The Role of Intermediaries in Dynamic Auction Markets

Kenneth Hendricks
University of Wisconsin & NBER
hendrick@ssc.wisc.edu

Alan Sorensen
University of Wisconsin & NBER
sorensen@ssc.wisc.edu

September 2015

PRELIMINARY DRAFT: PLEASE DO NOT CIRCULATE OR CITE

1 Introduction

Many goods and assets are traded in dynamic markets where sellers and buyers arrive, find trading partners, bargain to determine a price, and—if they fail to trade—continue in the market to be rematched. The main questions that arise in these kinds of markets are: what are the trading frictions and how do they affect the allocation and prices? how do they vary with the thickness of the market? how are they affected by the trading mechanism? and finally, how can an intermediary alleviate these frictions? There is a large theoretical literature that addresses these questions but the empirical literature is sparse. This paper develops and estimates a structural model of a dynamic market with an intermediary to provide quantitative answers to these questions.

Our data is from eBay, so we tailor the model to capture the main features of this market. Buyers and sellers arrive randomly over time to trade one unit of a homogenous good. Buyers have heterogeneous values that do not change over time. Sellers have essentially no value for the good. Every seller contracts with the intermediary to sell its unit in a second-price auction of fixed duration. At every moment in time, the market consists of a number of overlapping auctions with different closing times. Every buyer chooses one auction in which to bid and submits her bid to the intermediary. The intermediary receives the bids and posts in real time the highest losing bid in each auction. Buyers use these bids and the closing times of the auctions to sort across the auctions. Low value buyers tend to choose auctions with later closing times because the posted prices in earlier auctions typically exceed their values. High value buyers tend to choose auctions with earlier closing times, because the probability of being outbid by a subsequent buyer is lower in these auctions. The buyer with the highest bid at the close of the auction wins and pays the second highest bid to the seller. Buyers who submit losing bids either exit for exogenous reasons or enter a pool of losers and return to bid again at some future time. Sellers almost always sell in our model, but they exit the market if they do not sell. The trading frictions in the market are on the buyer side: private information about values, strategic behavior, and bidding costs. The main role of the intermediary is to run the auctions, collect the bids, and post information about the bids.

The option to bid again implies that each buyer's willingness to pay in an auction is her private value of the good less the option value of losing. In thin markets, this option value depends on the number and values of the buyers who lost in prior auctions, so the strategies of the buyers could in principle depend on the entire history of the game. Since rivals can exploit this information,

bidders would also have to take into account the effect of their decisions on the option value of losing. We make assumptions about the size of the buyer population and about the rate at which losing buyers return to the market such that (i) the benefit from tracking the bids of past losers is negligible and (ii) her bid has a negligible impact on the option value of losing. These assumptions imply that each buyer has a dominant strategy to bid her willingness to pay, and this strategy is monotone increasing in value and independent of the state of play. The latter property is crucial to our estimation strategy: it allows us to estimate the distribution of values of buyers without having to solve for the equilibrium auction choice rule, which (unlike the bid function) can depend on the history of play.

We show that the primitives of the model are identified using data on bids and buyer identities. We estimate the model using data from eBay auctions for iPads and compute the total surplus generated by the trading mechanism. We examine how surplus varies with the thickness of the market by estimating the model separately for the 16GB and 64GB iPad models. The market for the 16GB version has more sellers and buyers arriving per day than the 64GB market. This comparison also provides evidence on whether buyers' bids indeed reflect continuation values, since the option value of losing is higher in the thicker market.

We use the estimates to simulate two counterfactuals that serve as benchmarks for evaluating the efficiency of the eBay trading mechanism. The first counterfactual is a centralized exchange market in which all sellers and buyers who arrive during a fixed period of time participate in a uniform price auction. Buyers submit their bids to the intermediary, who orders the bids from highest to lowest and finds the price that clears the market. The intermediary then allocates the available units to buyers who bid at least the market-clearing price, and they pay the market-clearing price to the sellers. If enough buyers and sellers arrive during the period, then this price is essentially the Walrasian price. In our application, running the uniform auction every day or two would be sufficient to achieve the competitive outcome, in which case the uniform auction would be essentially equivalent to a posted price mechanism.

The second counterfactual is intended to represent a fully decentralized market in which each seller runs her own second-price auction. The distinguishing feature of such a market is that the matching of buyers to sellers is random. In the specific context of our model, it means that the intermediary does not post standing high bids. Random matching then results from an equilibrium in which

each arriving buyer chooses to bid in the auction that is next to close. We simulate the market outcomes under this choice rule and compute total surplus and revenues. We find that the eBay mechanism generates substantially higher surplus and revenue than the decentralized market.

In the theoretical literature, Satterthwaite and Shneyerov (2008a, 2008b) study models of dynamic matching and bargaining markets that are quite similar to our model of a decentralized market. Their focus is on whether the allocation converges to the competitive outcome as the number of trading opportunities per unit of time gets large. In the relevant empirical literature on dynamic models of trade, Gavazza (2011) studies a decentralized market for used aircraft where the trading mechanism is pair-wise bargaining rather than an auction. The model generates a rich set of testable predictions on how market outcomes change with the thickness of the market. He tests these predictions using cross-sectional variation in the stocks of different models of aircraft. Gavazza (2013) introduces dealers into the model and estimates its primitives in order to quantify the effects of trading frictions on allocations and prices, and to measure the extent to which dealers alleviate these frictions. Backus and Lewis (2012) study identification and estimation issues in a dynamic model of sequential auctions in a differentiated good market. In the structural empirical literature on eBay auctions, Bajari and Hortascu (2003), Gonzalez, Hasker, and Sickles (2004), Canals-Cerda and Pearcy (2006), Ackerberg, Hirano, and Shahriar (2006) and Lewis (2007) have estimated static models of bidding. Adachi (2014) estimates a dynamic model of the eBay mechanism that is closely related to ours, although the estimation strategy is quite different.

2   A Dynamic Model of Bidding

In this section, we develop a theoretical model of bidding in a dynamic environment like eBay. We provide a partial characterization of the equilibrium and show that this is sufficient to identify and estimate the primitives of the model from data on bids.

An infinite number of sellers arrive over time to sell a single unit of a homogenous good. Each seller contracts with an intermediary to sell her unit in a second-price auction. It will be convenient to assume the arrival times of the sellers are equally spaced and normalize the time between arrivals to be one unit. The duration of each auction is set to J units of time, where J is a large integer.[1]

---

[1] eBay allows the seller to choose auction length; typically auctions run for 1, 3, 5, 7, or 10 days.

Thus, at any time t, the market consists of a set of J overlapping auctions. We index the auctions by j and order them by their closing times with $j = 1$ denoting the next-to-close auction and $j = J$ denoting the auction with the most distant closing time. If at time t, the time remaining in auction 1 is $d \in (0, 1)$, then the closing time of auction $j = 2, .., J$ is equal to $d + j - 1$. At closing time, the item is awarded to the buyer who submitted the highest bid at a price equal to the second-highest bid or, if no one else bids, the seller's starting bid. Since virtually all of the sellers in the market that we study choose starting bids that are low enough to ensure the sale of the product, we normalize the starting bid to zero.[2]

On the demand side, new buyers randomly arrive over time to buy a single unit of the good. Buyers are risk-neutral and their utility functions are quasi-linear in income. Each new buyer draws a value x for the good from a distribution F_E with density f_E and support $X = [0, \bar{x}]$. The values of the buyers are independently and identically distributed, and they do not change over time. If a buyer arrives and the standing bid in at least one auction is less than her willingness to pay, then she chooses an auction in which to bid and submits a bid at that time.[3] If her bid is the highest bid at the closing time of the auction, then she wins the auction and exits. If she does not bid or her bid is not the highest bid at closing time, then she is a loser and either exits the market with positive probability or bids again in some future auction. Thus, each auction can have two types of buyers: new buyers who have never bid before and returning buyers who have lost in a previous auction.

