

A Demand System for a Dynamic Auction Market with Directed Search*

Matthew Backus

Gregory Lewis

Cornell University and eBay Research

Harvard University and NBER

October 8, 2012

Abstract

We develop a demand system for a dynamic auction market with directed search. In each period, heterogeneous goods are exogenously supplied and sold by second-price auction, and incumbent bidders choose which good to bid on and how much to bid. Bidder valuations are multidimensional, private and perfectly persistent, and the population of bidders evolves according to an exogenous entry and endogenous exit process. We prove that the state of the market — which includes active bidders' types and information sets — evolves as a geometrically ergodic Markov process. We characterize best responses as solutions to a partially observed Markov decision problem and provide conditions under which the econometrician can identify equilibrium strategies from time series data. We provide additional conditions under which this allows nonparametric identification of preferences. When the market is large so that each bidder's actions are informationally small, we show bidderwise identification: the valuations of bidders whose individual time series includes a bid on every product are identified. Two-stage nonparametric and semiparametric estimation procedures are proposed, and shown to work well in Monte Carlo and counterfactual simulations.

*This revision is a substantial departure from our earlier working paper, motivated by many thoughtful comments. We are grateful to Dan Akerberg, Chris Adams, Susan Athey, Lanier Benkard, Steve Berry, Gary Chamberlain, Matt Gentzkow, Phil Haile, Ken Hendricks, Hugo Hopenhayn, Ariel Pakes, Maher Said and Robert Zeithammer for their input, and to conference and seminar participants at the FTC Applied Micro, Haas, Harvard/MIT, IOOC, Johns Hopkins, LSE, Michigan, NBER-CEME, Stanford/SITE/SIEPR, Toulouse, UBC, UCLA, UCSC, UCSD and Yale. Danyang Su provided excellent research assistance. All remaining errors are our own.

1 Introduction

In many markets, goods are traded by an auction mechanism. Important examples include financial assets, housing, procurement contracts and used goods. The goods being sold at any point in time are typically heterogeneous, and buyers have heterogeneous preferences. As a result, these markets exhibit directed search: families look for homes with enough bedrooms for their children; construction firms look for projects for which they have the relevant experience and capacity; and car dealerships attend car auctions looking to fill specific gaps in their inventory.

Buyers also often get multiple purchase opportunities over time. This is most obvious in the case of today's online auction markets: eBay has auctions for physical goods closing every second, while Google's Doubleclick and Microsoft's Advertising Exchange trade online advertisements at a much faster rate. This creates the possibility for buyers to inter-temporally substitute, adjusting their current bids to account for the option value of waiting for future purchasing opportunities. When preferences are persistent, so-called "leakage effects" can also arise: a bidder who bids aggressively on a particular product may be revealing their preferences to the other bidders, which has strategic implications.

All of these effects are well understood by theorists, although seldom combined in a single model. However they have not received much attention in the empirical auctions literature, which typically considers the identification and econometric analysis of a repeated cross-section of observations of a static auction game. This is problematic for the markets cited above, in which directed search and dynamics may be important for counterfactuals.

For example, one question often emphasized is the design of optimal reserves. In a static model of a second-price auction, changing the reserve price has no effect on bidding. But when that auction is embedded in a market, changing the reserve price affects participation. And when we allow for dynamics, one can distinguish between a transitory increase in the reserve price (which has no strategic implications in a second price auction) and a permanent increase in the reserve price (which affects strategies by changing continuation values).

There are many more counterfactuals that only make sense in a richer model. For example, consider trying to predict what would happen to the bids on small state-funded construction projects when the government increases the supply of large federal contracts. This question fits naturally into a framework with heterogeneous goods and directed search. Or what about the effects of introducing of a new good? In fixed price markets, the usual approach

is to estimate demand as a function of product characteristics and predict the market share and consumer surplus of a new product. In our model we can follow the exact same logic by projecting multidimensional types down to characteristic space. In order to model substitution to the new good, however, we need a model of how bidders choose which auctions to participate in.

By ignoring the option value of participating in future auctions and how this affects the bidding function, the static auction model also mischaracterizes buyer valuations and therefore consumer surplus, in a way that is sharpest for high-type bidders.¹ Moreover, the static auction model throws away one of the most interesting features of auction data, which is our ability to link bidders across time. We believe that exploiting this feature is particularly fruitful for identifying cross-elasticities, similar in spirit to the use of second-choice data in demand estimation.

This paper aims to fill this gap in the empirical auctions literature. We introduce a particular model of a dynamic auctions market with directed search. We are guided in our modeling choices by the discrete choice demand system literature that has been employed in the analysis of durable goods markets. Our agents have unit demand and perfectly persistent preferences, described by their valuations for each of a discrete set of available goods. Each period, a set of goods is exogenously supplied, and bidders choose which good to bid on and how much to bid. Search frictions arise because the matching of buyers to sellers is random conditional on the choice of product, so that the number of bidders participating in the auction of any given object varies across auctions.

Dynamics are kept as simple as possible. Winning bidders exit with certainty and losing bidders exit exogenously with a fixed probability. Entry is governed by a distribution over the number of new entrants. Each new entrant draws their valuations independently and privately. The recent history of the game is publicly observable.

Despite these modeling choices, the dynamic environment remains complicated. Buyers have strategic incentives to learn each other's private valuations, making inferences from the observed history of the game. In order to characterize best response functions, we show that it is possible to rewrite the bidder's problem as a partially observable Markov Decision Problem (POMDP). We solve the POMDP to show that in a pure strategy Bayesian equilibrium, bidders participate on the good that gives them the highest expected surplus,

¹Surplus estimation is more convincing in auction models because unlike fixed price markets, it is possible to invert individual buyer behavior and nonparametrically identify the valuations of the highest types.

and then bid their valuation, less their continuation value when pivotal, plus a term that reflects leakage effects.

We then return to the questions of identification and estimation. The decision problem faced by an individual type is nonparametrically identified from the data; this means that the econometrician can look at the data and deduce the best response functions for each type. Since the equilibrium strategies are mutual best responses, this identifies them. This allows us to provide conditions under which the entire model is nonparametrically identified.

But these conditions are not constructive. To make progress, we consider an equilibrium in which leakage effects are small, as would be the case in a large market. Then we can go much further, showing that if a bidder is observed bidding once on each product, their type is identified by applying a contraction mapping. This “bidderwise” identification result provides part of the rationale for our decision to introduce both directed search and dynamics in the same model, since in a static model the econometrician only sees each bidder once.

The last part of the paper is concerned with estimation. We offer two approaches. One follows the nonparametric identification logic directly, taking the set of bidders who are observed bidding on every product and inverting back to their type using the contraction map developed in identification. But the set of bidders observed bidding on every product is a selected sample, and to get the true distribution of valuations, it is necessary to re-weight the estimated density. We show how to do this.

A disadvantage of the nonparametric method is that it is data intensive, since if the product space is large, the set of bidders who bid on every product may be quite small. Alternatively, by making a parametric assumption on the type distribution, we can use all the data. We show how to simulate moments given any parameter vector. We can thus apply simulated GMM to consistently estimate the model parameters.

This approach extends easily to a random coefficients model in which “types” are now idiosyncratic preferences over product characteristics, as in Berry, Levinsohn, and Pakes (1995). Monte Carlo simulations show the estimation approaches perform well in moderately sized samples, and a counterfactual simulation shows that taking account of dynamics is important in determining the optimal sales policy for a monopolist seller.

1.1 Literature Review

This paper builds on existing work in auction theory, search theory, empirical auctions and fixed-price demand systems. On the theory side, Milgrom and Weber (2000) introduced a model of sequential auctions over a finite horizon with a finite number of players. They showed that under certain information structures there are no learning effects, so that the strategies in each period are not history-dependent. The literature on dynamic matching and bargaining games went in a different direction, working with a continuum of players and an infinite horizon, allowing analysis of a steady-state and steady-state strategies (see Gale (2000) for a fantastic overview). Some recent work has focused on equilibria in which there is only temporary asymmetry in information, as between entrants and incumbents (Hendricks, Onur, and Wiseman 2008) or players use memoryless strategies (Said 2009). All of these models are cleverly designed to avoid “leakage effects” in which atomistic players must take account of how their actions today will affect their opponent’s future beliefs. By contrast, our model allows for this possibility, at the cost of significant modeling complexity. Our contribution here is to link this problem with the computer science and operations research literatures on partially observed Markov decision processes. This may enable computational analysis of these difficult problems.

Burdett, Shi, and Wright (2001) introduced a directed search model in which sellers posted prices, and buyers selected which seller to purchase from. They showed that in the symmetric equilibrium sellers all offer the same price and buyers randomize, creating frictions due to inefficient rationing. Subsequent papers explored simultaneous search in labor markets (Albrecht, Gautier, and Vroman 2006, Galenianos and Kircher 2009) and college admissions (Chade, Lewis, and Smith 2011). Peters and Serevinov (2006) develop an analogous model of an internet auction market where sellers post identical reserves and buyers randomize. In our paper, sellers are differentiated by the good they sell, buyers choose sellers, and search frictions arise within each buyer-product set due to random matching of buyers to auctions.

The existing empirical auctions literature is large; for excellent overviews, see Paarsch and Hong (2006), Athey and Haile (2007) and Hendricks and Porter (2007). Relative to this literature, we innovate in two distinct ways. The first is to allow for multiple kinds of goods and multidimensional bidder preferences, and model the directed search process by which bidders decide what good to bid on. Even in a purely static model, modeling search allows for realistic substitution patterns when the set of available objects being auctioned is changed.

Indeed, the participation decision is entirely analogous to discrete choice over products. A number of recent papers in the literature have also focused on the participation decision, but in a model where agents get some signal affiliated with their type before entry and must decide whether or not to enter (Marmor, Shneyerov, and Xu 2007, Li and Zheng 2009, Roberts and Sweeting 2012). Endogenizing participation leads to selection problems that must be addressed in identification and estimation.

The second innovation is that we allow for perfectly persistent private types in a dynamic environment. This is in contrast to the recent empirical literature on dynamic games, which considers the case with transient private information shocks (Aguirregabiria and Mira (2007), Bajari, Benkard, and Levin (2007), Pakes, Ostrovsky, and Berry (2007); Fershtman and Pakes (2012) is a recent exception). Jofre-Bonet and Pesendorfer (2003) is in a paper in this spirit, analyzing a dynamic auction game in highway procurement where the contractors have transient private information. Zeithammer (2006) found empirical evidence of dynamic behavior on eBay, showing that bidders shade down current bids in response to the presence of upcoming auctions of similar objects. Other related papers are Ingster (2009), who develops a dynamic demand model for identical objects, ignoring leakage effects; Sailer (2006) who estimates a dynamic model with participation costs; and Nekipelov (2007), who estimates a model where bidders attempt to prevent learning by late bidding.

We emphasize the issue of nonparametric identification in the paper, following the pioneering work of Athey and Haile (2002) for auction markets and the subsequent work on fixed price markets (Berry, Gandhi, and Haile 2012). Adams (2012) has analyzed a static version of our model with directed search, and provided partial identification results. Finally, our work is related to the dynamic demand literature for fixed price markets (e.g. Hendel and Nevo (2006), Gowrisankaran and Rysman (2009)).

The next section introduces the theoretical framework, while section 3 analyzes nonparametric identification. Section 4 describes our two different estimation approaches, while section 5 performs simulation exercises. Section 6 concludes.

2 Theory

In this section we introduce a stylized model of a large auction model with two novel features. The first is that there are many different products which are auctioned simultaneously in each

period, and buyers have distinct preferences for these products. In equilibrium there will be directed search: sellers are offering different products, and buyers will pick an auction for a good they value relatively highly. Assuming simultaneous auctions is a good way of approximating participation frictions in online markets that are truly sequential.² An important implication of directed search is that the set of bidders in any auction is selected, an issue that has been emphasized in many recent empirical auctions papers (Marmer, Shneyerov, and Xu 2007, Li and Zheng 2009, Roberts and Sweeting 2012).

The second is that buyers are long-lived and forward looking, with perfectly persistent private valuations. This introduces dynamic incentives: buyers may shade their bids to reflect the option value of buying later, or garble their bids to avoid revealing their valuations to their opponents (“leakage effects”). The option value effect will be important in markets where bidders have unit demand, as in the purchase of durable goods. Leakage effects will matter in markets where there are a small number of strategically inclined players, such as financial markets (e.g. auctions for distressed assets). Both of these are present in our analysis.

The stage game is complex, combining a matching model with a dynamic auction model. Accordingly, we make the dynamics as simple as possible. We assume that supply and the entry of new bidders is exogenous; that the exit of losing bidders is exogenous; and that bidders have unit demand. It would have been even simpler to avoid dynamic issues completely by assuming that all bidders exit after every period, generating a static auction game with directed search. We analyze this special case below, and show that the static model is not identified: Observations of the same bidder over time are the key to our identification strategy. We begin by describing the stage game, and then describe these dynamics.

