of highest bids is not continuous in some state. Then either the invariant type density is continuous but the bid function yields discontinuities; or the invariant type density is not continuous. By assumption $\beta(x, s) = x_a - v_\beta(x, s)$ and so $\partial \beta(x, s) \partial x_j = -\frac{\partial (v_\beta(x,s)}{\partial x_j} < 0$ if $a \neq j$ and $1 - \frac{\partial (v_\beta(x,s)}{\partial x_j} > 0$, since $\frac{\partial(v_\beta(x,s)}{\partial x_j} < \delta < 1$. So the bidding function is never constant in type, in any state, which implies there are no discontinuities when the type density is continuous.

So there must be a discontinuity in the invariant type density. Now all Borel sets in the type-space receive positive measure, from Lemma 1. Moreover, there cannot be any atoms, since the entry process places zero measure on a single point and exit from that point is almost sure, so no point can be recurrent. This implies that in any state, the probability of winning $G_1(b|s)$ is continuous and increasing in b. But this implies the probability of exit by winning changes continuously in type. Combined with the continuous entry density and random loser exit, we have that for any sequence of types $x_n \to x$ the transition kernels K under β converge. This implies that the invariant density over those types converges, yielding a contradiction. □

The lemma gives us a way to work around the standard problem that in auction games, the payoff function is discontinuous in actions. By considering only strategies that generate value functions with bounded derivatives (as any equilibrium strategies must do), we obtain the property that the highest bid density is always continuous, implying smooth payoffs in type when agents employ these strategies.

For an equilibrium, we require (a) that equilibrium strategies maximize payoffs, and thus take the form $\beta(x,s) = x_s - v(x,s)$ (Lemma 2), and that (b) bidder's calculations of the continuation value are consistent with the long-run distribution of play conditional on states. Taken together, an equilibrium requires a fixed point of the mapping:

$$\Gamma(v(x,s)) = G_1(\beta(x,s)|s)(x_a - \mathbb{E}[B_1|B_1 \leq \beta(x,s),s]) + (1 - G_1(\beta(x,s)|s))\delta \sum_{j\in S} Q_{st} v(x,j)$$

where $\beta(x,s) = x_s - v(x,s)$, and the distribution $G_1|s$ of B_1 is the distribution of highest bids in state s when the type measure is μ^x_β, the invariant measure when the strategies are β. Such an invariant measure exists by Lemma 1.

Based on Lemma 3, consider the space V of C^1 functions on $[0,\bar{x}]^J \times S$ with range on $[0,\bar{x}]$ and first partial derivatives in the first J arguments everywhere bounded on $[0,\delta]$ (i.e. $\frac{\partial v(x,s)}{\partial x_j} \in [0,\delta]\ \forall x\ \forall s\ , j = 1\cdots J$). Under the norm $\|v\| = \sup_{x\in\mathcal{X}} \max_{s\in S} \left[v(x,s) + \sum_{j=1}^{J} \frac{\partial v(x,s)}{\partial x_j}\right]$, V is a Banach space. To see this, notice first that is a normed vector space. Also, the space of C^1 functions on a bounded set is complete. So given a Cauchy sequence of functions on V, the limit point is C^1. Under the norm, the first partials must converge and since all members of the sequence have first partials in the closed set $[0, \delta]$, so does the limit point. The limit point is thus in the function-space, proving Vcomplete.

Now, we would like to apply the Schauder fixed point theorem to the operator Γ. Since V is compact and convex, it will suffice to show (i) that $\Gamma(V) \subseteq V$ and (ii) Γ is continuous. For (i), notice that for any $v \in V$, $\Gamma(v)$ is the value function induced by playing the β associated with v. Therefore increasing the valuation of any type along a single dimension must (weakly) raise the value function, but not by more that δ, since the additional payoff is realized at best next period. Also, to show that $\Gamma(v)$ is C^1, consider the partial derivatives:

$$\begin{aligned}\frac{\partial \Gamma(v(x,s))}{\partial x_j} &= g_1(\beta(x,s)|s)\frac{\partial \beta(x,s)}{\partial x_j}(x_a - \mathbb{E}[B_1|B_1 \leq \beta(x,s), s]) \\ &+ G_1(\beta(x,s)|s)\left(\frac{\partial x_a}{\partial x_j} + \frac{\partial \mathbb{E}[B_1|B_1 \leq \beta(x,s), s]}{\partial x_j}\right) \\ &- g_1(\beta(x,s)|s))\delta \sum_{j \in S} Q_{st} v(x,j) + (1 - G_1(\beta(x,s)|s))\delta \sum_{j \in S} Q_{st} \frac{\partial v(x,j)}{\partial x_j}\end{aligned}$$

We need to show this is continuous in x. It will suffice to show that the stationary distribution of highest bids in any state has continuous density, since the continuous differentiability of v implies the same property for β, and all the other terms are either constant or will be continuous if the density is continuous. But this is immediately true by Lemma 3. This suffices to show that $\Gamma(V) \subseteq V$.

