

NBER WORKING PAPER SERIES

DYNAMICS AND EFFICIENCY IN DECENTRALIZED ONLINE AUCTION MARKETS

Kenneth Hendricks
Alan Sorensen

Working Paper 25002
<http://www.nber.org/papers/w25002>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2018

We are grateful to Neil Sundareshan and eBay for providing data for this project, and to Matt Backus, Hugo Hopenhayn, Brad Larsen, Rob Porter, and Dan Quint for helpful suggestions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Kenneth Hendricks and Alan Sorensen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Dynamics and Efficiency in Decentralized Online Auction Markets
Kenneth Hendricks and Alan Sorensen
NBER Working Paper No. 25002
September 2018
JEL No. D10,D4,D44,D47,L0

ABSTRACT

Economic theory suggests that decentralized markets can achieve efficient outcomes if buyers and sellers have many opportunities to trade. We examine this idea empirically by developing a tractable dynamic model of bidding in an overlapping, sequential auction environment and estimating the model with detailed data from eBay. Bidders in the model discount their bids to reflect the option value of losing – if they lose, they can come back to try again – and the structure of the model makes it so they effectively bid against a stationary distribution of rivals. We find that dynamic participation makes the market meaningfully more efficient than a benchmark in which buyers have only one opportunity to bid, but the observed outcomes still fall well short of the fully efficient competitive equilibrium.

Kenneth Hendricks
Department of Economics
University of Wisconsin, Madison
1180 Observatory Drive
Madison, WI 53706
and NBER
hendrick@ssc.wisc.edu

Alan Sorensen
Department of Economics
University of Wisconsin, Madison
6454 W.H. Sewell Social Science Bldg
1180 Observatory Drive
Madison, WI 53706
and NBER
sorensen@ssc.wisc.edu

1 Introduction

Many goods, especially used goods, are sold in decentralized, dynamic auction markets. For instance, online platforms like eBay create virtual markets in which a large number of sellers and buyers arrive over time, get matched, and bargain over price. These markets are dynamic because buyers who fail to purchase, and sellers who fail to sell, can return to the market to try again. They are decentralized because individual buyers or sellers decide with whom to match. This matching may be inefficient, and frictions resulting from private information and strategic behavior can cause trade within a matched set of traders to be inefficient. However, the opportunity to trade many times makes the market for each trader large over time, and this feature can mitigate the effects of matching and trading frictions. This raises two questions. The positive question is: How does the option to trade again affect prices and efficiency? The normative question is: Can a centralized mechanism—in which the intermediary matches buyers and sellers—achieve significantly better outcomes?

[Satterthwaite and Shneyerov \(2007, 2008\)](#) have addressed these questions from a theoretical perspective. They study a stylized environment with two-sided private information in which many buyers and sellers arrive each period to trade one unit of a homogeneous good. Every seller holds a first-price auction with a reservation price for her unit, and each buyer is randomly matched to one seller. Buyers and sellers who fail to trade either exit or return the following period to be rematched. The authors focus on what happens to prices and allocations when the number of trading opportunities for each trader is large. They show that all steady state, Bayesian equilibria converge to the competitive outcome. More precisely, the distribution of prices converges to the Walrasian price and, in the limit, only sellers with values below the Walrasian price sell the good, and only buyers with values above the Walrasian price get the good. The authors conclude that simple selling mechanisms like individual auctions can allocate supply almost efficiently in a decentralized, dynamic auction, and that the gains from running a centralized mechanism—like a double auction for all buyers and sellers in a period—may be quite limited.

This paper studies these questions from an empirical perspective. The goal is to quantify the impact of dynamic competition on prices and efficiency in a real-world market, and evaluate how close they come to the Walrasian outcome. A simplifying feature of our application is that sellers do not appear to value the good or the option to sell again: most sellers choose reserve prices at which they are certain to sell. Of the sellers who set binding reserve prices, only a small fraction fail to sell, and an even smaller fraction return to sell again. By contrast, most buyers lose, and half of them return to bid again. Therefore, in our model, we treat the sellers as myopic players whose values for the good are zero and focus on the dynamic incentives of the buyers. In this setting, a

buyer’s option to return after a failed bid has two effects on prices and efficiency. It increases the level of competition in the market because the number of bidders includes both new and returning buyers. We refer to this as the *dynamic participation* effect. It also changes bidding behavior. Anticipating the possibility of returning at a later time, buyers shade their bids below their values to reflect the option value of losing. We refer to this as the *dynamic bidding* effect. Our objective is to measure the impact of these two effects on the stationary distribution of prices and on allocative efficiency.

Our model is tailored to capture the main features of the eBay marketplace. As in the Satterthwaite and Shneyerov model, each seller has only one unit to sell and each buyer wants only one unit, their private values do not change over time, and buyers who fail to win return to the market with positive probability. However, the arrival process and auction mechanism are different. In our model, new buyers and sellers arrive randomly in continuous time, and the return times of losing buyers are also stochastic. Each seller offers her unit in an ascending, second-price auction of fixed duration. Thus, at any moment in time, the market consists of a large number of overlapping second-price auctions with different closing times. The matching process is also different. Auctions on eBay are *open*, meaning that buyers observe the state of play in each available auction—the highest losing bid and time remaining—and use this information to decide in which auction to bid and how much to bid. For example, low value buyers do not participate in auctions where the posted bid exceeds their willingness-to-pay. As a result, buyers are sorted across auctions based on their values.

We provide a partial characterization of the equilibrium. When a buyer has the option to bid again after losing, her optimal bid depends on her beliefs about the state of the market if and when she returns to try again. We assume that she believes this state will be a random draw from the stationary distribution generated by equilibrium play. The assumption is plausible in large markets where the arrival rate of buyers is sufficiently high that the state is likely to undergo many transitions before a losing buyer returns. This is the case for our application. Given this assumption, we show that a buyer’s equilibrium bid is equal to her value less a continuation value that depends only on her type. Consequently, her equilibrium bid is invariant to the state of play or choice of auction. We do not solve for the equilibrium auction choice strategy, which (unlike the bid function) does depend on the state of play, and is consequently difficult to derive. While this means we can only partially characterize the equilibrium, we show that the primitives of our model can be empirically identified without specifying a model of auction choice.

We estimate our model using detailed data from eBay auctions for used iPads. The main primitives are the entry rate of new buyers, the exit and return rates of losing buyers, and the distribution of values of new buyers. We observe all bids, including the highest bid, and the identities of

the bidders (and sellers). The latter allows us to identify exits and distinguish between new and returning buyers. Given this information, we can estimate the entry, exit and return rates of buyers directly from the data. The primary empirical task is to then identify and estimate the unknown distribution of new buyers' values.

To achieve this task, we exploit the monotonicity and invariance properties of the equilibrium bid strategy. As noted by [Backus and Lewis \(2017\)](#), these properties imply that a bidder's type can equivalently be represented by her equilibrium bid, or "pseudotype." By applying the transformation of variables that [Elyakime, Laffont, Loisel and Vuong \(1994\)](#) and [Guerre, Perrigne and Vuong \(2000\)](#) introduced for static, first-price auctions, we show that the unobserved value of a buyer can be expressed as a function of her bid b and $G_{M|B}$, the distribution of the maximum rival bid M conditional on b . The dependence between a buyer's bid and her maximum rival bid arises from two sources. One is that her choice of auction is based on the highest losing bid that is posted in each auction; the second is that her bid may affect the decisions of later buyers who arrive before the close of the auction. We account for this correlation by estimating $G_{M|B}$ using the set of auctions chosen (in the different states) by buyers who bid b . Thus, the distribution of values of new buyers can be identified and estimated based on the choice rule that buyers actually use, instead of deriving the equilibrium choice rule or imposing assumptions about that choice rule.