2.1 The Stage Game

Figure 1 depicts the stage game. There are J distinct kinds of goods sold in a market, indexed by $j = 1 \dots J$. In period t the supply of good j is $N_{j,t}$, and overall supply is $N_t = (N_{1,t}, N_{2,t} \dots N_{J,t})$. Supply is exogenous, and the vector N_t is drawn iid each period from a distribution F_N with support $\{0, 1, 2 \dots \bar{N}\}^J$. Each good is simultaneously auctioned

²For example, one can think of eBay as a sequence of second-price sealed bid auctions ordered by closing time (Bajari and Hortacsu 2003). Without participation frictions, every incumbent bidder should participate in every auction, but in fact it is rare for the average bidder to participate in more than a single auction a day. Our model makes a pair of assumptions — multiple auctions per period and single auction participation — that may be thought of as a reduced form for some more complete model of search frictions.

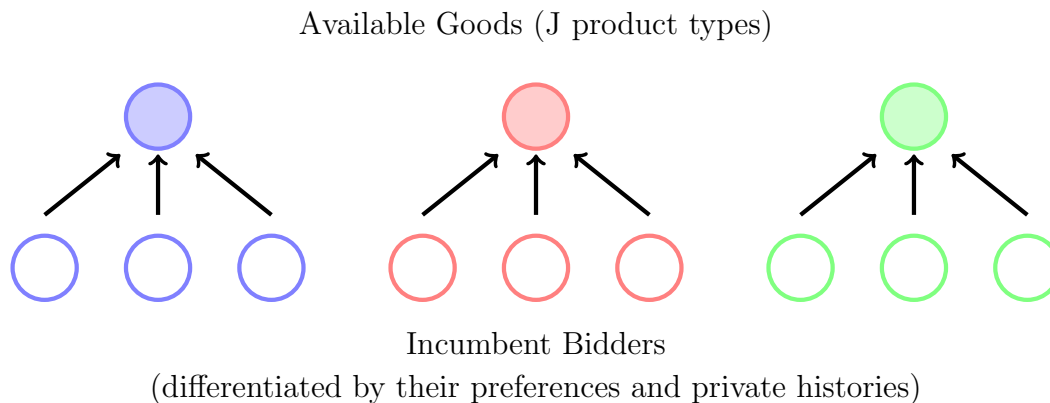


Figure 1: **The Stage Game.** Bidders decide which of the available goods to bid on and then how much to bid. There is heterogeneity on both sides of the market: bidders are differentiated by their preferences and private histories, and there are J product types.

by second-price sealed-bid auction. In each time period, there is some set of incumbent bidders, differentiated by their private information (which describes both preferences and knowledge). They choose a good to bid on and are randomly matched to an auction of that good.³ Without observing the number nor identities of the rival bidders in the auction, they decide how much to bid. Bidders are risk neutral and if they win good j they get a payoff equal to their valuation for that good x_j less the price paid.⁴

2.2 Market Dynamics

Bidders are assumed to have unit demand, and so winning bidders exit at the end of a period with certainty. Losing bidders exit randomly and independently at some rate $1 - r$, with a zero payoff on exit. Bidders are assumed to maximize their lifetime utility on the platform. The survival rate r plays the role of a discount rate, ensuring that lifetimes are almost surely finite and lifetime utility is bounded.⁵

³Nothing in the model gives a bidder a reason to pick one auction for good j over another, and so random matching is a sensible assumption. This sort of matching technology is commonly assumed in the theory literature, dating back to Rubinstein and Wolinsky (1985) (in bargaining) and Wolinsky (1988) (in auctions); and arises as an equilibrium phenomenon in the directed search model of Burdett, Shi, and Wright (2001).

⁴If no one participates in an auction or if the highest bid is zero, the good is not sold; if only one bidder participates and makes a positive bid, they pay zero. Losing bidders get a zero payoff.

⁵We could also allow for a separate discount rate, but this adds nothing to the analysis except that for the identification results we would need to assume that this discount rate is known by the econometrician.

At the end of each period, E_t buyers enter the market, independently drawing J -dimensional valuations x from some distribution F_X , where E_t is sampled independently over time from a distribution F_E with support $\{0, 1, 2 \dots \bar{E}\}$. So that the size of the market doesn't explode, we assume that whenever the number of incumbent bidders strictly exceeds $\bar{I} - \bar{E}$, no-one enters. It follows that the total number of bidders in the market never exceeds \bar{I} .⁶

We assume that F_X has support \mathcal{X} a connected subset of the hypercube $[0, \bar{x}]^J$ and strictly positive density on its support. Valuations are perfectly persistent and remain fixed while the bidder is in the market. Bidders are also differentiated by what they know. All information from the past t_P periods is publicly available. Specifically, at the end of the period and for each bidder the platform publishes an auction ID for the auction they participated in, the type of object they were bidding on, their bidder ID and their bid. Thus a bidder bidding for the k th time will have seen bid observations from periods $t - 1, t - 2 \dots t - t_P - k + 1$. We call this their history $h_{i,t}$ and let the set of all histories be \mathcal{H} .⁷

Bidder i 's information set at time t consists of their valuation x_i , the realized supply n_t and their history $h_{i,t}$. Throughout the analysis, we restrict attention to anonymous, stationary and symmetric pure strategies γ and β .⁸ Stationary symmetric pure strategies are a mapping from information sets to participation and bidding decisions: $\gamma : \mathcal{X} \times \mathcal{H} \times \mathcal{N} \rightarrow \mathcal{J}$ and $\beta : \mathcal{X} \times \mathcal{H} \times \mathcal{J} \times \mathcal{N} \rightarrow \mathbb{R}$.⁹ We will say that these strategies are anonymous if they are invariant under relabeling of auction numbers and rival bidder identities.

2.3 Long-Run Properties

We begin by analyzing the behavior of the dynamic system when the agents employ any (possibly sub-optimal) strategies γ and β . Recall that at any point in time each bidder's information set consists of their valuation, their history and the current supply. Taking as a state variable the collection of all incumbent valuations and all incumbent histories (i.e. everything privately known) and the public history and current supply (i.e. everything publicly known), all actions are fully determined by the state. These actions, along with the random matching outcomes, determine the endogenous exit of agents; and since entry and

⁶With endogenous entry, one would expect a condition like this to hold: as the number of participants goes to infinity, expected surplus falls to zero and entry stops. With exogenous entry, we must impose it.

⁷Since bidders can be arbitrarily long-lived, this set will be infinite dimensional.

⁸This restriction is common in the dynamic matching and bargaining game literature (Gale 2000).

⁹Stationarity and symmetry imply the mappings do not depend on identity i and time period t directly.

supply are exogenous, the system evolves as a discrete-time Markov process.

We can therefore analyze the long-run behavior of the dynamic system using standard techniques. We would like the system to be recursive in the sense that it is possible to revisit previous states. To do this, we assume that auction and player identities are “recycled”: whenever an identity is no longer needed, in the sense that it is in no incumbent bidder’s history, it will be assigned to the next auction or player that requires an identity. This is without loss of generality whenever strategies are anonymous.¹⁰

Letting \mathcal{I}_t be the set of incumbent bidders and h_t^P be the public history, we define the state as outlined above, anonymizing the histories: $s_t = (\{x_i, h_{i,t}\}_{i \in \mathcal{I}_t}, h_t^P, N_t) \in \mathcal{S}$. Let \mathcal{S} be metrized by the sup norm, and let $\mathcal{B}(\mathcal{S})$ denote the Borel sigma-algebra on \mathcal{S} .¹¹ Then there exists a Markov transition function $P_{\gamma, \beta}$ on $(\mathcal{S}, \mathcal{B}(\mathcal{S}))$ that describes the state transitions. What happens to this Markov process in the long-run? It cannot literally have a steady-state, as there is continuous fluctuation in the number of players and their valuations. But there is a unique ergodic measure over the state space, which gives the long-run probability of any measurable set of states occurring:

Lemma 1 (Long-Run Evolution). *Given anonymous, symmetric and stationary strategies γ and β , there is a unique ergodic measure $\mu_{\gamma, \beta}$ on $(\mathcal{S}, \mathcal{B}(\mathcal{S}))$, converged to at geometric rate.*

In the proof, we need to rule out two extreme cases: that there is no ergodic measure (the chain somehow drifts away over time); or multiple ergodic measures (the chain splits at some point). We show this by arguing that the market periodically collapses when everyone spontaneously exits. So the chain cannot drift away as it is anchored by these collapses; nor can it have two disjoint non-communicating classes of states, as both communicate through collapses. In fact, using the techniques of Stokey, Lucas, and Prescott (1989), we prove a stronger result: that the chain converges geometrically to the ergodic measure.

2.4 The Single-Agent Decision Problem

Now that we understand that the market dynamics are Markovian, we turn to the strategic analysis. We start with the decision problem of a single agent. Each period he must choose

¹⁰This transformation makes the system recursive, but any individual bidder faces a non-recursive decision problem: since their private history grows each period, they will never revisit a previous state.

¹¹Throughout the paper we will use $\mathcal{B}(A)$ to denote the Borel sigma-algebra on a space A and $\mathcal{P}(A)$ the associated space of probability measures on $(A, \mathcal{B}(A))$.

which product to bid on and how much to bid, knowing his own valuation and history. In equilibrium, he understands what his rivals will bid on and what they will bid given their information sets. What he doesn't know is what is in their information sets: their valuations, as well as any history they may have privately observed. Some part of this is impossible to know, as in each period there is a new cohort of entrants whose valuations are independently drawn from F . But the history may be informative as to the valuations held by incumbent bidders at $t - 1$, some of which will survive and play again in period t .

Because both the state and the observations (innovations in his private history) evolve as a controlled Markov process, and because the state is only partially observed, his problem can be described as a *partially observable Markov decision problem* (POMDP) (Sondik 1978).¹² From his perspective, the state transitions are described by a transition kernel $P_{\gamma_{-i}, \beta_{-i}}(s'|s, j, b)$ where the prime denotes that s' is tomorrow's state, j is his chosen good and b is his bid. The transition rule depends on the strategies of the other bidders γ_{-i} and β_{-i} , as well as the current state s and his own action (j, b) . Since the state is not observable, he has to form beliefs over the state on the basis of what he does see: his private history.¹³

A well-known result in computer science and operations research is that POMDP's can be transformed back into Markov Decision Problems (MDP's). This is achieved by a transformation of the state space. To solve a POMDP the agent has to form beliefs about the current state using the observations he has made. The so-called "belief MDP" implements this idea, creating a state space from the set of all beliefs π about the underlying state s . These beliefs evolve according to Bayes rule as a controlled Markov process, because the agent's observations of the system are themselves Markov. So instead of solving the POMDP on the partially observed state space, one may equivalently solve the MDP on the belief space, finding strategies $\gamma(\pi)$ and $\beta(\pi, j)$ that depend on current beliefs π and chosen object j .

Let us make this transformation, and consider the agent's decision problem.¹⁴ Let $G_{j, \pi}^M$ be the distribution of the maximum rival bid B_j^M in a random auction for item j when the state

¹²A quick taxonomy: autonomous (uncontrolled) Markovian systems with hidden state variables are *hidden Markov models*; controlled models with observable states are *Markov decision problems*.

¹³Because valuations are iid, his valuation is not informative after conditioning on his history.

¹⁴For the moment we defer the question of where the agent's initial beliefs — their prior — comes from.

is distributed according to π . The value function satisfies the following Bellman equation

$$V(\pi) = \max_{j,b} \left[G_{j,\pi}^M(b) (x_j - E_G[B_j^M | B_j^M < b]) + r(1 - G_{j,\pi}^M(b)) \int V(\pi') dQ(\pi' | \pi, j, b, B_j^M > b) \right] \quad (1)$$

The left hand side is the expected value of the game given beliefs π . The intuition for the right hand side is as follows. If an agent bids b when bidding on an item-type j , they win with probability $G_{j,\pi}^M(b)$ and pay $E_G[B_j^M | B_j^M < b]$, where the probabilities are assessed with respect to their beliefs π . At this point the game ends.

However, if they lose and survive — which happens with probability $r(1 - G_{j,\pi}^M(b))$ — then they will update their beliefs tomorrow and get a payoff of $V(\pi')$. Belief updating follows Bayes rule, and is described by a conditional probability measure $Q(\pi' | \pi, j, b, B_j^M > b)$ that gives the ex-ante distribution of the beliefs the agent may hold tomorrow after they observe the events of the current period. It is ex-ante in the sense that the agent has not yet observed the history of this period, and makes forecasts based on what currently believe, what they themselves will do, and noting that if they play tomorrow they must have lost today.