2.1   Entry and Exit Dynamics

Each auction is a continuous-time, ascending price auction with proxy bidding. At any time t, the intermediary posts the highest losing bids in the J auctions but does not report the highest bids. We will refer to the highest losing bids as standing bids. Let $r_j(t)$ and $w_j(t)$ denote respectively the standing bid and the high bid in auction j at time t and define $r(t) = (r_1(t), .., r_J(t))$ and $w(t) = (w_1(t), .., w_J(t))$ as the vectors of standing bids and high bids respectively in the market at time t. A buyer who decides to bid in auction j enters a maximum bid b and the intermediary bids on her behalf up to that level. We will refer to such bids as proxy bids. If b is less than $w_j(t)$, then the buyer loses and either exits with probability α or enters the pool of losers with probability $1 - \alpha$.

---

[2]eBay also allows sellers to set a secret reserve price. We ignore this choice since most sellers in our market choose not to exercise this option.

[3]This is a strong assumption since it rules out strategic delay (e.g., "sniping"). However, it appears consistent with the data in our application since most of the buyers in the ipad market do not wait until the last minute to bid.

The standing bid in auction j increases to b (plus a small bid increment). If b exceeds $w_j(t)$, then the buyer becomes the new high bidder and the previous high buyer (assuming $w_j(t) > 0$) exits with probability α or enters a pool of losers with probability $1 - \alpha$. The standing bid in auction j increases to $w_j(t)$. When the auction ends, the item is awarded to the buyer with the highest proxy bid at the standing bid and she exits for certain. The exit rates are exogenous and the same for all buyers.

The probability laws determining entry are also exogenous. Let $E(t)$ denote the (random) number of new buyers who arrive in an interval $[0, t]$. The family $E = \{E(t); t \geq 0\}$ of random variables is an arrival process with continuous time parameter and discrete state space consisting of the non-negative integers. The arrival process is a Poisson process with parameter λ. The number of new buyers who arrive in a time period $[t, t + \Delta)$ is defined as $N(t) = E(t + \Delta) - E(t)$. The probability that n new buyers arrive in that period is

$$\Pr\{N(t) = n\} = \frac{(\lambda\Delta)^n e^{-\lambda\Delta}}{n!}$$

and the expected number of new buyers is $\lambda\Delta$.

The arrival of returning buyers depends upon the number and values in the pool of losers. This set can be represented by a counting measure defined on X and its evolution studied using the theory of point processes.[4] Let $k(t) \in K$ denote the counting measure at time t where K is the space of counting measures and let $l(t)$ denote the number of buyers in the pool. The number of buyers who leave the pool to bid during period $[t, t + \Delta)$ is denoted by $M(t)$. We assume that the time that a buyer spends in the pool of losers before returning to the market is independent of her value and distributed exponential with parameter β. Thus, the probability that a buyer who loses at time t and returns in some future period $[t + s, t + s + \Delta)$ conditional on not exiting or arriving earlier than $t + s$ is

$$1 - e^{-\beta\Delta} \approx \beta\Delta.$$

This hazard rate does not depend upon when she entered the pool nor on how long she has been in the pool. Hence, the buyers in the pool are equally likely to leave, so $M(t)$ is a binomial random variable with parameters $l(t)$ and $\beta\Delta$. The expected number of buyers who leave the pool during the period $[t, t + \Delta)$ is $\beta\Delta l(t)$ and the expected number of buyers who arrive to bid during this period

---
[4]See Matsuki (2013).

is $\beta\Delta l(t) + \lambda\Delta$. Since every buyer who arrives either loses or displaces another buyer, the number of buyers who enter the pool of losers during the period is on average equal to $(1 - \alpha)(\beta\Delta l(t) + \lambda\Delta)$.[5] In the long run, the number of buyers who leave the pool must on average equal the number of buyers who enter the pool, which implies that the size of the pool fluctuates around

$$\bar{l} = \frac{\lambda(1 - \alpha)}{\alpha\beta}.$$

To summarize: the state of the market at any time t consists of the set of values in the pool of losers, $k(t)$, standing bids, $r(t)$, high bids, $w(t)$, and d, the time remaining in the next-to-close auction. We have described how the state changes when a buyer arrives and bids. The probability law determining the transitions of the state and the long-run behavior of the stochastic process depend on the equilibrium rules that govern how new and returning buyers choose an auction and bid in that auction.

2.2   Bidding Behavior

When a buyer arrives to the market, she has to decide in which auction to bid and how much to bid. These decisions depend on the information that the intermediary provides about the market. In our application, the intermediatry posts the bid history of individual auctions, including auctions that ended prior to arrival. The bid history of an auction consists of the highest losing bids, the times at which they were submitted, and the identity of the buyers submitting them. The winning bids in auctions that have closed are not posted. The buyer can learn the bid history of an auction by first clicking on the auction and then clicking on its bid history. By combining these individual bid histories, a buyer could re-construct the real-time history of the market (minus the winning bids), although it may be somewhat costly to do so. Let $h(t)$ denote the history of the market prior to time t and let $H(t)$ denote the set of all histories.

Since exit is private information, buyers do not observe the pool of losers. The buyer could use the market histories to learn about the pool of losers but, if they did so, then their choice and bidding strategies are unlikely to be stationary. It would also give rise to a dynamic incentive problem

---

[5]If a buyer is the first to bid in an auction then she does not displace anyone but, in every auction this buyer is offset by the buyer who wins the auction and exits for certain.

known as the "leakage" problem (see also Backus and Lewis (2012)). Buyers may lose and return to bid again, so they would have to take into account the impact of their decisions on the beliefs of future rivals. This issue is likely to be important for "thin" markets in which the number of sellers and buyers is small. Our application is a large market: the number of buyers in the pool of losers is on average over three hundred.[6]. Furthermore, return times are exponential. Thus, the benefit of building a market history in order to track bidder identities and losing bids is negligible. As a result, we will assume that buyers do not use the bid histories to learn about about the pool of losers. Instead, their beliefs about the probability law determining the number and values of return buyers is determined by the invariant distribution on K generated by their participation and bid strategies. Let F_L denote the distribution of values in the pool of losers induced by this invariant distribution.

In what follows, it will be convenient to approximate the distribution of the number of returning buyers by a Poisson distribution with parameter $\gamma = \beta \bar{l}$. This approximation is useful because it allows us to treat returning and new buyers symmetrically. A buyer who returns to the market has to condition her beliefs about $M(t)$ on her arrival. In the case of the Poisson distribution, a returning buyer's belief about the number of returning rivals conditional on her own arrival is the same as a new buyer's unconditional belief—namely, that it is distributed Poisson with parameter γ.[7] Thus, each buyer believes that the probability of m buyers returning to the market in any period $[t, t + \Delta)$ is given by

$$\Pr\{M(t) = m\} = \frac{(\gamma \Delta)^m e^{-\gamma \Delta}}{m!}$$

and the total number of buyers who arrive in that period is a Poisson random variable with parameter $(\gamma + \lambda)\Delta$. As a result, in any symmetric equilibrium, returning buyers and new buyers that have the same value and arrive in same state make the same decisions.

When buyers arrive, they do not observe the high bids in the auctions open for bidding at that time. They need to form beliefs about these bids based on the market history. Let $\omega(t) \in \Omega$ denote the information set that a buyer who arrives at time t uses to construct her beliefs about $w(t)$. It certainly includes the standing bids $r(t)$ and closing times $d(t)$ but it is also could include statistics such as the difference between t and the times at which the high bids in the J auctions were

---

[6]We compute this number using estimates of the arrival rate of new buyers, and exit and return rates of losers and plugging them into.the above formula for $\bar{l}$.

[7]Myerson (2000) refers to this property of Poisson games as *environmental equivalence*.

submitted. The crucial restriction here is that Ω is finite and time-invariant. A pure (stationary) strategy is a choice function $\rho : X \times \Omega \to \{0, 1, .., J\}$, where 0 represents the option of not bidding and a bid function $\sigma : X \times \Omega \to \Re_+$. These two rules, together with the exogenous arrival and exit of buyers and sellers, determine the continuous time transitions of the state of the market. An equilibrium consists of a triple $(\rho^*, \sigma^*, \mu^*)$ such that (i) given μ^*, ρ^* selects the auction with the highest expected payoffs and σ^* is the optimal bid in that auction; (ii) given (ρ^*, σ^*), μ^* is the invariant distribution that gives the long-run probabilities over any set of states. Let Ψ denote the stationary distribution induced by μ on Ω that gives the long-run probabilities of any subset of Ω.