Our model generates several testable implications. For the data to be consistent with our model, the estimated inverse bid function needs to be increasing in b . The invariance property implies that buyers who lose and return should bid the same amount. Since we observe bidder identities, we can track the bids of buyers who lose and return, and directly test this implication. In a stationary equilibrium, the number of returning bidders per auction depends only on the arrival rate of new buyers and the exit rate of losers. More generally, the flow of losing buyer types that return to the market must on average be equal to the flow of buyer types (new and returning) that leave the market, either by winning or by exiting. The latter restriction implies that the stationary density of losers' values is proportional to the density of new bidders' values. Since we can estimate the distribution of values of returning buyers directly from the data, we use the restrictions to test the model. We find that the data do not reject any of these implications.

Given estimates of the model primitives, we simulate a number of counterfactuals. The first set focuses on the allocative efficiency of the eBay trading mechanism, which we compare to two hypothetical benchmarks. The first is the fully efficient benchmark, which we compute by finding the price that would clear the market if all buyers and sellers were pooled. The second is the inefficient benchmark achieved by a *static*, decentralized auction market in which each buyer gets only one chance to win and is randomly matched to a seller. We find that the actual eBay mechanism meaningfully increases efficiency relative to this second benchmark, but falls well short of the fully

efficient outcome.

A second set of simulations evaluates the impact of the auction selection rule on efficiency and prices. In particular, we simulate outcomes when buyers bid in the soonest-to-close auction, which randomly matches buyers to auctions based on their arrival times. This is a dynamic version of the random matching that occurs in a static, decentralized market, except that in this case buyers are able to return to try again if they lose. We find that this selection rule would increase allocative efficiency relative to the outcomes we observe in the data, because the selection rule we observe in the data tends to disproportionately match high-value buyers with other high-value buyers. This matching of high-value buyers also generates significantly more price dispersion than would be observed under the counterfactual selection rule.

In the third set of simulations, we measure the magnitude of the dynamic bidding effect for different values of the exit rate. In theory, the market should converge to the competitive outcome as the exit rate goes to zero. Buyers with values above the market-clearing price should bid that price and win almost surely (although it may take many tries), and buyers with values lower than the market-clearing price should bid their value and lose almost surely. Thus, as the exit rate goes to zero so that buyers have many opportunities to bid, the dynamic bidding effect should eliminate prices above the market-clearing price, and the dynamic participation effect should eliminate prices below the market-clearing price. Our simulations suggest that the latter effect is more substantial: when the exit rate is small, low prices are mostly eliminated, but a surprising amount of dispersion above the market-clearing price still remains.

2 Related Literature

Our structural analysis assumes that buyers are forward-looking and bid optimally. There is an extensive literature that tests implications of this assumption in a variety of settings. The early papers focused on environments with a fixed number of buyers and a finite number of objects. [Milgrom and Weber \(1999\)](#) showed that, when only the winning bid is revealed after each auction, equilibrium prices in sealed bid, sequential auctions of k identical objects are a martingale; the price of each object is on average the same and equal to the expected value of the $(k + 1)$ -st order statistic. [Engelbrecht-Wiggans \(1994\)](#) extended this analysis to stochastically equivalent objects, a setting in which a buyer draws an independent private value for each object and does not learn that value until it is sold. He shows that if the number of objects is sufficiently large, then prices will on average decline. A large number of empirical studies have tested these predictions (see the survey by [Ashenfelter and Graddy \(2000\)](#)), and the evidence suggests that buyers are often forward-looking. For example, [Ashenfelter \(1989\)](#) examined prices of pairs of identical wines sold in the same lot size—an environment that matches the one described by [Milgrom and Weber \(1999\)](#),

in which prices are predicted to be the same on average. He found that prices were twice as likely to decrease as increase, but they were identical for the majority of the 4615 comparisons.

The more recent literature considers settings where buyers arrive (and exit) randomly over time and face an infinite sequence of sealed-bid, second-price auctions. [Zeithammer \(2006\)](#) studies a market for differentiated goods in which each buyer values only one type of good but can observe the types of goods that will be for sale in the next few auctions.¹ He exploits exogenous variation in this state to show that buyers in eBay auctions of MP3 players and DVD movies bid more aggressively when their preferred product is in lower supply. [Coey, Larsen, and Platt \(2015\)](#) study a consumer search setting in which all buyers have the same value for the good but face different deadlines for procuring it. In their model, a buyer increases her bid and is more likely to buy from a fixed-price listing as her deadline approaches. They show that bidder behavior in eBay markets is consistent with these predictions.²

Our paper provides a framework for evaluating the extensive structural literature that models eBay auctions as a sequence of independent, static games with exogenous participation (e.g., [Bajari and Hortascu \(2003\)](#), [Gonzalez, Hasker, and Sickles \(2004\)](#), [Canals-Cerda and Percy \(2006\)](#), [Ackerman, Hirano, and Shahriar \(2006\)](#), and [Lewis \(2007\)](#)). Our results suggest this may not be a good approximation. First, by ignoring the distinction between new and returning buyers, these papers implicitly treat the stationary distribution of buyer values as the primitive rather than the distribution of new buyer values. This matters for counterfactuals (e.g., higher reserve prices) since changes in the auction mechanism will lead to a different stationary distribution of buyer values. Second, by ignoring the overlapping structure of the auctions, they fail to account for buyers choosing an auction based on the state of play. We find that the selection effects from endogenous participation have a significant impact on our estimates of the distributions of buyer values. Finally, by ignoring the option to bid again, they overestimate the values of the bidders, although our results suggest that this effect may be second-order.

Our paper is most closely related to the nascent structural literature ([Adachi \(2016\)](#), [Backus and Lewis \(2017\)](#), [Bodoh-Creed et al \(2017\)](#)) that models online auctions like eBay as a dynamic game in which bidders know their values for the goods and those values are perfectly persistent over time. This setting is very different from the one studied by [Jofre-Bonet and Pesendorfer \(2003\)](#) in their

¹[Said \(2011\)](#) examines bidding for stochastically equivalent objects and shows that a unique symmetric Markov equilibrium exists in which buyers bid their value less a continuation value that depends on the level of competition. [Budish and Zeithammer \(2017\)](#) also study this setting but consider the possibility that bidders know their values of the objects and examine how this information affects equilibrium bidding behavior in the special case of two objects.

²These papers focus on inter-auction dynamics. [Hopenhayn and Saeedi \(2017\)](#) study intra-auction dynamics of equilibrium bidding in second-price auctions when bidding opportunities and values follow a joint Markov process. They show that their model can explain repeat bidding and sniping and estimate the parameters of the stochastic process by matching moments.

seminal paper on estimation of repeated first-price auctions. They assume that each bidder gets an independent draw for each object and does not learn that value until the object is auctioned. This assumption simplifies the analysis, because it implies that a bidder cannot learn anything about her rivals' types from their previous bids, and as a result she does not have to worry about how her bids may affect her rivals' behavior in future auctions. The eBay papers deal with this issue by assuming that the market is sufficiently large that bidders' beliefs about the distributions of the maximum rival bid in the current and future auctions are determined by the stationary distributions implied by equilibrium play. Given this assumption, they show (as we do) that the equilibrium bid of a buyer is her type less a continuation value that depends only on her type, and that the latter can be estimated from bid data. Of course, the continuation value function and how it is estimated depends on the model.