This dynamic programming formulation allows a clean characterization of the optimal policy. Write the conditional continuation value when the agent bids b on product j and the highest rival bid is B as $v(\pi, j, b, B) = \int V(\pi') dQ(\pi' | \pi, j, b, B_j^M = B)$. Taking a first order condition, we obtain the following characterization of the bidding function:

Theorem 1 (Optimal Strategies). *The participation strategy $\gamma(\pi)$ must satisfy:*

$$\gamma(\pi) \in \arg \max_{j \in J(\pi)} \max_b \left[G_{j,\pi}^M(b) (x_j - E_G[B_j^M | B_j^M < b]) + r(1 - G_{j,\pi}^M(b)) \int V(\pi') dQ(\pi' | \pi, j, b, B_j^M > b) \right] \quad (2)$$

where $J(\pi)$ denotes the set of products available at state π . Let $j^* = \gamma(\pi)$ and $\beta(\pi, j^*) = b^*$. Then whenever $G_{j^*,\pi}^M$ has a density $g_{j^*,\pi}^M$ at b^* , b^* satisfies the implicit equation:

$$b^* = x_{j^*} - rv(\pi, j^*, b^*, b^*) + r \left(\frac{1 - G_{j^*,\pi}^M(b^*)}{g_{j^*,\pi}^M(b^*)} \right) E_G \left[\frac{\partial}{\partial b} v(\pi, j^*, b, B) \Big| B > b^* \right] \quad (3)$$

The first part is intuitive: the object bid on must be the best of the available options, in the sense of maximizing expected payoff. The second part describes the optimal bid. Ignoring the third term in (3) momentarily, we can use the usual intuition from second price auctions to explain the result. Bidders must bid in such a way that they experience no ex-post regret. If a bidder marginally wins, they get a surplus of $x_j - b$. If on the other hand they marginally lose, they get the value of the game tomorrow, which is the second term on the right hand side. This takes into account the fact that their bid was pivotal, which is informative about future competition.¹⁵ These two terms must be equal, otherwise they would regret having not bid marginally higher.

The extra third term in (3) reflects the fact that in a dynamic environment bidders must account for the effects of their bid on future competition. These are termed “leakage effects” in the auction literature. Leakage only matters upon losing (i.e. when $B > b^*$), and is captured by the partial derivative in these states. We expect leakage effects to lead to more bid shading: high bids convince rivals that competition is tough, lowering their perceived continuation value and therefore increasing their optimal bids, which hurts the agent.

2.5 Equilibrium

Our equilibrium notion is Bayes-Nash equilibrium, slightly modified to take into account the fact that we are interested in a market in long-run equilibrium. We place three restrictions on equilibrium play. First, agents must form beliefs according to Bayes Rule whenever possible. To make sense of this, we endow agents with an initial set of beliefs that are consistent with long-run play: if the equilibrium strategies are (γ, β) and the ergodic measure over states is thus $\mu_{\gamma, \beta}$, their initial beliefs are $\pi_0 = \mu_{\gamma, \beta}$.

So an agent entering the market will start with beliefs $\mu_{\gamma, \beta}$, observe the last t_P periods of data, and immediately update to new beliefs π'_0 before their first auction. This process corresponds to what a sensible analyst might do upon observing a long time series of data from a market in steady-state equilibrium: work out what the conditional distribution of various statistics of interest is, given the recent history. We also require that agents optimize given beliefs, so that the strategies (γ, β) are a solution to the POMDP defined by (1).

¹⁵Ignoring leakage, a bidder who bids their valuation less their unconditional continuation value falls prey to a “winner’s curse”: when they win, they learn that their rivals had low valuations for the object, implying that they may have been able to win it even cheaper tomorrow.

Finally, off-path beliefs must sustain equilibrium play.¹⁶

A brief discussion of how this compares to other equilibrium notions is useful. Our restrictions are identical to those of Bayes-Nash equilibrium, except that the initial beliefs come from long-run play rather than a prior over the initial state of the market at time zero and Bayesian updating to form a posterior at time t . Whether this is reasonable or not is arguable, but by Lemma 1, for large t (i.e. in the long-run), the initial distribution of types should be irrelevant, so that the two approaches to the prior coincide.¹⁷ Any equilibrium that obeys our restrictions is also an experience-based equilibrium (Fershtman and Pakes 2012). This is because our agents must have assessments of expected returns to on-path actions that are consistent with long-run play.

We will not attempt to prove existence of a pure strategy equilibrium, as this involves challenges beyond the ambition of this paper.¹⁸ Instead, we will assume that such an equilibrium exists and proceed to ask whether the primitives can be identified from observable data under equilibrium play.

3 Identification

3.1 Observable Data

We assume that the econometrician observes an arbitrarily long time series of the data collected by the platform: all bids, bidder identities, auction numbers and outcomes. The resulting data set is shown graphically in Figure 2. Each bidder’s private history and actions are observed and collected in the vector y_t^i . Following a single bidder over time, we get

¹⁶We ignore off-path beliefs and refinements in what follows. They are perhaps less of a force for multiplicity here than in other games: because of the private information, agents have different payoff functions; and because winning leads immediately to exit, it is hard to get different types to pool based on the threat of future punishment.

¹⁷Assuming that initial beliefs come from long-run play is theoretically unattractive in that they are then themselves equilibrium objects. On the other hand, assuming a common prior at the start of the market (i.e. before any play) begs the question of where the prior came from.

¹⁸Indeed, even in a static auction game it can be hard to prove existence (Athey 2001, McAdams 2003). We are dubious that a pure strategy equilibrium will exist without adding some noise or discretizing the bid observations so that a player’s beliefs about rival’s types never converge to atoms; although as Bergemann and Hörner (2010) note, one may be able to restore existence even in that case by enriching the strategy space in various ways (see the discussion in their paper). For simplicity, we simply assume existence and thus avoid this additional modeling complexity.

		Bidder Identity						Aggregate	
		...	i	i+1	i+2	i+3	i+4	...	Data
Time Period	⋮								
	t		y_t^i	y_t^{i+1}	\emptyset	\emptyset	\emptyset		y_t
	t+1		y_{t+1}^i	y_{t+1}^{i+1}	y_{t+1}^{i+2}	\emptyset	\emptyset		y_{t+1}
	t+2		y_{t+2}^i	\emptyset	\emptyset	y_{t+2}^{i+3}	\emptyset		y_{t+2}
	t+3		\emptyset	\emptyset	\emptyset	y_{t+3}^{i+3}	y_{t+3}^{i+4}		y_{t+3}
	t+4		\emptyset	\emptyset	\emptyset	y_{t+4}^{i+3}	y_{t+4}^{i+4}		y_{t+4}
	⋮								
Individual Time Series			y^i	y^{i+1}	y^{i+2}	y^{i+3}	y^{i+4}		

Figure 2: **Observable Data.** This is the dataset that we assume the econometrician observes. Each entry y_t^i consists of the private history $h_{i,t}$, object chosen $j_{i,t}$ and bid $b_{i,t}$ of bidder i in time period t . A notation of \emptyset means the agent was not present in the market at that time period, either because they had not yet entered, or had already exited. Aggregating all the individual data from a given time period gives us the aggregate data y_t ; aggregating all the data from a given individual across time periods gives us the individual time series y^i .

an individual time series y^i ; aggregating all the data across individuals from a single time period together, we get a vector of actions and histories y_t . Both arrangements of the data are useful to us. By Lemma 1, the aggregate time series $\{y_t\}$ is geometrically ergodic, and we can apply time series econometrics techniques in its analysis. But since types are sampled iid on entry, for the identification arguments it will be easier to look at the collection of individual time series $\{y^i\}$.

The primitives of the model are the distribution of the number of entrants F_E , the distribution of the number of goods listed of each type F_N , the distribution of types F_X and the survival rate r . Notice that everything except the type distribution is trivially identified directly from the aggregate data (e.g. the number of entrants each period is observable, and therefore so is its distribution). So the challenge is to identify the bidder's valuations F_X .

We make our identification argument in two steps. First we show that the econometrician can deduce the equilibrium strategies $\sigma^e = (\gamma^e, \beta^e)$ from the observables. Second, if different types take different actions (at least along some paths), we can use our knowledge of σ^e to invert from the distribution of individual bidder time series to the type distribution.

3.2 Identifying Equilibrium Strategies

To start, we rewrite the Bellman equation (1) in terms of observables, rather than the bidder's beliefs. Beliefs enter through the bidder's perception of the distribution of the highest rival bid $G_{j,\pi}^M$, and the transition kernel $Q(\pi'|\pi, j, b, B_j^M > b)$. In equilibrium, bidders form their beliefs by starting with the ergodic distribution of states and updating by Bayes rule based on their history h and the supply realization n .¹⁹ Their beliefs are thus consistent with the ergodic measure over the state space μ , conditional on h and n , and we can define $G_j^M(b|h, n) \equiv E_\mu[B_j^M < b|h, n] = G_{j,\pi}^M(b)$. Similarly since future beliefs are a function of future private history and supply, their perceived transition kernel is $Q(h', n'|h, n, j, b, B_j^M > b)$. Because supply is exogenous and iid, this can be decomposed further, allowing us to write the following observable counterpart of the original Bellman equation:

$$V(x, h, n) = \max_{j, b} \left[G_j^M(b|h, n) (x_j - E_G[B_j^M | B_j^M < b, h, n]) + r(1 - G_j^M(b|h, n)) \int \int V(x, h', n') dF_N(n') dQ(h'|h, n, j, b, B_j^M > b) \right] \quad (4)$$

where $V(x, h, n)$ is the value of a bidder of type x with private history h facing a supply vector n . The objects in (4) are identified as the asymptotic limits of their empirical analogs.²⁰

Letting \mathcal{G} be the set of conditional bid distributions $\{G_j^M(b|h, n)\}$ and \mathcal{Q} be the set of conditional transitions $\{Q(h'|h, n, j, b, B_j^M > b)\}$, define the set of empirical best responses $\Sigma(\mathcal{G}, \mathcal{Q}, F_N)$ as the set of policies that solve (4) (i.e. the set of best responses to equilibrium actions). Each element σ of this set is a best response function $\sigma(x)$. One can construct all the elements of Σ by solving the MDP defined by (4) for each type x , and then taking the cartesian product of the sets of solutions for each type. This implies the empirical best responses are identified. Moreover, we know that the equilibrium strategies σ^e must be in $\Sigma(\mathcal{G}, \mathcal{Q}, F_N)$ otherwise some type could deviate and improve their payoff, contradicting the definition of equilibrium. Summarizing:

Lemma 2 (Identification of Optimal Policies). *The empirical best responses $\Sigma(\mathcal{G}, \mathcal{Q}, F_N)$ are identified. If $\Sigma(\mathcal{G}, \mathcal{Q}, F_N)$ is a singleton, equilibrium policies $\sigma^e = (\gamma^e, \beta^e)$ are identified.*

¹⁹Recall that because valuations are iid and have no impact on state transitions, they are redundant.

²⁰For example, let $G_j^{M,T}(b|H, n)$ be the empirical probability in a dataset of T auctions that the highest rival bid faced by a bidder bidding on good j is less than b when supply is n and his history is in some open set of histories H . Since the system is ergodic, we may apply the Birkhoff ergodic theorem to conclude that $\lim_{T \rightarrow \infty} G_j^{M,T}(b|H, n) = G_j^M(b|H, n)$ almost everywhere.

3.3 Inversion

Suppose that $\Sigma(\mathcal{G}, \mathcal{Q}, F_N)$ is indeed a singleton. Lemma 2 suggests a natural identification strategy: since the optimal policies are identified, why not invert from the observed actions to the types? For example, in the case of a static symmetric first-price auction, there is a unique monotone equilibrium bidding function β . Given this optimal policy and its inverse β^{-1} , one can invert the bid distribution pointwise to get the type distribution. This is basically the identification strategy of Guerre, Perrigne, and Vuong (2000).²¹

This approach is complicated by the fact that the type and action spaces are multidimensional, and we may observe repeated observations from a single bidder. To simplify this problem, we focus our analysis on the individual bidder time series $y^i \in \mathcal{Y}^{\mathcal{I}}$. For each type there is a distribution over these time series outcomes $F_{Y|X}$ implied by the primitives and equilibrium play. The econometrician can simulate these paths by drawing an initial public history at random, applying the equilibrium strategies, and then sampling rival high bids, transitions and supply according to $(\mathcal{G}, \mathcal{Q}, F_N)$ to generate the next state in the path. Therefore the distributions $F_{Y|X}$ are identified. The data can be summarized as a distribution $F_Y(y)$ over individual time series, generated as an (infinite) mixture of the time series from different types:

$$F_Y(y) = \int F_{Y|X}(y|x) dF_X(x) \quad (5)$$

This defines an inversion problem: we want to deduce F_X from F_Y given $F_{Y|X}$. Intuitively, this is only possible if different types have different strategies; that is, if the equilibrium is *separating*. We make a stronger assumption: that for each type there is some path — perhaps a long sequence of bids and participation decisions — that is separating in the sense that only that type could have played it.

Assumption 1 (Dynamic Separability). *Let $\mu_{Y|X}$ be the conditional measures on $(\mathcal{Y}^{\mathcal{I}}, \mathcal{B}(\mathcal{Y}^{\mathcal{I}}))$ corresponding to the distributions $F_{Y|X}$. For each type x , there is some set of outcomes $A \in \mathcal{B}(\mathcal{Y}^{\mathcal{I}})$ with $\mu_{Y|X}(A|x) > 0$ and $\mu_{Y|X}(A|x') = 0 \forall x' \neq x$.*

Dynamic separability condition is a weak condition in the common case when types are assumed to be one dimensional.²² Then it would suffice to assume a monotone equilibrium. In higher dimensions the condition becomes harder to satisfy and we may only be able to

²¹In their case, the optimal policy is identified from the first order condition for the bidding problem.