We turn next to the buyer's decision problem. Our second large population assumption concerns the option value of losing. In calculating this value, a buyer has to form beliefs about what the state of the market will be when she returns. Let $v(x, \omega)$ denote the expected payoff to a buyer with value x who behaves optimally in state ω. The ex ante value function is defined as

$$V(x) = \int v(x, \omega) d\Psi(\omega).$$

Key Assumption: A buyer with value x believes that her continuation value from losing and returning at some future time is $V(x)$.

The above assumption implies that each buyer views the state of the market when she returns as a random draw from a stationary distribution. It plays a critical role in the analysis. In thin markets, if a buyer loses and does not exit, relatively few buyers are likely to arrive before she returns. In that case, she will want to use the current state to forecast the state upon her return. She will also need to take into account how her bid affects the transition of the state. Recall that a buyer's bid only becomes part of the state whenever her bid is the high bid in an auction and a subsequent buyer submits a higher bid. At that point, her bid will influence the decisions of future buyers and hence her continuation value. Both of these effects, the information effect and the dynamic incentive effect, become negligible as the arrival rate of buyers gets large relative to the return time of a buyer. This is the case in our application. As we show later, the median number of auctions that finish in the time it takes a losing bidder to return is 8, and at an average of 9 buyers per auction, it implies that roughly 72 other buyers arrive before a losing bidder returns. In any case, it is a testable assumption.

We now characterize the optimal bidding decision of a buyer. The probability that she wins an auction is simply the probability that her bid exceeds the bid of the current high bid and that of any other buyer who decides to bid in the auction before it closes. Since her bid only affects the decisions of subsequent buyers if she loses, her expectations of winning depend only upon the current state. Let $G_j(p; \omega)$ denote her belief that the highest bid submitted by her rivals in auction j is p when the state of her arrival is ω. Note that the high bid in auction j is included in ω so the lower bound of the support of $G_j(p; \omega)$ is r_j. Ignoring any discounting and integrating over the random return time, her optimization problem can be expressed as

$$\max_b \int_0^b (x - p)dG_j(p; \omega) + (1 - \alpha)(1 - G_j(b; \omega))V(x).$$

Differentiating and solving for the optimal bid yields the optimal bid function.

Proposition 1 *The equilibrium bid function is given by*

$$\sigma(x) = x - (1 - \alpha(x))V(x). \tag{1}$$

The striking feature of this bidding strategy is that it does *not* depend upon the auction choice or bid history or the bidding strategy of rivals. We exploit this feature to derive a closed form solution for the value function.

Substituting $b = \sigma^*(x)$, the expected payoff to a bidder with type x who bids optimally in auction j in state ω is

$$v_j(x; \omega) = \int_0^{\sigma^*(x)} (x - p)dG_j(p; \omega) + (1 - \alpha(x))(1 - G_j(\sigma^*(x); \omega))V(x)$$

Her optimal choice of auction is to bid in the auction with highest positive expected payoff. Assuming for convenience that there is always at least one auction that has no bid (i.e., J is large),

$$v(x; \omega) = \max\{v_1(x; \omega), ..., v_J(x; \omega)\}.$$

10

Let $\rho^*(x;\omega)$ denote the equilibrium choice of a buyer with value x in state ω, and define

$$G(p) = \int G_{\rho^*(x;\omega)}(p;\omega)d\Psi(\omega).$$

Here G is the stationary distribution of the highest bid submitted by rivals when buyers use the equilibrium choice rule $\rho^*(x;\omega)$ and bid rule $\sigma^*(x)$. Taking expectations over the state and changing the order of integration, the ex ante value function can be expressed as

$$
\begin{aligned}
V(x) &= \int\left(\int_0^{\sigma^*(x)}(x-p)dG_{\rho^*(x;\omega)}(p;\omega) + (1-\alpha)(1-G_{\rho^*(x;\omega)}(\sigma^*(x);\omega))V(x)\right)d\Psi(\omega) \quad (2)\\
&= \int_0^{\sigma^*(x)}(x-p)dG(p) + (1-\alpha)(1-G(\sigma^*(x)))V(x)\\
&= \frac{\int_0^{\sigma^*(x)}(x-p)dG(p)}{[1-(1-\alpha)(1-G(\sigma^*(x)))]}
\end{aligned}
$$

Equation (2) reveals that the ex ante value function has a closed form solution that depends only upon G, the unconditional equilibrium probability that a buyer wins the good. We can now check that σ^* is increasing, a necessary condition for existence, and establish the properties of the ex ante value function.

Proposition 2 σ^* *is monotone increasing in x and*

$$V(x) = \int_0^x \frac{G(\sigma^*(s))ds}{[1-(1-\alpha)(1-G(\sigma^*(s)))]} \quad (3)$$

and convex in x.

Proof. Taking expectations over the state and changing the order of integration,

$$
\begin{aligned}
V(x) &= \int\left(\int_0^{\sigma(x)}(x-p)dG_{\rho(x;\omega)}(p;\omega) + (1-\alpha)(1-G_{\rho(x;\omega)}(\sigma(x);\omega))V(x)\right)d\Psi(\omega)\\
&= \int_0^{\sigma(x)}(x-p)dG(p) + (1-\alpha)(1-G(\sigma(x)))V(x)
\end{aligned}
$$

11

Differentiating with respect to x yields

$$
\begin{aligned}
V'(x) &= G(\sigma(x)) + [(x - \sigma(x)]g(\sigma(x)\sigma'(x) - (1-\alpha)g(\sigma(x))V(x) + (1-\alpha)(1 - G(\sigma(x)))V'(x) \\
&= \frac{G(\sigma(x))}{[1 - (1-\alpha)(1 - G(\sigma(x)))]}
\end{aligned}
$$

where the last statement follows from the envelope theorem. Thus,

$$
\begin{aligned}
\sigma'(x) &= 1 - (1-\alpha)V'(x) > 0 \\
&\Longleftrightarrow 1 - \frac{(1-\alpha)G(\sigma(x))}{[1 - (1-\alpha)(1 - G(\sigma(x)))]} > 0 \\
&\Longleftrightarrow \alpha > 0.
\end{aligned}
$$

Integrating $V'(x)$ and imposing the boundary condition $V(0) = 0$ yields the solution for $V(x)$ given above. Differentiating $V'(x)$ yields

$$
V''(x) = \frac{\sigma'(x)\alpha}{[1 - (1-\alpha)(1 - G(\sigma(x)))]^2} > 0.
$$

Q.E.D. ∎

Proposition 2 is important for computing counterfactuals. Given model primitives and a choice rule, it allows us to easily solve the model through simulation. Monotonicity of the bid function implies that, given any simulation in space of bids, there is an equivalent simulation in the space of values. Furthermore, because a buyer's optimal bid does not depend upon closing times, the standing bid in an auction corresponds to the second-highest value among the buyers who arrived and choose to bid in that auction. The choice rule assigns each buyer who arrives to an auction based on her value and the state of the market at the time of her arrival. This assignment in turn determines the state transitions. The auction then eliminates the highest value, and lower values either exit or enter the pool of losers. The pool of losers grows stochastically over time, and distribution of loser values approaches a stationary distribution, F_L. In the stationary state, we can simulate the distribution $H(x) \equiv G(\sigma(x))$ from the fraction of times that a bidder of type x wins. Given H, we use equation (2) to compute the continuation values, $V(x)$, and equation (1) to compute equilibrium bids, $\sigma(x)$.

2.3 Identification

In what follows, we assume that the econometrician has access to data on the bid history and the identity of the bidders. The task is to show that the primitives of the model $(\alpha, \beta, \lambda, F_E)$ can be identified from these data.

The first step is to use equation (2) to obtain a closed form solution for the inverse bid function. Define $\eta(b) \equiv \sigma^{-1}(b)$.

Proposition 3: *The equilibrium inverse bid function is given by*

$$\eta^*(b) = b \left(1 + \frac{(1-\alpha)G(b)}{\alpha} \right) - \frac{(1-\alpha)\int_0^b p\,dG(p)}{\alpha}. \tag{4}$$

Proof: Use equation (2) to substitute for $V(x)$ in equation (1). Evaluating this equation at $x = \eta(b)$ and solving for $\eta(b)$ yields the closed form solution given above. Q.E.D.