Backus and Lewis (2017) treat eBay auctions as a sequence of sealed bid, second-price auctions of heterogeneous goods. A random number of buyers enter each period and, if they lose or choose not to bid, either exit or go into hibernation for a random number of periods.³ Bodoh-Creed et al (2017) propose a model in which a continuum of sellers and buyers enter each period, are randomly matched in sealed bid, second-price auctions, and exit only when they trade. Adachi (2016) develops a model like ours, but she assumes that when bidders arrive they are matched to the soonest-to-close auction in which the posted price is less than their optimal bid, and that they bid in that auction as if it is a sealed bid auction. In each of these models, the price or winning bid in an auction can be treated as a second or first order statistic from a random set of buyers. As a result, the models in these papers can be identified and estimated using only data on prices or winning bids.⁴

Our contribution consists of developing a model that is arguably a better approximation to the data-generating mechanism, providing conditions under which it is identified, and showing how it can be estimated. The distinguishing features of our model are that the auctions are *overlapping* and *open*. These features imply that the matching of buyers to sellers is endogenous. When bidders arrive, they observe the posted prices and closing times of the auctions and bid in the auction that is the best match for them.⁵ This sorting means that the price and winning bid in an auction cannot be treated as realizations of second or first order statistics, since bidders with the highest values among a set of bidders who arrive during a period often choose to bid in different auctions. The sorting can also generate dependence between a buyer's bid and maximum rival bid. Consequently,

³Buyers observe the current product for sale but not future products, so the payoff relevant state for each bidder is the product and her (multidimensional) type.

⁴If the number of potential bidders is unobservable and not all bids are observed, then methods proposed by Song (2004) and Freyberger and Larsen (2017) can still be used to estimate the underlying distribution of buyers' valuations.

⁵We assume that all bidders find a match since there is always an auction that closes within a day that has no bids and a low start price.

the continuation value in our model is a function of $G_{M|B}$, and identifying and estimating our model requires observing all bidders, the auctions in which they bid, and the value of their bids (including the winning bid). This data requirement is analogous to the result of [Athey and Haile \(2002\)](#) that the symmetric, affiliated private value model is not identified unless all bids are observed. The difference here is that the dependence is induced through bidders' selection of auctions rather than through their values.

Our model nests random matching as a special case. Random matching can occur in our model if buyers wait until the last minute to submit their bids or if the auctions are sealed bid. In these instances, the set of auctions in which buyers submitted $B=b$ is a random sample. This would imply that the fraction of auctions in which b is the winning bid among auctions selected by pseudotype b is equal to the fraction of *all* auctions in which the winning bid is less than b , in which case the value function is determined by G_M , the stationary distribution of the highest rival bid. This distribution can be estimated non-parametrically from the distribution of winning bids. On the other hand, a significant difference between these two fractions would indicate that the matching of buyers to auctions is non-random. In that case, estimates of the distribution of buyer values obtained from applying the inverse bid function using G_M rather than $G_{M|B}$ would be biased. In our empirical analysis, we estimate both distributions and find that they differ substantially.

More generally, our paper contributes to the theoretical literature on dynamic models of trade with private information. [McAfee \(1993\)](#) considers a discrete time model in which a large number of buyers and sellers are matched randomly each period, and return in the next period if they fail to trade. He shows that, in steady state, there is an equilibrium in which all sellers choose to sell via second-price sealed-bid auctions. As noted above, [Satterthwaite and Shneyerov \(2007, 2008\)](#) show that the allocation in these kinds of environments converges to the competitive outcome when traders are forward-looking and the number of trading opportunities per unit of time gets large. [Loertscher, Muir, and Taylor \(2018\)](#) use the tools of dynamic mechanism design to characterize the optimal market clearing policy in thin markets. Their focus is on the tradeoff between efficiency and costly delay. The designer can clear compatible trades when they occur, as in the eBay mechanism, or accumulate buyers and sellers and clear the market at fixed intervals, as in the centralized exchange mechanism we consider in Section 7. Our main contribution to this literature consists of empirically measuring the efficiency gains that result from letting the market thicken and clearing it at fixed intervals.⁶

Finally, our paper contributes to the empirical literature on search-and-bargaining models of trade, such as in [Gavazza's \(2011, 2016\)](#) studies of the market for used aircraft, [Brancaccio et al's \(2018\)](#)

⁶In our application, buyers and sellers arrive at sufficiently high rates that most of the gains can be achieved by clearing the market every two or three days. Consequently, delay costs are likely to be negligible for most buyers.

study of global shipping, and Buchholz’s (2017) study of the New York City taxi market. These papers approximate markets with finite numbers of buyers and sellers with a continuum of agents. They focus on the steady state of the dynamic market and use the restrictions on entry and exit flows as the basis for estimating the models’ primitives.⁷ By contrast, we work with the stationary state of a finite market, and we use the restrictions on flows as over-identifying tests of our model.

3 A Dynamic Model of Trade

An infinite number of sellers arrive over time to sell a single unit of a homogenous good. Each seller values the good at zero. Upon arrival, each seller contracts with the platform to sell her unit in an ascending, second price auction. Time is indexed by t . It will be convenient to assume that sellers’ arrival times are equally spaced and normalize the time between arrivals to be one. The duration of each auction is set to J units of time, where J is a large integer. Thus, at any time t , the market consists of a set of J overlapping auctions.⁸ 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. Note that, if the time remaining in the next-to-close auction is $d \in (0, 1)$, then the closing time of auction $j = 2, \dots, J$ is equal to $d + j - 1$.

Each seller chooses a start price, which we will refer to as the seller’s *reserve* price.⁹ At any time t , the platform posts the current highest *losing* bid in each auction or, if no one has submitted a bid yet, the seller’s reserve price. The platform does not reveal the current high bids in the auctions. We will refer to the highest losing bid (or reserve price if there are no bids) in an auction as the *posted* bid in that auction. Let $r_j(t)$ denote the posted bid in auction j at time t and define $r(t) = (r_1(t), \dots, r_J(t))$ as the vector of posted bids in the market at that time. In an auction that receives bids, the unit is awarded to the highest bidder at the second-highest bid or, if there is only one bid, the reserve price. Sellers who sell their units exit the market. If an auction fails to attract any bids, then the seller either exits or returns at some future time to sell the unit again.

On the demand side, buyers arrive randomly over time to buy a single unit of the good. Each buyer is endowed with a private value x for the good that is an independent, random draw from a distribution F_E with density f_E and support $X = [0, \bar{x}]$. The values of the buyers do not change over time. A buyer who bids in auction j at time t 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. Let $w_j(t)$ denote the high bid in auction j at time t . If a buyer’s proxy bid exceeds $w_j(t)$, then the posted bid $r_j(t)$

⁷Bodoh-Creed et al (2017) also use the steady state restriction that the flow of new buyer types into the market has to equal the flow of types that leave the market to help identify and estimate their model.

⁸In our application, sellers arrive at a rate of roughly one per hour and auction duration is typically seven days.

⁹eBay also allows sellers to set a secret reserve price. We ignore this choice since most sellers in our application choose not to exercise this option.

increases to $w_j(t)$ and she becomes the high bidder. If her bid remains the high bid until closing time, then she wins the auction and exits. If her proxy bid is less than the current high bid in the auction, or is less than the bid of a subsequent bidder, then she loses the auction and either exits the market with probability α or enters the pool of losers with probability $1 - \alpha$ and returns at some future time (which can be before auction j closes) to bid again in a *different* auction. An important feature of our model is that losing buyers do not return immediately. The exit rate α is exogenous and the same for all buyers, although later we show how this assumption can be relaxed.

We impose two restrictions on bidding behavior in our theoretical model. First, a buyer bids when she arrives. This restriction rules out strategic delay (e.g., last-minute bidding or “sniping”). Second, the bid is a proxy bid and therefore her only and final bid in that auction. This restriction rules out bidding strategies such as bidding incrementally more than the posted bid until the bid exceeds the high bid. As we show later, these modeling restrictions allow us to focus on inter-auction dynamics, the main concern of this paper. Of course, the option to bid again in another auction reduces the value of behaving strategically within an auction. Nevertheless, in our application we do observe some buyers bidding incrementally; we examine the prevalence of this kind of bidding behavior and discuss how we address it in our empirical analysis below.

In the theoretical analysis that follows, we focus on buyers and treat sellers as non-strategic players who set their reserve prices to zero. The main reason is that, in our application, sellers seem primarily interested in selling their item and do not appear to value the good or the option to sell again. Most sellers choose very low reserve prices at which they are certain to sell. Of the sellers who set binding reserve prices, only a small fraction fail to sell, and an even smaller fraction return to sell again. By contrast, most buyers lose, and half of them return to bid again.

