

Dynamic Demand Estimation in Auction Markets*

Matthew Backus
Columbia and NBER

Gregory Lewis
Microsoft Research and NBER

June 15, 2016

Abstract

Economists have developed empirically tractable demand systems for fixed price markets. In contrast, empirical auction techniques treat each auction in isolation, ignoring market interactions. We provide a framework for estimating demand in a large auction market with a dynamic population of buyers with unit demand and heterogeneous preferences over a finite set of differentiated products. We offer an empirically tractable equilibrium concept under which bidders behave as though they are in a steady-state, characterize bidding, and prove existence of equilibrium. Having developed a demand system, we show that it is non-parametrically identified from panel data, and that this result is robust to typical data limitations, reserve prices, random coefficient demand, public signals that refine beliefs about market conditions, unobserved heterogeneity, idiosyncratic preferences, and random latent outside options. We apply the model to estimate demand and measure consumer surplus in the market for compact cameras on eBay. Our analysis highlights the importance of both dynamic bidding strategies and panel data sample selection issues when analyzing these markets.

*This revision is a substantial departure from our earlier working papers, 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 comments, and to conference and seminar participants at Columbia, FTC Applied Micro, Haas, Harvard/MIT, IOOC, Johns Hopkins, LSE, Michigan, Microsoft Research, 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

Most goods and services are sold at fixed prices. Yet auctions are used as the allocation mechanism in a wide variety of contexts, including procurement contracts, treasury bills and the granting of oil drilling and spectrum rights. Technology companies such as Facebook, Google and Microsoft sell advertisers access to online consumers through display and search advertising auctions. And though the majority of eBay’s revenue now comes from fixed price listings of new goods, they still sell a large number of goods (both in absolute and dollar terms) by auction.

These markets share many common features. Buyers get multiple purchase opportunities over time, either for exactly the same product (e.g. in search advertising where a product is a keyword), or for close substitutes (e.g. in treasury bill auctions). This allows bidders to intertemporally substitute, making participation and bidding decisions in light of the option value of waiting for future purchasing opportunities. Both bids and participation choices will reflect individual-specific preferences over the heterogeneous products available. For example, in highway procurement, Lewis and Bajari (2011) document matching between contractors and contract based primarily on distance and contract size, while in online labor markets, employers are more likely to award contracts to workers from their own country (Krasnokutskaya et al., 2016).

In the fixed price context, an influential discrete choice demand estimation approach has been developed that models how buyers with unit demand and heterogeneous preferences decide which goods to purchase (e.g. Berry (1994), Berry et al. (1995), Gowrisankaran and Rysman (2009)). However, until recently, there has been no analogous work in the empirical auctions literature, which has generally focused on the analysis of a repeated cross-section of static auctions for identical products. The goal of this paper is to fill this gap.

A “whole-market” approach that incorporates dynamics and product differentiation can be important for policy counterfactuals. For example, in setting optimal reserves, a static model of second-price sealed bid auctions would suggest that when a monopolist changes her reserve price policy, she should expect no change in buyer behavior (it is weakly dominant for a buyer to bid their valuation regardless). By contrast, in our dynamic setting, a new reserve price policy will affect buyer’s continuation values, and therefore bids throughout the marketplace. Reserve prices will also have an effect on the endogenous set of incumbent bidders. And without taking product differentiation seriously, we cannot use auction data to

say much about cross-price elasticities, or how an increase in the supply of one good would affect the transaction prices of other goods (see Newberry (2015) for evidence that increased supply depresses prices on eBay).

The present paper analyzes a model that is very closely related to the above work on fixed price markets, but with a data-generating process that is given by equilibrium play in a dynamic auction market. Vis-à-vis that literature, our approach is analogous to the second-choice demand estimation approach of Berry et al. (2004) because we use bids placed by the same bidder on different products. This is informative about the *joint* distribution of valuations or, in the parlance of the fixed-price market literature, cross-elasticities. The auction format also allows us to credibly identify a region of the distribution of valuations that we often struggle with: the right tail. This region can be hard to identify in fixed-price markets because of limited variation in characteristics (in particular, prices); for the subset of buyers whose choice behavior doesn't change over the support of observed prices, identification is entirely parametric. To put it more succinctly, we can't invert an inequality, but we can invert bids.

Here we offer what we hope is a comprehensive analysis of our model: we characterize participation and bidding decisions, prove existence of equilibrium, offer a number of non-parametric identification results, and apply the model to measure consumer surplus in the market for compact cameras on eBay. Along the way, we deal with many of the methodological issues that are likely to arise in similar models of auction markets.

Specifically, we consider a model of repeated second-price auctions of differentiated products, in which buyers with unit demand and heterogeneous, perfectly persistent and privately known multidimensional valuations for a finite set of types of good must work out when to participate and how much to bid. Persistence of preferences is a realistic feature of many auction markets, but causes game theoretic problems that do not arise in fixed price markets. So-called "leakage effects" are a particular concern, whereby a bidder's bid today may reveal their valuation to other bidders, affecting the evolution of future play (this remains true even when other bidders cannot observe rival bids; less information just implies a harder filtration problem for the bidders and the outside analyst).

When markets are "small" (in the sense that there are few bidders in the market), the approach of the dynamic games literature has been to simply assume that everything persistent is commonly known, and that all private shocks are iid (see e.g. Jofre-Bonet and Pesendorfer

(2003), Aguirregabiria and Mira (2007), Bajari et al. (2007), Pakes et al. (2007); Fershtman and Pakes (2012) is a recent exception). But in “large” markets (i.e. many anonymous bidders), it is attractive to instead assume that players form beliefs about their rivals in a simpler way. Here we assume that bidders simply take the distribution of rival bids for each product as given and optimize against this, forming neither more refined beliefs (by conditioning on other information available to them) nor considering the possibility of trying to manipulate the evolution of this distribution with their bid. An equilibrium is a fixed point in which bidders optimize given their beliefs about rival bids, and those beliefs are consistent with the long run time-averages of equilibrium play, i.e. the stationary distribution. This is similar to ideas elsewhere in the literature, such as the model of belief formation in Krusell and Smith (1998) and the oblivious / mean-field equilibrium concept of Weintraub et al. (2008) and Iyer et al. (2014).

In equilibrium, losing bidders “win” the opportunity to play the game again tomorrow. The optimal strategy requires that bidders bid their valuation less their continuation value whenever this is positive, and otherwise not participate. A bidder’s valuation is thus identified by adding their bid to their continuation value. In other words, to estimate demand, one must estimate continuation values. We show that these can be learned from observing a bidder’s “pseudo-type” — their optimal bid on each product — as well as the distributions of competing bids and the exit and supply processes.¹

The existence of an inversion from bids to types is often sufficient to prove identification in static auctions (e.g. Guerre et al. (2000), Athey and Haile (2007)): when each bidder bids exactly once, the distribution of pseudo-types is just the distribution of bids, which is observed. But when bidders bid multiple times, they may win (and therefore exit) before recording a “complete,” or invertible, history consisting of a bid on every product (i.e. before their pseudo-type is known). The probability of reaching a complete history varies with type, so we need to correct for selection in measuring the distribution of pseudo-types. By combining this new selection correction idea with the familiar inversion approach, we are able to bound demand. Our identification result is pointwise up to the set of bidders who prefer not to participate in some auctions.

In subsequent sections, we extend the model to cases with limited data, reserve prices,

¹A common refrain in the empirical auctions literature is that losers’ payoffs are “normalized to zero.” We view this as a strong assumption, rather than a normalization. In our model the continuation value that bidders anticipate when they lose is not constant across bidders, but instead depends on their private type.

random coefficient demand, public signals that refine beliefs about market conditions, unobserved heterogeneity, idiosyncratic preferences and random latent outside options. Most of these extensions are reasonably straightforward given the tools developed earlier, as well as results on deconvolution used in Li and Vuong (1998) and Krasnokutskaya (2011). The main exception is the case with outside options and endogenous exit, which is harder. There we show that observable differences in the exit rate across bidders imply a first order separable differential equation that can be used to identify differences in their outside options.

The final section of the paper applies this framework to data from eBay’s compact camera market. We document that substitution over time and across cameras with different levels of resolution occurs in this market, and that there is both observable and unobservable heterogeneity within cameras of a given resolution. We model preferences as a linear combination of a random taste for a camera and a random taste for camera resolution, and estimate the distribution of these random coefficients. We show that our estimated model is able to match the data well, and then use it to calculate consumer surplus. We find that consumer surplus is twice as large as we would have estimated had we ignored continuation values, \$6.04 per bidder rather than \$2.73. We chose the application for its expositional value rather than any policy implications: it gives us the opportunity to show how our methodology can be implemented with limited data; how it fares under unobserved heterogeneity, and to illustrate that dynamics and substitution are important for understanding auction markets.

Literature Review. The paper is related to various strands of literature. Jofre-Bonet and Pesendorfer (2003) was the first paper to attack estimation in a dynamic auction game. Subsequent to this, a number of papers have looked at dynamics on the eBay platform specifically. Budish (2008) examines the optimality of eBay’s market design with respect to the sequencing of sales and information revelation. In Said (2009), the author investigates efficiency and revenue maximization in a similar setup through the lens of dynamic mechanism design. Zeithammer (2006) developed a model with forward-looking bidders, and showed both theoretically and empirically that bidders shade down current bids in response to the presence of upcoming auctions of similar objects. Ingster (2009) develops a dynamic model of auctions of identical objects, and provides equilibrium characterization and identification results. Sailer (2006) estimates participation costs for bidders facing an infinite sequence of identical auctions. Nekipelov (2007) analyzes a model where bidders attempt to prevent learning by late bidding, while Hopenhayn and Saeedi (2016) develop and estimate a model

in which bidder's bidding opportunities and valuations evolve dynamically within an auction according to a Markov process. Bodoh-Creed et al. (2013) employ our methodology in analyzing the optimal fee structure for the eBay platform. Hendricks and Sorensen (2014) offer a model, similar to our own but in continuous time and with a single product type, and analyze the efficiency of the eBay trading mechanism.

Relative to this literature, our main contribution is the focus on sequential auctions of heterogeneous objects, where bidders have multidimensional persistent private valuations. In short, we are focused on developing a demand system. This is also a topic of interest in other recent papers. Adams (2009) examines the problem of nonparametric identification when differentiated products are sold by simultaneous auction, and Krasnokutskaya et al. (2014) develop a model of participation and bidding in an online labor market. Our paper differs from these papers in allowing buyers to participate repeatedly and dealing with the dynamic issues this creates. One of the advantages of using a dynamic model is that we are able to make use of panel data in estimation, rather than treating it as a repeated cross-section, which allows identification of individual bidder preferences.

Paper Structure. The paper proceeds as follows. In the next section we introduce the basic model, and prove existence of and characterize equilibrium. Sections 3 and 4 discuss non-parametric identification of the basic model and a series of extensions. Section 5 presents our application, while Section 6 concludes.

2 Model and Equilibrium Analysis

We consider a market in which competing products are sold by second-price sealed bid auctions. These auctions are held in discrete time, with either zero or one good auctioned per period over an infinite horizon. Since our focus is on demand, we assume for simplicity that supply is random and exogenous. Bidders have unit demand, and enter the market with private and perfectly persistent valuations for each of the objects. They are inattentive, and are active in any particular period with constant probability. When active, they choose whether or not to participate in the current auction, and how much to bid. We show that they bid their valuation less their continuation value, and assume they assess the latter based on the steady-state distribution of supply and competing bids, rather than on current market

conditions (e.g. the number of competing bidders in the current auction). Winning bidders exit the market with certainty, while losing bidders exit with constant probability.

We have chosen this set of assumptions to match some features of the market for digital cameras on eBay, which is our empirical application. In any eBay category (such as digital cameras), there are many different products sold by auction to a large number of anonymous buyers.² Although these auctions typically last for many days, and thus overlap — so that at any given point in time there are many auctions occurring simultaneously — they finish at different ending times, in sequence. As Bajari and Hortacısu (2004) and Hendricks and Porter (2007) have noted, this timing, combined with the way the proxy bidding system works, imply that eBay is well approximated as sequence of second-price sealed bid auctions.

2.1 Environment

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

Supply. There are J distinct kinds of goods sold in a market, indexed by $j = 1 \dots J$. We denote the set of products by \mathcal{J} . In each period t , a good may be available for purchase. Supply is exogenous and Markov, with the current product j_t drawn from a stationary multinomial distribution conditional on the lagged product j_{t-1} . We allow for the possibility that no product is available in a given period, and so supply can be summarized by a square transition matrix Q of size $J + 1 \times J + 1$, where the entry $Q_{j,k}$ gives the probability product k will be supplied next when j is currently offered (and the last row and column give the cases where no product is offered now and will be offered later, respectively). We assume moreover that the multinomials have full support, so that regardless of what was supplied at $t - 1$, it is possible that any of the J products (or nothing) is supplied at t . When a good is available, it is sold by second-price sealed-bid auction.

Demand. At the beginning of each period, E_t buyers enter the market, where E_t is sampled independently over time from a distribution F_E with support $\{0, 1, 2 \dots \bar{E}\}$. Each buyer has unit demand, and has a *perfectly persistent* value for each of the goods summarized by a privately known vector of valuations $\mathbf{x} = (x_1, x_2 \dots x_J) \in \mathbb{R}^J$. This value is drawn iid across buyers on entry according to a distribution \mathbf{F} with strictly positive density over its support

²eBay hides the identity of the bidders by replacing parts of the username with asterisks.

$\mathcal{X} = [0, \bar{x}]^J$. The new entrants combine with the population of buyers from previous periods to form a cohort of bidders of size N_t . To ensure that the maximum number of buyers \bar{N} is bounded, we assume that when $N_t = \bar{N}$, no-one enters ($E_t = 0$).

New entrants are always active; the remaining incumbent members of the cohort are active in each period with iid chance τ . Active buyers participate in the current auction (if one is held), observing the product currently under auction, and placing a weakly positive bid $b \in \mathbb{R}^+$. We assume that the good sells only if the highest bid is *strictly* positive, so a buyer can always opt-out of an auction by bidding zero. Agents are risk-neutral and have quasi-linear utility, receiving a total payoff of $x_j - p$ for buying a good j at price p , and zero otherwise. If an auction is held and a bidder wins the auction, they exit the market. All other active bidders exit with probability $(1 - r)$ (i.e. $r \in (0, 1)$ is the survival rate). Inactive bidders do not exit. We assume that agents do not discount future payoffs, although the exogenous exit probability effectively leads to discounting. We discuss bidder beliefs and strategies in the next section.

2.2 Analysis

We begin our analysis by looking at the buyer’s incentives. In our application to the compact camera market on eBay, it seems reasonable to assume that buyers have simple models of the competition they face. We formalize this idea by assuming that bidders believe that the distribution of the highest competing bid is equal to the historical average and best respond accordingly (this assumption or some variant is implicitly used in the existing literature (Ingster, 2009; Hendricks and Sorensen, 2014)). Let B_j^1 be a random variable denoting the highest bid in an auction for good j , and let G_j^1 be its cumulative distribution function.

Assumption 1 (Beliefs about competing bids). *Following any history of play, bidders believe that the highest rival bid in an auction for good j has distribution G_j^1 .*

This is a sensible way for bidders to form beliefs, though it is non-standard. Agents in games are usually modeled as Bayesian, and so would condition on the recent bidding history and any other observables in their information set (e.g. the number of competing bidders in the current auction) to update their prior and form a posterior over the current distribution of types in the market.³ We have elected to model bidder beliefs in this way both because

³Even if this data was hidden a bidder could learn from their own experience over time (e.g. a bidder

it substantively simplifies the game theory, and because the gap between the behavior of sophisticated Bayesian bidders and those that obey assumption 1 is probably small.

Just as information about competitor valuations has no effect on the optimal bid in a static second-price auction, knowing the current distribution of types in the market is useful only to the extent that it predicts *future* competition, which determines continuation values. Since bidders participate only intermittently and eBay is a large market with high turnover of market participants, current market conditions have little predictive power for future competition. Assumption 1 is thus a good approximation to the behavior of Bayesian bidders in this marketplace, though the model can be extended to weaken this assumption and allow bidders to condition on other state variables (see section 4.3.2).

Under this assumption, the only payoff relevant-information in each period is the bidder’s type and the current object under auction (we show this below). So without loss of generality a pure strategy is just a vector $\beta : \mathcal{X} \rightarrow \mathbb{R}^J$, a bid on each product as a function of type, with j th coordinate β_j . Moreover, let $\mu_j^\beta(\cdot)$ denote the ergodic measure of the first-order statistic of bids, which depends on β .⁴ We can think of this measure as capturing the time-average we would compute if we picked a β and simulated the market for an arbitrarily long time. We are now in a position to define what it means for the market to be in equilibrium (subject to our restriction on beliefs):

Definition 1 (Equilibrium). *A pure strategy equilibrium is a tuple $(\beta^e, \{G_j^1\})$ such that*

(Optimality) $\beta^e(\mathbf{x})$ is a best response for type \mathbf{x} given beliefs $\{G_j^1(b)\}_{j \in \mathcal{J}}$, and

(Consistency) for every j , $G_j^1(b) = \mu_j^{\beta^e}([0, b])$.

The definition mirrors the requirements for a Bayes–Nash equilibrium: play must be optimal given beliefs, and beliefs must be consistent with play. The consistency requirement explicitly connects beliefs, as specified by Assumption 1, to the distribution of the highest bid on each product implied by rivals’ strategies.

who lost yesterday might reason that competition was fierce yesterday, and may still be so today). This gives rise to a winner’s curse effect (Budish, 2008). There is some evidence that this kind of learning occurs on eBay: Coey et al. (2015) show that bidders tend to (modestly) increase their bids in subsequent auctions for the same product, although they attribute this to deadlines rather than learning.

⁴In an earlier version of this paper, we proved that for any β , such an ergodic measure exists, is unique and converged to at geometric rate - details available on request. The speed of convergence makes us more confident that the data is drawn from the recurrent ergodic class.