Proposition 3 establishes that the buyers' private values can be estimated directly from data on bids. It extends the structural approach developed by Elyakime, Laffont, Loisel and Vuong (1994) and Guerre, Perrigne and Vuong (2000) for estimating static, private value auctions to a dynamic environment.

Proposition 3 *The model primitives, (λ, β, α) and F_E, are identified from data on bids if bidder identities are known.*

Proposition 4 is the main result of this section. The parameters are identified from bidder participation: $\widehat{\lambda}$ is the average number of new buyers per auction, $\widehat{\beta}$ is the average number of losers per auction, and $\widehat{\alpha}$ is the fraction of losing buyers who exit. A non-parametric estimate of $G(b)$ can be obtained by computing the fraction of times that a buyer wins when he submits a bid of b. We can then evaluate equation (3) on bids of new buyers to obtain a non-parametric estimate of F_E. An estimate of F_L can be obtained by evaluating equation (3) on bids of returning buyers.

Stationarity imposes an interesting over-identifying restriction on F_L. Given $l(t)$, the flow of x types in a period of length Δ from the pool of losers is given by the expected number of buyers that leave

during the period times the density at x :

$$\beta \Delta l(t) f_L(x).$$

The flow back into the pool over this period is

$$(1 - \alpha)[1 - G(\sigma^*(x))][\beta \Delta l(t) f_L(x) + \lambda \Delta f_E(x)].$$

On average, these two flows have to equal each other at every x. Recall that, in the stationary equilibrium, the expected number of bidders in the pool is

$$\bar{l} = \frac{(1 - \alpha)\lambda}{\alpha \beta}.$$

Equating flows at $l(t) = \bar{l}$, we obtain

$$f_L(x) = \frac{\alpha(1 - G(\sigma^*(x)))}{[1 - (1 - \alpha)(1 - G(\sigma^*(x)))]} f_E(x). \tag{5}$$

Equation (5) is a testable restriction. We can estimate α, f_L, f_E and G from the data and see if the estimates satisfy (5).

The remarkable aspect of Proposition 4 is that the primitives of our model are identified and estimated without solving for the equilibrium choice rule.[8] The intuition is that there is only one dimension of unobserved heterogeneity among the buyers and it can be inferred from bids because the optimal bid does not depend upon the auction choice. This result is important because solving for the equilibrium choice rule in dynamic bidding environments may not be feasible. The choice of auction is observable to subsequent buyers, so each buyer has to worry about the impact of her choice of auction on the choices of subsequent buyers, and hence upon the probability that she wins the auction. This dynamic incentive complicates an already difficult problem.

Two questions remain. First, does a stationary equilibrium exist? The bid rule is clearly stationary but we have not shown that the equilibrium choice rule is stationary. Second, assuming a stationary equilibrium exists, does a stationary distribution over states exist? The transition probabilities are clearly not time-invariant since transitions within a period are governed by the arrival of buyers

---

[8]The caveat is that we have not proved the existence of a stationary equilibrium choice rule and simply assume that the data are generated by one.

and their behavior whereas, at any $t \in \mathbb{N}$,.the transition is deterministic with one auction closing and another auction opening with posted and high bids equal to zero. However, it is relatively straightforward to show that the stochastic process is Markovian and we believe that arguments similar to those presented in Backus and Lewis (2012) can be applied to show that there is an invariant distribution that gives the long-run probabilities of any set of states.

2.4   Endogenous Exit

When a buyer learns that she has lost the auction, she has to decide whether or not to exit the market. Exit is costless but continuation may be costly. We denote the buyer's cost of staying in the market and bidding again by c.We assume that it is drawn from a distribution F_C with support $[0, \bar{c}]$. The cost is drawn after she bids and loses, and it is independently distributed across losses.

The probability that a buyer with type x exits is given by

$$\Pr\{c > V(x)\} \equiv 1 - F_C(V(x)).$$

The optimal bid function is

$$\sigma(x) = x - F_C(V(x))V(x).$$

The ex ante value function is given by the *implicit* function

$$V(x) = \frac{\int_0^{\sigma(x)} (x - p)dG(p)}{[1 - F_C(V(x))(1 - G(\sigma(x)))]}.$$

where $G(p)$ is the probability of winning at a bid of p. Therefore, given G, F_C and x, we have three equations to solve for three numbers: the bid $b = \sigma(x)$, the continuation value $v = V(x)$, and the probability of exit $\alpha = 1 - F_C(v)$.

However, F_C is not known. Can it be identified from the data? The answer is yes. To see why, we use the transformation $x = \eta(b)$ and express the above three equations in bid space as follows. The probability of exit becomes

$$\alpha(b) = 1 - F_C(V(\eta(b))).$$

The inverse bid function is

$$\eta(b) = b + (1 - \alpha(b))V(\eta(b)),$$

and the value equation becomes

$$V(\eta(b)) = \frac{\int_0^b (\eta(b) - p)dG(p)}{[1 - (1 - \alpha(b))(1 - G(b))]}.$$

Substituting $V(\eta(b))$ into the inverse bid function, we obtain

$$\eta(b) = b\left(1 + \frac{(1 - \alpha(b))G(b)}{\alpha(b)}\right) - \frac{(1 - \alpha(b))\int_0^b pdG(p)}{\alpha(b)}.$$

Once again, estimates of the private values can be obtained directly from data on bids and exits. Thus, F_E (and F_L) are identified. To identify F_C, we solve $v(b) = V(\eta(b))$.for each bid b and then plot $\alpha(b)$ against $v(b)$ to determine the distribution of costs.

3   Data

Our primary data consist of all eBay listings for iPads posted between February-September 2013. We focus on auctions for used 16GB iPad Minis. The data were obtained from eBay's internal data warehouse. For each listing, the data contain information about the seller (e.g. identity, feedback rating) and about the timing and characteristics of the listing (e.g. start date, end date, starting bid, reserve price, shipping options, etc.). We also observe all of the bids submitted for each listed item. Importantly, we observe the identities of all bidders and the amounts and times of all bids they submitted, which allows us to track bidders who lose an auction and return later to bid again in another auction. We also observe the bid submitted by the winning bidder—not just the sale price. As explained above, these winning bids are important for estimating G, the distribution of the maximum rival bid.

eBay sellers can choose whether to sell their items by auction or at a fixed price. Auctions are the most common (and commonly known) type of sale on eBay. Fixed-price listings remain on the site for a longer period of time—usually 30 days—and any buyer may purchase the item immediately if she is willing to pay the posted price. For sellers who choose the auction mechanism, they may also (for a small fee) give buyers the option of winning the item immediately by paying a "buy-it-now"

price specified by the seller. Of the used iPad Minis successfully sold in our data, 65% were sold by auction. Of those sold by auction, 45% offered a buy-it-now option, and 14% ended with a buyer exercising the buy-it-now option.

Over 70% of the listings in our data ended successfully with a sale. In the analyses below we focus exclusively on these listings, for two reasons. First, in examining the unsuccessful listings, we found that many of them were problematic or bogus: e.g., they were listings for broken iPads, or for hot tips about how to get a free iPad, etc. Second, legitimate auctions that ended without a sale typically had very high starting bids or very high reserve prices, so effectively these listings were more like fixed-price listings.

Our model assumes that buyers have unit demands. For iPads it seems reasonable that most buyers would be interested in buying only one unit, and in the data 92 percent of successful buyers bought exactly one iPad during the sample period. However, a small fraction of buyers appear to be systematically buying and reselling iPads on the site. There were 35 eBay users who purchased 10 or more iPad Minis in our data, and some of these users also appear in the data as sellers. In the analyses reported below we exclude buyers who purchased more than two iPad Minis during the sample period.

Table 1 shows summary statistics for the auctions that ended with a sale. The average sale price was \$295.29, and the average shipping fee was \$7.03. The retail price for a new unit of this particular model was \$329, so the used units on eBay were selling at an average 10% discount relative to the new retail price. The table indicates that over 5% of the auctions closed at prices above the retail price; this mostly results from bundled extras (such as cases or extra chargers) that are included in some auctions. The average number of bidders per auction is 9, but this number varies substantially across auctions.