3.1 Market Dynamics and Strategies

Over any interval of time, two kinds of buyers can arrive: new buyers and returning buyers. The arrival process of new buyers is an exogenous Poisson process with parameter λ . Thus, the expected number of buyers arriving in any period of length Δ is $\lambda\Delta$. The arrival process for returning buyers depends upon the number of buyers in the losers’ pool. The amount of time each buyer spends in the pool of losers before returning is distributed exponential with parameter β . Thus, in any interval $[t, t + \Delta)$, the probability that a losing buyer returns, conditional on not having returned earlier than t , is $1 - e^{-\beta\Delta}$, which for small Δ is approximately equal to $\beta\Delta$. This arrival rate does not depend on when the buyer entered the pool, on how long she has been in the pool, or on her value. The number of buyers who leave the pool in an interval $[t, t + \Delta)$ is therefore a Binomial random variable with parameters $k(t)$ and $\beta\Delta$, where $k(t)$ denotes the number of buyers in the pool at time t . The expected number of buyers who leave the pool during this period is $\beta\Delta k(t)$.

The size and composition of the pool evolves stochastically over time as new and returning buyers arrive and submit bids. When a buyer leaves the pool, she either (i) submits a losing bid and returns immediately with probability $1 - \alpha$; or (ii) submits a high bid and displaces another buyer who then enters the pool immediately with probability $1 - \alpha$; or (iii) bids in an auction that has not yet received any bids. Thus, the flow of buyers out of the pool causes it to shrink over time. Similarly, when a new buyer arrives and bids in an auction that already has at least one bid, she either submits a losing bid or displaces a current high bid, and the losing buyer then enters the pool with probability $1 - \alpha$. Thus, the flow of new buyers into the market causes the pool to grow over time. The probability laws determining the evolution of types in the losers' pool depend on the behavior of the buyers.

We are interested in studying the market in its stationary state, which requires the size of the losers' pool to satisfy the condition that the number of buyers leaving the pool is on average equal the number of buyers entering the pool. In any interval of time in which one auction closes, this condition implies that the number of buyers in the pool fluctuates around¹⁰

$$\bar{k} = \frac{(\lambda - 1)(1 - \alpha)}{\alpha\beta} \quad (1)$$

The distribution of types in the losers' pool is governed by a stationary distribution, which we denote by F_L . This distribution will be a rescaling of the primitive F_E , reflecting the fact that low-value buyers have a higher likelihood of being selected into the losers' pool. We discuss steady-state restrictions on the relationship between F_E and F_L in more detail below.

We turn next to the strategies of buyers. Upon arrival, a buyer observes the posted bids and closing times of the available auctions but not the losers' pool or the past history of the market.¹¹ In forming beliefs about the probability of winning an auction, we assume that a returning buyer ignores any information from previous arrival states. However, her beliefs can still differ from those of new buyers since she knows that the pool of losers has one less buyer in it. But, in our application, \bar{k} is quite large, so the number of returning buyers in any interval Δ is approximately distributed Poisson with parameter $\beta\bar{k}\Delta$. The Poisson distribution has the property that a returning buyer's belief about the number of returning rivals conditional on her arrival is the same as a new buyer's unconditional belief.¹² In what follows, we adopt the Poisson approximation and assume that both

¹⁰This expression follows from recognizing that the number of returning and new buyers arriving during the period is on average $\beta\bar{k}\Delta + \lambda\Delta$ and the number of buyers that enters the losers' pool is on average $(1 - \alpha)(\beta\bar{k}\Delta + \lambda\Delta - 1)$.

¹¹In our application, the platform reports the history of highest losing bids and partially masked identities of losing buyers for each auction. The assumption here is that buyers do not bother to use this information in forming beliefs about the pool of losers or about the high bids. The value of this information is likely to be quite small in thick markets where buyers have the option to bid again.

¹²Myerson (JET, 2000) refers to this property of Poisson games as *environmental equivalence*.

types of buyers believe that the number of returning buyers is distributed Poisson, and that the values of returning bidders are drawn independently from F_L . Since new and returning buyers have the same beliefs about the probability of winning, we assume that they also have the same strategies.

We focus on stationary strategies that map the observed state of the market, $\omega = (r, d) \in \Omega$, into actions. More formally, a buyer's strategy consists of a stochastic selection or matching rule $\rho : X \times \Omega \rightarrow \Delta \mathbf{J}$, which associates her type (x) and the observed state (ω) with a probability measure over the set of available auctions ($\Delta \mathbf{J}$, with $\mathbf{J} = \{1, \dots, J\}$), and a bidding function $\sigma : X \times \Omega \times \mathbf{J} \rightarrow \mathfrak{R}_+$ that specifies the bid she would submit in each of the auctions in her choice set. Given any stationary strategy (ρ, σ) , the state of the market is a Markov process. We assume this process is stationary and ergodic.

3.2 Optimal Bidding Behavior

The optimal bidding decision of a buyer in our model depends on her beliefs about what the state of the market will be if and when she returns. Letting Φ denote the stationary distribution of the market state, we impose the following assumption about those beliefs.

Assumption 1: Given any arrival state ω , a buyer believes that the state of the market if and when she returns is a random draw from Φ .

Assumption 1 is a strong assumption. It implies that a buyer believes that the level of competition when she returns does not depend on the current state or her actions. This is critical for our analysis. In thin markets, the current state may be informative about the return state, and a buyer may need to take into account how her losing bid can affect her continuation value through its impact on the decisions of subsequent buyers. This dynamic incentive problem would greatly complicate the analysis of the game. However, Assumption 1 is plausible in thick markets—or, to be more precise, if the arrival rate of buyers is sufficiently rapid relative to an individual buyer's return time. In our application, by the time a losing bidder returns to bid again, on average over fifty other buyers have arrived, and six auctions have closed. Thus, an individual bidder's decisions will have negligible impact on the state at the time of her return, since typically many other market transitions will have taken place. We also find that the correlation between the current state and future states is quite low, and essentially zero after six auctions have closed. Therefore, we view Assumption 1 as a good approximation for our application.

We now formalize the buyer's optimization problem. Let $M_j(\omega)$ be the random variable denoting the highest rival bid in auction j (including the unknown current high bid) in state ω and let $G_{M_j|B}$ denote the distribution of $M_j(\omega)$ conditional on the buyer choosing auction j and bidding $B = b$

in that auction. The distribution depends on B because, in an open ascending price auction, the buyer's bid influences the selection decisions of subsequent buyers. For example, suppose a buyer bids $b' > b$ in auction j and, as a result, the posted bid increases from r to r' . The buyers with values between r and r' are obviously not going to bid in auction j , and buyers with higher values may also decide to switch to other auctions.

Given Assumption 1, the optimization problem of a buyer with type x in state ω who chooses auction j can be expressed as

$$v_j(x, \omega) = \max_b \int_0^b (x - m) dG_{M_j|B}(m|\omega, b) + (1 - \alpha)(1 - G_{M_j|B}(b|\omega, b))V(x; \rho), \quad (2)$$

where m is the price she pays if she wins the auction and $V(x; \rho)$ is her continuation value if she loses the auction and survives, uses ρ to select an auction upon her return, and bids optimally in that auction. The first term is the expected payoff from winning the auction. The second term represents her expected payoff if she loses. The payoff from exit is normalized to zero.¹³

The following proposition characterizes the optimal bidding strategy.

Proposition 1 *The optimal bid of type x is given by*

$$\sigma(x) = x - (1 - \alpha)V(x; \rho). \quad (3)$$

Proposition 1 establishes that a buyer's optimal bidding strategy is to bid her value less the option value of losing.¹⁴ The optimal bid does not depend on the observed state, her beliefs about the behavior of her rivals in the auction, or the choice of auction. The proof of the proposition is in Appendix A. It extends the arguments of static second-price auctions to settings in which the participation of subsequent buyers is endogenous. The critical assumption here is that a buyer can bid only once in an auction.