We can now specify the buyer's decision problem. It is dynamic, since losers may have an opportunity to bid at a later date. The relevant state variables are just the buyer's type and the product-type currently being auctioned, since Assumption 1 implies that this is sufficient for beliefs about the highest competing bid (which together with the bid determines allocations and payments). These state variables are Markov, and so the buyer faces a Markov decision problem.

Define the state transition matrix $\tilde{Q} \equiv \sum_{s=1}^{\infty} \tau(1-\tau)^{s-1} Q^s$. This is the distribution over products a bidder bidding on j today expects to see when they are next active. Notice that as $\tau \rightarrow 0$ (so that bidders are infrequently active), each row of \tilde{Q} will converge to the steady-state distribution of supply, which we denote by π . Let $v_j(\mathbf{x})$ be the continuation value of a bidder of type \mathbf{x} who is active and bidding on product j . Then we can write down a Bellman equation:

$$v_j(\mathbf{x}) = \max_{b \in \mathbb{R}^+} \int \left(1(b \geq B_j^1 \vee b > 0)(x_j - B_j^1) + 1(b < B_j^1 \vee b = 0)r \sum_k \tilde{Q}_{j,k} v_k(\mathbf{x}) \right) dG_j^1(B_j^1). \quad (1)$$

The first term in the integral represents the case where the bidder submits the largest nonzero bid, winning the auction and obtaining surplus equal to current valuation less a payment given by the second-highest bid (potentially zero). The second term represents the case where the bidder loses and survives to bid another day, obtaining their continuation value for the next period in which they will be active. These events determined by the realization of the highest competing bid, which by Assumption 1 the bidder believes to be distributed according to G_j^1 . Solving the above maximization problem, we get:

Lemma 1 (Best Responses). *Suppose that beliefs satisfy Assumption 1. Then a bidder's expected payoff is maximized by bidding their valuation less continuation value if this is positive, and otherwise bidding zero:*

$$\beta_j(\mathbf{x}) = \max\{x_j - r \sum_k \tilde{Q}_{j,k} v_k(\mathbf{x}), 0\}. \quad (2)$$

When $\beta_j(\mathbf{x})$ is zero or interior to the support of G_j^1 , the best response is unique. Moreover, $\beta_j(\mathbf{x})$ is continuous, increasing in x_j (strictly when $\beta_j(\mathbf{x}) > 0$) and decreasing in x_k for $k \neq j$ (strictly when $\beta_j(\mathbf{x}) > 0$ and $\beta_k(\mathbf{x}) > 0$).




Figure 1: These figures offer an illustration of the bidding strategies described by Lemma 1 in a two-dimensional example. In panel 1(a) the two functions represent optimal bids on product 1 as a function of bidders' valuations for good 1, fixing their valuation for good 2. Panel 1(b) represents three exemplary bidders: their types $\{x, x', x''\}$ and their associated bids $\{b, b', b''\}$. See the text for discussion.

In static second price auctions, bidders bid their values so that they are indifferent between marginally winning (paying their bid, earning zero surplus) and marginally losing (earning zero surplus). The same logic applies here: bidders must be indifferent between marginally winning and losing, and since losers receive their continuation value, bids must be equal to value less continuation value. It is possible that a bidder prefers to wait for a product they value more highly rather than win the current object at a price of zero, in which case opting out of the auction by bidding zero is optimal.

The assumption that bids are either zero or interior to the support of $\{G_j^1\}_{j \in \mathcal{J}}$ is necessary for uniqueness, since if there is an interval $[b, b']$ on which rival bids are unsupported and the putatively unique best response falls in this range, any other bid in the range will also be a best response. Note that the characterization in (2) is implicit — i.e., $v_j(\mathbf{x})$ is defined recursively according to the Bellman equation in (1) — but since the value functions in Markov Decision Problems are unique (Stokey et al., 1989), the characterization of optimal bids as valuations less continuation values immediately delivers uniqueness of $\{\beta_j(\mathbf{x})\}$ (i.e. that they are functions). Continuity follows from the continuity of payoffs in types. Monotonicity is natural: if a bidder values object j more, they optimally increase their bid for it, and correspondingly shade their bids on substitute objects down.

We offer some graphical intuition for the bidding strategies of Lemma 1 in Figure 1 with a two-dimensional example. Panel 1(a) depicts the monotonicity properties of optimal bids. The more a bidder values product 1 the more she bids on it. However, bidders with higher valuations also shade more; a bidder's continuation value is monotone increasing in her value for product 1, which induces the concavity of the function. Increases in her valuation for product 2 cause a level shift, and may even drive non-participation if x_1 is small — or x_2 is large — enough. Panel 1(b) depicts three conjectured bidder types (x , x' , and x'') and their corresponding bids (b , b' , and b'' respectively). The bidder of type x has a high valuation for both products and therefore shades both of her bids substantially. The bidder of type x' gets the same high utility from product 1, but zero utility from product 2. Therefore she does not participate in auctions of product 2; she prefers to wait for auctions of product 1. Finally, the bidder of type x'' is unlikely to win in either state of the world, and therefore shades her bid very little, but participates in both auctions.

Theorem 1 (Existence). *A pure strategy equilibrium exists in which strategies are characterized by Lemma 1. When $J = 1$, there is a unique equilibrium within the class of strategies characterized by Lemma 1.*

The proof is non-trivial. We must find a fixed point in strategy-space: the equilibrium strategies β^e must be a best response to the distributions $\{G_j^1\}$, i.e. the ergodic distributions of the highest competing bid on each product generated by β^e and the entry, supply and exit processes. To show that such a strategy β^e exists, we apply Schauder's fixed point theorem on the space of continuous functions on the type space \mathcal{X} . Doing so requires some functional analysis. In the appendix we show that the set of best responses to any continuous bidding function is uniformly equicontinuous and bounded (a compactness condition), and that the best response varies smoothly with the bidding function (a continuity condition).

The case with a single product ($J = 1$) is easier, and we can prove uniqueness of equilibrium. This is because the types are totally ordered and monotone bidding preserves that ordering, so that the chance of winning an auction is entirely pinned down by a bidder's type. As a result who wins and who loses can be determined without knowing the equilibrium bidding function, so that the economy admits a unique ergodic distribution of types. Uniqueness of optimal bids follows from a contraction mapping argument.⁵

⁵That uniqueness result for the single product case is only within the class of strategies characterized by Lemma 1: somewhat trivially, it is possible to construct an equilibrium in which a point mass of bidders for

3 Identification

In this section we state and prove the main theoretical result of the paper, which is that the demand system outlined above is non-parametrically identified up to constraints imposed by non-participation (a limitation we make precise below). We think this is useful because it makes explicit the assumptions that are needed to identify the primitives of the dynamic game. Our result is constructive, so it also provides some guidance as to a sensible estimation strategy.

We begin with a scenario in which the econometrician has excellent data and wishes to estimate the baseline model of Section 2 (we investigate cases with limited data and more complex demand models in Section 4 below).

Assumption 2 (Observable Data). *In every auction, the econometrician observes the product-type being sold, all positive bids and the identities of the bidders that made them.*

Even in this case we assume that the econometrician does not observe zero bids, since those stand in for non-participation in our model (i.e. non-participation is non-observable). Despite this, identification of a number of components of the model follows directly. First, $\{G_j^1\}$ is directly observable from the data as the distribution of winning bids on each product. The supply transition matrix Q is identified by the probability that item k is auctioned following item j .

The activity rate τ and survival probability r cannot be directly measured from the sample analogs because some types will optimally choose to not participate even when active. To deal with this, we measure τ and r from a sample consisting of bidders who have been observed bidding on every product. Such bidders have revealed themselves to be types that will participate whenever active. Define for each of these bidders a qualifying time t_i^1 , the time at which they qualified to be in this subset (i.e. the time at which they bid on the last remaining product they had not previously bid on). Then for bidders in this sample

$$r = \mathbb{P}(\text{bidder } i \text{ is observed in any period } > t_i^1 | \text{bidder lost in period } t_i^1) \quad (3)$$

and,

whom, for some j , $\beta_j(\mathbf{x})$ is neither zero nor interior to G^1 — i.e. those who submit $\sup_{\mathcal{X}} \beta_j(\mathbf{x})$ — choose to bid something arbitrarily larger.

$$r\tau = \mathbb{P}(\text{bidder } i \text{ is observed in period } t_i^1 + 1 | \text{bidder lost in period } t_i^1). \quad (4)$$

Since the RHS of (3) and (4) are observed in the data, from this pair of equations one can solve for r and τ . Once we have r , Q and τ , we can compute the state transition matrix $\tilde{Q} = \sum_{s=1}^{\infty} \tau(1-\tau)^{s-1} Q^s$. Next, define the inverse bidding function $\xi : \mathbb{R}^J \rightarrow \mathbb{R}^J$ by

$$\xi(\mathbf{b}) \equiv \mathbf{b} + r\tilde{Q}(I - r\tilde{Q})^{-1}G^1(\mathbf{b}) (\mathbf{b} - E[\mathbf{B}^1 | \mathbf{B}^1 < \mathbf{b}]). \quad (5)$$

In view of the above analysis showing that r , \tilde{Q} and $\{G_j^1\}_{j \in \mathcal{J}}$ are identified, ξ is also identified. Notice that $\xi(\mathbf{b})$ is equal to the vector of a buyer's bids plus a second term that is equal to their vector of continuation values.⁶ There is a clear intuition for the second term: when a buyer wins with a bid b_j , he gets a payoff equal to his value less the expected payment, which we can decompose as $(b_j - \mathbb{E}[B_j^1 | B_j^1 < b_j]) + (x_j - b_j)$. The latter component is equal to his continuation value (Lemma 1 again, as he only wins if $b_j > 0$), and so the payoff can be written as $(b_j - \mathbb{E}[B_j^1 | B_j^1 < b_j])$ plus an opportunity to play the game again. The buyer's continuation value is thus equal in value to the stream of payments from a set of J annuities, where the j th pays $G^1(b_j)(b_j - \mathbb{E}[B_j^1 | B_j^1 < b_j])$ in expectation whenever j is auctioned. The pre-multiplication by $r\tilde{Q}(I - r\tilde{Q})^{-1}$ sums the appropriate discounted geometric series.

In view of Lemma 1, this sum in (5) equal to the bidder's value whenever the bid is positive, so ξ is an inverse bidding function on the region $\mathbf{b} > 0$. When components of the bid vector \mathbf{b} are zero, however, we still obtain an upper bound on the bidder's valuation. Combining this yields

Lemma 2 (Bidder-wise Inversion). *Let $\mathbf{b} = \beta(\mathbf{x})$. Then $\mathbf{x} \leq \xi(\mathbf{b})$ with equality on every dimension j such that $b_j > 0$ (i.e. $\mathbf{x} = \xi(\mathbf{b})$ if $b_j > 0$ for all j).*

This result characterizes the *bidder-wise* content of data from an auction market. Should we observe a complete, strictly positive bid vector, it implies that we can exactly identify the bidder's type. In the particular case where $J = 1$, the inverse bidding function takes the simple form:

$$\xi(b) = b + \frac{r}{1-r} G^1(b)(b - E[B^1 | B^1 < b])$$

⁶The diligent reader can derive this expression by stacking the Bellman equations from (1) in vector form and then substituting in the optimal bidding strategies from Lemma 1.

the discrete-time analogue to the inverse bidding function in Hendricks and Sorensen (2014). Given this inversion result, it will be helpful from now on to refer to the bid-vector $\mathbf{b} \equiv \beta(\mathbf{x})$ as type \mathbf{x} 's "pseudo-type". Let the distribution of pseudo-types be $\tilde{\mathbf{F}}$. It will suffice to identify $\tilde{\mathbf{F}}$ to bound demand \mathbf{F} . To see this, define a vector-valued indicator function $1(\mathbf{b} > 0)$ with $1_j(\mathbf{b} > 0) = 1$ if $b_j > 0$ and $1_j(\mathbf{b} > 0) = 0$ otherwise. Then by Lemma 2, we have $1(\mathbf{b})\xi(\mathbf{b}) \leq \mathbf{x} \leq \xi(\mathbf{b})$. These pointwise inequalities translate into bounds on \mathbf{F} :

$$\mathbb{P}(\{\mathbf{b} : \xi(\mathbf{b}) \leq \mathbf{x}\}) \leq F(\mathbf{x}) \leq \mathbb{P}(\{\mathbf{b} : 1(\mathbf{b} > 0)\xi(\mathbf{b}) \leq \mathbf{x}\}). \quad (6)$$

The missing part of our identification proof is a way of identifying the distribution of pseudo-types (and thus the probabilities on each side of the inequality in (6)). With a single product, this is a simple counting exercise: how many bidders are there who make each bid b (i.e. what is the bid distribution $G(b)$)? With multiple products this counting exercise is complicated by selection, because most bidders will exit before placing a bid on every available product.

We deal with sample selection by choosing subsets of the data for which selection can be explicitly modeled and accounted for. There are many such subsets one could construct in order to prove identification. For example, one possible subset is the set of all complete bid vectors. The (joint) distribution of complete bids is observable, but selected: pseudo-types with high b_1 (e.g.) will typically bid on product 1 and then win and exit, and therefore have low density in the selected sample of complete bid vectors.

One problem with this conditioning set is that there are many paths a bidder can take to make a complete bid vector, and so calculating the probability of this event requires enumerating and computing the probability of many paths. So instead we work with the set of bidders whose history falls into the following set H : they enter when good 1 is available, bid, lose and survive, then next period good 2 is available, and the bid, lose and survive ... J is available, they bid, lose and survive. Such bidders can be identified given the data available under Assumption 2. Every pseudo-type who participates on every product is almost surely represented in this set, as every type on the interior of $\tilde{\mathcal{X}}$ can lose an auction and survive to bid on a different product. We re-weight the observed frequency of pseudo-types observed in this subset by the inverse probability of inclusion— a selection correction that will allow us to learn about the distribution of pseudo-types and, from (6), the distribution of types. To compute the weights, we need to identify the probability $s_H(\mathbf{b})$ that a pseudo-type \mathbf{b} ends up in the set H , which is relatively straightforward for our choice of H :

Lemma 3 (Selection correction). *Let Assumption 2 hold. Then $s_H(\mathbf{b})$ is identified.*

Let the joint density of their bids be $g_H(\mathbf{b})$. Lemma 3 allows us to re-weight the density $g_H(\mathbf{b})$ to get the density of pseudo-types with $\mathbf{b} > 0$:

$$\tilde{f}(\mathbf{b}|\mathbf{b} > 0) \propto \frac{g_H(\mathbf{b})}{s_H(\mathbf{b})} \quad (7)$$

But this only characterizes bidders with $\mathbf{b} > 0$. Revisiting our example in Figure 1(b), bidders of types x and x'' are represented in H , but not x' , who never participates in auctions for product 1. Pseudo-types that don't participate on some products will never enter H (i.e. $s_H(\mathbf{b}) = 0$), and so to move from the conditional distribution to the unconditional distribution takes a little more work. The idea is to define less restrictive sets $H' \supset H$ (e.g. sets of bidders who are observed bidding on product 1) and look for “extra mass” in those sets generated by these non-participating pseudo-types. In the appendix, we make this argument formal, applying induction on the cardinality of the set of products each type participates in, starting from types who bid on every product. Combining these arguments:

Theorem 2 (Non-parametric identification). *Let Assumption 2 hold. Then $\tilde{\mathbf{F}}$ is point identified and \mathbf{F} is partially identified, according to (6).*

This result combines the bidder-wise inversion of Lemma 2 with the selection correction result of Lemma 3. The dynamics of the auction market are important in both steps: in the first, the inverse bidding function $\xi(\cdot)$ captures bidders shading their bids by the option value of losing and participating in future auctions; in the second, $s_H(\mathbf{b})$ accounts for selection into our sample, which is inherently dynamic when $J > 1$.

The reason that Theorem 2 yields partial rather than complete identification is that we allow for bidder non-participation. If bidders were required to participate (e.g. primary dealers in US Treasury auctions), then identification would be exact. Moreover, we note that identification of our model is only incomplete in the left tail of the type distribution, precisely because it is driven by non-participation. Our identification result is exact in the right tail, which is typically the relevant region for computing policy-relevant statistics, e.g. welfare or expected revenue.

Finally, the proof and our exposition above hinged on a specially selected subset of bidders with history H . We chose this subset for expositional clarity alone: though our identification

argument is constructive, in practice we expect that the econometrician will be thoughtful about which subsets — and therefore which variation in the data — to exploit in light of their particular application.⁷ In the same spirit, since the choice of H was expositional, so too is Lemma 3. One could construct a similar argument for alternative subsets.

4 Extensions

In this we provide characterization and identification results for variations on the basic model. We begin by showing that the model remains identified with data on only the winning and second highest bid in each auction, as well as in the presence of reserve prices. We also show, parallel to the the demand estimation literature for fixed price markets, that the model is amendable to projection onto product characteristics, as in the random coefficient utility models of Berry et al. (1995). We then generalize the model to a setting in which there are additional state variables — beyond the product-type — that may affect bidder valuations and beliefs. After characterizing bidding behavior in this case, we are also able to show that the model remains non-parametrically identified in cases with unobserved product heterogeneity and idiosyncratic payoff shocks. Lastly, we consider a version of the model in the presence of an unobserved, external fixed price market that agents may purchase from instead of buying at auction, generating endogenous exit.

4.1 Limited Data

In many real world applications, the data does not contain all bids and identities. Or even if it does — as in our eBay application here — one may be concerned that not all bids are serious (in the sense that they are equal to $\beta_j(\mathbf{x})$), and want to identify the model using data only on the winning and second highest bid. Haile and Tamer (2003) offer a behavioral model of ascending price auctions to motivate this concern, and a number of papers have taken identification under this or similar data limitations seriously (Athey and Haile, 2002; Song, 2004; Menzel and Morganti, 2013; Platt, 2015) We consider the case where the econometrician observes the two highest order statistics and show that the model remains identified.

⁷This argument anticipates the illustrative empirical application of Section 5, where we choose to use only data for bidders who are observed bidding multiple times. See in particular the discussion of the empirical design in Section 5.2.