²²Examples: models with a single kind of good; identical preferences over different goods up to a scalar unobservable (Haile, Hong, and Shum 2006); in characteristic space, models with a single random coefficient.

get partial identification (we discuss this later in the paper). As noted earlier, we will also need a second assumption on the uniqueness of the empirical best responses:

Assumption 2 (Unique Empirical Best Responses). *For any set of primitives F_E, F_N, F_X , r and equilibrium strategy σ^e which induce bid distributions \mathcal{G} and transitions \mathcal{Q} , the empirical best response set $\Sigma(\mathcal{G}, \mathcal{Q}, F_N)$ is a singleton.*

This is not the same as requiring a unique pure strategy equilibrium: multiple equilibria are fine because each equilibrium will generate different observables. Instead, it rules out the possibility that there are two different type distributions and strategies (F_X^1, σ_1^e) and (F_X^2, σ_2^e) that generate identical data, and therefore both σ_1^e and σ_2^e are best responses.

Theorem 2 (Non-Parametric Identification). *Under assumptions 1 and 2, the model is identified.*

The proof appeals to a result in the statistics literature on the identification of mixtures (Blum and Susarla 1977). Our dynamic separability assumption is stronger than necessary for the inversion. We use it because it is relatively transparent, whereas the necessary and sufficient condition is devoid of economic content.²³

Neither of the assumptions are on the primitives, and unfortunately we don't know enough about the full Bayesian equilibrium to verify that these assumptions always hold. What we can say is that they are both in principle testable by solving the MDP and then simulating paths. Testing these conditions would give a sense of whether identification is possible with any particular dataset.

We will show below that in large markets we can go further and dispense with the assumptions. First though, we will use our analysis to get a negative identification result for the static case in which bidders only live for a single period (i.e the survival rate is $r = 0$). Data from the static model essentially consists of repeated iid observations of the static stage game. In each stage game, agents pick an item j to bid on and then bid on it. There are no dynamic implications — no leakage, no continuation value — and so bidders bid their valuations in any auction they choose to participate in.

Corollary 1 (Static Model). *Fix $r = 0$ and any entry and supply distributions F_E and F_N . F_X is identified iff $J = 1$.*

²³Fox and Gandhi (2011) make a similar point about identification of finite mixture models, and introduce their own sufficient condition called reducibility. Dynamic separability and reducibility are closely related.

When $J = 1$, identification is obvious: valuations are one-dimensional and are directly identified from the bid distribution. But when $J > 1$, the bid distribution for product j is informative only about the marginal distributions of valuations of bidders who have selected into bidding on product j . As in Roy’s (1951) classic model of the labor market, unless there is sufficient variation in the choice sets, some bidders will never select into bidding on product j and therefore even the marginal distributions are not fully identified.

But even if F_N has full support, so that sometimes only a single product is available and everyone selects into bidding on it, it is still hard to infer the correlation structure of F_X . For example, in a two-good economy with many high bids for good 1, the challenge is to learn what the contribution is from types with high and low valuations for good 2, respectively. The supply variation may provide an experiment: there may be some supply vectors in which the bidders with high valuation for good 2 switch to bidding on good 2, making it clear what the relative contributions of the different bidder types are. But there are “not enough” supply vectors to separate out the types: since the set of possible supply vectors is finite and the set of types is infinite, there is always partial pooling in the participation decision. And since each bidder is observed exactly once, the correlation structure cannot be inferred by following the same bidder over time across different supply vectors. As we will see below, this possibility is the key to identification, and is the reason we chose to focus on a dynamic model.

3.4 Large Market Approximation

Although the full Bayesian model is identified under the conditions we have outlined above, it will be formidably difficult to estimate in markets with large numbers of players and products. Happily, those are also the markets in which it may be reasonable to use more tractable equilibrium notions.

We consider an equilibrium in which bidders ignore the effects of their own actions on evolution of other’s beliefs — that is, they ignore leakage effects. This is reasonable in large markets: as the market size approaches the continuum limit, the distribution of opposing types is close to the steady-state distribution with probability approaching one, and so it is natural to assume that agents mutually ignore the actions of individual players, eliminating leakage effects.²⁴ In that case the assumption acts as an equilibrium refinement, eliminating

²⁴Bodoh-Creed (2012) argues that under certain smoothness conditions, equilibrium strategies of contin-

undesirable equilibria where agents coordinate on the basis of *a priori* unreasonable grounds.

Notice that in this special case bidder’s actions have no dynamic implications apart from determining whether they win and exit, or lose and persist. Formalizing this requires a little care. Fixing rival strategies γ_{-i}, β_{-i} and a belief-state π , let $\tilde{j} = \gamma_{-i}(\pi)$ and $\tilde{b} = \beta_{-i}(\pi, \tilde{j})$. Then define an autonomous transition kernel $Q(\pi'|\pi) \equiv Q(\pi'|\pi, \tilde{j}, \tilde{b}, B_j^M > \tilde{b})$. This kernel is autonomous because we have substituted rival strategies for what would otherwise be the controlled part of the transition process.

Definition 1 (No Leakage Equilibrium). *Symmetric strategies σ^e constitute a no-leakage equilibrium if σ^e solves the belief MDP with autonomous transition kernel $Q(\pi'|\pi)$, where on-path beliefs and the prior are consistent with long-run equilibrium play under σ^e .*

A no leakage equilibrium has the properties we desire. The transition kernel is correct on-path, because in a symmetric equilibrium rival strategies coincide with the player’s own strategy (i.e. they do what they expected to do). But agents treat the transition kernel as autonomous, ignoring their ability to influence rival beliefs by changing their actions when solving their decision problem. The value of considering these equilibria is that equilibrium strategies are considerably simpler and better defined.²⁵

Definition 2 (Strategy Monotonicity). *Fix two types x^1 and x^2 with $x_j^1 < x_j^2$ and $x_k^1 = x_k^2$ for $k \neq j$. Strategies are monotone if at any (h, n) , whenever type 1 bids on product j , type 2 also bids on product j and bids strictly more; and whenever type 2 bids on product $k \neq j$, type 1 also bids on product k and bids weakly more.*

Strategy monotonicity seems like something one would desire from reasonable demand systems. It says that if one type values a product more than another otherwise identical type, they should be more likely to bid on it, and when they do, they should bid more. Conversely, they should be less likely to substitute away to another product, and if they do, they should bid less. No leakage equilibria have this property:

um games will be an epsilon-Bayes-Nash equilibrium of the finite game with sufficiently many players.

²⁵In the analysis of industry dynamics, Weintraub, Benkard, and Van Roy (2008) introduce the related notion of oblivious equilibrium, in which firms optimize against the ergodic distribution of rival firm types (instead of conditioning on the full state vector). The no leakage assumption is different: our bidders are sophisticated in how they form beliefs about rival play, conditioning on all information available; and are naïve only in predicting how their own actions will affect future play.

Lemma 3 (No Leakage Equilibrium). *In a no leakage equilibrium, strategies are monotone. When optimally bidding on product j at information set (h, n) , bidders bid their valuation for j less their (survival)-discounted continuation value in that state.*

Recall that if we ignore leakage effects, bidders bid their valuations less their continuation value discounted by the survival rate. So whenever their valuation for the good they are bidding on increases they bid more. since their continuation value increases at a slower rate than their valuation (a rate of at most $r < 1$). Similar arguments show that participation strategies must be monotone.

This has strong implications for identification. Solving the empirical MDP for a type x produces a value function $V(x, h, n)$, and by standard arguments this value function is *unique*. Consequently the discounted continuation value is also unique, and since bids are equal to valuation less discounted continuation value, so are the best responses. Bidders who are observed bidding on all J products will be distinguishable, by the monotonicity of the strategies. In fact, we can explicitly construct an inversion from a set of bids that includes a bid on each product back to the bidder’s type.²⁶

Theorem 3 (Bidderwise Identification). *In a no leakage equilibrium the type of any bidder is identified iff the bidder is observed bidding at least once on each of the J products.*

The proof strategy is quite novel. We first reduce the bid-vector of each bidder to a J -vector, by randomly selecting one of their bids on each product. We then conjecture a J -vector of discounted continuation values v at each of those points, and infer the type $\hat{x} = b + v$ who would have made those bids. Solving the MDP for that conjectured type \hat{x} we get a new continuation value for each of those points v . This algorithm defines a mapping from the space of continuation values \mathbb{R}^J into itself, and we show that this map is a contraction map.²⁷ Thus it converges to a unique value function, identifying the true type according to $x = b + v$.

This result allows us to prove an explicit and constructive result regarding identification of the model. The idea is to apply the bidderwise identification result to every bidder observed bidding on all products. However, this set of bidders is selected: in particular, bidders with high valuations are unlikely to be observed bidding on all products since they will win early

²⁶An earlier version of this paper required that a bidder was observed making bids in *every* state, which in this model is a zero probability event.

²⁷In the context of fixed-price demand systems, Berry (1994) constructs a similar inversion from market shares to product/market level unobservables.

and exit. We can correct for this by re-weighting: if $\tilde{F}_{\mathcal{J}}$ is the distribution of types recovered by applying the contraction mapping of Theorem 3 to this set of bidders, then we can divide through by the probability that each type would be observed bidding on all products to recover $F_{\mathcal{J}}$, the distribution of types who bid on all products *with positive probability*. The following lemma establishes that these weights are recoverable.

Lemma 4 (Selection Correction). *Let \mathcal{A} be the set of all subsets of \mathcal{J} . In a no leakage equilibrium the probability of observing a type x bid on any $A \in \mathcal{A}$ is identified.*

A more serious problem for identification stems from directed search: some types may never bid on some products, regardless of the state of the world (i.e. for those types the probability of bidding on all products is *zero*). Then dynamic separability may fail and the model may be only partially identified. Still, Theorem 3 can be applied to deduce the marginal distribution of valuations for the set of goods on which these bidders actually bid. This may be enough for certain counterfactuals. To formalize this, define the “product support” of type x as

$$J(x) = \{j : \mathbb{P}(\gamma(x, h, n) = j) > 0\}$$

This is the set of all products that type bids on with positive probability. Let X_A be the random variable X restricted to only the products in $A \in \mathcal{A} \setminus \emptyset$ (i.e. defined by an identity map from \mathcal{X} to $[0, \bar{x}]^{|A|}$) and define the conditional distribution function $F_A(x_A) \equiv P(X_A \leq x_A | J(X) = A)$. This is the joint distribution of valuations for products in A of types with product support A .

Theorem 4 (Identification in Large Markets). *In a no-leakage equilibrium, the distribution F_A is nonparametrically identified for all $A \in \mathcal{A} \setminus \emptyset$.*

Notice how much stronger this is than the earlier identification Theorem 2. Not only do we dispense with assumptions 1 and 2, we also sharply define the limits of the data. The proof is in two parts. First we show that for any product set A , we can apply the bidderwise identification result plus a selection correction approach to get the type distribution for types with that product support. Then we proceed by induction from the maximal set down, getting the type distribution and then subtracting the contribution of higher-level sets. For example, when there are two products, we get the density of types who bid on both goods, then get the marginal valuations for good 1 and remove the contributions of those

who bid on both to get the marginal for those who only bid on good 1. The nice thing about this identification approach is that it may be directly employed in estimation.²⁸

4 Estimation

In this section we explore estimation of the full Bayesian model, as well as the model where we assume a no leakage equilibrium. There are two basic approaches. The first approach is nonparametric, following the logic of the identification section by inverting from the distribution of bidder time series to types. In the no leakage case, this is relatively straightforward: we simply implement the constructive identification argument. The disadvantage — and it is a serious one — is that it is impractical in markets with many products or high turnover in participants, since then it will be unusual to see many bidders bidding in multiple states.

Our second approach is a semiparametric approach based on simulated generalized method of moments (an approach often used for demand estimation). There we assume a parametric structure on the distribution of types, and choose parameters to match moments implied by the structural model with those observed in the data.

Even in that case, if there are a large number of products the model quickly becomes unwieldy. So a third approach is to follow the IO literature and project product valuations onto characteristics (McFadden 1974). Instead of valuations for products, types are now random coefficients indicating the marginal value of product characteristics. In a standard specification, this implies a linear structure for valuations in characteristics. Regardless of the approach — nonparametric, semiparametric or characteristic-based — we apply a two-step estimation approach, as is common in the estimation of dynamic games (Aguirregabiria and Mira 2007, Bajari, Benkard, and Levin 2007, Pakes, Ostrovsky, and Berry 2007).