We are interested in characterizing equilibria in which σ is strictly increasing. In such an equilibrium, it follows from Proposition 1 that there is a one-to-one relationship between a buyer's type and her bid. Consequently, her type x can equivalently be represented by her bid. We follow Backus and Lewis (2016) and refer to $b \equiv \sigma(x)$ as type x 's "pseudotype." Abusing notation slightly,

¹³If exit means not buying the good, then the value of the outside option is zero and x is a buyer's willingness-to-pay for the good. If exit involves buying the good at a fixed price (e.g., retail market), then the value of the outside option is the consumer surplus from this purchase and x needs to be interpreted as net of this surplus.

¹⁴Adachi (2016), Backus and Lewis (2016) and Bodoh-Creed, Boehnke, and Hickman (2017) have also shown that this strategy is an equilibrium in their models.

let $G_{M_{\rho(x,\omega)}|B}(\sigma(x)|\omega, \sigma(x))$ denote the probability that type x wins when she uses $\rho(x, \omega)$ to select an auction in state ω . Taking expectations, we exploit the equivalence between x and b and define

$$G_{M|B}(m|\sigma(x)) = \int G_{M_{\rho(x;\omega)}|B}(m|\omega, \sigma(x))d\Phi(\omega). \quad (4)$$

as the probability that the highest bid submitted by rivals is less than m in the auctions chosen by type x . Here B plays two roles: it accounts for the set of auctions that type x selects and the bid she submits in those auctions. This probability is easily computed from bid data. It corresponds to the frequency of the highest rival bid being less than m in the set of auctions where buyers submit a bid of b .

Integrating (2) over ω , changing the order of integration, and solving for $V(x; \rho)$ yields

$$\begin{aligned} V(x; \rho) &= \int \left(\int_0^{\sigma(x)} (x - m)dG_{M_{\rho(x;\omega)}|B}(m|\omega, \sigma(x)) + (1 - \alpha)(1 - G_{M_{\rho(x;\omega)}|B}(\sigma(x)|\omega, \sigma(x)))V(x; \rho) \right) d\Phi(\omega) \\ &= \int_0^{\sigma(x)} (x - m)dG_{M|B}(m|\sigma(x)) + (1 - \alpha)(1 - G_{M|B}(\sigma(x)|\sigma(x)))V(x; \rho) \\ &= \frac{\int_0^{\sigma(x)} (x - m)dG_{M|B}(p|\sigma(x))}{[1 - (1 - \alpha)(1 - G_{M|B}(\sigma(x)|\sigma(x)))]} \end{aligned} \quad (5)$$

The numerator is the expected surplus of a buyer of type x in the set of auctions that she selects with positive probability. The denominator is the proportionality factor that accounts for the possibility that she can lose and return many times.

Equations (4) and (5) provide a partial characterization of a monotone equilibrium. Given a selection rule ρ , a monotone equilibrium consists of a pair $(\sigma^*, G_{M|B}^*)$ such that (i) given $G_{M|B}^*$, $\sigma^*(x)$ is strictly increasing and satisfies equation (5) and (ii) $G_{M|B}^*$ is the stationary distribution generated by (ρ, σ^*) . Condition (i) requires buyers to bid optimally given their beliefs about their continuation value, and condition (ii) requires those beliefs to be consistent with the distribution generated by their behavior. We demonstrate in the next section that if the data are generated by a monotone equilibrium, then these two conditions are sufficient to identify the distribution of private values, F_E . In other words, it is not necessary to solve the dynamic model to identify and estimate F_E . An advantage of this approach is that our estimate of F_E is based on the selection rule (ρ) that buyers actually use. The cost of not solving for a selection rule or specifying one is that we cannot prove existence of a monotone equilibrium. We examine this issue next.

3.3 Selective vs Random Matching

In our model, the sample of auctions in which a buyer bids is a selected sample that depends on her type and the observed state. Selective matching can induce dependence between a bidder's pseudotype B and her highest rival pseudotype M . By contrast, the theoretical and empirical literature on dynamic trade make modeling assumptions under which the matching of buyers to sellers is random. In this simpler case, B and M are independently distributed, and we can show that a monotone equilibrium exists.

One way for matching to be random in our model is to assume that the auctions are closed: the platform posts auctions' closing times but not their start prices or highest losing bids.¹⁵ A buyer arrives, selects and bids in an auction, and then waits until the auction closes to learn whether she has won or lost, in which case she either exits or enters the losers' pool. In this information environment, there is a symmetric equilibrium in which buyers always choose to bid in the soonest-to-close auction. The argument is straightforward: a buyer cannot gain by deviating and bidding in a later auction since, given the soonest-to-close strategy of rivals, the level of competition in later auctions is the same as in the current auction. The random arrival times of buyers then generates a random assignment of buyers to sellers.¹⁶

In this equilibrium, a buyer's decision has no impact on the decisions of rival buyers, and every auction is ex ante identical. Let G_M denote the marginal distribution of the highest rival bid in the auction. Applying the above analysis, the buyer's continuation value is

$$V(x) = \frac{\int_0^{\sigma(x)} (x - m) dG_M(m)}{[1 - (1 - \alpha)(1 - G_M(\sigma(x)))]}. \quad (6)$$

where

$$\sigma(x) = x - (1 - \alpha)V(x). \quad (7)$$

Equations (6) and (7) and the soonest-to-close selection rule fully characterize the equilibrium. Note that, in this setting, Assumption 1 is not required.

¹⁵In a later section, we use this model as a benchmark for evaluating the effect of the selection rule generating the data on efficiency and revenues.

¹⁶This situation can also arise in open auctions if buyers wait until the last minute to submit their bids. In that case, posted prices are largely uninformative and the set of auctions in which buyer type b bids is essentially a random sample. Several papers (e.g., Ockenfels and Roth (2006), Bajari and Hortascu (2003)) have argued for this model of eBay auctions.

A revealed preference argument implies that V is increasing: a type x' can always earn a higher surplus than a lower type x by bidding like x . However, to establish that σ is increasing, we need to show that $(1 - \alpha)V'(x)$ is bounded from above by one.

Proposition 2 *Suppose the platform posts the closing times of the auctions but not their reserve prices or losing bids. Then a unique symmetric equilibrium exists and it is monotone increasing.*

Proof: Differentiating V with respect to x yields

$$\begin{aligned} V'(x) &= G_P((x)) + [(x - \sigma(x))g_M(\sigma(x))\sigma'(x) - (1 - \alpha)g_M(\sigma(x))V(x)\sigma'(x) + (1 - \alpha)(1 - G_M(\sigma(x)))V'(x)] \\ &= \frac{G_M(\sigma(x))}{[\alpha + (1 - \alpha)G_M(\sigma(x))]} \end{aligned} \quad (8)$$

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

$$\sigma'(x) = 1 - (1 - \alpha)V'(x) > 0 \quad (9)$$

$$\Leftrightarrow 1 - \frac{(1 - \alpha)G_M(\sigma(x))}{[\alpha + (1 - \alpha)G_M(\sigma(x))]} \quad (10)$$

$$\Leftrightarrow \alpha > 0. \quad (11)$$

Let H denote the distribution of the highest rival type. By definition, $H(x) = G_M(\sigma(x))$. Integrating $V'(x)$ and imposing the boundary condition $V(0) = 0$ yields

$$\begin{aligned} V(x) &= \int_0^x \frac{G_M(\sigma(s))ds}{[\alpha + (1 - \alpha)G_M(\sigma(s))]} \\ &= \int_0^x \frac{H(s)ds}{[\alpha + (1 - \alpha)H(s)]}. \end{aligned}$$

Consequently, equations (6) and (7) can be solved recursively. Derive $H(x)$ from the primitives,¹⁷ compute $V(x)$ as above, and then substitute it into equation (7) to obtain the equilibrium bid function. Q.E.D.