Assumption 3 (Limited Data). *In each auction, the econometrician observes the product auctioned, the top two bids and the identities of the bidders placing them.*

Lemma 2 still applies, so identifying $(\tilde{\mathbf{F}}, \{G_j^1\}_{j \in \mathcal{J}}, Q, r, \tau)$ is sufficient to identify $\xi(\cdot)$ and bound \mathbf{F} according to (6). The primitives $\{G_j^1\}_{j \in \mathcal{J}}$ and Q are still directly observable in the data. But now the difficulties in identifying r and τ we had earlier are exacerbated by the fact that a bidder will have (unobserved) gaps in their history whenever they were active but did not place one of the top two bids. Our data is thus censored. But since the probability of being censored depends on the distribution of the second highest bids, which is observed, the parameters τ and r remain point identified (see the Appendix for a proof).

Thus the remaining problem is identifying the distribution of pseudo-types $\tilde{\mathbf{F}}$ from the observed bidding data. Once again, we can think of this as a selection problem. Consider the subset of bidders whose history falls into the set H' (a modified version of H): they enter when good 1 is available, *make the second highest bid (and are thus observed)*, then next period good 2 is available, they *make the second highest bid*... good J is available and they *make the second highest bid*. Since the event “make the second highest bid bidding on product j ” is measurable with the data promised by Assumption 3, these bidders can be identified in the data. And by similar arguments to those offered before, we can identify the probability $s_{H'}(\mathbf{b})$ that we need for selection correction. This gives us:

Theorem 3 (Identification under Limited Data). *Let Assumption 3 hold. Then \mathbf{F} is partially identified, according to (6).*

For further intuition on this result, recall the discussion of our selection of the history H at the end of Section 3. There are a number of possible histories on which we could have built our identification result — H was an expository choice, but H' is also feasible. As long as the process by which bidders are selected into the sample can be recovered from observables, our identification result is robust subject to the emendation of the selection correction.

4.2 Random Coefficients

Let us suppose that the products have fixed (non-time-varying) characteristics, observable to the econometrician. These are summarized in a $J \times K$ matrix Z , where the number of characteristics K is less than the number of products J . We make two assumptions on Z :

that it is of full rank (this is necessary for identification), and that all entries are positive (this is unnecessary but simplifies the analysis by allowing the inequality in (8) below).

Let bidder i have valuations of the form $\mathbf{x}_i = Z\boldsymbol{\alpha}_i$, so that their preferences are summarized by a random coefficient $\boldsymbol{\alpha}_i$ (a $K \times 1$ column vector). Instead of sampling valuations \mathbf{x} upon entry from \mathbf{F} , we assume that bidders sample their random coefficient $\boldsymbol{\alpha}$ from \mathbf{F}_α instead. We would like to identify the distribution \mathbf{F}_α .

This demand structure is a special case of the model in which valuations \mathbf{x} come from some higher-dimensional distribution \mathbf{F} , and as a result optimal bidding and the inversion result in Lemma 2 remain unchanged. We can bound \mathbf{F}_α using the more general result in (6):

$$\mathbb{P}(\{\mathbf{b} : \xi(\mathbf{b}) \leq Z\boldsymbol{\alpha}\}) \leq \mathbf{F}_\alpha(\boldsymbol{\alpha}) \leq \mathbb{P}(\{\mathbf{b} : 1(\mathbf{b} > 0)\xi(\mathbf{b}) \leq Z\boldsymbol{\alpha}\}) \quad (8)$$

where we have simply replaced \mathbf{x} in (6) by $Z\boldsymbol{\alpha}$. Intuitively, the probability that a random coefficient is lower than some level $\boldsymbol{\alpha}$ is the same as the probability that the valuations are lower than the corresponding level $Z\boldsymbol{\alpha}$, and can be bounded in the same way.

Corollary 1 (Identification under random coefficients). *Suppose Assumption 3 holds. Then \mathbf{F}_α is partially identified, according to (8).*

Though the result is straightforward, it highlights the generality of the type space of our model from Section 2. As in demand systems for fixed-price markets, projection down to preferences over characteristic space yields dimension reduction. Since we can partially identify a high-dimensional model with random preferences over products, we can also partially identify a lower-dimensional model of random preferences for characteristics too.

4.3 Demand with Richer Signals

We now analyze a more substantial variation of the basic model in which we allow for both public and private signals that affect both beliefs and valuations. This framework is the basis for two extensions: public signals that affect beliefs about the state of the market and unobserved heterogeneity in product attributes.

To accommodate this, we need a richer state space. Assume that in each period, a state variable $s \in \mathcal{S}$ is publicly observed. This state variable may affect bidder valuations, so we model a bidder's valuation as a random function $x_i(s)$. Exactly how $x_i(s)$ is sampled will

depend on the specifics of the model. For example, in the original model, the state variable was the product-type under auction (i.e. $s \in \mathcal{J}$), so valuations were $x_i(j) \equiv x_{i,j}$ where $x_{i,j}$ was the j th element of \mathbf{x}_i , and \mathbf{x}_i was sampled iid across bidders i from \mathbf{F} . The state variable may also affect beliefs about competing bids, so we modify Assumption 1 accordingly:

Assumption 4 (Beliefs with Public States). *When the public state is s , bidders believe that the highest rival bid in the current auction has distribution $G^1(\cdot|s)$.*

For example, bidders may believe that when there is a lot of supply of a given good on the market, the highest competing bid will be lower. Assume also that bidders believe the state transitions according to an exogenous Markov transition kernel $P(s'|s)$ (exogenous in that it is unaffected by the actions in the current auction).

Definition 2 (Equilibrium with Public States). *A pure strategy equilibrium is a pair $(\beta^e(\cdot, s), G^1(\cdot|s))$ such that:*

(Optimality) $\beta^e(\mathbf{x}|s)$ is a best response for type \mathbf{x} given beliefs $G^1(b|s) \forall s \in \mathcal{S}$.

(Consistency) $G^1(b|s) = \mu([0, b]|s) \forall s \in \mathcal{S}$.

Then by the same logic that we offered for Lemma 1, it is optimal for bidders to bid their valuation less their discounted continuation value, and we obtain:

Lemma 4 (Best responses, general state space). *The equilibrium bidding function $\beta(s)$ satisfies:*

$$\beta(s) = \max\{x(s) - r \int v(s') dP(s'|s), 0\} \quad (9)$$

where $v(s)$ is the value function, defined according to the Bellman equation:

$$v(s) = \max_b G^1(b|s)(x(s) - E[B^1|B^1 < b, s]) + r(1 - G^1(b|s)) \int v(s') dP(s'|s) \quad (10)$$

4.3.1 Public States and Bidder Beliefs

A special case of particular interest is when $\mathcal{S} = \mathcal{J} \times \{s_1, \dots, s_S\}$ and $x_i(s) = x_{i,j}$, i.e. when the state informs beliefs about competitor play, but does not encode valuation-relevant information. We call attention to this not because it is technically demanding, but because it remedies a limitation of the baseline model: that it misses the intuition of a dynamic

marketplace in which bidders respond to changes in market fundamentals. We can encode countably sophisticated information into the public state space \mathcal{S} , including histories, changing primitives, and foresight.

For instance, suppose that bidders know not only what is being auctioned today, but the sequence of products that are being auctioned $k - 1$ periods into the future. Then, $\mathcal{S} = \mathcal{J}^k$, and bidders' continuation values will now depend upon realizations of future supply. Alternatively, suppose that τ , the re-arrival rate of bidders, evolves over $\{\tau_L, \tau_H\}$ in a public and Markov process. Now, $\mathcal{S} = \mathcal{J} \times \{\tau_L, \tau_H\}$. In the τ_H world bidders face more competition from more bidders, and so it is useful to let this be reflected in $G^1(\cdot|s)$.

It is straightforward to see that this presents no theoretical difficulty for identification. The simplest approach is by brute force application of Theorem 2: simply define $\tilde{\mathcal{J}}$ to contain $|S|$ duplicates of the product space \mathcal{J} and proceed as before. It is also possible to show, using a contraction mapping argument, that the full $J \times |S|$ bid vector is identified from any J -vector that includes a bid for each product type.⁸

4.3.2 Unobserved Heterogeneity and Idiosyncratic Shocks

We now introduce a variant of the model with unobserved product heterogeneity and idiosyncratic taste shocks (this will be useful in the application). In this setting unobserved product heterogeneity (ξ_t) is public and transient, and idiosyncratic taste shocks ($\varepsilon_{i,t}$) are private and transient. We assume that valuations are:

$$x_i(j, \xi_t, \varepsilon_{i,t}) = x_{i,j} + \xi_t + \varepsilon_{i,t} \tag{11}$$

where $x_{i,j}$ is sampled as in the basic model, ξ_t is sampled iid across auctions from F_ξ and $\varepsilon_{i,t}$ is sampled iid across bidders and auctions from F_ε . We assume in addition that $x_{i,j}$, ξ_t and $\varepsilon_{i,t}$ are mutually independent and that ξ_t and $\varepsilon_{i,t}$ are mean zero. We make a full support assumption on $\varepsilon_{i,t}$ so that every type \mathbf{x} bids on every product with positive probability (i.e. they participate when they get a sufficiently large idiosyncratic shock). Notice that a random coefficients model with unobserved heterogeneity and idiosyncratic shocks in the spirit of Berry (1994) and Berry et al. (1995) is a special case of this model, in the sense

⁸This implies that the requirements for a “complete” bid vector are not made more taxing as the econometrician adds public states. Details available from authors upon request. However, the dimension of the first-stage objects the econometrician must estimate will grow as public states are added.

that it projects the valuations $x_{i,j}$ onto characteristics.

We split the identification analysis into two parts: a first step in which we define pseudo-types and show an inversion from pseudo-types to types; and a second step in which we show that the distribution of pseudo-types is identified.

Inversion. For valuations of the form in (11), Lemma 4 implies we have bids⁹:

$$b_{i,j,t} = \max\{x_{i,j} + \xi_t + \varepsilon_{i,t} - r \sum_k \tilde{Q}_{j,k} v_k(\mathbf{x}), 0\}$$

where the continuation value in the first term depends only on the permanent part of the state \mathbf{x} rather than the transitory iid shocks ξ_t and $\varepsilon_{i,t}$ (though the presence of idiosyncratic shocks creates option value that is included in $v_k(\mathbf{x})$).

Define a pseudo-type vector $\check{\mathbf{b}}(\mathbf{x})$ elementwise according to $\check{b}_{i,j}(\mathbf{x}) = x_{i,j} - r \sum_k \tilde{Q}_{j,k} v_k(\mathbf{x})$. We use the check notation to distinguish the pseudo-types from the bids. The pseudo-type is the permanent part of the bid vector; positive bids can be written as pseudo-type plus the two transitory shocks. In contrast to bids, pseudo-types can be negative. We can then write the bids as:

$$b_{i,j,t} = \max\{\check{b}_{i,j} + \xi_t + \varepsilon_{i,t}, 0\}$$

Let $B_{j,t}^1$ be the winning bid in each auction, and define \check{G}_j^1 as the distribution of $B_{j,t}^1 - \xi_t$. In the presence of unobserved heterogeneity, all bidders will adjust their bids by ξ_t (which they all value equally), and so the raw distribution of winning bids G_j^1 will have additional variance relative to the distribution \check{G}_j^1 .

Lemma 5 (Inversion, general demand). *Suppose $\{\check{G}_j^1\}_{j \in \mathcal{J}}$, \tilde{Q} , r and F_ε are known. Then $\mathbf{x} = \check{\xi}(\check{\mathbf{b}})$, where $\check{\xi}(\check{\mathbf{b}}) = \check{\mathbf{b}} + r\tilde{Q}(I - r\tilde{Q})^{-1}\check{\mathbf{u}}(\check{\mathbf{b}})$ and*

$$\check{u}_j(\check{\mathbf{b}}) = \int \check{G}_j^1(\check{b}_j + \varepsilon)(\check{b}_j + \varepsilon - E_{\check{G}_j^1}[B_j^1 | B_j^1 < \check{b}_j + \varepsilon]) dF_\varepsilon(\varepsilon)$$

As in the basic model, each pseudo-type can be mapped back to an underlying type. But there are a number of important differences between this result and the earlier one. First, since pseudo-types can now be negative, this mapping is exact (rather than holding only for

⁹Since $\varepsilon_{i,t}$ is private rather than public information, strictly speaking this is not implied by Lemma 4. But since $\varepsilon_{i,t}$ only affects the bidder's current valuation and nothing else, adding it to the bid is optimal.

strictly positive pseudo-types, as in the basic model). Second, the expected utility from the optimal bid in any period now includes the idiosyncratic shock, both in the continuation value and in adjusting the bids. The unobserved heterogeneity does not enter expected utility, as it shifts all bids in an auction up by ξ_t , and thus the extra valuation of ξ_t is exactly canceled out by an additional payment of ξ_t conditional on winning. Finally, this inversion cannot be applied bidder-wise to bid vectors, since bid vectors are not pseudo-types (they now include unobserved heterogeneity and idiosyncratic shocks). Instead, we will apply the inversion $\xi(\cdot)$ to the distribution of pseudo-types as a whole.

But first we need to identify both F_ε and $\{\check{G}_j^1\}_{j \in \mathcal{J}}$, as these are required by Lemma 5. Identifying these objects will require making some deconvolution arguments (Li and Vuong, 1998; Krasnokutskaya, 2011). Let us start with identifying $\{\check{G}_j^1\}_{j \in \mathcal{J}}$. Take pairs of bids placed in the same auction, by bidders that have just entered the data. The pair can be written as a common part ξ_t plus a vector of independent shocks (sums of the pseudo-type and the idiosyncratic shocks). Since all the terms are mutually independent, and $E[\xi] = 0$, by Kotlarski's lemma, the distribution F_ξ is identified (Kotlarski, 1967). Then, since each observed G_j^1 is a convolution of F_ξ and the corresponding \check{G}_j^1 , and F_ξ is identified, we can again apply deconvolution to get $\{\check{G}_j^1\}_{j \in \mathcal{J}}$.

Next consider identification of F_ε . Take a very specific sample: the pair of bids of bidders who entered the sample, bid on product 1, lost, and bid on product 1 again. So we have:

$$\begin{aligned} b_{i,1,t_1} &= \check{b}_{i,1} + \xi_{t_1} + \varepsilon_{i,t_1} \\ b_{i,2,t_2} &= \check{b}_{i,1} + \xi_{t_2} + \varepsilon_{i,t_2} \end{aligned}$$

with the following statistical structure: all the distinct variables are mutually independent; ξ_{t_1} and ξ_{t_2} are drawn from F_ξ ; ε_{i,t_2} is drawn from F_ε ; but ε_{i,t_1} is drawn from some other distribution \tilde{F}_ε because the sample is selected: only losing bidders in the first auction (more likely to be those who drew low ε 's and placed low bids) are in the sample.

This can be written in the form $Y_1 = M + U_1$ and $Y_2 = M + U_2$ for $M = \check{b}_{i,1}$, $U_1 = \xi_{t_1} + \varepsilon_{i,t_1}$, $U_2 = \xi_{t_2} + \varepsilon_{i,t_2}$ with M, U_1, U_2 mutually independent and $E[U_2] = 0$. Applying Kotlarski's lemma again, the distribution of $U_2 = \xi_{t_2} + \varepsilon_{i,t_2}$ is identified (see also Evdokimov and White (2012)). This is itself a convolution of the identified F_ξ and the unknown F_ε , and so by applying deconvolution again, we identify F_ε .

Identification of Pseudo-Type Distribution. In view of Lemma 5 and the above deconvolution analysis we know that given the distribution of pseudo-types $\tilde{\mathbf{F}}$ we can identify \mathbf{F} . But $\tilde{\mathbf{F}}$ is latent. So let us instead work with the set of bidders with history H , where H is defined as before (i.e. these are bidders who enter when good 1 is available, are active, bid, lose, survive...). Let G_H be the distribution of their bid vectors in those J auctions. Define the convolution operator $*$ acting on distributions F and G of independent random variables X and Y by $P_{F*G}(A) = \int_X \int_Y 1(X + Y \in A) dF(X) dG(Y)$ for any measurable set A . Note that F is identified from $F * G$ whenever G is known. Then we have:

$$\mathbf{G}_H = \left((\tilde{\mathbf{F}} * \mathbf{F}_\varepsilon) | \text{complete} \right) * \mathbf{F}_\xi$$

where \mathbf{F}_ε is the distribution of idiosyncratic error vectors and \mathbf{F}_ξ is the distribution of unobserved heterogeneity vectors (both J -vectors). We indicate that the distribution of the convolution of the pseudo-types and the idiosyncratic errors is selected by the process of looking for bid vectors that are complete in the first J auctions.

Then since \mathbf{G}_H is observed and \mathbf{F}_ξ is identified by previous arguments (since the ξ_t are iid, the vector has distribution equal to the J -product of the marginals F_ξ), the object $\left((\tilde{\mathbf{F}} * \mathbf{F}_\varepsilon) | \text{complete} \right)$ is identified. We would next like to get the unconditional distribution $\tilde{\mathbf{F}} * \mathbf{F}_\varepsilon$. Following the logic of lemma 3, this requires identifying for each draw from $\tilde{\mathbf{F}} * \mathbf{F}_\varepsilon$ (i.e. for a pseudo-type plus an idiosyncratic shock vector) the probability that such a bidder ends up in the set H . But since $\{\check{G}_j^1\}_{j \in \mathcal{J}}$, Q, r and τ are all identified, the proof goes through exactly as before. So by applying selection correction, $\tilde{\mathbf{F}} * \mathbf{F}_\varepsilon$ is identified. One last deconvolution suffices to identify $\tilde{\mathbf{F}}$ separately. Summarizing:

Theorem 4 (Identification with Unobserved Heterogeneity and Idiosyncratic Shocks).

Let Assumption 2 hold. For the demand system given by (11), \mathbf{F} , F_ξ and F_ε are all non-parametrically identified.

In view of the many deconvolution operations, and the slow convergence rates of estimators based on deconvolution, this identification argument cannot be taken directly to data unless the data set is very large indeed. In our application, where we use a similar demand system, we impose a number of parametric assumptions to be practical.