Even though we are looking only at auctions for a specific model (16GB WiFi), sale prices exhibit considerable variation. Some of this variation reflects heterogeneity in item or seller characteristics, such as color (white vs. black), included extras (like a case), and seller feedback ratings. Even after controlling for observable characteristics, however, much of the variance in prices remains. This underscores the relevance of our model: bidders in these markets use dynamic bidding strategies in part because they know that if they bid in enough auctions, eventually they might get lucky and be matched against fewer and/or weaker bidders.

17

Table 1: Summary statistics: Successful Auctions for 16GB iPad Mini ($N = 12,704$)

	Mean	Std. Dev.	Percentiles		
			0.10	0.50	0.90
Start price	128.56	110.03	0.99	139.99	290.00
Positive reserve price (0/1)	0.10	0.30	0.00	0.00	1.00
Reserve price (if positive)	261.83	39.57	200.00	270.00	320.00
Sale price	295.29	32.17	236.80	296.60	344.00
Shipping fee	7.03	5.50	0.00	6.60	15.00
Number of bids	20.94	17.51	1.00	17.00	54.00
Number of unique bidders	9.20	6.19	1.00	9.00	20.00
Minutes since last auction	63.76	108.57	1.53	29.58	245.07
# of auctions in next hour	1.54	1.49	0.00	1.00	4.00
Cover included (0/1)	0.19	0.39	0.00	0.00	1.00
Seller feedback (#)	6860.62	42921.54	5.00	119.00	11355.00
Seller feedback (% positive)	98.92	6.71	96.97	100.00	100.00

Repeat Bidding

Conditional on losing an auction, bidders come back to bid again in a subsequent auction 50.2% of the time. Return times are fairly short: conditional on returning to bid again, 47% of bidders return within an hour, and 8% return within a minute. The full distribution of return times is highly skewed, however, since there is a long right tail reflecting bidders who take 24 hours or more to come back. We use the observed return times to estimate the re-arrival rate, β. This parameter is not needed for estimating the distribution of valuations—it drops out of the inverse bid function because we are assuming no time discounting—but it is needed for simulating counterfactuals.

When buyers return to bid on a new item after having lost in a previous auction, they do not always bid on the exact same model (storage capacity and connectivity). However, for the specific model we are looking at (16GB WiFi), 83% of returning bidders choose to bid again on the same model. Among those who switch to bidding on a different model, most either bid on the 32GB WiFi version (8%) or on the 16GB WiFi+4G version (5%). While clearly there is some substitutability between the different models, we believe it is a reasonable approximation to treat the 16GB WiFi market as its own separate market.

We also believe it is a reasonable approximation to treat the used market as separate from the market for new items. While both new and used items are offered on Ebay, most buyers do not

appear to view new and used items as substitutes. Of the buyers who lost the bidding on a used item and returned to bid again, 79% chose to bid on another used item. Of those who bid on a new item when they returned, only 6% won. For buyers who bid on three or more items, the modal pattern was to bid exclusively on used items, and the next most common pattern was to bid exclusively on new items. Treating the used market as separate also avoids the issue of how to model buyers' willingness to pay for new vs. used items. In the empirical analysis we use normalized bids that adjust for item characteristics like color, added extras, and seller feedback ratings—an approach that implicitly assumes these are vertical characteristics (for example, all buyers have the same willingness to pay for an extra charger). We suspect this assumption is unlikely to hold with respect to item condition: some buyers probably care a lot more than others about whether the item is new vs. used.

Auction Choice

As discussed above, bidders do not always bid in the next-to-close auction when they arrive: some choose to bid in auctions that won't be ending for several hours, even if dozens of similar auctions are ending sooner. Table 2 describes the distribution of chosen auction. In the left panel, the auctions are numbered in ascending order starting from the soonest to close after the buyer's arrival. The buyer's arrival time is defined as the time of her first submitted bid. The right panel shows how distant (in minutes) the chosen auction was from the buyer's arrival. So, for example, 23.8 percent of arriving buyers chose to bid in auctions that were set to close within 5 minutes of the buyer's arrival, and approximately 54 percent chose auctions that ended within the hour. One reason for this is that buyers prefer to learn the outcome sooner rather than later, and conditional on winning an auction they want to receive the item as soon as possible. Another reason is that when a buyer searches for 'iPad' on eBay, the default sort order for the search results heavily favors auctions that are soon to close.[9]

An important assumption we make in our model is that the current vector of standing high bids is not informative about future vectors of standing high bids—i.e., the conditional distribution of the state at some re-arrival date $t + r$, conditional on the state at time t, is the same as the unconditional distribution of the state. This assumption is approximately true in the data. Over

---

[9]Unless the user specifies otherwise, the sort order is based on eBay's "Best Match" algorithm. For a straightforward search like 'iPad Mini' the Best Match sort index is highly correlated with the "ending soonest" index.

Table 2: Timing of chosen auctions, relative to buyer's arrival

Auction number*	% choosing	Cumul. %	Minutes to close	% choosing	Cumul. %
1	16.7	16.7	<5	23.8	23.8
2-10	28.8	45.5	5-30	18.3	42.1
11-30	26.2	71.7	30-60	11.8	53.9
31-50	10.7	82.4	60-180	25.4	79.3
51+	17.6	100.0	180+	20.8	100.0

* For each arriving bidder, auctions are numbered in ascending order based on time to close. So, for example, the table says that 16.7% of bidders chose to bid in the auction that was ending soonest when they arrived.

short time horizons, we do observe some persistence in the vector of standing high bids. To be precise, if r_t^o is the standing bid in the o^{th} next-to-close auction at time t, we find some evidence that the r_t^o's are serially correlated. However, the correlation is weak. For example, after conditioning on item characteristics the correlation between r_t^1 and r_{t+1}^1 (i.e., the correlation between the standing bid in this period's next-to-close auction and next period's next-to-close auction) is only 0.03. Moreover, the median number of auctions that finish in the time it takes a losing bidder to re-arrive is 6, and there is virtually no correlation between r_t^o and r_{t+6}^o for any o. Therefore it appears reasonable to assume that the current vector of standing high bids is not informative about future vectors of standing high bids.

As we argue below, an auction's "next-to-close interval" is important because the bidding decisions of buyers who arrive during that interval are the most useful for econometric analysis. The length of the interval depends on how closely the auction follows the previous one, and the number of bidders who arrive (stochastically) during the interval naturally depends on its length. Table 3 shows some summary statistics describing variation across auctions in the next-to-close intervals. The exact timing of an auction's close relative to the previous auction is difficult for a seller to control, and will typically just be a function of the (random) time at which the sellers submit their listings. In the empirical analysis below we therefore assume that variation in the next-to-close interval is exogenous.

Table 3: Next-to-close intervals

|	Mean	Std. Dev.	Percentiles		
			0.10	0.50	0.90
Length (minutes)	63.8	108.6	1.5	29.6	245.1
Number of bidder arrivals	9.2	10.2	0	6	29

4 Estimation

The primary objective of our empirical analysis is to recover F_E, the distribution of buyers' valuations. We observe bids and bidder identities, so we can distinguish in the data between bidders who are bidding for the first time and bidders who are returning to bid after losing in a previous auction. For convenience we will refer to these two types as "new bidders" and "returning losers."

As explained above, the option to bid again after losing causes bidders to submit bids (b) below their true valuations (x). To estimate F_E, we must invert the observed bids using the inverse bid function given in equation 4. Doing so requires estimates of the exit rate (α) and the distribution of the maximum rival bid (G). Our approach is to estimate these in a first step, and then plug the estimates into the inverse bid function to estimate F_E in a second step.

An important detail is that the items auctioned in our data are not perfectly identical. We adopt the conventional approach in the empirical auctions literature of working with normalized bids. We regress observed bids on item characteristics, Z, and then use the estimated coefficients from this regression to normalize bids as $\hat{b} = b - Z\hat{\gamma}$. These normalized bids then reflect the bids that would have been submitted if all auctions were for exactly identical items. The normalizing regression includes indicators for color (white vs. black); indicators for whether the auction included a cover, keyboard, screen protector, stylus, headphones, and/or extra charger; seller feedback ratings; shipping fee; month dummies (to control for a steady downward trend in prices over time); and day-of-week and hour-of-day dummies.