¹⁷For example, assuming $R = 0$ and using the Poisson approximation for returning buyers,

$$H(x) = \sum_{n=0}^{\infty} \sum_{m=0}^{\infty} \frac{e^{-\lambda} (\lambda F_E(x))^n}{n!} \frac{e^{-\gamma} (\gamma F_L(x))^m}{m!}$$

The intuition here is that, when matching is random, a small increase in type from x to x' has two effects on the continuation value: the direct effect due to x' earning a higher expected surplus than x when she bids $b = \sigma(x)$, and the indirect effect of x' optimally changing her bid from b . The latter effect is zero if type x is bidding optimally. Therefore, under random matching, V' consists only of the direct effect. This intuition also makes clear why existence can be a problem with selective matching. In this case, there is a third effect that needs to be taken into account: x' may select a different set of auctions than x or affect the decisions of subsequent buyers differently than x . This selection effect on V could cause σ to be non-monotone, especially when α is small.

We close this section by observing that the existence argument of Proposition 2 extends to our model when the selection rule does not induce any dependence between B and M . Though beyond the scope of this paper, an interesting theoretical challenge is to characterize the class of matching rules that do not generate this dependence. When running counterfactuals, we experimented with a number of simple rules (e.g., bid in the soonest-to-close auction in which the posted price is less than the optimal bid) that generated different distributions of M but little or no correlation between M and B .

4 Identification and Estimation

In this section, we outline a strategy for identifying and estimating the primitives of the model, and discuss tests of the model. These tests are especially important in our case given the assumptions we make.

Our data on buyers consist of their identities, their bids (including winning bids), the times at which the bids are submitted, and the auctions in which the bids are submitted. Buyer identities are crucial because they allow us to distinguish between new and returning buyers and to observe who exited. We assume that every buyer who arrives at the platform submits a bid in some auction. The justification for this assumption is that, in our application, there is always an auction available that has not yet received any bids and has a zero start price. The task is to show that the unobserved primitives of the model $(\alpha, \beta, \lambda, F_E)$ can be identified from these data.

We use equation (5) 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*

$$\begin{aligned}\eta(b) &= b + \frac{(1-\alpha)}{\alpha} \int_0^b (b-m) dG_{M|B}(m|b) \\ &= b + \frac{(1-\alpha)}{\alpha} G_{M|B}(b|b) [b - E(M|M < b, b)]\end{aligned}\tag{12}$$

Proof: Use equation (5) to substitute for $V(x; \rho)$ in equation (3). 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 private values of bidders 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, first-price auctions to a dynamic environment in which objects are identical and bidder values are perfectly persistent.

Proposition 4 *The model primitives, α, β, λ , and F_E , are identified.*

Proposition 4 is the main result of this section. Identification of the entry, return, and exit parameters is straightforward given the data on bidder identities and participation: $\hat{\lambda}$ is the average number of new buyers arriving per period, $\hat{\beta}$ is the mean return time of a loser who does not exit, and $\hat{\alpha}$ is the fraction of losing buyers who exit. A non-parametric estimate of $G_{M|B}(b|b)$ can be obtained by computing the fraction of times that b is the winning bid in auctions bid by pseudotype b . Similarly, a non-parametric estimate of $E[M|M < b, b]$ is the average price the pseudotype b pays when she wins. Given these estimates, we evaluate equation (12) on the sample of bids of new buyers, and use the inverted bids to obtain a non-parametric estimate of F_E . We can also use equation (12) to derive estimates of the private values of returning buyers, and use the inverted bids to obtain a non-parametric estimate of F_L , the stationary distribution of values in the losers' pool.

The remarkable aspect of Proposition 4 is that F_E is identified without solving for the equilibrium selection rule. This result is due to the invariance property of the bid function. This property allows the econometrician to use each buyer's bid to directly infer her type, effectively conditioning on the set of auctions she chooses in the data. However, this convenience comes at a cost: the econometrician needs to observe all bids, not just the prices, in order to estimate $G_{M|b}(b|b)$. The econometrician also needs to assume that all the bids are realizations of pseudotypes. This is a strong assumption for ascending auctions like eBay, since in practice some buyers pursue bidding strategies in which they appear to submit bids that are less than their true pseudotypes. In our empirical work, we try to deal with this issue by focusing only on the maximum bid submitted by a

buyer in an auction and interpreting it as her pseudotype. Nevertheless, bid censoring may still be a problem. Below we describe tests of the model that can provide some indication of the severity of this issue.

Our model generates several testable implications. First, buyers who lose and return should bid the same amount. The data on bidder identities allows us to track the bids of buyers who lose and return and to directly test whether a buyer's maximum bid is the same across auctions. Bid censoring is one reason why this prediction can fail. Second, $\sigma(x)$ needs to be strictly increasing. Since $\sigma(x)$ is increasing if and only if $\eta(b)$ is increasing, we can test this prediction from bid data using equation (12). Third, in a stationary equilibrium, the number of buyers flowing out of the loser's pool must on average be equal to the flow entering the pool. This condition implies that the expected number of returning buyers in any period is

$$\beta\bar{k} = \frac{(1 - \alpha)(\lambda - 1)}{\alpha}. \quad (13)$$

We test this condition below using the data on sale rates and bidder identities. Fourth, in a stationary equilibrium, the flow of x types out of the pool of losers must equal the flow of x types entering the pool. On average, the flow of x types that leave the pool during a period of length Δ is

$$\beta\Delta\bar{k}f_L(x), \quad (14)$$

and the flow back into the pool over this period is on average

$$(1 - \alpha)[1 - G_{M|B}(\sigma(x)|\sigma(b))][\beta\Delta\bar{k}f_L(x) + \lambda\Delta f_E(x)]. \quad (15)$$

Equating these two flows yields

$$f_L(x) = \frac{\lambda\alpha(1 - G_{M|B}(\sigma(x)|\sigma(x)))}{(\lambda - 1)[1 - (1 - \alpha)(1 - G_{M|B}(\sigma(x)|\sigma(x)))]} f_E(x). \quad (16)$$

Equation (16) expresses the density of values in the losers' pool in terms of the density of values of new buyers. The proportionality factor reflects the natural censoring that results from the auction outcomes: buyers with high values are more likely to win than buyers with low values. As a result, the density f_L is greater than f_E at low values and smaller at high values. Given estimates of the

parameters α and λ , along with estimates of the distributions F_E , F_L , and $G_{M|B}$, we can test this implication of stationarity. It is highly informative because it involves a continuum of constraints.

As a final note before we turn to the empirical application, the previous analysis assumes the buyer’s exit rate does not depend on her type. We provide empirical support for this assumption in the data section below, but in Appendix B we also show that the model can be extended to allow for endogenous exit.

5 Data

Our primary data consist of all eBay listings for iPads posted between February-September 2013, 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 bids submitted by winning bidders, which are important for estimating G , the distribution of the maximum rival bid.

We focus on the used market for a specific model: the 16GB WiFi-only iPad Mini. Since there is some substitution between models (e.g. 16GB vs. 32GB) and between new vs. used items, one might be concerned this definition of the market is too narrow. Substitution is indeed evident in the bidding data: 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. However, among bidders who lost an auction for a 16GB WiFi model, 83% of returning bidders chose to bid again on the same model. Among those who switched to bidding on a different model, most either bid on the 32GB WiFi version (8%) or on the 16GB WiFi+4G version (5%). Similarly, most buyers did 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. Thus, while there is obviously some substitutability between models and item conditions, we believe it is a reasonable approximation to treat the used 16GB WiFi market as its own separate market.