4.4 Reserve Prices

In many markets, sellers set reserve prices on their items. They do this to avoid selling them at a price below their best outside use (either retained by the seller, sold in a different market, or sold in this market at a later date), and to extract more money from buyers. Sellers may set reserves non-strategically (e.g. at cost, or the \$1 starting bid recommended for many products on eBay), or to maximize expected revenue. They will generally know everything commonly known by the bidders, which may include characteristics latent to the econometrician.¹⁰ Reserves are sometimes implemented as minimum bids, which restrict participations to types with sufficiently high values (e.g. on eBay). In other settings (e.g. online advertising), the reserve may only be applied after the bids have been recorded (“secret reserves”), in order to adjust allocation and payment.

Reserve prices can be easily accounted for in the model. Suppose that sellers charge reserve prices on each item sold, drawn iid from distributions $\{G_j^R\}$. We assume that buyers know the distributions of reserve prices. Re-define G_j^1 as the conditional distribution of the maximum of the highest competing bid and the reserve price on product j , and let assumption 1 hold as before. Bids are equal to valuation less continuation value, where the continuation value is dependent on the distribution of reserve prices. High reserve prices directly depress continuation values by lowering the likelihood that a bidder is successful in future auctions. They also mean fewer bidders exiting the market, toughening competition.

For identification, the function $\xi(\cdot)$ still offers a valid inversion, though the distributions $\{G_j^1\}$ are now amended to reflect the reserves. So if we can identify the pseudo-type distribution $\tilde{\mathbf{F}}$, we can bound \mathbf{F} . With secret reserves, we can pointwise identify $\tilde{\mathbf{F}}$ just as before.

On the other hand, with minimum bids, we again face a selection problem: only bids that clear the minimum bid will be observed. So in the selection correction step, we need to take into account the reserves in calculating $s_H(\mathbf{b})$. Moreover — in contrast to all the prior analysis — some types might never be observed.¹¹ For example, suppose there is only one product, and it always gets a \$100 reserve. Then types with values below \$100 will never bid. The mass of such types cannot be inferred from the data. This implies that we can only

¹⁰For example, Roberts (2013) explores a model in which an unobserved heterogeneity term ξ is observed by sellers and used to set reserve prices. He shows that the reserve prices can then be used as controls in estimating demand.

¹¹In the baseline model, every type bids on at least one product with positive probability. This remains true in the limited data case, since there is always some chance of being one of the top two bidders (e.g. when no other bidders are active that period).

identify $\tilde{\mathbf{F}}$ for the subset of types who bid on at least one product with positive probability. And even conditioning on types that bid with positive probability, we can only partially identify $\tilde{\mathbf{F}}$. This translates into weaker bounds on \mathbf{F} , though the right tail of valuations will still be well-identified, and this is often what is needed for policy counterfactuals.

4.5 Outside Options

We conclude our extensions with a variation of the model in which there are latent outside options. We introduce outside options by describing a more general environment with *passive search*. In each period t , both new entrants and active incumbents are presented with either an auction of one of the goods in \mathcal{J} (as in the basic model), or an opportunity to buy one of the goods in \mathcal{J} at a fixed price in another market (observed by the bidder, unobserved to the econometrician). Beliefs about opposing bids in the auction market follow Assumption 1, and are captured by the distributions G_j^A (we switch superscript to indicate “auction”). Prices in the outside market are stochastic, sampled independently and identically over time according to product-specific continuous distributions G_j^O (where the “O” stands for outside) with strictly positive density over their support. We make a support inclusion assumption: $\text{supp}(G_j^O) \subseteq \text{supp}(G_j^A)$, so that the highest bid in the auction market would always suffice to buy the corresponding good in the fixed price market. This assumption seems weak given our intuitions about the size of price fluctuations in auction versus fixed price markets. We make use of it in our arguments below.

Search in this environment is “passive” in the sense that active bidders do not actively choose which product to bid on, nor whether to go to the fixed price or auction market. This is consistent with our initial model.¹² But bidders are sophisticated and take into account future opportunities when making decisions today.

The transition matrix \tilde{Q} is now of size $2J \times 2J$, indicating the distribution over options that a bidder participating on some product j in one of the two markets {Auction, Fixed Price} will encounter when next active. We simplify by assuming that this distribution is independent of the past state, so that it is multinomial (π^A, π^O) , where the J -vector π^A is the distribution over the J products offered at auction, π^O is the distribution over outside options, and $\sum_j \pi_j^A + \sum_j \pi_j^O = 1$. This simplification is substantive: without it we would need to make inferences about the Markovian transitions between unobserved states in the outside market.

¹²See a previous version of this paper for a model with optimal auction entry.

The buyer's decision whether or not to buy at price p_j in the fixed price market for good j takes the form of a reservation rule: buy iff $p_j < b_j^O$, for b_j^O a reservation value. We can therefore think of the buyer as "bidding" b_j^O in the fixed price market, winning when b_j^O exceeds the price p_j and paying p_j . This allows us to shoehorn the fixed price market into our original auction market analysis:

Lemma 6 (Best responses, outside option). *The equilibrium bidding function $\beta(\mathbf{x})$ satisfies:*

$$\beta_j^k(\mathbf{x}) = \max\{x_j - rv(\mathbf{x}), 0\}, \quad k = \{A, O\} \quad (12)$$

where $v(\mathbf{x})$ is the ex-ante value function, defined according to the Bellman equation:

$$v(\mathbf{x}) = \sum_{j=1}^J \sum_{k=\{A, O\}} \pi_j^k \left(\max_b G_j^k(b) (x_j - E_{G_j^k}[B^1 | B^1 < b]) + (1 - G_j^k(b))rv(\mathbf{x}) \right) \quad (13)$$

Here we take advantage of the iid multinomial state transitions to work instead with an "ex-ante" Bellman equation, where $v(\mathbf{x})$ is the value of an active bidder prior to the realization of the good/market they're participating in. As a result, all bids take the form of valuation less continuation value (if positive), where that continuation value is a scalar function of type. An immediate implication of (12) is that the buyer's bid in the auction market and their reservation value in the fixed price market are identical. Knowing a bidder's pseudo-type \mathbf{b} (still a J -vector) is therefore enough information to characterize their behavior in both markets. This insight will help in the identification argument that follows.

Identification. We consider identification where the econometrician has excellent data on the auction market (i.e. Assumption 2 holds), but knows nothing about the outside market, nor observes bidder participation in the outside market. As a result, our identification arguments are based on an analysis of how exit rates vary across pseudo-types, which is informative as to their outside options.

Let a pseudo-type \mathbf{b} be the J -vector of bids on each product (equal across the two markets), and define the inverse bidding function $\xi : \mathbb{R}^J \rightarrow \mathbb{R}^J$:

$$\xi(\mathbf{b}) = \mathbf{b} + \frac{r}{1-r} \sum_{j=1}^J \sum_{k=\{A, O\}} \pi_j^k \left(G_j^k(b_j) \left(b_j - E_{G_j^k}[B^1 | B^1 < b_j] \right) \right)$$

where once again $\xi(\mathbf{b})$ takes the form of bid plus continuation value, where the continuation value is equal to a stream of discounted payoffs in both auction and fixed price markets. Given the distribution of pseudo-types $\tilde{\mathbf{F}}$ (partially identifiable using the exact same argument as in Lemma 3), we can partially identify \mathbf{F} .

So the only new part of the analysis is showing that $\xi(\mathbf{b})$ remains identified in the presence of latent outside options. The distribution of winning bids $\{G_j^A\}$ can be identified from the data, as can $\{\pi_j^A\}$ (up to a normalizing factor), but $\{G_j^O, \pi_j^O\}$ are both latent.

To identify these, we look at how exit rates change with pseudo-types. Define $\Pi^O = \sum_j \pi_j^O$, the probability of entering the fixed price market when active, and $\tilde{\pi}_j^O = \frac{\pi_j^O}{\Pi^O}$, the probability of drawing good j when active in the fixed price market. The probability that pseudo-type \mathbf{b} exits after participating for a single period in the fixed price market is $k(\mathbf{b}) = (1 - r) + r \sum_j \tilde{\pi}_j^O G_j^O(b_j)$, where the first term comes from exogenous exit after losing, and the second from the chance that they buy. So after losing an auction, the probability that a bidder exits our data is given by:

$$e(\mathbf{b}) = \underbrace{(1 - r)}_{\text{exogenously exits}} + \underbrace{r}_{\text{survives}} \underbrace{(\Pi^O k(\mathbf{b}) + (\Pi^O)^2 (1 - k(\mathbf{b})) k(\mathbf{b}) \dots)}_{\text{exits in the outside market}} = (1 - r) + r \frac{\Pi^O k(\mathbf{b})}{1 - \Pi^O (1 - k(\mathbf{b}))}$$

We can learn the function $e(\mathbf{b})$ by observing the subsequent exit rates of bidders who have made complete bid vectors \mathbf{b} . The exit rates of “extremal” pseudo-types are particularly informative. One can show $e(\mathbf{0}) = (1 - r) + r \frac{\Pi^O (1 - r)}{1 - r \Pi^O}$. Let $\mathbf{b}_j = (0, 0 \dots b_j \dots 0)$ denote a pseudo-type who bids b_j on good j and zero on everything else, and let $\bar{\mathbf{b}}_j$ be the type who makes the highest such bid \bar{b}_j . By our earlier support inclusion assumption, $\bar{\mathbf{b}}_j$ purchases j in the fixed price market with certainty. Then we have

$$e(\bar{\mathbf{b}}_j) = (1 - r) + r \frac{\Pi^O (1 - r) + r \pi_j^O}{1 - r \Pi_{-j}^O}$$

So the $J + 1$ exit rates $e(\mathbf{0})$ and $\{e(\bar{\mathbf{b}}_j)\}_{j=1}^J$, can be written as a function of $J + 1$ latent variables: r and $\{\pi_j^O\}$, since Π^O is just $\sum_{j=1}^J \pi_j^O$. In the appendix we show that this system of equations has a unique solution, so that r and $\{\pi_j^O\}$ are identified.

Finally, consider type \mathbf{b}_j , who “bids” b_j on product j and zero otherwise. Their per-period probability of exit in the fixed price market is $k(\mathbf{b}_j) = k(b_j) = (1 - r) + r \tilde{\pi}_j^O G_j^O(b_j)$ and

consequently their exit probability upon losing in the auction market is:

$$e(b_j) = (1 - r) + r \frac{\Pi^O(1 - r) + r\pi_j^O G_j^O(b_j)}{1 - r (\Pi^O - \pi_j^O G_j^O(b_j))}$$

Take a derivative in b_j on both sides of the above equation:

$$e'(b_j) = r \frac{r\pi_j^O G_j^{O'}(b)(1 - \Pi^O)}{(1 - r (\Pi^O - \pi_j^O G_j^O(b_j)))^2}$$

which is a separable first-order ordinary differential equation in $G_j^O(b)$. For fixed r , $\{\pi_j\}$ and $e'(b_j)$, this equation has a unique solution for $G_j^O(b)$ (see appendix). Then since r , $\{\pi_j\}$ and $e'(b_j)$ are identified, so is $G_j^O(b)$ for each j , which by earlier arguments suffices to identify $\xi(\cdot)$. To use this function to bound the distribution of types \mathbf{F} , we need to identify the distribution of pseudo-types $\tilde{\mathbf{F}}$, which requires a selection correction argument. But the argument in Lemma 3 suffices for this, since the presence of an outside market doesn't affect the probability of the particular sequence of events H considered in that proof. Summarizing:

Theorem 5 (Identification with an outside option). *Suppose Assumption 2 holds, and that $\text{supp}(G_j^O) \subseteq \text{supp}(G_j^A)$ for every j . Then \mathbf{F} is partially identified, according to (6).*

5 Empirical Application

We present an application of our identification result to the auction market for compact cameras on eBay.com. We selected this market as an illustration for several reasons. First, as noted earlier, the eBay auction design is strategically similar to a sequence of second-price sealed bid auctions. Second, most consumers only purchase a single camera, so the unit demand assumption seems reasonable. Third, compact cameras are measurably differentiated in a salient characteristic — namely resolution — which consumers may value differently, affording us an opportunity to highlight a random coefficients variation of our model. The demand model we estimate allows for random consumer preferences for resolution in the spirit of Berry et al. (1995), but in the “pure characteristics” setting of Berry and Pakes (2007). We estimate the parameters of this model using a maximum likelihood approach and use those estimates to characterize consumer surplus and optimal bidding strategies.

5.1 Data and Market Overview

The eBay Marketplace. eBay is widely considered to be the world leader in online auctions. Various elements of its platform design, such as the use of proxy bidding agents, feedback scores and “buy-it-now” offers have been widely copied. At any time, eBay hosts a large number of object listings from a variety of sellers. Buyers can browse these, either by navigating through categories delineated by the site, or by directly searching for key phrases. For example, a search for “digital compact camera” will typically bring up thousands of listings. Some of these are offered at “buy-it-now,” i.e. fixed, prices. As Einav et al. (2016) have documented, over 60% of listings on eBay are now at fixed prices (though auctions were more common at the time our data was collected). We will ignore the presence of a fixed price market in what follows, effectively ignoring substitution opportunities between the two markets in favor of focusing on substitution within the auction market.

Restricting to auctions yields a list of items, ordered by time until auction end. These auctions all end at different times, so bidders face a set of *sequential* auctions. Unsurprisingly, there is substantial heterogeneity in the cameras offered, in terms of brand, resolution, zoom and accessories (to name some of the most salient features). Bidders can research past prices for different items by searching the completed listings.

Bids are placed through eBay’s proprietary “proxy bidding” system. Bidders enter the maximum they are willing to pay for the item, and then eBay’s proxy bidding system will bid up from the current standing price in standardized increments on their behalf until either their bid is the highest yet entered in the system, or an additional increment would take them over their maximum. For example, if bidder A enters a bid of \$8000 on a camera where the standing price is \$6000 and the highest bid placed in the system by a rival is \$7000, then the system will update the standing price to \$7100 ($\$7000 + \100 increment), and will record this bidder as the currently high bidder. Under unit demand high bidders become “committed” to the auctions they enter, in the sense that if they bid in another auction, there is a risk of winning a second object they don’t need.

Fear of premature commitment is perhaps one reason why bidders tend to bid late in auctions. For example, on eBay Motors, 75% of bids are placed in the last day of the auction (Lewis, 2007). Similar results have been found elsewhere on eBay and in other auction markets, and a wide range of alternative explanations for late bidding have been offered (see e.g. Roth and Ockenfels (2002)). The combination of late and proxy bidding suggest that eBay’s auction

market is well approximated as a sequence of second-price sealed bid auctions, so that our model can be applied to this setting.

Data. We purchased a dataset concerning all sales of digital compact cameras over a 2-year period from TeraPeak, a data analytics company. The data includes attributes of the camera auctioned (resolution, zoom, brand, product name, bundling of a tripod, extra battery etc), attributes of the listing (starting price, secret reserve, listing title), and the outcome of the auction. Each listing may be associated with several bids, all of which we observe — including the highest bid, which is not visible on the website and typically unavailable in “scraped” auction datasets from the platform. Market participants are persistent in our dataset — as in our model — and we observe their attributes (feedback, location) and construct measures of experience and activity from observed behavior.

We work with a subset of the data, consisting only of *new* compact cameras, sold in the 3-month period between February 5th and May 6th of 2007. We restrict attention to new cameras to limit the influence of unobserved heterogeneity, though as we will see, this is still a substantial problem. We analyze this particular time period because supply was relatively stable over those 3 months, so the stationarity assumptions implicit in our calculations of the continuation values are reasonable. We pick cameras with the most common resolution levels: those with (rounded) resolution between 5 megapixels and 10 megapixels (MP). We clean the data by excluding auctions with missing data, potential shill bidding, outlying bids and auctions that are terminated by the first bidder exercising a buy-it-now option (this is unusual). See the data appendix for further details on sample construction.

This leaves us with a dataset of 19160 auctions, 4387 sellers and 74375 unique bidders. Summary statistics are presented in Table 1. In the top panel, we summarize the data at an auction level (i.e. an observation is an auction). The average gap between the winning bid and the closing price is \$10.60. This is a direct estimate of the average consumer surplus of winning bidders in a static auction model, but we will show that it is a substantial underestimate once continuation value is taken into account. The probability of sale is quite high, at 0.95. Those products that do not sell tend to be among the small handful which employ “secret reserve prices”, whereby the item only sells if the highest bid meets a reserve price set by the seller but hidden from bidders. On average an auction attracts 7 – 8 unique bidders, and 16 – 17 bids, though there is substantial variance.

Table 1 also summarizes the data at the seller and bidder level. There are just over four

listings for every seller in our marketplace. The distribution of seller shares is skewed, with large, experienced sellers making up the bulk of listings. However there are many such large sellers and the seller concentration measure (HHI) remains low at 0.04. On the bidders' side we find some motivation for the assumptions of our model. Most bidders are unsuccessful in acquiring an item (the purchase rate is 24.3%), but they are active in the market for nearly a week and on average are observed bidding in two different auctions, with some participating in many more than that.¹³ This repeated participation is not driven by multi-unit demand—99% of our bidders make one or fewer compact camera purchases.

5.2 Descriptive evidence

Having introduced the data, we now offer some descriptive evidence that indicates our model is a reasonable approximation to the behavior we see in the market.

Do bidders substitute across products? If bidders generally bid on the same products repeatedly (or have identical preferences over product characteristics), one might think that the single-good models familiar from the existing empirical auctions literature may suffice. So to get some evidence on this, we look for evidence of substitution across products. There are many distinct products sold in our dataset, but in the demand system below we allow random coefficients over camera resolution, and so this is the relevant definition of a product for the purposes of our analysis. We therefore calculate a transition matrix across cameras of different resolution, shown in Table 2. Each entry in the matrix is the probability (expressed as a percentage) that a bidder who bids on a camera of the resolution in a given row subsequently bids on a camera with the column resolution. We find that most buyers that bid on multiple cameras tend to bid on a product with the same resolution next, but the probability of this is far away from 100%, ranging from 45.5% to 80.5%. When they substitute, they tend to pick a product of similar resolution (the biggest off-diagonal elements are generally close to the diagonal).