4.1 First-Stage Estimation

The first step is to estimate the conditional bid distributions, the empirical transitions, the entry and supply distributions, and the exit rate. The entry and supply distributions are

²⁸A disadvantage is that this result requires a non-standard equilibrium concept. To address this, one could try to prove that leakage effects are uniformly bounded in sufficiently large markets. The contraction mapping arguments above would then produce partial identification results, where each type is identified up to the error produced by ignoring leakage.

discrete, and so can be estimated from simple averages. The conditional bid distributions and empirical transitions can be estimated nonparametrically. In principle data limitations can be overcome by appropriate choice of bandwidth, but since the history vectors are high dimensional, estimating objects like $G_j^M(b|h, n)$ can be difficult. For this reason researchers may choose to estimate these distributions within pre-defined blocks, treating certain histories and supply vectors as equivalent.²⁹

4.2 Second-Stage: Nonparametric Approach

To estimate the distribution of valuations nonparametrically, we follow the main identification argument given in Section 3. For each type in the type-space, we solve the empirical MDP to recover optimal strategies. Using those strategies, we may simulate the distribution of time series outcomes for each type, and then invert from the observed distribution of time series to the implied distribution of types.

One of the practical difficulties of this approach is that it requires solving the empirical MDP for each type in \mathcal{X} , possibly necessitating a rather coarse discretization of \mathcal{X} for computational reasons. In large markets, it will be attractive to assume that bidders are playing a no leakage equilibrium, because then we can apply the bidderwise identification approach, solving the bid inversion problem only for the bid vectors actually observed in the data. Following this inversion, we smooth the estimated type density and then re-weight according to the selection correction probabilities that we derive in the proof of Theorem 4.

We omit a formal analysis of the asymptotic properties of this estimator, both because it takes us into the realm of nonparametric estimation with dependent data — which we understand poorly at best — and because we suspect that the semiparametric approach outlined below is more likely to be used in practice. Yet intuitively Lemma 1 guarantees that the data generating process converges geometrically to an ergodic distribution, and so the asymptotics should be well-behaved. This is supported by the estimator’s performance in our Monte Carlo experiments (see below).

²⁹A previous version of the paper suggested formalizing such “coarsening” as part of the equilibrium concept. This is a good idea if the modeler has a strong prior on what the bidders condition on in making their decisions, but otherwise is problematic.

4.3 Second-Stage: Semiparametric Approach

The semiparametric approach proceeds in the opposite direction. Instead of inverting bids to valuations, we take draws from a parametrized type distribution and match moments of the observed data with those generated by simulation. This is simulated GMM, and so inference is relatively standard, which is an advantage relative to the nonparametric case.

We give a brief overview of the approach, following Duffie and Singleton (1993). Let the type distribution F be finitely parameterized by some vector of parameters $\theta \in \Theta$, with θ_0 the true parameter vector. In contrast to our identification arguments, let us think of the data as a long time series, with all the data for each period summarized in a random vector Y_t . Let $Z_t = (Y_t, Y_{t-1}, Y_{t-2} \cdots Y_{t-l})$ for some positive integer $l < \infty$. Let $g_t^* \equiv g(z_t, \theta_0)$ be a known (vector-valued) function of the data, defined for all possible realizations of Z_t . The simulated GMM estimator proceeds by minimizing the distance between the sample average $\frac{1}{T} \sum_{t=1}^T g_t^*$ and the moments predicted by the model $E[g(z_t, \theta)]$:

$$\hat{\theta}_T = \arg \min_{\theta \in \Theta} \left(\frac{1}{T} \sum_{t=1}^T g_t^* - E[g(z_t, \theta)] \right)' W \left(\frac{1}{T} \sum_{t=1}^T g_t^* - E[g(z_t, \theta)] \right) \quad (6)$$

These moments $E[g(z_t, \theta)]$ are constructed by simulating the economy forward. In these simulations it is not necessary to solve for a new set of equilibrium strategies for each evaluation of θ , since these have been pinned down by the first-stage. Thus the moments depend on θ only through the type distribution.³⁰

As Duffie and Singleton (1993) argue, the simulated GMM approach will only have good asymptotic properties under certain conditions. The main challenge is that the moments depend on the parameter vector indirectly through the simulated DGP.³¹ These indirect dependencies may be substantial in our application: perturbations of the parameters describing the distribution of valuations, even holding random simulation draws fixed, may induce large discrete changes in outcomes: e.g., rank reversals in bids or changes in bidders' participation

³⁰In this two-stage approach the econometrician is restricted to moment conditions in which the policy functions are evaluated at θ_0 , although the type distribution is evaluated at θ . This is computationally very convenient and avoids the problems of equilibrium multiplicity that would occur in re-solving for new equilibrium policies at some other θ . But by holding the policy functions fixed and focusing on the “partial derivative” of the objective function in θ , the objective function may become flatter.

³¹Another problem is the *initial conditions problem*: every time the econometrician simulates moments for a particular parameter vector, they must pick an initial condition for the state of the system. This can be overcome by a “burn in” period under Duffie and Singleton’s (1993) stronger ergodicity condition.

choices, both of which may have dynamic effects via the set of survivors in future periods. The solution to this problem is to strengthen the standard ergodicity assumption of Hansen (1982) to geometric ergodicity of the DGP for all $\theta \in \Theta$. Such a condition holds by our Lemma 1. We suspect therefore that the simulated GMM estimator, with thoughtfully chosen moments, should have good asymptotic properties in our environment.³²

4.4 Moment Choice

This leaves the question of which moments to choose. In view of our earlier identification arguments, it may be reasonable to focus on (averages of) individual bidder outcomes; and in particular to look at the behavior of the same bidder over time (since this is what pins down the correlation structure of F_X). For example in each time period one could construct the fraction of bidders bidding on each good and the mean and mean squared bids on each product. Given two successive time periods t and $t - 1$ one could look at the substitution matrix across products, and the correlation in bids by the same bidder, conditioning on product identity.

Many of these functions have a particular form. They can be written as:

$$g(z_t, \theta) = \frac{1}{I_t} \sum_{i \in \mathcal{I}_t} \tilde{g}(z_t^i, \theta)$$

where \mathcal{I}_t is the set of bidders who were present in period t , with cardinality I_t ; and \tilde{g} is an identity-symmetric function of the individual time series outcomes $z_t^i = \{y_t^i, y_{t-1}^i \cdots y_{t-l}^i\}$.³³

An advantage of this formulation is that we can re-write the moments more conveniently:

$$\begin{aligned} E[g(z_t, \theta)] &= \frac{1}{I_t} \sum_{i \in \mathcal{I}_t} \int \tilde{g}(z_t^i, \theta) d\tilde{F}(z_t^i | \theta) \\ &= \int \tilde{g}(z_t^i, \theta) d\tilde{F}(z_t^i | \theta) \quad (\text{by symmetry}) \\ &= \frac{\int \int \int \tilde{g}(z_t^i, \theta) dF_{Z|Y}(z_t^i | y^i) w(y^i) dF_{Y|X}(y^i | x) dF(x | \theta)}{\int \int w(y^i) dF_{Y|X}(y^i | x) dF(x | \theta)} \end{aligned}$$

The distribution \tilde{F} is the ergodic distribution of individual observations, which can be ob-

³²An additional complication is error in the first-stage estimation: the arguments of Andrews (1994) or Ai and Chen (2003) may help here to argue in favor of \sqrt{N} -consistency.

³³If $l > 1$, the function \tilde{g} needs to be well-defined when a bidder wasn't present in some period $t - l$.

tained by simulation of the full economy. The last step separates the problem of simulating moments into three pieces. First, we sample a type from the distribution $F(x|\theta)$. Then we sample a time series y^i for that type using the conditional distribution of time series outcomes $F(y^i|x)$. As argued above, the distribution of time series outcomes for each type is identified from the empirical MDP, and so this is possible without θ . Finally, given an individual time series y^i , we sample z_t^i by drawing a period t according to weights $\psi(z_t^i) = E \left[\frac{1}{I_t} | z_t^i \right]$ and then getting z_t^i from y_t^i (i.e. $F_{Z|Y}$ is a discrete distribution with masses given by ψ).

In performing this decomposition, we have to carefully weight how different individual observations are sampled. Conditional on any bidder's time series, the probability that an observation from period t is chosen is proportional to the inverse of the number of incumbent bidders in period t . This explains the form of $F_{Z|Y}$. Similarly, the probability of sampling a bidder should be proportional to how often their observations would be chosen, with weights $w(y^i) = \sum_{y_t^i \in y^i} \psi(z_t^i(y_t^i))$. Dividing through by the sum of these weights in the denominator, we get a well-defined distribution. The weights are identified directly from the data.³⁴ For example, suppose that z_t^i is a product-bid pair for bidder i at period t . An estimate of $E \left[\frac{1}{I_t} | z_t^i \right]$ is given by averaging the inverse of the number of incumbent players across all instances in which the product-bid pair z_t^i was observed in the data.

What is useful about this decomposition is that only the first piece depends on θ . So one can fix a large sample of types x , simulate a large set of time series outcomes $\{y^i(x)\}$ and corresponding weights $\{w\}$, and calculate $\int \tilde{g}(z_t^i, \theta) dF_{Z|Y}(z_t^i | y^i)$ for each time series observation y^i . Then for each new parameter draw of θ during the optimization, sample those mean outcomes with weights $f(x|\theta)$ to evaluate the objective function. The advantage of this estimation approach is its computational simplicity and the fact that the simulated objective function will be continuous in θ , which will not generally be the case for moments that require simulating the entire auction marketplace.³⁵ This will make numerical optimization easier. This trick may be usefully applied elsewhere: for example, in macro-econometric models with heterogeneous types.

³⁴To see that these are the right weights, consider Figure 2 and the relative probability of sampling any two entries, say y_1^1 and y_4^4 . These are equally likely to be sampled in the ergodic distribution, since periods 1 and 4 are equally likely to be sampled, and each entry in those periods is sampled with probability 0.5. Now go the other way. The probability of sampling bidder 1 is $1/5$; the bidder-weight is $1/2 + 1/3 + 1/2$ and the probability of sampling the first entry is $3/8$, multiplied for a total weight of $1/10$. Similarly for bidder 4 we get corresponding weights of $1/5$, $1/2 + 1/2 + 1/2$, $1/3$, multiplied for a total weight of $1/10$.

³⁵When $J > 1$, it is impossible to "hold the draws fixed": as θ changes, different bidders will win and exit, resulting in discrete changes in simulated behavior for future states. So the simulated objective function will be discontinuous in θ .

4.5 Characteristic Space Approach

Demand estimation in product space is difficult when the product space is large and so we may want to project valuations down onto product characteristics (e.g. McFadden (1974)). To do this, we assume that valuations depend purely on the characteristics of the goods z_j :

$$x_j = z_j \beta_i \tag{7}$$

where $\beta_i \sim F_\beta$. This is the pure characteristics model of Berry and Pakes (2007). Bidders differ in their tastes for the characteristics, and we would like to recover the distribution of the random coefficients F_β .

We assume that the characteristic space Z has dimension K , with $K \leq J$ (i.e. we are projecting down onto a lower dimensional space). An implication of this is that the characteristic-space restriction is testable: given any bidder observed bidding on all J products, we can solve for their valuation x and then ask if there is any β_i that rationalizes their imputed valuations. More practically, it simplifies parametric estimation. Rather than optimizing over a high-dimensional type space, we may instead simulate moments by sampling random coefficients from a known parametric distribution and matching as before. Optimization is now over a lower dimensional vector that parameterizes the distribution F_β over \mathbb{R}^K .

The characteristic-based specification is readily adaptable for applied work. For example, researchers may want to add idiosyncratic iid demand shocks $\varepsilon_{i,j,t}$ or unobserved product heterogeneity $\xi_{j,t}$ to the model. Idiosyncratic demand can be easily accommodated into the simulated GMM approach, as it will integrate out of the mean bids (although it will create “mixing” in the participation decision). Unobserved heterogeneity is more challenging as it will create difficulties in correctly measuring the bid distributions \mathcal{G} and state transitions \mathcal{Q} . It may be possible to introduce these features into our framework and retain nonparametric identification using modern measurement error techniques (e.g. Li and Vuong (1998), Krasnokutskaya (2011) and Hu and Shum (2012)), but this is beyond the scope of this paper.

5 Simulations

In this section we generate a simulated economy and use it in two ways. First, we take the simulated data and run a Monte Carlo simulation to see how our various proposed estimation

approaches work in finite samples. Second, we run a counterfactual to assess the optimal level of supply for a monopolist who controls the supply of one of the products in a two-product economy. We compare the results from the full structural model with a “straw man” reduced-form approach to determining the profit-maximizing supply.

5.1 The Simulated Economy

We consider a simulated economy in which there are just two products. Supply of both products is random, with an equal probability of either 10 or 15 units being supplied each period. There are exactly 120 entrants each period (i.e. E_t is deterministic), and they draw their valuations from a bivariate normal distribution. In the base case their valuations for each good are independent, with mean 200 and standard deviation 100. Winners exit with certainty, and losing bidders survive with probability 0.75, so that there is quite a lot of persistence in the economy.