If we assume the exit rate, α, is constant, then we can estimate it simply as the fraction of losing bidders who do not return to bid again. Over a time horizon of three weeks, the fraction of bidders who exit is 0.406. Changing the time horizon, e.g. to two weeks or four weeks, has little impact on this estimate, since most bidders return relatively quickly if they are going to return at all. If we

Figure 1: Exit rate as a function of bid
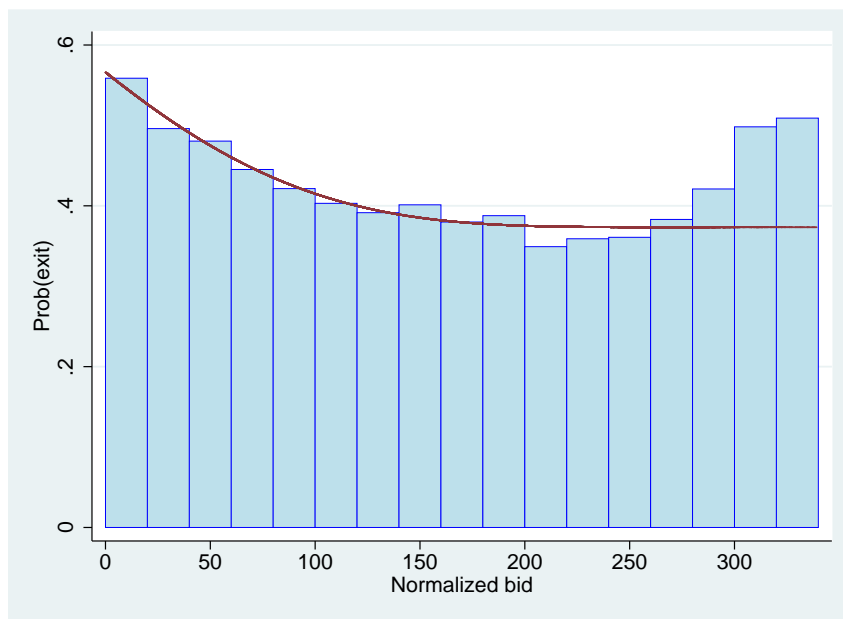
allow for endogenous exit, as discussed in Section 2.4, then we must estimate a bidder's probability of exit as a function of her bid. Since we observe bidders' bids and also whether they exit (which we define as not coming back within three weeks), we can estimate the exit function $\alpha(b)$ directly. Figure 1 shows binned exit frequencies along with an estimate of $\alpha(b)$ computed using the semi-nonparametric method proposed by Gallant and Nychka (1987). Exit rates are relatively flat with respect to bidders' types—at least over the relevant range. Bidders who submit bids below 100 are somewhat more likely to exit, but adjusting their value functions is unimportant: their value functions are essentially zero since they virtually never win the auction. There is a slight uptick in exit rates for high-value bidders, perhaps because these bidders elect to purchase at retail (or at a fixed price on ebay) when they lose an auction. However, in what follows we focus on the simpler model with exogenous exit, since the differences in exit rates are small, and allowing for endogenous exit makes simulating the model meaningfully more difficult.[10]

We estimate the re-arrival rate, β, as the inverse of the average return time. Since we assume re-arrival times are exponentially distributed, this is the maximum likelihood estimate of β. In the

---

[10]Since we can estimate the $\alpha(b)$ function, estimating the model with endogenous exit is not much more difficult than with exogenous exit. However, when simulating counterfactuals in a model with endogenous exit, we must compute a new equilibrium in which value functions accurately reflect exit functions and exit functions are optimal given the value functions.

Figure 2: Estimates of $G(m|b)$ at different levels of b
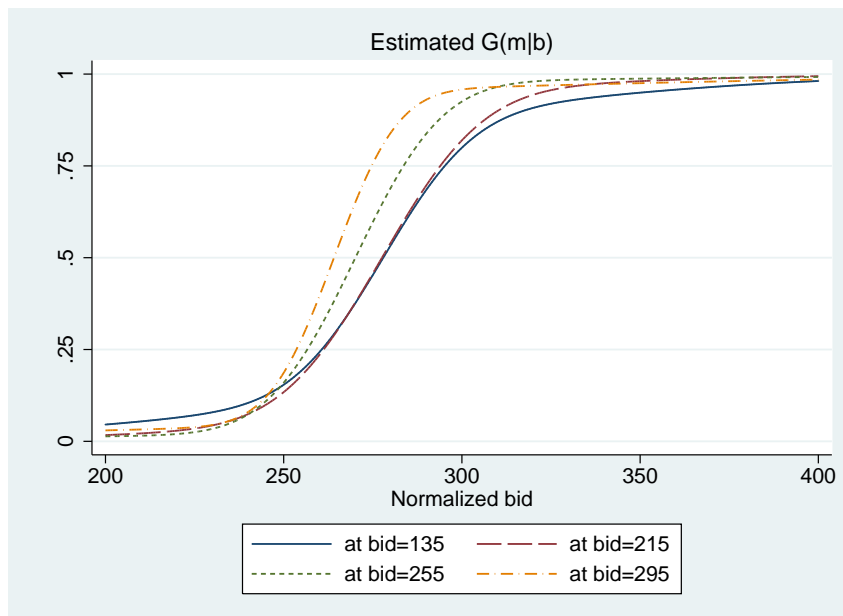
data, the distribution of re-arrival times does not perfectly match the exponential distribution—in particular, there is too much density at very short return times—but the approximation is reasonably good. Note that an estimate of β is not needed for estimating F_E: it drops out of the bid function because we are assuming no time discounting. However, an estimate of β is needed to simulate the model, which we do in the counterfactual analyses in Section 4 below.

We use the data on winning bids to estimate G, the distribution of the maximum rival bid. As noted above, this distribution may itself depend on the bidder's type, since a bidder's choice of which auction to bid in (i.e., the $\rho(x; \omega)$ function) depends on the bidder's type. To account for this dependency, we group bids into 45 bins of width 10, and compute a separate estimate of G for each of these bins. The estimate in each bin is computed by fitting a two-component normal mixture to the empirical distribution of winning bids corresponding to the bids in that bin.[11] Figure 2 shows the estimated G for four of the bins, chosen to roughly correspond to the 25^{th}, 50^{th}, 75^{th}, and 95^{th} percentiles of the bid distribution. The estimates show that high types tend to bid against stronger rivals, presumably because of the sorting induced by the ebay mechanism (lower types select out of auctions with high standing bids).

---

[11] For every submitted bid there is a corresponding winning bid in the auction; hence the data used for this exercise can be thought of as a set of (b, m) pairs.

Our primary empirical objective is to estimate F_E, which requires inverting the bids using equation (4). The dependence of G on type is important in computing $\int_0^b p\, dG(p)$, since here the shape of the distribution matters. (Elsewhere in equation (4) we only require the value of G at a given bid, which we could compute without allowing the shape of G itself to depend on bids.) Since we estimate G's in separate bins, when calculating the integral we use a linearly weighted average of the estimated G's from adjacent cells in order to avoid jumps in the value of the integral as we move from one bin to the next.[12]

With estimates of α and the function G, we can invert the observed bids to determine the bidder's underlying valuation. However, in estimating the distribution of these valuations, there are two reasons why we do not want to use all observed bids. One is the usual concern that much of the early bidding in eBay auctions does not appear to be serious bidding: early bids are often very low, and early bidders sometimes end up incrementing their bids upward. More importantly, our model implies that bidders who choose to bid in later (not soon-to-close) auctions are a selected sample: they may have opted not to bid in the soon-to-close auctions because the standing bids in those auctions exceeded their valuations. This means, for example, that to construct the likelihood for a bid submitted in an auction that was third in line to close, we would have to calculate the probability of the bid *conditional* on the bidder choosing not to bid in either of the previous two auctions—and this conditioning would depend in non-obvious ways on the standing bids in those two auctions.