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 to adjust for item characteristics like color, added extras, and seller feedback ratings—an approach that implicitly assumes these are characteristics that are valued uniformly across buyers (for example, all buyers

have the same willingness to pay for an extra charger). We doubt this assumption would hold with respect to item condition: some buyers probably care a lot more than others about whether the item is new vs. used.

Our model assumes that buyers have unit demands. For iPads it seems reasonable that most buyers would be interested in buying only one unit. However, a small fraction (less than 6%) of buyers bought two or more units during the sample period. In the analysis below, we treat these buyers as new bidders in the first auction they bid in, and as returning bidders in all subsequent auctions, even if they had previously won an auction.

When a seller posts an item for auction on eBay, she chooses the starting price of the auction. This starting price serves as a public reserve price, since the system only accepts bids above the starting price. The seller also has the option of setting a secret reserve for a small additional fee, but this option is rarely used—in our data only 10% of listed items had secret reserve prices. Many sellers choose effective reserve prices that are clearly intended to be non-binding: 22% of listed items had effective reserve prices below \$1, and 41% had reserve prices below \$180, which is the first percentile of the distribution of final sale prices. Sellers also have the option to create a fixed price listing, in which case the price is fixed and the listing remains active on the site for up to 30 days until the item is sold. For the specific product we are studying, auctions are the most common form of sale: 65% of successfully sold items were sold by auction. In the analyses below we focus on auction listings only.

Table 1 shows summary statistics for the 5,622 auction listings in our sample. The majority of these listings ended successfully with a sale,¹⁸ and the average sale price (conditional on sale) was \$288.86 with an average shipping fee of \$7.33. The retail price for a new unit of this particular model was \$329, not including tax and shipping, so the used units on eBay were selling at an average discount of at least 10% relative to the new retail price. The average number of bidders per auction is 9.27, 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.

One simplification in our theoretical model is that the interval between seller arrivals (or equivalently, auction closings) is constant. We then measured the arrival rate of buyers relative to this period by normalizing it to one. In reality, the arrival rates of sellers and buyers differ by time of

¹⁸eBay requested that we not report the exact conversion rate, but it is higher than 85%.

Table 1: Summary statistics for auction listings ($N=5,622$)

	Mean	Std. Dev.	Percentiles		
			0.10	0.50	0.90
Start price	141.89	119.48	0.99	150.00	295.95
Positive reserve price (0/1)	0.09	0.28	0.00	0.00	0.00
Reserve price (if positive)	274.53	42.75	220.00	280.00	325.00
Sale price (if sold)	288.86	31.32	255.60	290.00	325.00
Shipping fee	7.33	5.62	0.00	6.60	15.00
Number of bids	21.12	17.73	0.00	18.00	46.00
Number of unique bidders	9.27	6.40	0.00	9.00	18.00
Minutes since last auction	61.95	103.19	3.82	28.82	137.52
Cover included (0/1)	0.19	0.39	0.00	0.00	1.00
Seller feedback (#)	6,781.88	42,526.86	12.00	124.00	3,032.00
Seller feedback (% positive)	99.02	5.55	98.28	100.00	100.00

day in predictable ways, so one possible concern is that these differences influence bidding behavior. However, we find that even though the number of auctions that close varies substantially by time of day, the number of bidders per auction closing is approximately the same, as shown in Table 2. The implication is that sellers' and buyers' arrival rates vary proportionally by time of day, so the assumption of constant arrival rates is a harmless normalization. This also means that thinking of time in terms of auction closures (i.e., one period equals one auction closing) is approximately correct. We therefore estimate λ , the arrival rate of new buyers, as the average number of new buyers per auction closing, which is 5.47.

Table 2: Bidders per auction closing, by time of day

Time block	Mean	Std. Dev.	Percentiles		
			0.10	0.50	0.90
00:00-06:00	8.31	6.13	0	8	18
06:00-12:00	9.23	6.40	0	9	18
12:00-18:00	9.40	6.43	1	9	18
18:00-24:00	9.29	6.41	1	9	18

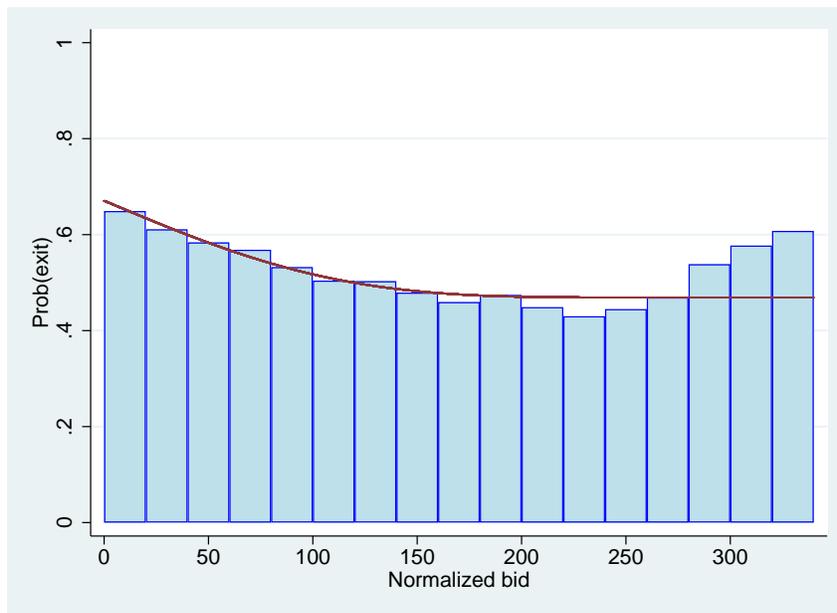
Conditional on losing an auction, 49.8% of bidders come back to bid again in a subsequent auction. Our estimate of the exit rate, α , is thus 0.502.¹⁹ Return times are fairly short: conditional on returning to bid again, 21% of bidders return within an hour, and 10% return within 5 minutes. The full distribution of return times is highly skewed, however, since there is a long right tail

¹⁹We say a bidder returned if she comes back to bid again within three weeks. Changing the time horizon, e.g. to two weeks or four weeks, has little impact on our estimate of α , since most bidders return relatively quickly if they are going to return at all.

reflecting bidders who take 24 hours or more to come back.

Although our model can be extended to allow for endogenous exit, as shown in Appendix B, our baseline model assumes exit is independent of the bidder’s type. Since we observe bidders’ bids and also whether they exit, we can estimate an exit function $\alpha(b)$ to see if exit rates appear to depend on bidders’ types. Figure 1 shows binned exit frequencies along with a semi-nonparametric estimate of $\alpha(b)$. 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 these bidders are not especially relevant in the model. Adjusting their value functions to reflect a higher exit rate is unimportant: their value functions are essentially zero anyway, since they have virtually no chance of winning an auction. There is a slight uptick in exit rates for high-value bidders, perhaps because these bidders elect to purchase at retail when they lose an auction, and adjusting for this difference might be more important. However, the differences in exit rates are small, and allowing for endogenous exit makes computing counterfactuals meaningfully more difficult, so we use the inverse bid function from the simpler model with exogenous exit when we estimate the distributions of bidders’ valuations below.²⁰

Figure 1: Exit rate as a function of bid



Items in our data rarely fail to sell, but in that event sellers also have the option to come back

²⁰Since 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.

and try again. Unsold items can be re-listed, typically without having to pay any additional fees to eBay. Among the items in our data that failed to sell, 63 percent appear to have been relisted, based on subsequent appearance of an item with the same seller ID and the exact same product title. Because sellers in our data typically set low reserve prices and the majority of auctions end successfully with a sale, in what follows we focus on dynamics among bidders and largely ignore the seller dynamics.