Do bidders have option value? Another focus of our approach has been to emphasize dynamic bidding and option values. Option values arise from price fluctuations — a bid that loses today may win tomorrow. Many economists have the intuition that in large markets

¹³We say a bidder is “active” from the start of the first auction they bid in to the end of the last auction.

Table 1: Summary Statistics: Online Auctions of Compact Cameras

	Mean	Std. Dev.	Min	Max
Auction-Level Data				
Winning bid	203.0	108.3	0.01000	2900
Closing price	192.4	97.08	0.01000	2025
Shipping cost	18.49	9.482	0	190
Starting price	49.83	86.66	0.01000	600
Secret reserve?	0.0872	0.282	0	1
Item sold?	0.949	0.220	0	1
Bid count	16.72	10.17	1	95
Number of unique bidders	7.615	3.910	1	34
Camera resolution (megapixels)	6.912	1.279	5	10.10
Optical zoom	4.520	2.767	1	18
Digital zoom	4.183	1.088	2	10
Comes with accessories	0.477	0.499	0	1
Number of auctions	19160			
Seller-Level Data				
Number of listings	4.367	51.34	1	2677
Number of sales	4.120	50.99	0	2664
Seller feedback	1001.9	7145.5	-1	252861.2
Number of sellers	4387			
Seller HHI	0.0351			
Bidder-Level Data				
Auctions participated in	1.962	3.441	1	186
Incumbent period (days)	2.706	5.537	-0.000746	77.82
Time between bids (days)	4.208	9.640	0	93.06
Number of purchases	0.243	0.475	0	22
≤ 1 purchase	0.990	0.101	0	1
Bidder feedback	82.88	463.8	-999	60013
Winning bidder's age	36.27	16.02	1.742	95.98
Number of bidders	74375			
Buyer HHI	0.0000636			

Notes: Summary statistics for the full dataset, which consists of all auctions of compact cameras auctions on a large online platform that ended during the period Feb 5 – May 6, 2007. Observations with missing product characteristics have been dropped. The “incumbent period” for a buyer is measured as the time from the start-date of the first auction bid in to the end of the last auction bid in. “Time between bids” is the gap between the first and second bids by a buyer on different objects, measured only for the subsample of bidders with multiple bids. “Seller HHI” is the Herfindahl–Hirschman index for sellers, based on their share of items sold. “Buyer HHI” is the analogous measure for buyers, based on the share of items bought.

Table 2: Substitution Patterns for Repeat Bidders

	5MP	6MP	7MP	8MP	10MP
5MP	64.8	17.3	12.3	4.2	1.4
6MP	5.9	75.3	13.7	3.5	1.7
7MP	2.9	8.8	80.8	5.2	2.4
8MP	2.6	5.4	10.2	78.3	3.6
10MP	1.9	6.1	13.0	8.0	71.0

Notes: Each entry in the matrix gives the observed frequency with which a bidder who is observed bidding on the row product is next observed bidding on the column product.

the law of one price holds, and so price fluctuations should be minimal. This is not true here. The top left panel of Figure 2 shows the median daily transaction prices (including shipping) on different models of camera, over the sample period. There is quite a lot of variability — an typical change is 10% — and since it is a median price, this is not driven entirely by outliers.

To rule out variation based on compositional effects, we drilled down to look at the highest volume seller’s most popular product (a Kodak Easyshare Z710), and plotted the price series against time (fluctuations could be within a day), shown in the top right panel of Figure 2. Prices vary from below \$200 to over \$250 dollars. Doing some rough back of the envelope calculations based on the observed price distribution, a bidder with a valuation of \$250 (red dashed line), bidding once a day, with a daily exit hazard rate of 0.5, should optimally shade their bid down to \$238 (green dashed line), because of the option value (and should expect to pay closer to \$225, yellow dashed line). Indeed this option value should be present in many online markets: Einav et al. (2011) have found the standard deviation in price to be on the order of 10% of the transaction price in most eBay categories.

We also observe that the average time between bids is just over four days, with many auctions closing in the meantime. This suggests that bidders are inattentive (i.e. τ close to zero). In view of this, Assumption 1 seems quite reasonable: information from the current auction is probably a poor predictor of what conditions will be like the next time the bidder bids.

Is there unobserved product heterogeneity? The way we specify valuations below incorporates random coefficients over camera resolution, as well as common preferences for observable camera attributes (e.g. brand) and an unobservable component. To motivate the




Figure 2: Market Dynamics. The left panel shows the median daily price (including shipping) over the sample period, separately by camera resolution (4MP cameras omitted for clarity). The right panel shows the closing transaction prices on auctions of a Kodak Easyshare Z710 from a single seller over the sample period, superimposing over this the valuation, optimal bid and expected payment of a bidder with \$250 valuation (see text for more details).

inclusion of this unobserved product heterogeneity, we perform a simple heuristic estimation exercise. We take bids that have been “normalized” to account for observable heterogeneity (we explain this process below), and look at bidders who have bid exactly twice on cameras at two different resolution levels (e.g. 7MP and 8MP cameras). Dividing the difference in their normalized bids by the difference in the camera resolutions gives us a crude estimate of this bidder’s willingness to pay for megapixels, and the distribution of differences across such bidders gives us an idea of the diversity in bidders’ willingness to pay.

Figure 3 presents the density of that statistic computed at the bidder level for this restricted subset, without any kind of selection correction. While the mean of this distribution is positive (reassuringly) and the variance is large (motivating heterogeneous preferences for resolution), there is a troublesome and substantial mass to the left of zero. It is implausible that many consumers have negative marginal utility from higher resolution cameras, and we instead interpret this as evidence of unobserved product heterogeneity in our data that confounds this simple approach (we will offer more direct evidence below).




Figure 3: Density of $\frac{b_{i,j1} - b_{i,j2}}{\text{res}_{i,j1} - \text{res}_{i,j2}}$. This figure presents a kernel density plot of the ratio between the difference in bids across cameras of two resolution levels and the difference in resolution levels, across bidders.

5.3 Estimation

We construct a demand system in which consumers obtain value from the purchase of a single camera, a value which has an idiosyncratic component and a common component. Formally, bidder i 's valuation for product j offered in auction t is given by:

$$x_{i,j,t} = \underbrace{\alpha_{i,c} + \text{res}_j \alpha_{i,r}}_{\text{idiosyncratic}} + \underbrace{Z_{j,t} \gamma + \xi_{j,t}}_{\text{common}} \quad (14)$$

This combines the random coefficients of section 4.2, and the unobserved heterogeneity of 4.3.2 and adds some observable common demand shifters $Z_{j,t}$. Bidder i 's type in our model, α_i , is a double: their fixed utility draw for obtaining any camera ($\alpha_{i,c}$) as well as an idiosyncratic preference shock for resolution ($\alpha_{i,r}$). The auction-specific term $\xi_{j,t}$ captures unobserved heterogeneity that is observable to all bidders but not the econometrician. We assume that it is distributed normally with mean zero and variance $\sigma_{\xi,j}$ that varies freely with the resolution type of the camera. The bidder type α_i is drawn, iid upon bidder entry, from the distribution \mathbf{F}_α , which is our main estimation target. We assume mutual independence of $(\alpha_i, Z_{j,t}, \xi_j)$. On the supply side, we assume that the distribution of arriving auctions is

iid multinomial over the product space with a J -vector of probability weights π .

Estimation of supply and exit rates. Estimation begins with the recovery of the survival rate r , and the multinomial supply π . Note that with multinomial supply, the activity rate τ plays no role in optimal bidding and need not be estimated (whereas with Markov supply transitions Q , the activity rate drives a wedge between supply transitions and the state transitions \tilde{Q}).

Survival Rate of Bidders (r): We estimate this parameter using a censored negative binomial model fit to the likelihood of bidder exit using the full sample of bids. A bidder exits if they do not return to bid again. However we treat this outcome as censored if the last bid event was in the final six weeks of our dataset.¹⁴

Supply (π): $\hat{\pi}$ is the observed market share of each of the six resolution types.

Bid normalization. Cameras differ in both observable and unobservable features. The first step in working with the bids is to normalize out the contribution of observables for which bidders have common preferences (Haile et al., 2006). Applying our earlier theory to the valuations in (14), the bidding equation takes the form:

$$b_{i,j,t} = \max\{0, \alpha_{i,c} + \text{res}_j \alpha_{i,r} + Z_{j,t} \gamma + \xi_{j,t} - v(\alpha_{i,c}, \alpha_{i,r})\} \quad (15)$$

i.e. valuation less continuation value. Restricting attention to positive bids, and adding and subtracting $\psi_j \equiv E_{\mathbf{F}_\alpha} [\alpha_{i,c} + \text{res}_j \alpha_{i,r} - v(\alpha_{i,c}, \alpha_{i,r})]$, we get an estimating equation in reduced form:

$$b_{i,j,t} = \psi_j + Z_{j,t} \gamma + e_{i,j,t}$$

where the error term $e_{i,j,t}$ combines an individual-specific deviation in bid from that of the average type and the common unobserved heterogeneity. Since these are both independent of $Z_{j,t}$, we can estimate γ consistently by OLS with product fixed effects.

¹⁴We maintain the assumption of exogenous exit on the part of bidders, though it could be weakened using the methods of Section 4.5. While the assumption might seem somewhat unrealistic, we found no evidence in reduced-form regressions of a systematic correlation between exit of losing bidders and the level of their bid. In the end, we chose to highlight the extensions that were most necessary to fit the basic features of the data discussed above.

We only run this regression on a subsample of our data, consisting of the two highest bids in every auction. This is motivated by the behavioral intuition of Haile and Tamer (2003). Because eBay’s proxy bidding system is formally an ascending auction, the bid of the third, fourth, or n -th highest bidder may not reflect their intended final bid (what they would have bid in a sealed-bid auction i.e. valuation less continuation value). It is common for bidders to “test out” a sequence of ascending bids to see if they can become the standing high bidder, but if in this process the price comes to exceed their intended final bid then it will be censored from the dataset. For this reason we restrict attention to the two highest bids in each auction, who are never censored in this way.

As controls $Z_{j,t}$ we include product line fixed effects, listing attributes including shipping options, seller feedback, and optional listing features (e.g., sellers may pay a fee for their results to be highlighted in search results), as well as a set of dummies for resolution, optical zoom, and digital zoom levels.¹⁵ We report results for the main controls of interest in Table 3. They are generally sensible, with bids increasing in seller feedback, decreasing in shipping costs, and particularly high for listings with free shipping, consistent with the findings of Einav et al. (2011).

Define the normalized bids $\tilde{b}_{i,j,t} = b_{i,j,t} - Z_{j,t}\gamma = \alpha_{i,c} + \text{res}_j\alpha_{i,r} - v(\alpha_{i,c}, \alpha_{i,r}) + \xi_{j,t}$. We will work with our estimate of these bids $b_{i,j,t} - Z_{j,t}\hat{\gamma}$ in what follows. Consider the difference between two normalized bids on the same product, by the same bidder in two different auctions. These should differ only by $\xi_{j,t_1} - \xi_{j,t_2}$, and since these random variables are independent of each other, we can estimate the variance of $\xi_{j,t}$ as $\sigma_{\xi,j}^2 = \text{Var}(b_{i,j,t} - b_{i,j,t'})/2$ where the RHS variance is estimated by pooling over all $t < t'$ pairs available.

Opposing bids. Recall from the identification analysis of section 4.3.2 that the relevant strategic object for bidders to form beliefs about is $\{\check{G}_j^1\}$, the distribution of the highest competing bid net of unobserved heterogeneity (since $\xi_{j,t}$ just acts as a common bid shifter). So we need to estimate $\{\check{G}_j^1\}$. Also, for the reasons outlined above, we only want to use the top two bids in each auction in estimation. This places us in the “limited data” case, and will require correcting for selection into the top two bids. So we also need to estimate $\{\check{G}_j^2\}$.

We observe the raw distributions of the first and second highest normalized bids $\{G_j^1, G_j^2\}$ (i.e. the convolution of $\check{b}_{i,j}$ and $\xi_{j,t}$). Based on the shape of these, we assume a Gamma

¹⁵Note that we include resolution dummies to avoid omitted variable bias, however we exclude these coefficients when predicting the normalized bids below.

Table 3: First-Stage Normalization Regression

	(1) Winning bid
Free shipping	24.21*** (1.913)
Shipping cost	0.405*** (0.0446)
Seller feedback (thousands)	0.0665*** (0.00764)
R^2	0.721
N	18432

Notes: This table presents selected coefficients from the first-stage normalization regression. Unreported here, the regression also includes dummies for all (rounded) resolution, optical zoom, and digital zoom values, as well as product-level fixed effects. It also includes a large array of listing attributes such as featured listing status, whether the seller paid for a scheduled end-time, and dummies for bundled accessories. Note that the coefficients on the resolution dummies are set to zero before predicting the normalized bids for subsequent analysis. Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

distribution for $\{\check{G}_j^1, \check{G}_j^2\}$ with separate shape and scale parameters for each product and order statistic. We estimate these parameters by method of moments (i.e. we pick these parameters to match the mean and variance of the distribution of the normalized highest bid, taking our estimates of $\sigma_{\xi,j}^2$ in the earlier step as given).

We summarize the results of all these preliminary estimation steps in Table 4. We find that bidders are quite likely to return and bid again upon a loss ($r = 0.57$). Unobserved heterogeneity varies substantially with resolution type and is particularly important for 10MP cameras, where we anticipate there is more differentiation of high-end products.

Aside: consumer surplus. Even without estimating the full demand system, at this point we can already say something useful about consumer surplus. In the static model consumers bid their valuations whereas in our dynamic setting we have shown that they shade their bids substantially. This shading will bias static estimates of consumer surplus.¹⁶

In the static framework where $b_j = x_j$, the difference between the highest bid and the

¹⁶Coey et al. (2015) also make a point of computing consumer surplus in a dynamic framework, although theirs is based on valuations with “deadlines.”

Table 4: First-Stage Parameters

r						0.5646 (0.0014)					
π		5mp	6mp	7mp	8mp	10mp	0.1382 (0.0033)	0.3020 (0.0044)	0.3816 (0.0046)	0.0934 (0.0028)	0.0847 (0.0027)
σ_ξ		5mp	6mp	7mp	8mp	10mp	15.43 (0.1474)	12.75 (0.0729)	14.16 (0.0539)	11.16 (0.1581)	22.24 (0.1667)
$G^{(1)}$		5mp	6mp	7mp	8mp	10mp					
	k	12.57 (0.5235)	19.21 (0.5097)	17.06 (0.4253)	17.68 (1.0151)	38.98 (1.7325)					
	θ	10.04 (0.3798)	8.17 (0.2020)	10.83 (0.2522)	11.86 (0.6243)	7.12 (0.4260)					
$G^{(2)}$		5mp	6mp	7mp	8mp	10mp					
	k	11.63 (0.5355)	18.57 (0.5348)	16.33 (0.4240)	15.80 (0.9675)	44.71 (1.9974)					
	θ	9.91 (0.4035)	7.81 (0.2070)	10.59 (0.2544)	12.36 (0.6959)	5.83 (0.3753)					

Notes: This table presents estimates of the first-stage parameters of the model. See text for estimation details. Standard errors are presented in parentheses.

second-highest bid is a direct measure of surplus per auction. So for a psuedotype \mathbf{b} entering

a random auction according to π ,

$$\mathbb{E}[CS_{static}(\mathbf{b})] = \sum_j \pi_j \{G_j^1(b_j)[b_j - \mathbb{E}[B^{(1)}|b^{(1)} < b_j]]\}$$

To compute this object we take total surplus to be the total sum of $\{B^{(1)} - B^{(2)}\}_i$, i.e. the valuation of the first-highest bidder minus the price paid, across all auctions, and then divide by the total number of bidders. From Table 1 the average difference between the first- and second-highest bid in our dataset is \$10.60, which yields an expected consumer surplus of $(\$10.60 \times (19160 \text{ auctions} \div 74375 \text{ bidders})) = \2.63 per bidder in our dataset.

In the dynamic model, however, bidders shade their bids: $b_j = x_j - rv(\mathbf{b})$. Note moreover that, if we take bidders' beliefs to be correct, $v(\mathbf{b}) = \mathbb{E}[CS_{dynamic}(\mathbf{b})]$.¹⁷ In this setting,

$$\begin{aligned} \mathbb{E}[CS_{dynamic}(\mathbf{b})] &= \sum_j \pi_j \{G_j^1(b_j)[\underbrace{x_j}_{=b_j + \sum_k \tilde{Q}_{j,k} v_k(\mathbf{b})} - \mathbb{E}[B^{(1)}|b^{(1)} < b_j]] + (1 - G_j^1(b_j))r \sum_k \tilde{Q}_{j,k} v_k(\mathbf{b})\} \\ &= \underbrace{\sum_j \pi_j \{G_j^1(b_j)[b_j - \mathbb{E}[B^{(1)}|b^{(1)} < b_j]]\}}_{=\mathbb{E}[CS_{static}(\mathbf{b})]} + \underbrace{\sum_j \pi_j r \sum_k \tilde{Q}_{j,k} v_k(\mathbf{b})}_{=r\mathbb{E}[CS_{dynamic}(\mathbf{b})]} \\ &= \frac{1}{1-r} \mathbb{E}[CS_{static}(\mathbf{b})]. \end{aligned}$$

Note that in the second equality, we take advantage of the fact that π is the ergodic distribution of Q , and therefore $\tilde{Q}\pi = \pi$ for any τ .

This computation yields a substantially larger estimate of consumer surplus per bidder: $((0.4354)^{-1} \times \$2.73 =) \6.04 , or \$24.35 per auction. This large correction reflects the fact that the bidders who are most likely to win also shade the most, and therefore it is for these that the static model most underestimates valuations.

Estimation of Demand. At this stage, we have estimates of π , r and the distributions $\{G_j^1, G_j^2\}$ of first and second order statistics, after adjusting for observed and unobserved heterogeneity. We also have a dataset of normalized bids that were themselves first or second highest bids in auctions, and therefore plausibly equal to valuation less continuation value.