To avoid such complications, we base our estimation on bidding that occurs during each auction's next-to-close interval: the period of time in which the auction is next to close. On average in our sample this period lasts 64 minutes (median $=$ 30 minutes), during which roughly 9 bidders arrive on average (median $=$ 6). We say a bidder "arrives" during auction t's next-to-close interval if she submits a bid in *any* auction during that time window. (We don't observe when the bidder literally arrives and starts browsing through auctions; we only observe when she submits a bid.) We construct the likelihood function based on arriving bidders' decisions about whether and what to bid in the next-to-close auction. Let b_{it} be the bid submitted by bidder i during the next-to-close interval of the t^{th} auction, and let x_{it} be the bidder's underlying valuation.[13] Let r_{it} be the

---

[12]For example, for a bid equal to 300, we use a 50/50 mixture of the estimated G's in the 290-300 and 300-310 bins. For a bid equal to 298, we use a 70/30 mixture of G's from the 290-300 and 300-310 bins.

[13]Bidders' valuations are indexed by t to reflect the fact that there may be observable differences in the items auctioned. As explained above, we control for these differences by working with normalized bids. We regress bids on item characteristics, and then compute the normalized bid as the baseline predicted bid plus the residual from this

standing high bid of the t^{th} auction when bidder i arrives, and p_t the closing price. We then make the following inferences from the bidding decisions:

- If the bidder chose not to bid in the next-to-close auction, we infer that $x \leq \sigma^{-1}(r_{it})$, where σ^{-1} is the inverse bid function.

- If the bidder bids in the next-to-close auction but loses with a bid below the final price, we infer that $\sigma^{-1}(b_{it}) \leq x_{it} < \sigma^{-1}(p_t)$.

- If the bidder bids in the next-to-close auction and wins, or loses but has the second-highest bid, we infer that $x_{it} = \sigma^{-1}(b_{it})$.

The likelihood for an observed bid b_{it} is therefore

$$
L_{it} = \begin{cases}
F_{it}(\sigma^{-1}(r_{it})) & \text{if did not bid} \\[2em]
F_{it}(\sigma^{-1}(p_t)) - F_{it}(\sigma^{-1}(b_{it})) & \text{if bid, but below final price} \\[2em]
f_{it}(\sigma^{-1}(b_{it})) & \text{if bid and was highest or second-highest bidder}
\end{cases}
\tag{6}
$$

where $F_{it} \in \{F_E, F_L\}$ depending on whether b_{it} was submitted by a new bidder (E) or a bidder who was returning after having lost in a previous auction (L).

With parametric assumptions on F_E and F_L, estimation of the distributions' parameters by maximum likelihood is straightforward. In the interest of flexibility, we instead construct nonparametric estimates of the two distributions using a simple binning method. As described above, each observed bid leads to an inference about the bidder's underlying valuation. We use these inference rules to fill in a histogram of valuations: that is, we separate the space of values into bins, and assign each bidder's valuation to the bin (or bins) implied by her bid. So, for example, if a bidder wins her auction, we count her in the bin containing her inverted bid. If a bidder bids in the

---

regression.

Figure 3: Estimated densities of valuations
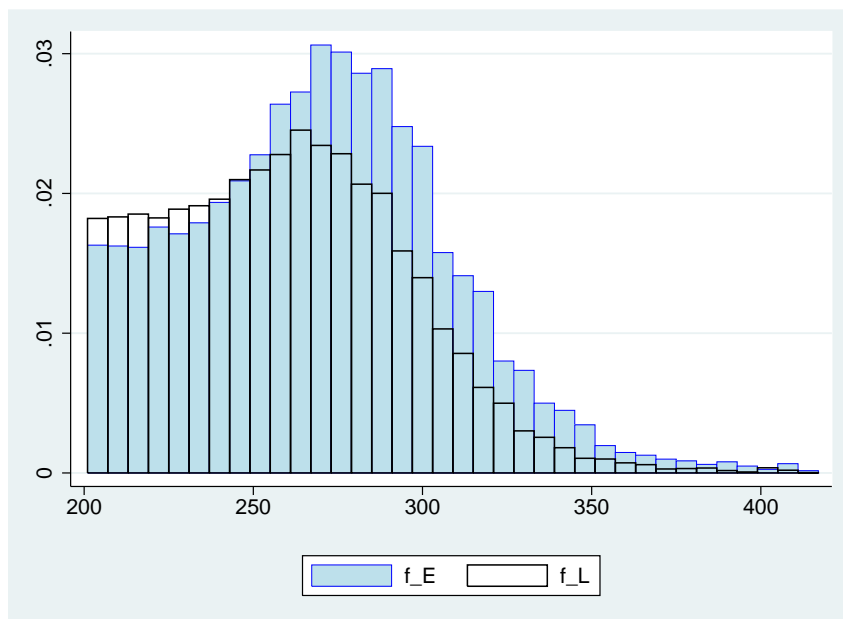
next-to-close auction but her bid ends up being below the second-highest bid, we assign her to the bins bracketing the inverted standing bid and the inverted price, giving equal weight to each bin (i.e., if the difference between the standing bid and the price spans n bins, we assign $1/n$ bidders to each of those n bins). For bidders who choose not to bid in the next-to-close auction, we assign them to bins below the inverted standing bid (again dividing the weights equally across bins).

Figure 3 shows the histogram estimates of the densities, f_E and f_L. We plot the two histograms only for values above 200, because the data are essentially uninformative about the shape of the distributions below that point. We observe very few auctions with standing bids or prices below 200, so our procedure yields estimated histograms that are almost perfectly flat between 0 and 200. As should be expected, f_L has less density on high valuations, since bidders with high valuations are less likely to show up as returning losers. This difference between f_E and f_L is driven entirely by the data, not the model; we are not imposing any restrictions on how the two distributions differ.

Before moving on to counterfactual analyses of the model, we note that our estimates imply dynamic incentives have a quantitatively meaningful impact on bidding. Most previous studies using eBay data have implicitly assumed that buyers are bidding myopically, interpreting the auction price

as a realization of a second-order statistic from the distribution of valuations. But in a dynamic framework buyers submit bids *below* their true values, due to the option value of losing. Since this option value is largest for buyers with high values—the buyers whose bids determine the final prices—estimates based on an assumption of static bidding may substantially understate both the level and the dispersion of buyers' true values. Based on our estimates, it appears that a static model of bidding would especially mis-estimate the upper tail of the distribution of bidder values. For example, we estimate that the winning bidder's true value (x) is on average \$31.07 higher than the bid she submitted.

Efficiency gains from bidder sorting

[Note: The simulations described below are preliminary. They are based on estimates with exogenous exit (i.e., α is estimated as a constant) and with F_E and F_L parameterized as lognormal distributions.]

Having recovered the primitives of the model—the distribution of buyers' valuations (F_E) and the exit rate (α) and re-arrival rate (β)—we can simulate market outcomes under alternative selling mechanisms. Note that while we are able to estimate these primitives without imposing a particular auction choice rule (the rule bidders use to choose which auction to bid in upon arrival), simulating the model requires specifying this choice rule. Obtaining prices from the simulation also requires computing the bid function, which requires solving for the continuation value $v(x)$. We do this by first simulating the auctions in type space (x), then determining the distribution of the highest rival type (H) as the average probability in the simulations that a buyer with type x wins, and then computing the value function as

$$v(x) = \int_0^x \frac{H(s)}{1 - (1 - \alpha)(1 - H(s))} ds \ .$$

We simulate prices and allocations (of items to buyers) under three regimes:

1. *Fully decentralized mechanism*: buyers are randomly assigned to auctions, which we implement by having each buyer bid in the next-to-close auction when she arrives. Bidding in the next-to-close auction is an equilibrium auction choice rule if the auction mechanism is closed

instead of open (i.e., if buyers do not observe standing bids and know nothing about previous bids in the auctions).

2. *Open auctions with sorting*: upon arrival, buyers bid in the soonest-to-close auction in which their optimal bid exceeds the current standing bid. Effectively, buyers go to the next auction in which their value is at least the second-highest among bidders who have already bid in that auction. We consider this to be a reasonable guess at the sorting rule that buyers actually use on eBay.

3. *Market-clearing with a uniform auction*: rather than having buyers bid in separate auctions depending on when they arrive, we pool all arriving buyers and pool all sellers, and find the price at which the market clears. If there are n sellers, the market-clearing price is the n^{th}-highest value among the buyers. This is an interesting benchmark because it yields the allocatively efficient outcome, and in principle it could be implemented (at least approximately) by eBay—for example, by holding daily uniform auctions.