Next, we need to show Γ continuous. Since sequential continuity coincides with continuity on Banach spaces, it will suffice to show $v_n \to v$ implies $\Gamma(v_n) \to \Gamma(v)$. For each (v, v_n) write (β, β_n) for the strategies implied by these continuation values according to $\beta(x,s) = x_s - v(x,s)$. Since the value functions are geometric series and are one-to-one with the strategies, one can alternately prove $U(\beta_n) \to U(\beta)$ for $U(\beta(x,s)) = G_1^{\beta}(\beta(x,s)|s)(x_a - E_\beta[B^1|B^1 < \beta(x,s), s])$. Let $\mu^n|s$ be the conditional invariant measure over types generated when strategies are β_n and the state is s, and similarly define $\mu|s$ for β.

Let $U^n(\beta(x,s)) = G_1^{\mu_n|s}(\beta(x,s))(x_a - E_{\mu_n|s}[B^1|B^1 < \beta(x,s)])$ (i.e. the beliefs are consistent with the invariant measure induced by strategy β^n, but β is actually played). Then by the triangle inequality, $\|U(\beta^n) - U(\beta)\| \leq \|U^n(\beta^n) - U^n(\beta)\| + \|U^n(\beta) - U(\beta)\|$. It will suffice to show each of the RHS terms converges. The first term converges by arguments similar to those already made: G and g are continuous in b, as is the expected payment, and so if the strategies converge, so do the levels and slopes of the one-period payoffs.

Next, for the second term, we need to show that $\mu_n|s$ converges weakly to $\mu|s$ for any $s \in S$. It suffices to show μ_n converges weakly to μ. Let us denote the Markov transition kernel on types generated by β_n in state s by K_n^s. Let h be any uniformly continuous function on $X^N \equiv [0, \bar{x}]^{\bar{N}}$. Then Karr (1975) gives a set of three conditions for convergence: (i) $K^s h$ is continuous for all s; (ii) $K_n^s h \to K^s h$ uniformly on every compact subset of X^N for all s; (iii) the family of measures $\{K_n^s(x, \cdot)\}$ is tight $\forall n \geq 1$ $\forall x \in X^N$ $\forall s \in S$. The last condition is immediate, since X^N is compact. For (i), notice that if $x, x' \in X^N$ are "close enough", then β will assign the same winner in both cases. Thus $K^s h(x')$ and $K^s h(x)$ can differ only at the mass-points assigned to losing bidders, and the uniform continuity of g implies that we can choose x, x' "close-enough" so the values at these mass-points differ by no more than ϵ, for any $\epsilon > 0$.

It remains to show property (ii). Fix a state s, and take any compact set A of X^n, and on that set define $\Delta = \min_{x \in A} B_\beta^1(x, s) - B_\beta^2(x, s)$, the smallest difference between the highest and second highest bids when the strategies are β. This is well-defined because A is compact and the strategies are bounded. Now, if we can show that for all $x \in A$ the transition kernels K_n^s and K^s are the same, we are done. Notice that if K_n^s and K^s imply the same winner, the transition kernels will be identical. So choose a N_ϵ such that $\beta_n(x, s)$ and $\beta(x, s)$ differ by no more than $\Delta/2$. Then $B_\beta^1(x, s) - \Delta \geq B_\beta^2(x, s) \Rightarrow B_{\beta_n}^1(x, s) + \Delta/2 > B_{\beta_n}^2(x, s) + \Delta/2$, which implies that they predict the same winner on the set A. Property (ii), weak convergence of the invariant distributions and the theorem follow.