¹⁷This relies on an assumption that τ is sufficiently small that each time a bidder participates, sufficient time has passed that the market is once again in steady-state i.e. $v(\mathbf{b})$ can be calculated without sophisticated conditioning on histories. We could be more mathematically precise — as we were in the earlier selection correction argument — but as we argued when discussing Table 1, τ appears to be small.

In principle, we could continue by applying our nonparametric identification strategy from Theorem 4 directly in estimation, for bidders who are observed bidding on each least two different products (since the random coefficient is two-dimensional, two different products suffices for identification). But there are less than 300 such bidders in our dataset, and this makes the necessary deconvolution analysis unattractive in view of the slow convergence properties of such estimators (Carroll and Hall, 2004).

So instead we take a parametric approach, assuming that α_c and α_r are distributed $\Gamma(k_c, \theta_c)$ and $\Gamma(k_r, \theta_r)$, respectively. We chose the gamma distribution because our estimation sample consists of the right tail of bids, i.e. first- and second-highest bids, and so we want a parametric model with flexibility to fit the shape of that tail. Therefore $\theta \equiv \{k_c, \theta_c, k_r, \theta_r\}$ makes up the set of parameters of the demand system we ultimately hope to recover. Following the logic of our non-parametric identification argument, we need to see bidders bid multiple times in order to identify demand. So we estimate the model using a sample of bidders who submit either the first or second-highest bid in auctions on exactly two distinct resolution types, a sample of 264 bidders (very few bidders come first or second on three or more distinct products, so we simplify the analysis by excluding their data). Under our parametric assumptions, we have the following likelihood of each observation:

$$\mathcal{L}(\mathbf{b}_i|\theta) = \int \underbrace{\frac{\mathbb{P}\{\mathcal{B}_i|\beta(\boldsymbol{\alpha})\}}{\mathbb{P}\{|\mathcal{B}_i|=2|\beta(\boldsymbol{\alpha})\}}}_{\text{selection probability}} \left(\underbrace{\prod_{j \in \{\mathcal{B}_i\}} f_{\xi,j}(b_j - \beta_j(\boldsymbol{\alpha}))}_{\text{unobserved heterogeneity}} \right) \underbrace{dF(\boldsymbol{\alpha}|\theta)}_{\text{demand}}. \quad (16)$$

where \mathbf{b}_i is a two-vector, indicating an observed pair of bids by i ; \mathcal{B}_i is the set of products bidder i bids on; $\beta(\boldsymbol{\alpha})$ is type $\boldsymbol{\alpha}$'s pseudo-type (i.e. what they would bid on each product in the absence of unobserved heterogeneity); $f_{\xi,j}$ is the density of ξ_j , assumed normal with mean zero and variance $\sigma_{\xi,j}^2$ and $F(\boldsymbol{\alpha}|\theta)$ is the distribution of $\boldsymbol{\alpha}$ at parameter vector θ .

This likelihood function has three components: the first, $\mathbb{P}\{\mathcal{B}_i|\beta(\boldsymbol{\alpha})\}/\mathbb{P}\{|\mathcal{B}_i|=2|\beta(\boldsymbol{\alpha})\}$ is the selection probability; the likelihood that a bidder of type $\boldsymbol{\alpha}$ is observed in a subset of the product space \mathcal{B}_i conditional on the event $|\mathcal{B}_i|=2$, i.e. our sample construction. The second component of the likelihood function is the deviation of the observed bid \mathbf{b} from the predicted bid $\beta(\boldsymbol{\alpha})$ on the components $j \in \mathcal{B}_i$, which can be accounted for by unobserved heterogeneity. Finally, we integrate with respect to the type distribution $F(\boldsymbol{\alpha}|\theta)$, the only

point at which the parameter vector θ enters.¹⁸

Estimating bid functions and selection probabilities. In order to compute the likelihood of any observation at a parameter vector θ , we will need to compute the optimal bidding function $\beta(\boldsymbol{\alpha})$ and the selection probability $\mathbb{P}\{\mathcal{B}_i|\beta(\boldsymbol{\alpha})\}$. We start with the optimal bid function. In the case of iid supply, this is given in Lemma 6 by (12), where the continuation value is defined recursively according to the Bellman equation (13), taking $\pi_j^O = 0$ for all j . Fix a type $\boldsymbol{\alpha}$ and an associated J -vector of valuations \mathbf{x} . We solve for their optimal bids and associated continuation value by value iteration: given a (scalar) continuation value $v(\mathbf{x})$, define the following mapping from \mathbb{R}^+ into itself:

$$T_{\mathbf{x}}(v) = \frac{\pi \cdot (G^1(\mathbf{x} - v)(\mathbf{x} - \mathbb{E}_{G^1}[B^1|B^1 \leq \mathbf{x} - v]))}{1 - r((1 - G^1(\mathbf{x} - v)) \cdot \pi)}$$

where \cdot denotes dot product, and G^1 and E_{G^1} are now J -vectors. The continuation value for type $\mathbf{x}(\boldsymbol{\alpha})$ satisfies $T_{\mathbf{x}}(v) = v$, and from this we obtain the bidding function $\beta(\boldsymbol{\alpha}) = \max\{0, \mathbf{x}(\boldsymbol{\alpha}) - v(\boldsymbol{\alpha})\}$.

Next, given a bid \mathbf{b} associated with the type $\boldsymbol{\alpha}$, we work out the probability that a particular bidder is observed in a subset $\mathcal{B} \subseteq \mathcal{J}$ of the set of possible auctions. The building block for this is the function $\mathbb{P}(\mathcal{A}, \mathbf{b})$, defined as the probability that a bidder with bid vector \mathbf{b} , entering in a randomly sampled period, exits before they are “observed” in any auction outside of the set $\mathcal{A} \subseteq \mathcal{J}$. Define $\mathbb{P}(\mathcal{A}, \mathbf{b}, j)$ in the same way, but additionally conditioning on the bidder entering in state j , so that $\mathbb{P}(\mathcal{A}, \mathbf{b}) = \sum_j \pi_j \mathbb{P}(\mathcal{A}, \mathbf{b}, j)$. Recalling that we treat bidders as “observed” only when they make bids that qualify for our estimation sample (i.e. one of the top two bids in the auction), we can write the latter probability recursively as:

$$\begin{aligned} \mathbb{P}(\mathcal{A}, \mathbf{b}, j) = & 1(j \in \mathcal{A}) (G_j^1(b_j) + (1 - G_j^1(b_j))(1 - r + r\mathbb{P}(\mathcal{A}, \mathbf{b}))) \\ & + 1(j \notin \mathcal{A}) ((1 - G_j^2(b))((1 - r) + r\mathbb{P}(\mathcal{A}, \mathbf{b}))) \end{aligned} \quad (17)$$

On the first line of the RHS we consider the case where $j \in \mathcal{A}$: that bidder will be observed only in \mathcal{A} if they exit immediately (either by winning or exogenously exiting), or survive and then transition to a new random state, in which case the chance is $\mathbb{P}(\mathcal{A}, \mathbf{b})$. The second line

¹⁸In principle there is a conceptual problem here if the observations overlap in the sense that some of the bidders in our estimation sample participate in the same auctions, as this places restrictions on the value of ξ . In practice this overlap problem does not occur.




Figure 4: Estimated Bid Functions and Selection Probabilities. This figure presents bid functions and selection probabilities as a function of α_c , holding fixed α_r at two levels: 0 and 20.

is for the case where $j \notin \mathcal{A}$, and then the bidder must *not* be observed in j , so they must make a bid below the second highest. Given that this is the case, again they must either exit immediately or survive and we get $\mathbb{P}(\mathcal{A}, \mathbf{b})$ again.¹⁹

Now G_j^1 and G_j^2 are known, and $\mathbb{P}(\mathcal{A}, \mathbf{b})$ is a weighted sum of the $\mathbb{P}(\mathcal{A}, \mathbf{b}, j)$, so for any (\mathcal{A}, b) we can get $\mathbb{P}(\mathcal{A}, \mathbf{b})$ as the solution to a linear system with unknowns $\{\mathbb{P}(\mathcal{A}, \mathbf{b}, j)\}_{j \in \mathcal{J}}$. Then we can use these objects to compute the probability that a bidder is observed in *exactly* the set \mathcal{B} . For instance, if we are interested in the probability that a bidder of type α is observed in a subset $\mathcal{B} = \{2, 4\}$ we can compute $\mathbb{P}(\{2, 4\} | \beta(\alpha)) = \mathbb{P}(\{2, 4\}, \beta(\alpha)) - \mathbb{P}(\{2\}, \beta(\alpha)) - \mathbb{P}(\{4\}, \beta(\alpha)) + \mathbb{P}(\{\phi\} | \beta(\alpha))$.

¹⁹This calculation assumes that the true state space coincides with the one that bidders use in forming beliefs. This is not necessary (we could condition on history), but is practical given our data constraints.

Figure 4 illustrates the resulting bid functions and selection probabilities when these procedures are applied to our data. The selection probability is for the event that a bidder bids on exactly two different products ($|\mathcal{B}| = 2$), i.e. that a bidder is selected into our estimation sample. Panel (a) shows the bid function for a bidder with $\alpha_r = 0$, i.e. a bidder who places no value on resolution, and consequently bids the same on every product. Their bid function “peels away” from the 45-degree line. This makes intuitive sense: as α_c increases, their continuation value increases as well, causing them to shade their bids. Panel (b) repeats this exercise for the case where $\alpha_r = 20$. They now bid differently on each product, with the five dotted lines showing their bids on the lowest resolution camera (bottom line) up to the top resolution camera (top line), as α_c varies. The shape is largely preserved, except that now there is a positive intercept because valuations are positive even with $\alpha_c = 0$.

In panels (c) and (d) we show the selection probabilities for these two types of bidders ($\alpha_r = 0$ in (c) and $\alpha_r = 20$ in (d)) as α_c varies. The “hump-shape” arises because the probability of $|\mathcal{B}| = 2$ is increasing at first as the likelihood that the bidder is ever first or second rises, but later declines as the probability that they win their first auction and exit before bidding again goes to one.

Optimization. We now proceed to the problem of optimizing the likelihood. Usually one would proceed by applying some optimization routine in an outer loop to the likelihood (16), evaluating it by Monte Carlo integration, holding the randomness in the draws of α fixed across evaluations k (Pakes and Pollard, 1989). But this is computationally costly, as for many different α draws we will have to evaluate the bidding and selection functions.

Notice that the parameter vector θ enters the likelihood (16) only through the distribution of random coefficients $F(\alpha|\theta)$. This allows us to employ an alternative approach based on re-weighting. Fixing α , define

$$h_i(\alpha) \equiv \frac{\mathbb{P}\{\mathcal{B}_i|b(\alpha)\}}{\mathbb{P}\{\phi|b(\alpha)\}} \left(\prod_{j \in \mathcal{B}_i} \phi \left(\frac{b_i - b_j(\alpha)}{\sigma_{\xi,j}} \right) \right), \quad (18)$$

and note that this object is computable for any α from the results of the first-stage estimation. Therefore, taking as our objective function the log-likelihood of the dataset, we can write

$$L(\theta, \{b_i\}) = \sum_i \log \int h_i(\alpha) dF(\alpha|\theta). \quad (19)$$

One natural and computationally efficient way to proceed would then be to sample α uniformly from the type space S times (for S the number of draws in the Monte Carlo integration), and then optimize the simulated criterion function by re-weighting their likelihood contribution according to $f(\alpha|\theta)$. Instead, we choose to use importance sampling, i.e. sampling from and re-weighting relative to a user-chosen distribution, denoted $G(\alpha)$. We do this for two reasons: first, for efficiency — we would like to sample more points from the center of the true distribution if possible, since this is where the gradient of the likelihood is sharpest. Second, we would like the sampled distribution to have unbounded support, mirroring our parametric assumption on $F(\alpha|\theta)$.

Let the set of points be drawn from $G(\alpha)$; then we have:

$$\hat{\theta} = \arg \max_{\theta} L(\theta, \{b_i\}) = \sum_i \log \sum_{s=1}^S h_i(\alpha) \frac{f(\alpha_s|\theta)}{g(\alpha_s)}. \quad (20)$$

We implemented the estimator in two steps: first, we choose G_1 to be a normal with very large variance, and obtained a first estimate $\tilde{\theta}$. In a second step, we drew from a distribution $G_2(\alpha) = F(\alpha|\tilde{\theta})$ and re-optimized. This two-step procedure should improve the efficiency of our estimator, as the second sample should be more centered and thus provide a better approximation to the integral.

5.4 Results

Estimates from this ML exercise are presented in Table 5. Our results suggest that camera resolution is the main determinant of consumer utility, but that there is substantial heterogeneity in how consumers value this attribute. They predict a mean taste for cameras of 22.26 with a standard deviation of 51.06, and a mean taste for resolution of 22.75 and a standard deviation of 6.41. Recall that we have estimated \mathbf{F} , whereas the observed bids in the data (i.e., first- or second-highest bids) will tend to come from bidders sampled the right tail of the distribution.

As a measure of goodness of fit we also compute moments of the bid distribution that were not used directly in estimation. In Table 6 we present means and standard errors of the observed bids by product type as well as means and standard errors as predicted by our model at $\hat{\theta}$. In order to compute the latter we simulated 10,000 draws from $F(\alpha|\hat{\theta})$ and then

Table 5: Demand System Estimates

	α_c	α_r
\hat{k}	0.19 (0.06)	13.73 (1.74)
$\hat{\theta}$	117.14 (31.70)	1.73 (0.22)

Notes: This table presents estimates for $\{k_c, \theta_c, k_r, \theta_r\}$, the parameters governing the distribution of random coefficient preferences for compact cameras according to the utility function (14).

Table 6: Sample and Simulated Moments from the Bid Distribution

	5mp	6mp	7mp	8mp	10mp
Sample mean bid	122.82	152.51	181.59	202.43	272.07
Simulated mean bid	120.61	144.56	168.25	192.05	239.68
Sample standard deviation	38.56	38.21	47.23	51.18	48.93
Simulated standard deviation	36.29	38.96	43.86	47.77	58.23

Notes: This table compares moments from the sample bid distribution to moments that are predicted from the estimates in Table 5. Simulated moments are computed by taking 100,000 draws from $F(\alpha|\hat{\theta})$, computing bid vectors for each type, simulating unobserved heterogeneity ξ , and taking a weighted average according to the probability that each bid is selected into the sample of observed bids.

weighted their bids according to the probability of being observed in each state.

The results in Table 6 suggests that our results fit the data reasonably well, though there is some divergence, in particular for the higher-end cameras in the market. We believe this to be largely a function of the parametric compromises we made in order to focus on the sharply identified set of 264 bidders who bid on multiple different resolution types. However for 6MP to 8MP cameras, which make up the bulk of our sample, we match these moments quite closely. Note that the moments are out-of-sample in that they incorporate bidders who are observed once but then exit; these bidders are not used in estimation.

6 Conclusion

This paper offers a flexible demand system for the study of auction markets *as markets*. We developed a notion of equilibrium in such markets and proved its existence, and in turn were able to characterize the conditions under which bidders' actions can be inverted to

infer their private type. While this is sufficient for identification if we treat auctions as independent, isolated draws from a distribution, in an auction marketplace we also need to account for the selection of bidders into the observed and identified set. Subject to the constraints of non-participation, we are able to partially identify the distribution of types by explicitly modeling this selection as a function of observable equilibrium objects. We show that this result extends to accommodate familiar features of auction data and demand models in fixed-price markets, including data limitations, random coefficient demand, unobserved heterogeneity, and latent outside options.

The selection correction turns out to be an important source of flexibility in the model. It allows us to accommodate standard limitations of auction data. It also gives the econometrician flexibility to pick and choose sources of variation in the data — in this case particular histories — that they trust for identification.

A second important source of flexibility in the model is that we have made all the substitution between products occur *inter-temporally*. High-dimensional preferences are thus projected down to a valuation for the current product and a continuation value which, thanks to the Markov dynamics of the game, we can model in a straightforward way. These Markov dynamics also give us the flexibility to extend the state space of the game to incorporate arbitrary *public* signals about the state of the market on which bidder behavior may depend.

We illustrated much of this flexibility in an application to the auction market for compact cameras on eBay. There we estimated the distribution of types for a utility function that looked a lot like the kind you might estimate in a fixed-price market using standard methods in industrial organization. While we were able to document a substantial bias in the estimates of consumer surplus that would come out of an alternative, static model, as well as match out-of sample moments of the bid distribution, the main objective was to illustrate the kinds of judgment calls faced by an econometrician applying our method, which can accommodate a wide array of data limitations, utility functions, information structures, and assumptions about outside options.

Our contribution — the development of a demand system for auction markets — is meant to mirror similar work in fixed-price markets, and to a similar end: the structural estimation of demand allows us to do counterfactuals of both private and public interest. Though we have developed a flexible framework, there remain several open directions for future work: we believe that our framework would generalize naturally to multi-unit demand, in which

bidders may shade against the opportunity cost of moving further down their marginal utility curve. This is an important direction for the modeling of ad auctions and treasury auctions. There are also unmet challenges in the modeling of substitution across auctions within-period rather than inter-temporally. Auction markets are a pervasive mechanism for the allocation of goods and services, and beyond our contribution there remains much work to be done to understand competition between and substitution among them.

References

- ACKERBERG, D. A., K. HIRANO, AND Q. SHAHRIAR (2009): “Identification of Time and Risk Preferences in Buy Price Auctions,” Working paper, UCLA.
- ADAMS, C. P. (2009): “Identifying Demand from Online Auction Data,” Working Paper, Federal Trade Commission.
- AGUIRREGABIRIA, V. AND P. MIRA (2007): “Sequential Estimation of Dynamic Discrete Games,” *Econometrica*, 75, 1–53.
- 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, 2107–2140.
- BAJARI, P., L. BENKARD, AND J. LEVIN (2007): “Estimating Dynamic Models of Imperfect Competition,” *Econometrica*, 75, 1331–1370.
- BAJARI, P. AND A. HORTAÇSU (2004): “Economic Insights from Internet Auctions,” *Journal of Economic Literature*, 42, 457–486.
- BERRY, S. (1994): “Estimating Discrete Choice Models of Product Differentiation,” *Rand Journal of Economics*, 25, 242–262.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63, 841–890.