To find equilibrium participation and bidding strategies, we look for a fixed point so that these strategies are mutual best responses. In theory we would like to allow strategies to be maps from type, the supply level and private history to the reals, but in practice this is computationally infeasible.³⁶ So instead we consider simpler strategies $\gamma(x, n)$ and $\beta(x, n, j)$ that condition only on the supply vector, which has four discrete points of support. We find a fixed point in this strategy space by best response iteration: taking an initial strategy, simulating out an economy governed by these strategies, computing a best response for every type, and repeating until convergence. While in principle there could be multiple equilibria, regardless of our starting point we find only a single equilibrium set of strategies. This gives us a strategy vector (γ^e, β^e) and a set of primitives F_N, F_E, F_X and r as described above.

Table 1 offers summary statistics for the simulated economy at the bid-, auction-, and bidder-level. Average bids are marginally higher than valuations, but only marginally because continuation values are only large for the small population of high-valuation bidders who are likely to win auctions. Conditional on participating in good 1 auctions, one might expect that the average valuation of bidders would be substantially higher than the underlying population mean of 200. This is offset by the difference between the entry distribution and the ergodic distribution; low-valuation bidders are more likely to lose auctions and persist,

³⁶While estimation on the full state space is difficult, finding an equilibrium is vastly more so, as it requires solving for a fixed point in a high-dimensional space.

as evident from the bidder-level data, and therefore push the mean observed bid downwards. The auction-level summary statistics highlight the relationship between supply and directed search. Holding fixed the supply of good 1, an increase in the supply of good 2 draws bidders away seeking lower prices.

5.2 Monte Carlo

We generate 100 Monte Carlo datasets by simulating 300 periods of data from the above data generating process. Since there are on average 25 auctions each period, this implies that the average dataset consists of 7500 auctions. We then apply both the nonparametric and semiparametric estimation approaches to recover the type distribution F_X , estimating the other primitives F_N and r in a first-stage (we assume F_E known since it is deterministic).

The model is only partially nonparametrically identified: even with the supply variation, many types will always choose to participate on the same good. For this reason, the nonparametric estimation approach follows the partial identification proof of Theorem 4, starting with the bidders observed bidding on both goods, identifying the distribution of types of those bidders and then proceeding to get the marginal distributions for bidders who bid only on a single good.

The results are shown in Table 2. We present the result for three groups separately: those who bid on both goods, those who bid on good 1 only and those who bid on good 2 only. For each group, we show their mean μ_i and standard deviations μ_{ii} of their valuations for good i , as well as the covariance in valuations σ_{12} . These statistics θ_0 are generated from the simulated economy; they differ from the parameters of the full joint distribution because they condition on a particular product support.

For types that bid on both items, the bidder-wise identification result applies, and we do a great job of backing out the underlying valuations. The mean point estimates are extremely accurate, and the standard deviations are small. Although this is not shown in the table, we can confirm that the estimates for each individual bidder are also good. We do less well in the estimates of the marginals. Although the mean estimates remain quite good, there is substantially more variation in the parameter estimates. Overall though, this approach performs extremely well without imposing any parametric assumptions.

The semiparametric approach assumes, correctly, that the underlying distribution of types

is multivariate normal. We estimate the five parameters $(\mu_1, \mu_2, \sigma_{11}, \sigma_{22}, \sigma_{12})$ by simulated GMM. The moments g_t we use in estimation are the mean bid on each good in period t , the variance in bids for each good in period t , and the mean covariance in bids between bidders who bid on one product in $t - 1$ and another product at t . Table 3 shows the results of not only the benchmark case but also other simulated economies with different parameters.³⁷

The mean parameter estimates are once again good, suggesting that the estimators are unbiased. But there is considerably more noise in the parameter estimates, particularly in the variance estimates. In some sense this is to be expected. Both the nonparametric and semiparametric estimators will inherit error from the first stage estimates of the conditional distributions \mathcal{G} and transitions \mathcal{Q} . But the semiparametric estimator has the additional problems of simulation and optimization error, which will add noise to the estimates. Nonetheless the results are good: the estimator performs well even when there is positive or negative correlation in valuations, or with higher levels of persistence.

5.3 Counterfactual

One might wonder if it is really necessary to estimate this complicated structural model in order to run counterfactuals. We consider one particular counterfactual and show that the full model is needed. In this counterfactual, an economist is asked by the firm who produces to good 1 to work out their long-run profit maximizing level of per period supply, given that their per unit cost is $c = 207$. That is, they want to solve the problem of maximizing per period profit:

$$\max_{q \in \mathbb{Z}^+} (E_{\mu_q}[p] - c) q$$

where q is the per period supply, p is the price (a random variable since this is an auction), and the expectation is taken with respect to the ergodic measure for an economy in which per period supply is q . To calculate the long-run price distribution for any supply level q , we have to re-solve for the equilibrium strategies and then simulate out the economy. We do this for a range of integer supply levels. Again there is no guarantee that the equilibrium strategies are unique, and therefore in principle for any particular supply level we could have a set of possible counterfactual outcomes, one for each equilibrium. In practice though we find only a single equilibrium at each supply level.

³⁷We have repeated this exercise for the nonparametric estimation procedure, and it performs comparably — we omit those results to conserve space.




Figure 3: **Counterfactual demand and revenues.** The left panel shows the estimated inverse demand curve (i.e. mean price as function of quantity supplied) using a linear model (solid) and the full structural model (dashed). The linear model passes through the mean prices in the observed data (dots). The linear estimates vastly underestimate the quantity elasticity of price. The right panel shows the projected profits according to the linear model (solid) and the full structural model (dashed).

As a straw man, we consider taking a simpler “reduced-form” approach to this optimization problem. Taking advantage of the current exogenous variation in supply, the economist estimates a linear demand system of the form:

$$p_{n,t} = \alpha_0 + \alpha_1 q_t^1 + \alpha_2 q_t^2 + \varepsilon_{n,t}$$

where n indexes auctions of good 1 within period t , and q_t^1 and q_t^2 are the supply of goods 1 and 2 respectively. Using the estimated inverse demand curve $\hat{p}(q) = \hat{E}[p|q, q^2 = 10]P(q^2 = 10) + \hat{E}[p|q, q^2 = 15]P(q^2 = 15)$, one can also solve the optimization problem above.

The results are displayed in Figure 3. In the left panel we plot the inverse demand curve (i.e. mean price as function of quantity supplied) estimated from the linear model and the structural model. The linear model significantly underestimates the quantity elasticity of price. The main problem is not the linear specification, because, as we can see, the true relationship is approximately linear for a large range of quantity levels (between 8 and 17 or so). Instead, the issue is that the linear model gives the average price response to a *transitory* change in quantity, whereas the structural model gives the impact of a *permanent* change in quantity, which is in fact the counterfactual of interest.

Transitory supply changes have no effect on equilibrium strategies. On the other hand, a permanent decrease in supply will increase bids, because bidders will correctly perceive that

their continuation values have fallen and thus bid more aggressively. The converse is true for an increase in quantity supplied. As a result the long-run inverse demand curve is steeper than the short-run curve.

The right panel shows the implications of this. The optimal supply level is between the two previous levels, and is relatively flat between 12 and 14. But using a reduced form approach, the economist would instead advise an increase to a supply of 19 units per period, which in fact would lead to a sharp fall in profits. This comparison makes clear the value of having a dynamic structural model.

6 Conclusion

Dynamics and market interactions are salient in real-world auction markets. Bidder participation choices can create limits for identification: if bidders never bid on a product, we can only hope to put an upper bound on their valuations. Moreover, this participation choice implies an endogenous distribution of competing bidders in any given auction, a distribution that responds to both short-term supply fluctuations and long-term structural changes. Because bidders are persistent, there is also an option value to losing; an option value that is increasing in a bidder's private valuation. Bidders shade their bids by that option value, which creates challenges for recovering their underlying valuation. In addition to these issues, we consider the problem of information leakage, whereby bidders seek to influence the beliefs of other bidders. Taking all of these features into account, we offer an model of demand in an auction platform market.

We show that for any strategies, the evolution of the state of the market is geometrically ergodic. Moreover, we make use of results from the computer science literature on partially observable Markov decision processes to characterize best responses. We combine these results in our general identification approach: Lemma 2 demonstrates that bidders' strategies are identified, and Theorem 2 relates general identification of the model to a set of intuitive and testable assumptions on bidders' strategies. If we further assume that leakage effects are negligible, we are able to offer constructive identification results in Theorem 3 and 4, which hinge on bidder-level invertibility. Theorem 4 in particular demonstrates the limits of identification in a market with directed search.

We believe the identification results here offer a robust framework for estimation in dynamic

auction markets. We offer three estimators that are designed to exploit different features of the results and the unique structure of auction data: a nonparametric estimator that directly implements the constructive identification result, a semiparametric estimator that transforms the problem into a simulated method of moments estimation exercise, and finally a variation on the latter that projects valuations down into characteristic space, an analogue to traditional methods for fixed-price markets.

Auction data offers a unique and rich environment for learning about consumer preferences. The choice to model a dynamic auction marketplace explicitly is driven not just by a desire to “get the model right,” but also because we want a framework that allows us to extract information from within-bidder variation in outcomes, across products and time, as the state of the market changes. We hope that this framework offers a first step in that direction.

References

- ADAMS, C. P. (2012): “Identification in the Roy Model with an Application to Auctions,” Working Paper, Federal Trade Commission.
- AGUIRREGABIRIA, V., AND P. MIRA (2007): “Sequential Estimation of Dynamic Discrete Games,” *Econometrica*, 75(1), 1–53.
- AI, C., AND X. CHEN (2003): “Efficient Estimation of Models with Conditional Moment Restrictions Containing Unknown Functions,” *Econometrica*, 71(6), 1795–1843.
- ALBRECHT, J., P. A. GAUTIER, AND S. VROMAN (2006): “Equilibrium Directed Search with Multiple Applications,” *Review of Economic Studies*, 73, 869–891.
- ANDREWS, D. W. K. (1994): “Asymptotics for Semiparametric Econometric Models Via Stochastic Equicontinuity,” *Econometrica*, 62(1), 43–72.
- ATHEY, S. (2001): “Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information,” *Econometrica*, 69(4), 861–889.
- ATHEY, S., AND P. HAILE (2007): “Empirical Models of Auctions,” in *Advances in Economics and Econometrics, Theory and Applications: Ninth World Congress of the Econometric Society*, ed. by R. Blundell, W. Newey, and T. Persson. Cambridge University Press.

- 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 (2003): “The Winner’s Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions,” *The RAND Journal of Economics*, 34(2), 329–355.
- BERGEMANN, D., AND J. HÖRNER (2010): “Should Auctions be Transparent?,” Discussion paper, Cowles Foundation Discussion Paper 1764.
- BERRY, S. (1994): “Estimating Discrete Choice Models of Product Differentiation,” *Rand Journal of Economics*, 25, 242–262.
- BERRY, S., A. GANDHI, AND P. HAILE (2012): “Connected Substitutes and the Invertibility of Demand,” working paper, Yale University.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63(4), 841–890.
- BERRY, S., AND A. PAKES (2007): “The Pure Characteristics Demand Model,” *International Economic Review*, 48(4), 1193–1225.
- BLUM, J., AND V. SUSARLA (1977): “Estimation of a Mixing Distribution Function,” *Annals of Probability*, 5(2), 200–209.
- BODOH-CREED, A. (2012): “Approximation of Large Dynamic Games,” working paper, Cornell University.
- BURDETT, K., S. SHI, AND R. WRIGHT (2001): “Pricing and Matching with Frictions,” *Journal of Political Economy*, 109(5), 1060–1085.
- CHADE, H., G. LEWIS, AND L. SMITH (2011): “Student Portfolios and the College Admissions Problem,” mimeo.
- DUFFIE, D., AND K. SINGLETON (1993): “Simulated Moments Estimation of Markov Models of Asset Prices,” *Econometrica*, 61(4), 929–952.

- FERSHTMAN, C., AND A. PAKES (2012): “Dynamic Games with Asymmetric Information: A Framework for Empirical Work,” forthcoming in the Quarterly Journal of Economics.
- FOX, J., AND A. GANDHI (2011): “Identifying Demand with Multidimensional Unobservables: A Random Functions Approach,” Working Paper, University of Michigan.
- GALE, D. (2000): *Strategic Foundations of General Equilibrium*. Cambridge University Press.
- GALENIANOS, M., AND P. KIRCHER (2009): “Directed search with multiple job applications,” *Journal of Economic Theory*, 144(2), 445–471.
- GOWRISANKARAN, G., AND M. RYSMAN (2009): “Dynamics of Consumer Demand for New Durable Goods,” Working paper, University of Arizona.
- GUERRE, E., I. PERRIGNE, AND Q. VUONG (2000): “Optimal Nonparametric Estimation of First-Price Auctions,” *Econometrica*, 68(3), 525–574.
- HAILE, P., H. HONG, AND M. SHUM (2006): “Nonparametric Tests for Common Values in First-Price Sealed-Bid Auctions,” NBER Working Paper 10105.
- HANSEN, L. P. (1982): “Large Sample Properties of Generalized Method of Moments Estimators,” *Econometrica*, 50, 1029–1054.
- HENDEL, I., AND A. NEVO (2006): “Measuring the implications of sales and consumer stockpiling behavior,” *Econometrica*, 74, 1637–1673.
- HENDRICKS, K., I. ONUR, AND T. WISEMAN (2008): “Last-Minute Bidding in Sequential Auctions with Unobserved, Stochastic Entry,” Working Paper, University of Texas.
- 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. North-Holland.
- HU, Y., AND M. SHUM (2012): “Nonparametric identification of dynamic models with unobserved state variables,” forthcoming in *Journal of Econometrics*.
- INGSTER, A. (2009): “A Structural Study of a Dynamic Auction Market,” Working Paper, NYU Stern.