The case where there is only one product ($|J| = 1$) is much simpler. Restrict attention to strategies $\beta(x, s)$ that are increasing in x for all $s \in S$. Then any two strategies β_0 and β_1 predict the same winner in every auction, and hence result in the same invariant measure. So take the operator $\Gamma(\beta) = x_a - v(x, s)$ defined on the space of continuous bounded increasing functions. It is easy to show Γ is a contraction mapping, and then the Banach fixed point theorem gives the existence of a unique fixed point. □

Proof of Theorem 2: Following precisely the logic of the identification example, we can write the vector of continuation values of a bidder who bids b as the solution to a linear system $V = (I - \delta(1-\rho)Q)^{-1}u$, where Q is the transition matrix, $\rho = \rho_1, \rho_2 \cdots \rho_S$ is the vector of exit probabilities and $u = [u_1, u_2 \cdots u_S]$ is the S-length vector with terms of the form $u_s = G_1(b_s|s)(b_s - \mathbb{E}[B_1|B_1 < b_s, s])$. Q and ρ are identified from the data, δ is assumed known and u is identified for any bidder observed bidding in every state. From Rust (1994), $(I - \delta Q)^{-1}$ exists for $\delta < 1$. Given these objects, the ex-ante continuation value and therefore the valuation x is identified for all complete observations. Since every bidder on the interior of $[0, \bar{\omega}]^J$ loses with positive probability, all types in X are observed bidding in every state. Thus the density of valuations for complete observations $g(x)$ has full support on X. Moreover,

$$f(x) = \frac{g(x)}{P(S,x) \int \frac{g(x)}{P(s,x} dx}$$

where $P(S, x)$ is identified from the data and defined in (6) below. □

6.1 Additional Estimation Details

Let A be a subset of S. We want to the probability that any type x ends up submitting bids in the states in A. Define $p(x, s) = G_1(\beta(x, s)|s) + (1 - G_1(\beta(x, s)|s))\, \rho(s)$, which is just the probability that a type x will exit the sample in state s, whether by winning or losing. Also define $P(B, x, s)$ to be the probability of a bidder x who enters the sample in state s being observed bidding only in states $B \subseteq S$. We can express this recursively:

$$P(B, x, s) = 1(s \in B) \left[p(x, s) + (1 - p(x, s)) \sum_{s' \in B} Q_{ss'} P(B, x, s') \right] \tag{5}$$

where Q is the Markov transition matrix (recall, states are assumed finite). Then the probability of observing a bidder x in group A can be defined implicitly as:

$$P(A, x) = \sum_{s \in A} \pi(s) P(A, x, s) - \sum_{B \subset A} P(B, x) \tag{6}$$

where π is the invariant measure over states. The idea is simply that the probability of seeing bids for every state s in A is equal to the probability that the bidder stays within A less the probability that he stays in a strict subset of A.

6.2 Monte Carlo Results

Table 1: Simple Model

		p	ρ	μ_1	μ_2	σ_1^2	σ_2^2	σ_{12}	MISE1	MISE2
A	true	0.50	0.25	100	100	100	100	0		
1	est	0.50	0.25	99.95	99.94	99.42	99.57	-0.33	4.79e-5	4.89e-5
	std	0.02	0.01	0.28	0.24	4.96	4.50	4.77	1.32e-5	1.24e-5
2	est			97.08	97.09	72.01	71.80		1.07e-3	1.06e-3
	std			0.38	0.32	4.16	4.03		3.92e-4	4.04e-4
B	true	0.50	0.125	100	100	100	100	0		
1	est	0.50	0.127	99.96	99.96	99.47	99.43	-1.06	9.55e-5	9.94e-5
	std	0.022	0.05	0.32	0.33	6.17	7.08	6.02	1.52e-5	1.75e-5
2	est			94.94	94.94	59.99	59.65		2.90e-3	2.93e-3
	std			0.46	0.43	4.30	4.74		4.20e-4	3.68e-4
C	true	0.50	0.25	100	100	100	100	50		
1	est	0.50	0.25	99.99	100.01	100.70	100.79	50.91	8.23e-5	8.10e-5
	std	0.02	0.01	0.26	0.24	5.26	3.85	4.59	1.67e-5	1.38e-5
2	est			96.46	96.48	66.34	65.63		1.05e-3	1.18e-3
	std			0.38	0.29	4.16	3.19		5.07e-4	4.00e-4
D	true	0.50	0.25	100	100	100	100	-50		
1	est	0.50	0.25	100.02	99.99	100.93	100.91	-49.78	3.37e-5	3.40e-5
	std	0.02	0.01	0.33	0.32	5.50	5.76	5.65	1.02e-5	1.16e-5
2	est			97.88	97.88	75.52	75.23		9.14e-4	9.39e-4
	std			0.42	0.39	3.92	4.45		4.64e-4	4.24e-4
E	true	0.50	0.25	100	150	100	400	100		
1	est	0.50	0.25	99.93	149.92	98.44	398.51	96.88	1.03e-4	1.63e-5
	std	0.02	0.01	0.27	0.41	4.25	18.94	9.14	1.59e-5	3.99e-6
2	est			95.78	143.83	61.73	277.45		1.24e-3	2.56e-4
	std			0.32	0.54	3.22	16.36		4.41e-4	1.23e-4