- (2004): “Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market,” *Journal of Political Economy*, 112, 68–105.
- BERRY, S. AND A. PAKES (2007): “The Pure Characteristics Demand Model,” *International Economic Review*, 48, 1193–1225.
- BODOH-CREED, A., J. BÖEHNKE, AND B. HICKMAN (2013): “Optimal Design of Two-Sided Market Platforms: An Empirical Case Study of Ebay,” Working Paper.
- BUDISH, E. (2008): “Sequencing and Information Revelation in Auctions for Imperfect Substitutes: Understanding eBay’s Market Design,” Working paper, Harvard University.
- CARROLL, R. J. AND P. HALL (2004): “Low order approximations in deconvolution and regression with errors in variables,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66, 31–46.
- COEY, D., B. LARSEN, AND B. PLATT (2015): “A Theory of Bidding Dynamics and Deadlines in Online Retail,” .
- EINAV, L., C. FARRONATO, J. LEVIN, AND N. SUNDARESAN (2016): “Auctions versus Posted Prices in Online Markets,” Working Paper.
- EINAV, L., T. KUCHLER, J. LEVIN, AND N. SUNDARESAN (2011): “Learning from Seller Experiments in Online Markets,” Working Paper, Stanford.
- EVDOKIMOV, K. AND H. WHITE (2012): “SOME EXTENSIONS OF A LEMMA OF KOTLARSKI,” *Econometric Theory*, 28, 925–932.
- FERSHTMAN, C. AND A. PAKES (2012): “Dynamic Games with Asymmetric Information: A Framework for Empirical Work,” Forthcoming in the *Quarterly Journal of Economics*.
- 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, 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.

- HAILE, P. A. AND E. TAMER (2003): “Inference with an Incomplete Model of English Auctions,” *The Journal of Political Economy*, 111, 1–51.
- 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, North-Holland, vol. III.
- HENDRICKS, K. AND A. SORENSEN (2014): “The Role of Intermediaries in Dynamic Auction Markets,” Working Paper, University of Wisconsin.
- HOPENHAYN, H. AND M. SAEEDI (2016): “Dynamic Bidding in Second Price Auctions,” Working Paper.
- INGSTER, A. (2009): “A Structural Study of a Dynamic Auction Market,” Working Paper, NYU Stern.
- IYER, K., R. JOHARI, AND M. SUNDARARAJAN (2014): “Mean Field Equilibria of Dynamic Auctions with Learning,” *Management Science*, 60, 2949–2970.
- JOFRE-BONET, M. AND M. PESENDORFER (2003): “Estimation of a Dynamic Auction Game,” *Econometrica*, 71, 1443–1489.
- KOTLARSKI, I. (1967): “On characterizing the gamma and the normal distribution.” *Pacific J. Math.*, 20, 69–76.
- KRASNOKUTSKAYA, E. (2011): “Identification and Estimation in Highway Procurement Auctions under Unobserved Auction Heterogeneity,” *Review of Economic Studies*, 71, 293–327.
- KRASNOKUTSKAYA, E., K. SONG, AND X. TANG (2014): “The Role of Quality in On-Line Service Markets,” Working paper, Johns Hopkins.
- KRASNOKUTSKAYA, E., C. TERWIESCH, AND L. TIEREROVA (2016): “Trading across Borders in Online Auctions,” Working Paper, Johns Hopkins.
- KRUSELL, P. AND A. A. SMITH (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106, 867–896.
- LEWIS, G. (2007): “Asymmetric Information, Online Disclosure and Adverse Selection: The Case of eBay Motors,” Working paper, Harvard.

- LEWIS, G. AND P. BAJARI (2011): “Procurement Contracting with Time Incentives: Theory and Evidence,” *Quarterly Journal of Economics*, 126.
- LI, T. AND Q. VUONG (1998): “Nonparametric Estimation of the Measurement Error Model Using Multiple Indicators,” *Journal of Multivariate Analysis*, 65, 135–165.
- MENZEL, K. AND P. MORGANTI (2013): “Large Sample Properties for Estimators Based on the Order Statistics Approach in Auctions,” *Quantitative Economics*, 4, 329–375.
- NEKIPELOV, D. (2007): “Entry Deterrence and Learning Prevention on eBay,” Mimeo, UC Berkeley.
- NEWBERRY, P. W. (2015): “The effect of competition on eBay,” *International Journal of Industrial Organization*, 40, 107 – 118.
- 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, 373–399.
- PAKES, A. AND D. POLLARD (1989): “Simulation and the Asymptotics of Optimization Estimators,” *Econometrica*, 57, 1027–1057.
- PLATT, B. (2015): “Inferring Ascending Auction Participation from Observed Bidders,” Working Paper.
- ROBERTS, J. W. (2013): “Unobserved heterogeneity and reserve prices in auctions,” *The RAND Journal of Economics*, 44, 712–732.
- ROTH, A. E. AND A. OCKENFELS (2002): “Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet,” *American Economic Review*, 92, 1093–1103.
- 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.
- SONG, U. (2004): “Nonparametric Estimation of an eBay Auction Model with an Unknown Number of Bidders,” Working Paper, University of British Columbia.

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, 1375–1411.

ZEITHAMMER, R. (2006): “Forward-looking bidding in online auctions,” *Journal of Marketing Research*, 43, 462–476.

A Proofs Appendix

Lemma 1.

Proof. Recall that G_j^1 is the CDF of the highest competing bid for good j . Then we have:

$$\beta_j(\mathbf{x}) = \arg \max_{b \in \mathbb{R}^+} \int (1(b \geq B_j^1 \wedge b > 0)(x_j - B_j^1) + 1(b < B_j^1 \vee b = 0)\tilde{v}_j(\mathbf{x})) dG_j^1(B_j^1) \quad (21)$$

for $\tilde{v}_j(\mathbf{x}) = r \sum_j \tilde{Q}_{j,k} v_k(\mathbf{x})$. Consider the RHS of (21). The integrand takes the value $x_j - B_j^1$ when $B_j^1 \leq b$ and $\tilde{v}_j(\mathbf{x})$ when $B_j^1 > b$ and so can be maximized by choosing b so that the event $B_j^1 \leq b$ occurs iff $x_j - B_j^1 - \tilde{v}_j(\mathbf{x}) \geq 0$. There are two relevant cases: first, if $\tilde{v}_j(\mathbf{x}) > x_j$, then the bidder strictly prefers to lose the auction. In this case they bid zero, equivalent to non-participation in our framework. Otherwise, $b = x_j - \tilde{v}_j(\mathbf{x})$ maximizes the RHS of (21) and thus $\beta_j(\mathbf{x}) = \max\{x_j - r \sum_j \tilde{Q}_{j,k} v_k(\mathbf{x}), 0\}$. If $\beta_j(\mathbf{x})$ is interior to the support of G_j^1 , then $\int 1(B_j^1 \leq b)(x_j - B_j^1) dG_j^1(B_j^1)$ is strictly decreasing, and so the crossing with $v_j(\mathbf{x})$ — and therefore the argmax of (21) — is unique.

For continuity, consider two types \mathbf{x}^1 and \mathbf{x}^2 . We know that $\max_j v_j(\mathbf{x}^2) - v_j(\mathbf{x}^1)$ is at most $\max_j |x_j^2 - x_j^1|$ because \mathbf{x}^1 can copy the strategy of \mathbf{x}^2 , getting the same allocations and making the same payments, so the only difference in value functions comes from the underlying valuations. Continuity of the value functions immediately follows. Bidding strategies take the form $\max\{x_j - \tilde{v}(\mathbf{x}), 0\}$, and since $x_j - \tilde{v}(\mathbf{x})$ is continuous in \mathbf{x} and the maximum of two continuous functions is itself continuous, we are done.

For monotonicity, consider two types \mathbf{x}^1 and \mathbf{x}^2 with $x_k^1 = x_k^2$ for all $k \neq j$ and $x_j^1 < x_j^2$. Now $v_k(\mathbf{x}^2) - v_k(\mathbf{x}^1) \leq x_j^2 - x_j^1$ for all k , since \mathbf{x}^1 can follow the strategy $\beta(\mathbf{x}^2)$ and

get the exact same payoff as type \mathbf{x}^2 except when bidding on product j , when they get a payoff at most $x_j^2 - x_j^1$ lower. Also $v_k(\mathbf{x}^2) - v_k(\mathbf{x}^1) \geq 0$ for all k since \mathbf{x}^2 can follow $\beta(\mathbf{x}^1)$ and get at least as high a payoff. Then if $\beta_j(\mathbf{x}^2) > 0$ and $\beta_j(\mathbf{x}^1) > 0$, we have: $\beta_j(\mathbf{x}^2) - \beta_j(\mathbf{x}^1) = x_j^2 - x_j^1 - r \sum_k \tilde{Q}_{j,k}(v_k(\mathbf{x}^2) - v_k(\mathbf{x}^1)) \geq x_j^2 - x_j^1 - r(x_j^2 - x_j^1) > 0$ proving the bid function is strictly increasing on this region; and examining the other cases (e.g. $\beta_j(\mathbf{x}^1) = 0$) yields that is weakly increasing. Finally, if $\beta_k(\mathbf{x}^2) > 0$ and $\beta_k(\mathbf{x}^1) > 0$, then $\beta_k(\mathbf{x}^2) - \beta_k(\mathbf{x}^1) = -r \sum_l \tilde{Q}_{k,l}(v_l(\mathbf{x}^2) - v_l(\mathbf{x}^1)) \leq 0$ which (together with analysis of other cases) shows the bid function is decreasing in the value of other products and strictly so if the $\beta_j(\mathbf{x}^2) > 0$. \square

Theorem 1.

Proof. We prove existence of a pure strategy equilibrium in continuous strategies. Let $C(\mathcal{X})$ be the space of continuous functions on \mathcal{X} , metrized by the sup norm. Since \mathcal{X} is compact, $C(\mathcal{X})$ is a Banach space. Moreover $C(\mathcal{X})$ is convex. Define the best response function $\Gamma(\beta)$ to any strategy $\beta \in C(\mathcal{X})$ as in Lemma 1, i.e. $\Gamma(\beta)(x) = \max\{x_j - \tilde{v}_j(\mathbf{x}), 0\}$, where \tilde{v} depends implicitly on the rival strategies β through the ergodic measure over the distribution of highest bids μ_j . Now $\Gamma(\beta)$ is uniformly bounded for every $\beta \in C(\mathcal{X})$, since 0 and \mathbf{x} respectively lower and upper bound all best responses for each \mathbf{x} , and \mathcal{X} is bounded. Moreover, every $\beta \in \Gamma(C(\mathcal{X}))$ admits a modulus of continuity of 1, since by the argument of Lemma 1, types that are ε -close make ε -close best responses to any opposing bid distribution. This establishes that the set of functions $\Gamma(C(\mathcal{X}))$ is uniformly equicontinuous, and it follows by Arzelà-Ascoli that $\Gamma(C(\mathcal{X}))$ is a compact subset of $C(\mathcal{X})$.

Next, we need to show that Γ is continuous in β when $\beta \in C(\mathcal{X})$. The first step is to show that the distribution of highest bids μ_j is continuous in β in the weak-* topology, because payoffs depend on rivals strategies only through the distribution of the highest competing bid. Now, this distribution arises from the composition of a sequence of functions: $\mu_j = g(\kappa_j(\beta(\mu_F)))$, where μ_F is the underlying type distribution and β is the bidding function, so that $\beta(\mu_F)$ is the distribution of pseudo-types $\tilde{\mu}_j$. The function κ_j maps the pseudo-type distribution into an ergodic distribution over the vector of bids in a auction for good j (a vector of bids of random length). In earlier working papers we have shown that κ_j is a continuous function (i.e. the ergodic distribution is well defined, unique, and smoothly varies with the bids of

entrants).²⁰ Finally g maps this distribution into the distribution of the highest bid by taking the max of each vector, itself a continuous operation. So if f is any continuous bounded real-valued function, $f \circ g \circ \xi_j$ is also continuous, implying that for β' close to β in the sup norm, $f \circ g \circ \xi_j \circ \beta(\mathbf{x})$ close to $f \circ g \circ \xi_j \circ \beta'(\mathbf{x})$ pointwise, and so certainly $E_{\mu_j(\beta)}f$ close to $E_{\mu_j(\beta')}f$, proving weak convergence.

The second step is to show that if the highest bid distributions are close in the sense of weak convergence, then so are the best responses. We first argue that the value functions must be close. Let $v_j(\mathbf{x}, \mu)$ be the value of the game for \mathbf{x} starting from an auction for product j , when the bid distributions are given by $\mu = \{\mu_j\}$. Since the player can only win once and exits exogenously at rate $(1 - r)$, this value is upper bounded for every j by the maximum expected single-period payoff when bidding optimally (over all j). The single-period payoff is $u_j(\mathbf{x}, \mu) = \int_0^{\beta_j(\mathbf{x}, \mu)} (x_j - B_j^1) d\mu_j(B_j^1)$. So under the sup norm $\|v_j(\mathbf{x})\| = \sup_j \sup_{\mathbf{x}} v_j(\mathbf{x})$:

$$\begin{aligned} \|v_j(\mathbf{x}, \mu) - v_j(\mathbf{x}, \mu')\| &\leq \|u_j(\mathbf{x}, \mu) - u_j(\mathbf{x}, \mu')\| \\ &= \left\| \int_0^{\beta_j(\mathbf{x}, \mu)} (x_j - B_j^1) d\mu_j(B_j^1) - \int_0^{\beta_j(\mathbf{x}, \mu')} (x_j - B_j^1) d\mu'_j(B_j^1) \right\| \\ &\leq \left\| \max \left\{ \int_0^{\beta_j(\mathbf{x}, \mu)} (x_j - B_j^1) d\mu_j(B_j^1) - \int_0^{\beta_j(\mathbf{x}, \mu)} (x_j - B_j^1) d\mu'_j(B_j^1), \right. \right. \\ &\quad \left. \left. \int_0^{\beta_j(\mathbf{x}, \mu')} (x_j - B_j^1) d\mu_j(B_j^1) - \int_0^{\beta_j(\mathbf{x}, \mu')} (x_j - B_j^1) d\mu'_j(B_j^1) \right\} \right\| \end{aligned}$$

where the last line follows since making the strategies the same across the two different distributions of rival bids must make one of the terms smaller and so the maximum payoff differential bigger. Now since the integrands in the final line are continuous, if μ is close to μ' in the weak-* topology, the final expression must be close to zero, which proves continuity by a sandwich argument. This immediately suffices to prove continuity of the best responses (from their functional form). Putting this all together, we get the required continuity of Γ in β . So Γ is a continuous mapping from a convex Banach space into a compact subset of that space, and thus has a fixed point by Schauder's fixed point theorem.

We prove uniqueness in the case $J = 1$ by constructing a contraction mapping. Let β and β' be two equilibrium bidding strategies, necessarily monotone and continuous by earlier results. By definition of equilibrium, $\Gamma(\beta) = \beta$ and $\Gamma(\beta') = \beta'$. Let $v(x, \beta)$ and $v(x, \beta')$ be the corresponding value functions (dropping the vector notation since there is a single

²⁰Details available from the authors on request.

product). Then Γ is a contraction:

$$\begin{aligned}
\|\Gamma(\beta) - \Gamma(\beta')\| &= \max_x (x - rv(x, \beta)) - (x - rv(x, \beta')) \\
&< \max_x |G_\beta^1(\beta(x)) (x - \mathbb{E}_\beta[B^1 | B^1 \leq \beta(x)]) \\
&\quad - G_{\beta'}^1(\beta'(x)) (x - \mathbb{E}_{\beta'}[B^1 | B^1 \leq \beta'(x)])| \\
&\leq \max_x G_\beta^1(\beta(x)) |\mathbb{E}_\beta[B^1 | B_1 \leq \beta(\mathbf{x})] - \mathbb{E}_{\beta'}[B_1 | B_1 \leq \beta'(x)]| \\
&\leq \|\beta - \beta'\|
\end{aligned}$$

where in the third line we use the fact that $G_\beta^1(\beta(x)) = G_{\beta'}^1(\beta'(x))$ for all x . This property holds because under both strategies the same winner is chosen in each auction (types are totally ordered), so that the distribution of highest bids is equal across strategies if evaluated at the bids of a fixed type x in any state s . Therefore Γ is a contraction which guarantees uniqueness of the fixed point. \square

Lemma 2.