- JOFRE-BONET, M., AND M. PESENDORFER (2003): “Estimation of a Dynamic Auction Game,” *Econometrica*, 71(5), 1443–1489.
- KRASNOKUTSKAYA, E. (2011): “Identification and Estimation in Highway Procurement Auctions under Unobserved Auction Heterogeneity,” *Review of Economic Studies*, 71(1), 293–327.
- LI, T., AND Q. VUONG (1998): “Nonparametric Estimation of the Measurement Error Model Using Multiple Indicators,” *Journal of Multivariate Analysis*, 65, 135–165.
- LI, T., AND X. ZHENG (2009): “Entry and Competition Effects in First-Price Auctions: Theory and Evidence from Procurement Auctions,” *Review of Economic Studies*, 76(4), 1397 – 1429.
- MARMER, V., A. SHNEYEROV, AND P. XU (2007): “What model for Entry in First-Price Auctions? A Nonparametric Approach,” Working Paper, Concordia University.
- MCADAMS, D. (2003): “Isotone Equilibrium in Games of Incomplete Information,” *Econometrica*, 71(4).
- MCFADDEN, D. (1974): “Conditional logit analysis of qualitative choice behavior,” in *Frontiers in Econometrics*, ed. by P. Zarembka, pp. 105–142. Academic Press.
- MILGROM, P., AND R. WEBER (2000): “A Theory of Auctions and Competitive Bidding II,” in *The Economic Theory of Auctions*, ed. by P. Klemperer. Edward Elgar Publishing Ltd.
- NEKIPELOV, D. (2007): “Entry Deterrence and Learning Prevention on eBay,” mimeo, UC Berkeley.
- PAARSCH, H. J., AND H. HONG (2006): *An Introduction to the Structural Econometrics of Auction Data*. MIT 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.
- PETERS, M., AND S. SEREVINOV (2006): “An Internet Auction with Many Traders,” *Journal of Economic Theory*, 130, 220–245.

- ROBERTS, J., AND A. SWEETING (2012): “Competition versus Auction Design,” working paper, Duke University.
- ROY, A. D. (1951): “Some Thoughts on the Distribution of Earnings,” *Oxford Economic Papers*, 3, 135–146.
- RUBINSTEIN, A., AND A. WOLINSKY (1985): “Equilibrium in a Market with Sequential Bargaining,” *Econometrica*, 53(5), pp. 1133–1150.
- 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.
- SAID, M. (2009): “Auctions with Dynamic Populations: Efficiency and Revenue Maximization,” Working paper, Yale University.
- SAILER, K. (2006): “Searching the eBay Marketplace,” Working Paper, University of Munich.
- SONDIK, E. (1978): “The Optimal Control of Partially Observable Markov Decision Processes over the Infinite Horizon: Discounted Costs,” *Operations Research*, 26(2), 282–304.
- STOKEY, N. L., R. E. LUCAS, AND E. C. PRESCOTT (1989): *Recursive Methods in Economic Dynamics*. Harvard University Press.
- WEINTRAUB, G., C. L. BENKARD, AND B. VAN 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

By Theorem 11.12 in Stokey, Lucas, and Prescott (1989), uniform geometric convergence in total variation norm will be achieved if their “condition M” holds: $\exists N \geq 1$ and $\varepsilon > 0$ such that for every $A \in \mathcal{B}(\mathcal{S})$, either $[P^N(s, A) > \varepsilon \forall s \in S]$ or $[P^N(s, A^c) > \varepsilon \forall s \in S]$ where $P(s, A) \equiv P_{\gamma, \beta}(s, A)$ is the probability of reaching the set A from state s in a single step, and $P^N(s, A)$ is the N -step transition kernel. We claim the following is sufficient for condition M: there exists some $s_0 \in S$ and some $N \geq 1$ and some $\varepsilon > 0$ such that $P^N(s, s_0) > \varepsilon \forall s \in S$. To prove this, notice that for any $A \in \mathcal{B}(\mathcal{S})$ either $s_0 \in A$ or $s_0 \in A^c$. If the former, then for any $s \in S$, $P^N(s, A) \geq P^N(s, s_0) > \varepsilon$. If the latter, then $P^N(s, A^c) \geq P^N(s, s_0) > \varepsilon$.

Now we must find such an s_0 . Let s_0 be the state where there are no bidders and no supply, and the history of the past t_P periods is null. This state is reachable in $t_P + 1$ steps: it occurs if for $t_P + 1$ periods there has been no supply and in the previous period everyone exited and no-one entered. The probability of this is at least $F_N(0)^{t_P+1}(1-r)^{\bar{I}}$. So the required condition holds with $R = t_P + 1$ and $\delta = F_N(0)^{t_P+1}(1-r)^{\bar{I}}$. \square

Proof of Theorem 1

Rewrite the payoff function at state π after choosing object j^* as:

$$\int_0^b (x_{j^*} - B_j^M) dG_{j^*, \pi}^M(B_j^M) + r \int_b^\infty \int V(\pi') dQ(\pi' | \pi, j^*, b, B_j^M) dG_{j^*, \pi}^M(B_j^M)$$

Take a first order condition in b wherever $G_{j^*, \pi}^M$ has a density $g_{j^*, \pi}^M$:

$$(x_{j^*} - b) g_{j^*, \pi}^M(b) - r g_{j^*, \pi}^M(b) \int V(\pi') dQ(\pi' | \pi, j^*, b, b) \\ + r \int_b^\infty \frac{\partial}{\partial b} \left(\int V(\pi') dQ(\pi' | \pi, j^*, b, B_j^M) \right) dG_{j^*, \pi}^M(B_j^M)$$

Substituting $v(x, j, b, B) = \int V(\pi') dQ(\pi' | \pi, j, b, B)$, setting the result equal to zero and rearranging terms gives the result. \square

Proof of Lemma 2

Fixing a type x and $(\mathcal{G}, \mathcal{Q}, F_N)$, the transitions and per period payoff in the empirical Bellman equation are identified for any actions j and b . This defines a MDP, so standard results (e.g. Rust (1994)) imply the existence of a unique value function $V(x, h, n)$, and a set of optimal policies $\Sigma(\mathcal{G}, \mathcal{Q}, F_N)$. By definition these are the empirical best responses. \square

Proof of Theorem 2

Let \mathcal{F} be the collection of conditional probability measures $\{\mu_{Y|X}(A|x)\}_{A \in \mathcal{B}(\mathcal{Y})}$, and let $\mathcal{L}(\mathcal{F})$ be the space generated by linear combinations of elements of \mathcal{F} . Let $C_0(\mathcal{X})$ be the Banach space of continuous functions on \mathcal{X} which vanish at infinity, under the sup norm. Blum and Susarla (1977) show that the inversion problem is solvable if $C_0(\mathcal{X}) \subseteq \mathcal{L}(\mathcal{F})$.

By assumption 2 and the argument in the main text, the collection of probability measures \mathcal{F} are identified. We must show that assumption 1 implies that $C_0(\mathcal{X}) \subseteq \mathcal{L}(\mathcal{F})$. For each x , pick a set $B(x) \in \mathcal{B}(\mathcal{Y})$ so that the collection of sets $\{B(x)\}_{x \in \mathcal{X}}$ is mutually disjoint and $\mu_{y|x}(B(x)|x) > 0$ and $\mu_{y|x}(B(x)|x') = 0$ if $x' \neq x$. Assumption 1 guarantees this is possible. Elements of $\mathcal{L}(\mathcal{F})$ can be written as $g(x) = \sum_{A \in \mathcal{B}(\mathcal{Y})} \mu_{Y|X}(A|x)w(A)$ for $w(A)$ a function indexed by the Borel sets A . Fix $f \in C_0(\mathcal{X})$. Write f as $f(x) = \sum_{A \in \mathcal{B}(\mathcal{Y})} 1(A = B(x))\mu_{Y|X}(B(x)|x)\frac{f(x)}{\mu_{Y|X}(B(x)|x)}$. But then $f \in \mathcal{L}(\mathcal{F})$ using the weight function $w(A) = 0$ whenever $A \neq B(x)$ for some x , and $w(A) = \frac{f(x)}{\mu_{Y|X}(B(x)|x)}$ otherwise (the weights are indexed by the set $B(x)$, not the type x). Since f was arbitrary, we have $C_0(\mathcal{X}) \subseteq \mathcal{L}(\mathcal{F})$. \square

Proof of Corollary 1

When $J = 1$, we have $b = v$ for every bidder, and so the distribution of bids $G(b) = F_X(x)$ immediately identifies demand. So suppose $J \geq 1$. We show that $C_0(\mathcal{X}) \not\subseteq \mathcal{L}(\mathcal{F})$ so the inversion problem is not solvable. For any type x , in the static game the surplus from (optimally) bidding their valuation for any product j when supply is n is:

$$S(x, j, n) = \int_0^{x_j} (x_j - B_j^M) dG_j^M(B_j^M|n)$$

Since the type distribution is assumed everywhere continuous, this surplus function is also continuous. Pick any type \hat{x} such that $S(\hat{x}, j^*, n) > S(\hat{x}, j, n)$ for $j \neq j^*$. Such a type exists: the set of types who are indifferent between two bidding on two products given any particular supply realization n is of lower dimension than the type space, and since the set of supply vectors is finite, so is their union, implying that the set of types with strict participation preferences is non-generic. Then by continuity, there is an open ball of types $B(\hat{x})$ who make the same participation choices as \hat{x} . The unique outcomes for these types $x \in B(\hat{x})$ are maximally a set of pairs (j, x_j) . So all functions in $\mathcal{L}(\mathcal{F})$, when restricted to the domain $B(\hat{x})$, have the form $\sum_j f_j(x_j)$ for some bounded measurable functions f_j . Continuous functions of the form $f(x) = \prod_{j=1}^J x_j$ are not in $\mathcal{L}(\mathcal{F})$. Thus following Blum and Susarla (1977), the model is not identified from the distribution of time series outcomes F_Y .

Finally, we argue that in the static case that if two type distributions generate the same distribution of time series outcomes F_Y , then there is no additional information in the data to tell them apart (i.e. non-identification by mixtures implies non-identification of the model). Suppose F_X^1 and F_X^2 are distributions over types that generate the same distribution of individual outcomes F_Y . Because all bidders are new entrants and drawn iid, the joint distribution of actions in every period conditional on the number of entrants is generated as an E_t fold repetition of F_Y ; implying that F_X^1 and F_X^2 will generate identical data.

Proof of Lemma 3

Assuming a no leakage equilibrium, taking a derivative in b in the empirical Bellman equation and setting it equal to zero gives the first order condition:

$$b^*(x, j, h, n) = x_j - rE[V(x, h', n')|h, n]$$

So $b^*(x^2, h, j, n) - b^*(x^1, h, j, n) = (x_j^2 - x_j^1) - r(E[V(x^2, h', n')|h, n] - E[V(x^1, h', n')|h, n])$. Since x^1 can copy x^2 and guarantee an identical payoff except when bidding on item j , we have $0 \leq E[V(x^2, h', n')|h, n] - E[V(x^1, h', n')|h, n] \leq x_j^2 - x_j^1$, implying the previous expression is positive. Similarly, $b^*(x^2, h, k, n) - b^*(x^1, h, k, n) = -r(E[V(x^2, h', n')|h, n] - E[V(x^1, h', n')|h, n]) \leq 0$. This proves bid monotonicity. For participation, we show that if $\gamma(x^1, h, n) = j$, then $\gamma(x^2, h, n) = j$. Suppose not. Then x^2 improves by bidding on some product k . But then x^1 has a profitable single-period deviation as he can bid on product k also, and return to his original strategy (under no leakage this is possible). Contradiction. \square

Proof of Theorem 3

Necessity is obvious: if a bidder is observed bidding on only a subset of the products, their valuation for the remaining products can only be bounded above (they chose not to bid on it), rather than point identified. For sufficiency, let b be a J -vector of bids with at least one bid on product j , where the j th element is a bid on a product of type j . Moreover, let (h_j, n_j) be the associated state of the world when that bid was observed. Under NL, bids are equal to valuations plus the discounted continuation value of the bidder:

$$x_j = b_j + \nu_j \quad (8)$$

– where ν is defined to be a J -vector of continuation values for the associated bids with $\nu_j = r \int V(x, h', n') dQ(h', n' | h_j, n_j)$. Define the following mapping $T : \mathbb{R}^J \rightarrow \mathbb{R}^J$,

$$T(\nu) = r \int V(b + \nu, h', n') dQ(h', n' | h, n) \quad (9)$$

This mapping is identified because the empirical best responses are identified. We use Blackwell's sufficient conditions to demonstrate that the mapping T is a contraction. First is monotonicity. Consider two vectors ν and ν' , with $\nu' \geq \nu$ componentwise. We need:

$$r \int V(b + \nu, h', n') dQ(h', n' | h, n) \leq r \int V(b + \nu', h', n') dQ(h', n' | h, n) \quad (10)$$