For each experiment A-E, panel 1 depicts structural estimates and panel 2 the naive estimates from treating bids as iid observations. Simulated datasets have 500 auctions, and 100 simulations are run.

Experiment A: baseline case, with $k = 3$, $\rho = 1/4$, and symmetric μ and Σ.

Experiment B: longer-lived bidders, with $k = 2$ and $\rho = 1/8$.

Experiment C: positive correlation in valuations, with $\sigma_{12} = 50$.

Experiment D: negative correlation in valuations, with $\sigma_{12} = -50$.

Experiment E: Asymmetric μ and Σ, as well as positive correlation in valuations.

Table 2: Forward-Looking Model

		p	ρ	μ_1	μ_2	σ_1^2	σ_2^2	σ_{12}	MISE1	MISE2
A	true	0.50	0.25	100	100	100	100	0		
1	est	0.50	0.25	99.92	99.90	94.59	94.90	-1.27	2.30e-4	2.26e-4
	std	0.02	0.01	0.31	0.30	6.69	6.47	5.18	4.60e-5	5.51e-5
2	est			97.17	97.15	72.80	72.25		6.63e-4	6.38e-4
	std			0.40	0.34	4.55	3.86		3.49e-4	3.60e-4
B	true	0.50	0.125	100	100	100	100	0		
1	est	0.50	0.125	99.90	99.97	97.67	98.22	-3.34	4.16e-4	4.18e-4
	std	0.02	0.01	0.37	0.45	9.64	8.68	5.85	6.24e-5	6.56e-5
2	est			94.87	94.90	60.82	58.95		1.98e-3	1.92e-3
	std			0.42	0.45	4.75	4.60		2.43e-4	2.60e-4
C	true	0.50	0.25	100	100	100	100	50		
1	est	0.50	0.25	99.84	99.88	94.62	92.94	46.01	3.03e-4	2.94e-4
	std	0.02	0.01	0.30	0.26	6.59	5.66	5.81	5.91e-5	6.08e-5
2	est			96.45	96.50	65.72	65.77		9.23e-4	8.26e-4
	std			0.28	0.33	3.53	3.69		3.11e-4	3.51e-4
D	true	0.50	0.25	100	100	100	100	-50		
1	est	0.50	0.25	99.91	100.20	99.77	97.14	-45.21	2.32e-4	2.33e-4
	std	0.02	0.01	0.41	0.36	7.38	7.97	6.38	5.57e-5	5.94e-5
2	est			97.84	97.95	76.04	74.97		5.44e-4	5.74e-4
	std			0.39	0.41	3.81	4.59		3.23e-4	2.91e-4
E	true	0.50	0.25	100	150	100	400	100		
1	est	0.50	0.25	99.79	149.78	92.20	372.42	88.56	3.46e-4	6.73e-5
	std	0.02	0.01	0.31	0.62	6.23	26.46	10.48	6.50e-5	1.43e-5
2	est			95.78	143.85	62.90	280.80		8.00e-4	2.03e-4
	std			0.29	0.71	3.64	15.64		4.02e-4	7.92e-5

For each experiment A-E, panel 1 depicts structural estimates and panel 2 the naive estimates from treating bids as iid observations. Simulated datasets have 500 auctions, and 100 simulations are run.

Experiment A: baseline case, with $k = 3$, $\rho = 1/4$, and symmetric μ and Σ.

Experiment B: longer-lived bidders, with $k = 2$ and $\rho = 1/8$.

Experiment C: positive correlation in valuations, with $\sigma_{12} = 50$.

Experiment D: negative correlation in valuations, with $\sigma_{12} = -50$.

Experiment E: Asymmetric μ and Σ, as well as positive correlation in valuations.