Proof. Let $G^1(\mathbf{b})$ be the J -vector giving the probability the pseudo-type \mathbf{b} will win on each product (i.e. it stacks $G_j^1(b_j)$); similarly let $E[\mathbf{B}^1 | \mathbf{B}^1 < \mathbf{b}]$ be J -vector of expected payments conditional on winning. Write the vector version of the Bellman equation in (1):

$$\begin{aligned}
v(\mathbf{x}) &= G^1(\mathbf{b}) (\mathbf{x} - E[\mathbf{B}^1 | \mathbf{B}^1 < \mathbf{b}]) + (1 - G^1(\mathbf{b}))r\tilde{Q}v(\mathbf{x}) \\
&= G^1(\mathbf{b}) (\mathbf{b} + r\tilde{Q}v(\mathbf{x}) - E[\mathbf{B}^1 | \mathbf{B}^1 < \mathbf{b}]) + (1 - G^1(\mathbf{b}))r\tilde{Q}v(\mathbf{x}) \\
&= G^1(\mathbf{b}) (\mathbf{b} - E[\mathbf{B}^1 | \mathbf{B}^1 < \mathbf{b}]) + r\tilde{Q}v(\mathbf{x})
\end{aligned}$$

where in the second line we use the bidding characterization of Lemma 1 to conclude that $G^1(\mathbf{b})\mathbf{x} = G^1(\mathbf{b})(\mathbf{b} + r\tilde{Q}v(\mathbf{x}))$. Re-arrange terms in the final line to get the value function as the solution to a linear system:

$$v(\mathbf{x}) = (I - r\tilde{Q})^{-1}G^1(\mathbf{b}) (\mathbf{b} - E[\mathbf{B}^1 | \mathbf{B}^1 < \mathbf{b}])$$

Since $r \in (0, 1)$ the matrix $(I - r\tilde{Q})$ is invertible (Stokey et al., 1989). Now by definition the

j -th component of ξ is

$$\xi_j(\mathbf{b}) = b_j + r \left[\tilde{Q}(I - r\tilde{Q})^{-1}G^1(\mathbf{b}) (\mathbf{b} - E[\mathbf{B}^1 | \mathbf{B}^1 < \mathbf{b}]) \right]_j = b_j + r \left[\tilde{Q}v(\mathbf{x}) \right]_j$$

where the last equality follows from the previous display (the j subscript indicates the j -th element of the vector). By Lemma 1, $b_j = \max\{x_j - r \left[\tilde{Q}v(\mathbf{x}) \right]_j, 0\}$ so if $b_j > 0$ we have $x_j = b_j + r \left[\tilde{Q}v(\mathbf{x}) \right]_j = \xi_j(\mathbf{b})$; and if $b_j = 0$ we have $x_j - r \left[\tilde{Q}v(\mathbf{x}) \right]_j \leq 0$ implying $x_j \leq r \left[\tilde{Q}v(\mathbf{x}) \right]_j = \xi_j(\mathbf{b})$. Together this gives $\mathbf{x} \leq \xi(\mathbf{b})$ with equality on the dimensions where $b_j > 0$. \square

Lemma 3.

Proof. Fix a pseudo-type \mathbf{b} . Let A_j be the event “ \mathbf{b} bids b_j on product j and loses, j periods after entry” (e.g. A_1 is bidding and losing on product 1 directly after entering). Then

$$s_H(\mathbf{b}) = \mathbb{P}(H|\mathbf{b}) = \mathbb{P}(\cap_{j=1}^J A_j) = \mathbb{P}(A_1)\mathbb{P}(A_2|A_1)\mathbb{P}(A_3|A_2, A_1) \dots \mathbb{P}(A_J|A_{J-1} \dots A_1)$$

Now, $\mathbb{P}(A_1) = r\tau\pi_1(1 - G_j^1(b_j))$ for π_1 the steady-state probability that good 1 is supplied. And $\mathbb{P}(A_j|A_{j-1} \dots A_1) = r\tau Q_{j-1,j}(1 - G_j^1(b_j|A_{j-1} \dots A_1))$. By Assumption 2, r , τ , and Q are identified, as are the conditional probabilities of winning $\{G_j^1(b|\cdot)\}$ following the (measurable) sequences of events $\{A_k\}_{k=1}^{j-1}$, so $s_H(\mathbf{b})$ is identified. \square

Theorem 2.

Proof. We will prove that $\tilde{\mathbf{F}}$ is point identified; the partial identification of \mathbf{F} via (6) follows immediately. To do this, we must find the unconditional density $\tilde{f}(\mathbf{b})$. We will do this iteratively, starting with pseudo-types who participate on all products (i.e. $\mathbf{b} > 0$), and proceeding with pseudo-types who participate on all but one product etc. Formally, we will show that $\tilde{f}(\mathbf{b})$ is identified by induction on the number of zero elements of \mathbf{b} , first showing that with no zero elements the density is identified (the base step), and then that if the density is identified for points with $n - 1$ zero elements, it is also identified for those with n zero elements (the induction step).

Base step. In the main text we show that $\tilde{f}(\mathbf{b}|\mathbf{b} > 0)$ is identified. We also observe the probability that a randomly chosen bidder will have a history that falls into H , $P(H)$.

Putting this together we get

$$\tilde{f}(\mathbf{b}) = \tilde{f}(\mathbf{b}|\mathbf{b} > 0) \frac{P(H)}{\int s_H(\mathbf{b}) \tilde{f}(\mathbf{b}|\mathbf{b} > 0) db}$$

where the second expression is the probability that a random pseudo-type has $\mathbf{b} > 0$.

Induction step. We must show that if $\tilde{f}(\mathbf{b})$ is identified for all points \mathbf{b} with $n - 1$ zero elements, then it is identified whenever \mathbf{b} has n zero elements. Fix a \mathbf{b} with n zero elements. Wlog, assume they bid \mathbf{b}^n on the first $J - n$ products (i.e. \mathbf{b}^n is a $J - n$ vector) and do not participate on the rest (i.e. bid zero). Let $f^n(\mathbf{b}^n)$ be the unconditional density of pseudo-types restricted to the first $J - n$ elements (i.e. $f^n(\mathbf{b}^n) = \int \tilde{f}(\mathbf{b}^n, s) ds$).

Consider the event H^n , that a bidder enters when good 1 is available, are active, lose and survive, then next period good 2 is available, they are active, lose and survive ... good $J - n$ is available, and they are active, bid, lose and survive. The joint density of the $(J - n)$ -dimensional bid vector conditional on this event, denoted $g_{H^n}(\mathbf{b}^n)$, is observed. By arguments similar to those in Lemma 3, the probability of ending up in that set for any \mathbf{b}^n , denoted $s_{H^n}(\mathbf{b}^n)$, is identified (note that evaluating this probability does not require a full J -vector of bids). And proceeding as in the base step, we can get the unconditional density of the $(J - n)$ -dimensional bid vectors, $f^n(\mathbf{b}^n)$.

Now, by definition $f^n(\mathbf{b}^n) = \int \tilde{f}(\mathbf{b}^n, s) ds$, which implies that $\tilde{f}(\mathbf{b}) = f^n(\mathbf{b}^n) - \int_{s \neq 0} \tilde{f}(\mathbf{b}^n, s) ds$. And by the induction assumption, $\tilde{f}(\mathbf{b}^n, s)$ is identified whenever s is not identically zero, implying that both $f^n(\mathbf{b}^n)$ and $\int_{s \neq 0} \tilde{f}(\mathbf{b}^n, s) ds$ are known, and so $\tilde{f}(\mathbf{b})$ is identified.

□

Theorem 3.

Proof. The text makes most of the required arguments. The missing step is to show that τ and r are point identified. Let H' be defined as in section 4.1; bidders in this set have been observed in J successive auctions, of products $1, 2, \dots, J$, coming second each time. Each such bidder has a known pseudo-type \mathbf{b} . Let t_i^1 be the time at which they entered H' (i.e. J periods after entry). Define the events $a \equiv$ “are observed in the auction at $t_i^1 + 1$ ” and $b \equiv$ “are observed in the auction at $t_i^1 + 2$ ”. Let G_j^2 be the distribution of the second-highest rival bid. Let j_1 be the type-of-good auctioned at period $t_i^1 + 1$ and let j_2 be the type-of-good

auctioned at $t_i^1 + 2$. We have:

$$P(a) \equiv p_a = \underbrace{r}_{\text{survived}} \underbrace{\tau}_{\text{active}} \underbrace{G_{j_1}^2(b_{j_1}|H')}_{\text{observed}}$$

and

$$P(b) \equiv p_b = \underbrace{r}_{\text{survived}} \underbrace{(1 - \tau(1 - rG_{j_1}^1(b_{j_1}|H')))}_{\text{did not exit in period } t_i^1 + 1} \underbrace{\tau}_{\text{active}} \underbrace{G_{j_2}^2(b_{j_2}|H', \text{ did not exit in period } t_i^1 + 1)}_{\text{observed}}$$

Let $k_1 = G_{j_1}^2(b_{j_1}|H')$, $k_2 = G_{j_1}^1(b_{j_1}|H')$ and $k_3 = G_{j_2}^2(b_{j_2}|H', \text{ did not exit in period } t_i^1 + 1)$. From the first equation, $r\tau = \frac{p_a}{k_1}$, and substituting into the second yields $p_b = \frac{p_a k_3}{k_1} \left(1 - \tau + \frac{p_a k_2}{k_1}\right)$. This is linear in τ and therefore has a unique solution. \square

Corollary 1.

Proof. The event $\mathbf{x} < Z\mathbf{a}$ is by definition $F(Z\mathbf{a})$, which by (6) is bounded below by $\mathbb{P}(\{\mathbf{b} : \xi(\mathbf{b}) \leq Z\mathbf{a}\})$ and above by $\mathbb{P}(\{\mathbf{b} : 1(\mathbf{b} > 0)\xi(\mathbf{b}) \leq Z\mathbf{a}\})$. Because Z has only positive entries, the event $\mathbf{x} < Z\mathbf{a}$ is both necessary and sufficient for the event $\boldsymbol{\alpha} < \mathbf{a}$, which in turn is by definition $\mathbf{F}_\alpha(\boldsymbol{\alpha})$, yielding the bounds result in (8). \square

Lemma 4.

Proof. Follows immediately from the first part of the proof of Lemma 1, replacing j subscripts with functions of s . \square

Lemma 5.

Proof. From the main text, we have that $\check{b}_{i,j}(\mathbf{x}) = x_{i,j} - r \sum_k \tilde{Q}_{j,k} v_k(\mathbf{x})$. Re-arranging and stacking as vectors, $\mathbf{x} = \check{\mathbf{b}} + r\tilde{Q}\mathbf{v}(\mathbf{x})$, for $\mathbf{v}(\mathbf{x})$ the J-vector of continuation values. By a similar argument to that in the proof of Lemma 2, we have that $\mathbf{v}(\mathbf{x}) = (I - r\tilde{Q})^{-1}\check{\mathbf{u}}(\check{\mathbf{b}})$ where $\mathbf{u}(\check{\mathbf{b}})$ is the vector of expected static surpluses of a type $\check{\mathbf{b}}$ when bidding on each of the products, prior to drawing the idiosyncratic shock ε and unobserved heterogeneity ξ . Now as indicated in the text, each element $u_j(\check{\mathbf{b}})$ is invariant to the realization of ξ , as it shifts

all bids in the auction up equally. But ε matters for static surplus. Integrating out:

$$u_j(\check{\mathbf{b}}) = E_\varepsilon \left[E_{\check{G}_j^1} \left[\max\{0, \check{b}_j + \varepsilon - B_j^1\} \right] \right] = \int \check{G}_j^1(\check{b}_j + \varepsilon) (\check{b}_j + \varepsilon - E_{\check{G}_j^1}[B_j^1 | B_j^1 < \check{b}_j + \varepsilon]) dF_\varepsilon(\varepsilon)$$

□

Theorem 4.

Proof. The result follows from the arguments made in the main text. □

Lemma 6.

Proof. The arguments given in the first part of the proof of Lemma 1 establish that $\beta_j^k(\mathbf{x}) = \max\{x_j - r \sum_{l=1}^J \sum_{m=\{A,O\}} \tilde{Q}_{(k,j),(m,l)} v_l^m(\mathbf{x}), 0\}$, where the subscript $(k, j), (m, l)$ gives the transition probability from market k , product j to market m product l . But since supply is by assumption multinomial, we can replace $\tilde{Q}_{(k,j),(m,l)}$ with π_l^m to get $\beta_j^k(\mathbf{x}) = \max\{x_j - r \sum_{l=1}^J \sum_{m=\{A,O\}} \pi_l^m v_l^m(\mathbf{x}), 0\}$. From the definition of the ex-ante value function in (13), $v(\mathbf{x}) = \sum_{l=1}^J \sum_{m=\{A,O\}} \pi_l^m v_l^m(\mathbf{x})$, and the result follows by making this substitution in the line above. □

Theorem 5.

Proof. The text makes most of the required arguments. There are two missing steps: (i) showing that the extremal exit rates $e(\mathbf{0})$ and $e(\bar{\mathbf{b}}_j)$ are sufficient to identify r and $\{\pi_j^O\}$; (ii) showing that the derivatives of the exit rates suffice to identify the price distributions $\{G_j^O\}$.

Step 1: To simplify notation, define $e_j \equiv e(\bar{\mathbf{b}}_j)$ and $e_0 \equiv e(\mathbf{0})$ and drop “O” superscripts. Recall that $e_j = (1 - r) + r \frac{\Pi(1-r) + r\pi_j}{1-r(\Pi-\pi_j)}$. Expand $\Pi = \sum_k \pi_k$, and multiply through by $1 - r(\sum_k \pi_k - \pi_j)$ to get:

$$e_j(1 - r \sum_{k \neq j} \pi_k) = (1 - r)(1 - r \sum_{k \neq j} \pi_k) + r\pi_j + r(1 - r) \sum_{k \neq j} \pi_k$$

Re-arranging and canceling terms:

$$re(j) \sum_{k \neq j} \pi_k + r\pi_j = e_j - 1 + r$$

Dividing by r :

$$e_j \sum_{k \neq j} \pi_k + \pi_j = \frac{1}{r} (e_j - 1 + r) \quad (22)$$

Fixing r and stacking the equations for each e_j ($j = 1 \dots J$), we have a linear system in $\pi_1 \dots \pi_J$, where the coefficient matrix is of full rank. It thus has either zero or one solution $\pi_1(r) \dots \pi_J(r)$ for each value of r . For the purposes of identification we know there must be at least one solution corresponding to the parameters of the true DGP, and need to show that we cannot find two different values of r and corresponding solutions of the linear system that are consistent with the data.

Let $\Pi(r) = \sum_k \pi_k(r)$, and use the remaining equation:

$$e_0 = (1 - r) + r \frac{\Pi(r)(1 - r)}{1 - r\Pi(r)}$$

Re-arranging and canceling terms:

$$1 - e_0 = r(1 - \Pi(r)e_0)$$

Solving for r :

$$r = \frac{1 - e_0}{(1 - \Pi(r)e_0)} \quad (23)$$

Let the RHS of (23) be $\phi(r)$. Since $0 \leq \Pi(r) \leq 1$, we have $1 - e_0 \leq \phi(r) \leq 1$, and $\phi(r)$ is continuous in r since $\Pi(r)$ is continuous in r . Moreover, $\phi'(r) = \frac{\Pi'(r)e_0(1-e_0)}{(1-\Pi(r)e_0)^2} > 0$, since $\Pi'(r)$ is positive (it is the sum of solutions to a matrix equation whose RHS, $\frac{1}{r}(e_j - 1 + r)$ is increasing in r). This implies that we have at least one solution to (23) (the RHS is below the LHS at $1 - e_0$ and above it at 1, both sides are increasing and continuous functions of r). Call any such solution r^* . We would like to show it is unique. Notice that we can rewrite the expression for $\phi'(r)$ as $\frac{\Pi'(r)e_0(1-e_0)}{(1-\Pi(r)e_0)^2} = \phi(r)^2 \frac{\Pi'(r)e_0}{1-e_0}$. At r^* we have $\phi(r^*) = r^*$, so $\phi'(r^*) = r^{*2} \frac{\Pi'(r^*)e_0}{1-e_0}$. Moreover, since $\Pi(r)$ is the sum of solutions to (22), it can be written as a weighted sum $\sum_j w_j \frac{1}{r} (e_j - 1 + r)$, where the weights depend on the inverse of a matrix whose entries are equal to 1 or elements of $\{e_j\}$ (i.e. the weights are not functions of r). It follows that $\Pi'(r)$ is equal to $\sum_j w_j \frac{1-e_j}{r^2}$. Using this, we get $\phi'(r^*) = \frac{e_0}{1-e_0} \sum_j w_j (1-e_j)$. This derivative is not a function of r , so all solutions r^* share the same derivative $\phi'(r^*)$. This implies a unique solution, since for two monotone continuous univariate functions to intersect multiple times it is necessary that when they cross, sometimes the one crosses the other from

above, and sometimes below, with necessarily different derivatives in these different cases.

Step 2: Fixing a j and dropping the j notation we have:

$$e'(b) = r \frac{r\pi G'(b)(1 - \Pi)}{(1 - r(\Pi - \pi G(b)))^2}$$

Solving for $G'(b)$ and grouping known functions we get:

$$\frac{dG}{db} = h(b)j(G)$$

for $h(b) = e'(b)/(r^2\pi(1 - \Pi))$ and $j(G) = (1 - r(\Pi - \pi G))^2$. This has solution $G(b) = J^{-1}(H(b)+c)$ for $H(b) = \int_0^b h(s)ds$ and $J(G) = \int_0^G j(s)ds$. The boundary condition $G(0) = 0$ implies $c = 0$ (via $J^{-1}(c) = 0$ and thus $J(0) = 0 = c$), giving final solution $G(b) = J^{-1}(H(b))$, where both $h(b)$ and $j(G)$ are identified by the argument in Step 1, and $e'(b)$ is observed for bidders who make complete bids (and moreover $e(b)$ is constant in bidder history, so that conditioning on bidders who have complete bid vectors does not introduce selection). \square

B Data Appendix

B.1 Data Source

We used data from eBay auctions for compact cameras that was purchased from Terapeak, a private company that uses eBay data to offer analytics tools for sellers on the platform. Our data included listing, seller, and bidder attributes as documented in Table 1 from the main text.

B.2 Sample Construction

We restricted attention to auctions for *new* compact cameras that ended between the dates of Feb 5 and May 6, 2007. By way of data cleaning we also imposed the following restrictions:

1. We were somewhat concerned that we may have been observing some “shill” bidding by sellers intent on raising revenues. We used the following rather coarse procedure to detect such bidding behavior: for any given bidder, if they won at least five auctions

and at 80% of them were from the same seller, then we flagged them as a shill bidder. We excluded all auctions in which a so-designated shill bidder was among the set of observed bidders.

2. There are many senses in which a bid may be an “outlier”. We excluded all winning bids that were more than twice the sale price of the listing (either the second-highest bid or the reserve price, whichever is greater).
3. Though we did not exclude auctions that used the Auction-Buy-it-Now feature, we did exclude auctions that ended with the execution of the BIN option. This option disappears once the first bid is made, so those listings function as regular auctions afterwards. See Akerberg et al. (2009) for a fuller treatment of Auction-Buy-it-Now auctions.
4. We drop listings for which we are missing data on attributes (e.g., zoom, resolution), for which those attributes are unreasonable or technically impossible, or where the product line is not indicated.