Since b is fixed, $b + \nu \leq b + \nu'$. Then, because a componentwise higher type can always do at least as well as a componentwise lower type by mimicking their strategy, we have $V(b + \nu, h, n) \leq V(b + \nu', h, n)$ everywhere on $\mathcal{H} \times \mathcal{N}$. Since Q is a distribution on that space and r is a positive constant, we have monotonicity. Next is discounting. We need:

$$r \int V(b + \nu + a, h', n') dQ(h', n' | h, n) \leq r \int V(b + \nu, h', n') dQ(h', n' | h, n) + \rho a \quad (11)$$

–for $\rho \in (0, 1)$ and $a > 0$ a J -vector with equal elements. However, a type with valuation raised by a can do no better than win immediately, implying an upper bound of $T(V + a) \leq TV + ra$ componentwise. Thus discounting holds with $\rho = r$. Applying the contraction mapping theorem we conclude that there is a unique fixed point ν . Combining this with the observable bid vector b , we use equation (8) to recover the bidder's type x . \square

Proof of Lemma 4

Define $P(A, x, h, n)$ to be the probability that a bidder who, starting from state (h, n) , is subsequently observed bidding *only* on a subset of products in A , and $q(x, h, n)$ to be the probability that a bidder of type x exits at (h, n) , whether by winning or losing. Let $Q(h', n'|x, h, n)$ be the transition operator induced by equilibrium play (i.e. substituting $\gamma(x, h, n)$ for j and $\beta(x, h, n, j)$ for b in $Q(h', n'|h, n, j, b, B_j^M > b)$). We can express $P(A, x, h, n)$ recursively:

$$P(A, x, h, n) = 1(\gamma(x, h, n) \in A) \cdot \left[q(x, h, n) + (1 - q(x, h, n)) \int P(A, x, h', n') dQ(h', n'|x, h, n) \right]$$

Since $q(x, h, n) \in (0, 1) \forall (x, h, n)$, it is easy to see that this recursion satisfies Blackwell's sufficient conditions for a contraction mapping, which we can exploit to obtain $P(A, x, h, n)$. Now define $\kappa(h, n)$ to be the ergodic measure with respect to histories and supply in which entry is possible (i.e. excluding states where the number of incumbents exceeds $\bar{I} - \bar{E}$). Then the probability of a bidder x being observed in all, and only in, auctions for products A can be written:

$$P(A|x) = \int P(A, x, h, n) d\kappa(h, n) - \sum_{B \subset A} P(B|x)$$

This follows since the probability of seeing bids for every product j 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 . One can construct the summation on the right-hand side of the final equation by beginning with singleton subsets of \mathcal{J} and proceeding iteratively. \square

Proof of Theorem 4

We proceed by induction on the cardinality a of the subsets $A \in \mathcal{A} \setminus \emptyset$. Our base step is the case $a = J$, so we prove that $F_{\mathcal{J}}$ is identified. The induction step is that if all of the distributions in $\{F_A\}_{\{A:|A|>a\}}$ are identified, then distributions in $\{F_A\}_{\{A:|A|=a\}}$ are identified.

Base Step:

By Theorem 3, we can identify the types of bidders who are observed bidding at least once on every product in \mathcal{J} , and get the distribution of types for these bidders. It will be easier to work with densities rather than distribution functions. Let the density of types for this

sample be $\tilde{f}_{\mathcal{J}}$. We have $\tilde{f}_{\mathcal{J}}(x) = P(\mathcal{J}|x)f_{\mathcal{J}}(x)$. By Lemma 4, $P(\mathcal{J}|x)$ is identified. So $f_{\mathcal{J}}(x) = \frac{\tilde{f}_{\mathcal{J}}(x)}{P(\mathcal{J}|x)}$ is identified (the denominator is non-zero since by assumption $J(x) = \mathcal{J}$).

Inductive Step:

Assume that all of the distributions in $\{F_A\}_{\{A:|A|>a\}}$ are identified. We want to show that distributions in $\{F_A\}_{\{A:|A|=a\}}$ are also identified. Fix a set A with cardinality a . Consider the set of bidders who are observed bidding only on products in A . This set consists both of types with product support A , and those with product support $J(x) \supset A$. An easy corollary of Theorem 3 is that the restricted types x_A are identified for those types with product support A . The argument follows the proof of Theorem 3, except that the mapping T defined in equation (9) is amended so that the vector $b + \nu$ on the RHS is equal to $b + \nu$ for products in A and 0 otherwise, so that T is now a mapping $T : \mathbb{R}^{|A|} \rightarrow \mathbb{R}^{|A|}$. This result correctly identifies x_A for these bidders, since the valuations of goods outside of the product support make no contribution to the continuation values (they never bid on these goods) and so setting them to zero has no effect on x_A .

In view of this, let the inversion mapping $\xi : (\mathbb{R} \times \mathcal{H} \times \mathcal{N})^{|A|} \rightarrow \mathbb{R}^{|A|}$ be implicitly defined as the outcome of the contraction mapping approach outlined in the paragraph above. Applying the inversion map ξ to the set of bidders observed bidding only in products in A , we get a density $\tilde{f}_A(x_A)$. For each A , partition x into x_A and x_{-A} with joint density $f(x_A, x_{-A})$. Then any type x with $J(x) = A$ makes a contribution of $\int P(A|x)f(x_A, x_{-A})dx_{-A}$ to this density. But some types with $J(x) \supset A$ are incorrectly classified as having restricted type $x_A = (\xi \circ \beta)(x)$ so that inversion of bids yields the density $\int P(A|x)1((\xi \circ \beta)(x) = x_A)dF(x)$. This can be decomposed

$$\begin{aligned} \tilde{f}_A(x_A) = & \frac{1}{s(A)} \left(\int P(A|x)1(J(x) = A)f(x_A, x_{-A})dx_{-A} \right. \\ & \left. + \sum_{B \supset A} \int P(A|x)1(J(x) = B)1((\xi \circ \beta)(x) = x_A)dF(x) \right) \end{aligned}$$

for $s(A)$ the share of bidders observed bidding only in set A . Re-arranging terms:

$$\begin{aligned} \int P(A|x)1(J(x) = A)f(x_A, x_{-A})dx_{-A} = & s(A)\tilde{f}_A(x_A) \\ & - \sum_{B \supset A} \int P(A|x)1(J(x) = B)1((\xi \circ \beta)(x) = x_A)dF(x) \end{aligned}$$

Now notice that $P(A|x)$ is constant in x_{-A} whenever $J(x) = A$, so write $P(A|x_A)$ for the probability of being seen bidding only on products in A with restricted type x_A and product support A . Then $\int P(A|x)1(J(x) = A)f(x_A, x_{-A})dx_{-A} = P(A|x_A) \int 1(J(x) = A)f(x_A, x_{-A})dx_{-A}$. And $f_A(x_A) \propto \int 1(J(x) = A)f(x_A, x_{-A})dx_{-A}$ where the normalizing constant comes from the constraint $\int f_A(x_A)dx_A = 1$. So we get:

$$f_A(x_A) \propto \frac{s(A)\tilde{f}_A(x_A) - \sum_{B \supset A} \int P(A|x_B)1(J(x) = B)1((\xi \circ \beta)(x) = x_A)dF(x)}{P(A|x_A)}$$

The denominator is identified by Lemma 4. By change of variables $\int P(A|x_B)1(J(x) = B)1((\xi \circ \beta)(x) = x_A)dF(x)$ is identified for all $B \supset A$ whenever $\{F_B\}_{\{B:|B|>a\}}$ are identified. \square

Table 1: Summary Statistics: Simulated Auctions of Two Goods

	Supply Vector			
	(15,15)	(15,10)	(10,15)	(10,10)
Bid-level data				
Avg valuation of bidders (good 1)	200.8351	200.7476	200.2437	200.4063
Avg bid (good 1)	200.4295	200.3533	199.8324	200.0033
Avg valuation of bidders (good 2)	200.7935	200.3346	200.6366	200.3604
Avg bid (good 2)	200.4026	199.9120	200.2481	199.9541
Auction-level data				
Avg bidders per auction (good 1)	13.4900	14.4689	18.7596	20.4310
Avg winning bid (good 1)	212.3141	212.3323	212.9482	213.0141
Avg bidders per auction (good 2)	13.4735	18.7511	14.4353	20.1094
Avg winning bid (good 2)	212.2328	213.0196	212.2590	213.0645
Bidder-level data				
Avg bidder life span	3.3630			
Avg bidder life span ($v_1 < \mu_1$)	3.6443			
Avg bidder life span ($v_1 > \mu_1$)	3.0822			

Summary statistics for the benchmark simulated economy where the mean valuations of entrants for good 1 and good 2 are 200, variances are 100, and covariance is 0. An observation in bid-level data is a simulated bid and underlying valuation. An observation in auction-level data is an individual auction. An observation in the bidder-level data is a bidder. The four columns in the first two panels correspond to data from periods with the four possible supply vectors: (15,15), (15,10), (10,15), (10,10).

Table 2: Monte Carlo Results: Nonparametric estimator

	μ_1	μ_2	σ_{11}	σ_{22}	σ_{12}
Bidders who bid on both items					
θ_0	197.6946	197.6906	66.6438	66.6442	4.9456
Average $\hat{\theta}$	198.0420	197.8758	67.4228	67.5580	2.7850
Std deviation of $\hat{\theta}$	0.2265	0.2199	1.6294	1.4711	1.1305
Bidders who bid only on good 1					
θ_0	215.8624	197.7556	17.3298	70.7754	5.2944
Average $\hat{\theta}$	215.4231		19.4245		
Std deviation of $\hat{\theta}$	0.1092		0.7490		
Bidders who bid only on good 2					
θ_0	197.7642	215.8568	71.0065	17.2577	5.1845
Average $\hat{\theta}$		215.4171		19.4129	
Std deviation of $\hat{\theta}$		0.0930		0.5608	

Monte Carlo simulations of the nonparametric estimation procedure. Each dataset consists of 300 periods of data from a simulated economy with 120 entrants each period, supply of each good identically and independently chosen at random from the set $\{10, 15\}$, survival rate $r = 0.75$ and valuations multivariate normal with parameters $\mu_1 = \mu_2 = 100$; $\sigma_{11} = \sigma_{22} = 100$; $\sigma_{12} = 0$. The nonparametric estimation procedure is used to estimate the joint distribution of bids for types who bid on both good; and the marginal distributions for those who bid only on a single good. Moments of these estimated distributions are recorded. The mean and standard deviations of these estimates across 100 repetitions are reported, as well as the truth for each group of types.

Table 3: Monte Carlo Results: Semiparametric estimator

	μ_1	μ_2	σ_{11}	σ_{22}	σ_{12}
Benchmark					
θ_0	200	200	100	100	0
Average $\hat{\theta}$	199.8262	199.7852	100.7446	100.1809	0.7720
Std deviation of $\hat{\theta}$	0.5490	0.4636	3.3925	3.2494	1.9159
Negative Correlation					
θ_0	200	200	100	100	-20
Average $\hat{\theta}$	199.5762	199.7759	99.3460	99.8421	-18.8666
Std deviation of $\hat{\theta}$	2.4388	0.7650	5.7142	5.8649	4.5680
Positive Correlation					
θ_0	200	200	100	100	20
Average $\hat{\theta}$	199.4154	198.5935	99.2190	98.1095	18.8579
Std deviation of $\hat{\theta}$	0.7029	3.4670	4.2294	5.0055	2.6607
Higher Persistence					
θ_0	200	200	100	100	0
Average $\hat{\theta}$	199.6549	199.6274	100.6775	100.7262	0.3364
Std deviation of $\hat{\theta}$	0.5059	0.4954	3.0815	2.9143	1.6271
Asymmetric Variance					
θ_0	200	200	100	80	0
Average $\hat{\theta}$	199.2110	199.2865	98.8734	79.9131	-0.0443
Std deviation of $\hat{\theta}$	2.1461	2.4859	4.1870	4.0724	1.8067

Monte Carlo simulations of the semiparametric estimation procedure. The benchmark dataset consists of 300 periods of data from a simulated economy with 120 entrants each period, supply of each good identically and independently chosen at random from the set $\{10, 15\}$, survival rate $r = 0.75$ and valuations multivariate normal with parameters $\mu_1 = \mu_2 = 100$; $\sigma_{11} = \sigma_{22} = 100$; $\sigma_{12} = 0$. In the negative (positive) correlation case, the covariance is $\sigma_{12} = \pm 20$. In the asymmetric variance case, we have $\sigma_{22} = 80$. In the higher persistence case, we assume $r = 0.8$. The parameters are estimated by simulated method of moments. The mean and standard deviations of these estimates across 100 repetitions are reported, as well as the truth for each group of types.