Table 4 summarizes the outcomes of the simulations. The eBay mechanism, which posts standing bids and allows buyers to sort themselves into auctions, yields higher auction revenues than a fully decentralized (closed) mechanism would. The sorting allocates bidders across auctions in a way that leads to more uniform competition: high-value bidders will rarely be alone in an auction, but once two high-value bidders are participating in an auction, other arriving high-value bidders will likely choose other auctions to bid in. Our simulations indicate that the average price in the sorting mechanism is quite close to the market-clearing price. However, there is still considerable dispersion in auction prices in the open mechanism, so buyers with relatively low-values (including values below the market-clearing price) still occasionally win auctions. This means that the open mechanism is meaningfully less efficient than the market-clearing mechanism from an allocative standpoint—as indicated by the numbers in Table 4 on gross bidder surplus—even though it delivers almost the same revenues.

Thick vs. thin markets

To examine how our results depend on market thickness, we compare the market for 16GB iPad Minis to the market for 64GB iPad Minis. The latter market is a much thinner market, with

Table 4: Counterfactual prices and bidder surplus

| | Mean | SD | Percentiles | | |
			.10	.50	.90
Prices:					
Closed auctions	321.37	45.97	268.53	317.36	381.04
Open with sorting	327.27	40.68	280.35	322.36	380.96
Market clearing	327.87	–	–	–	–
Gross bidder surplus:					
Closed auctions	376.37	69.65	302.28	363.94	464.53
Open with sorting	378.20	68.13	305.45	365.01	465.29
Market clearing	391.29	60.81	335.07	374.33	469.85

only about 1.2 auctions per hour (compared to 2.9 for the 16GB version). The average number of bidders per auction is similar in the two markets (8.7 for 16GB, 7.8 for 64GB).

We can estimate our model for the 64GB market in the same way as described above for the 16GB market. Figure 4 shows, in the left panel, the observed distributions of prices (demeaned) for 16GB vs. 64GB models. The right panel shows the estimated distributions of values (f_E) for the two different models. While the price distributions differ substantially—there is much more dispersion in prices for the 64GB model—our model can explain the difference as being driven partly by differences in the bidding as opposed to differences in the underlying values. Our estimates indicate that the underlying distributions of values are actually somewhat similar (after correcting for the difference in the means). But the prices in the 64GB market are more dispersed because in the 64GB market bidders' option values of losing are lower; the weaker dynamic incentives lead to less compression of the bids.

Figure 4: Prices and valuations in 16GB and 64GB markets
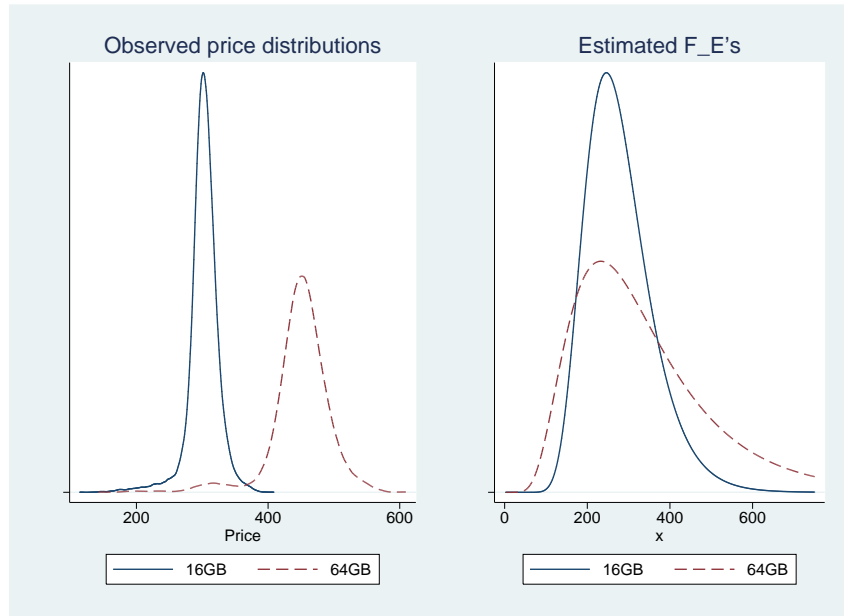
References

Ackerberg, D., Hirano, K., and Shahriar, Q. (2006). "The Buy-it-Now Option, Risk Aversion, and Impatience in an Empirical Model of eBay Bidding," working paper.

Adachi, A. (2014). "Competition in a Dynamic Auction Market: Identification, Structural Estimation, and Market Efficiency," forthcoming in *Journal of Industrial Economics*.

Ashenfelter, O. (1989). "How Auctions Work for Wine and Art," *Journal of Economic Persepectives*, 3, pp. 23-36.

Backus, M. and G. Lewis (2012). "A Demand System for a Dynamic Auction Market with Directed Search," working paper.

Bajari, P. and Hortacsu, A. (2003). "The Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions", *RAND Journal of Economics*, 34, pp. 329-355.

Barbaro, S. and Bracht, B. (2006). "Shilling, Squeezing, Sniping: Explaining late bidding in online second-price auctions," working paper.

Canals-Cera, J.J. and Pearcy, J. (2006). "Econometric Analysis of English Auctions: Applications to Art Auctions on eBay," working paper.

Elyakime, B., J.J. Laffont, P. Loisel and Q. Vuong (1994) "First-Price, Sealed-Bid Auctions with

Secret Reservation Prices," *Annales d'Economie et de Statistique,* 34: 115-41.

Gale, I. and Hausch, D. (1994). "Bottom-fishing and declining prices in sequential auctions," *Games and Economic Behavior*, 7, pp. 318-331.

Gallant, A. R. and D. W. Nychka (1987). "Semi-nonparametric maximum likelihood estimation." *Econometrica*, 363-390.

Gavazza, A. (2011). "The Role of Trading Frictions in Real Asset Markets," *American Economic Review*, 101, pp. 1106-1143.

Gavazza, A. (2013). "An Empirical Equilibrium Model of a Decentralized Asset Market,", working paper.

Gonzalez, R., Hasker, K., and Sickles, R. (2004). "An Analysis of Strategic Behavior in eBay Auctions," pp. 1535-1546.

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

Jeitschko, T. (1999). "Equilibrium Price Paths in Sequential Auctions with Stochastic Supply," *Economics Letters*, 64, pp. 67-72.

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

Katzman, B. (1999). "A Two Stage Sequential Auction with Multi-Unit Demands," *Journal of Economic Theory*, 86, pp. 77-99.

Lewis, G. (2011). "Asymmetric Information, Adverse Selection and Seller Disclosure: The Case of eBay Motors," *American Economic Review*, 101(4), pp. .

McAfee, R. and Vincent, D. (1993). "The Declining Price Anomaly," *Journal of Economic Theory*, 60, pp. 191-212.

Myerson, R. B. (2000). "Large poisson games." *Journal of Economic Theory*, 94(1), 7-45.

Ockenfels, A. and Roth, A. (2006). "Late and multiple bidding in second price Internet auctions: Theory and evidence concerning different rules for ending an auction," *Games and Economic Behavior*, 55, pp. 297-320.

Rasmusen, E. (2006). "Strategic Implications of Uncertainty over One's Own Private Value in Auctions," *B.E. Journals in Theoretical Economics: Advances in Theoretical Economics*, 6, pp. 1-24.

Roth, A. and Ockenfels, A. (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, pp. 1093-1103.

Satterthwaite, M. and A. Shneyerov (2008). "Convergence to Perfect Competition of a Dynamic Matching and Bargaining Model with Two-sided Incomplete Information and Exogenous Exit," *Games and Economic Behavior*, 63, pp. 435-467.

Satterthwaite, M. and A. Shneyerov (2008). "Dynamic Matching, Two-sided Incomplete Information, and Participation Costs: Existence and Convergence to Perfect Competition," *Econometrica*, 75, pp. 155-200.

Stryszowska, M. (2004). "Late and Multiple Bidding in Competing Second Price Internet Auctions," FEEM Working Paper No. 16.2004.

Wang, J. (2006). "Is Last Minute Bidding Bad," working paper.

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