5.1 Tests of model assumptions

Our model abstracts away from intra-auction dynamics, since buyers are assumed to bid when they arrive, and bid exactly once in whichever auction they choose. Of the various simplifying assumptions we make, this one is perhaps the most at odds with the data, since in reality “incremental bidding” (submitting multiple, increasing bids within a single auction) is relatively common. Roughly 44% of the bidders in our data submit multiple bids for the same item, but most of the incrementing happens before the auction nears its closing time: only 7% of bidders submit multiple bids in the last hour before the auction closes. The incremental bidding in the data could reflect within-auction strategic behavior: some bidders may be trying to learn about their rivals through incremental bidding, or even trying to influence the bidding decisions of subsequent bidders. Nevertheless, since incorporating these considerations would complicate the model considerably, and our goal is to keep the model as simple as possible, we estimate the model as though bidders submit only one bid, which we take to be the highest bid they submitted in the auction. More generally, the presence of incremental bidding raises the important question of which bids to take seriously when estimating the model. We return to this question in Section 6 below.

Even without incremental bidding, buyers in the the real-world marketplace might arrive, observe the bidding in several auctions of interest, and then make a strategic choice about *when* to submit their bids. While we cannot test for this directly, since we don’t observe users’ browsing behavior prior to their bid submissions, we can at least check for irregular bunching in the timing of bids. Contrary to what other studies using eBay data have found, we observe relatively little last-minute bidding in our data. Less than five percent of bids were submitted within five minutes of the auction’s close, and 58 percent of auctions were won by buyers who submitted their bids with more than an hour remaining in the auction. More directly, our model implies that the time between bids (across all auctions and bidders) should be exponentially distributed, and this appears to be approximately true in the data, as shown in Figure 7 in Appendix C. There is slightly more density near zero than would be consistent with an exponential distribution, but the difference is small.

One of the clearest implications of our model is that buyers’ bidding strategies are stationary: if a buyer loses an auction and returns to bid again in a subsequent auction, we expect her to submit the

same bid. This is approximately true in the data. Looking at bidders’ bids in successive auctions, there is a statistically significant upward trend, but it is small. That is, losing bidders tend to bid more aggressively when they return, but the increase in the bid is only 35 cents on average. Regressing bids on bidder fixed effects and the number of previous auctions lost, the seller fixed effects explain 87% of the variance in bids.

Another important assumption we make in our model is that bidders believe the current state of the market is uninformative about the future state—i.e., when a bidder loses she believes that when she returns the vector of posted bids will be a random draw from its stationary distribution. Over relatively short time horizons, we do observe some persistence in the vector of posted bids, so our assumption would be problematic if losing bidders typically returned very quickly to bid again. But most bidders take longer to come back: the median number of auctions that finish in the time it takes a losing bidder to re-arrive is six, and there is virtually no correlation between the posted bids of the current set of soon-to-close auctions and the posted bids six auctions later.²¹ Therefore it appears reasonable in our case to assume that bidders’ beliefs about the future state are not influenced by the current state.

6 Estimation

The primary objective of our empirical analysis is to recover F_E , the distribution of buyers’ valuations. Since 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, we can estimate F_E using the bids of new bidders. Monotonicity of the bid function $\sigma(x)$ (which we discuss below) means we can treat a bidder’s bid as her pseudotype, and recover her true type with the inverse bid function given by equation (12). This inversion requires estimates of the exit rate, α ; the probability of winning, $G_{M|B}(b|b)$; and the expected price conditional on winning, $E(M|M < b, b)$. Below we explain how we obtain these estimates from the data.

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 prices on item characteristics, Z , and then use the estimated coefficients $\hat{\gamma}$ 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 items with identical observed characteristics. 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; and month dummies (to control for a gradual downward trend

²¹For example, let r_t^o be the posted bid of the t^{th} auction when it becomes the o^{th} next to close. The correlation between r_t^1 and r_{t+6}^1 is -0.002.

in prices over time). In all that follows, when we refer to bids we mean normalized bids.

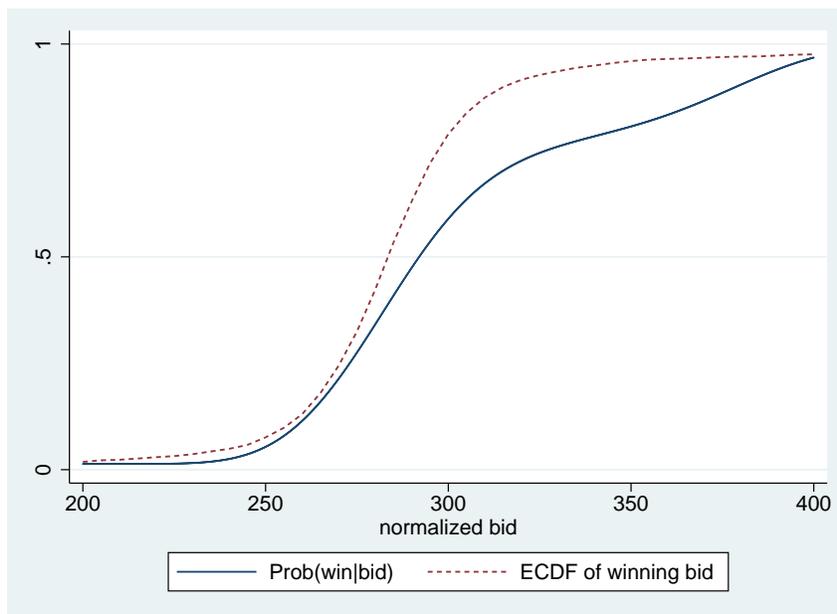
Estimating $G_{M|B}(b|b)$ is relatively straightforward, since it is simply the probability of winning at a bid equal to b . One could estimate this function by simply running a probit or logit regression of a win dummy on bids. To avoid the functional-form restrictions such an approach would impose, we instead use the semi-nonparametric maximum likelihood method of [Gallant and Nychka \(1987\)](#), approximating the latent density with a 6th-order Hermite polynomial. The estimated function is plotted in [Figure 2](#).

In outlining the bidding model, we were careful to note the possible dependence of G , the distribution of the maximum rival bid, on the buyer’s type. In the absence of any such dependence—i.e., if bidders assign themselves to auctions randomly—a natural estimate of $G_{M|B}(b|b)$ when the arrival processes for new and returns buyers are Poisson would be the empirical CDF of the winning bid across all auctions. [Figure 2](#) plots this empirical CDF alongside our semi-nonparametric estimate of $G_{M|B}(b|b)$ that accounts for bidder sorting, and shows that the two are meaningfully different. In particular, it is clear from these estimates that high-value bidders tend to sort into auctions where their probability of winning is *lower* than it would be if they chose an auction at random. Explaining this apparent puzzle is not the purpose of our analysis, but we return to discuss it in [Section 7](#) below. Even without knowing what causes the divergence between $G_{M|B}(b|b)$ and its unconditional counterpart, the figure highlights the role that auction selection plays in our analysis. On the one hand, our assumptions imply that bidders have stationary bidding strategies: regardless of the state or the auction they choose, they submit the same bid. On the other hand, in order to correctly apply the inverse bid function, we need to compute the correct $G_{M|B}(b|b)$, which must be the probability of winning conditional on the auction selection mechanism. Fortunately, we are able to do this, since computing it from the observed data effectively conditions on buyers’ actual auction selection rule.

The last component of the inverse bid function is the conditional price expectation $E(M|M < b, b)$. We estimate this by constructing a dataset of winning bids and the prices (second-highest bids) associated with those winning bids, and then running a local polynomial regression of the latter on the former. Note that by estimating the expected price conditional on winning with a bid equal to b , we are again implicitly accounting for the dependence of G on b .

As noted above, buyers’ strategic selection of which auctions to enter could in principle cause the bid function to be non-monotonic. With estimates of α and the functions $G_{M|B}(b|b)$ and $E(M|M < b, b)$, we can compute the bid function and directly check monotonicity. [Figure 3](#) shows the estimated bid function, which is indeed monotonic. This means we can invert the observed bids and estimate the distribution of bidders’ underlying valuations using a dynamic analogue to

Figure 2: Estimated probability of winning



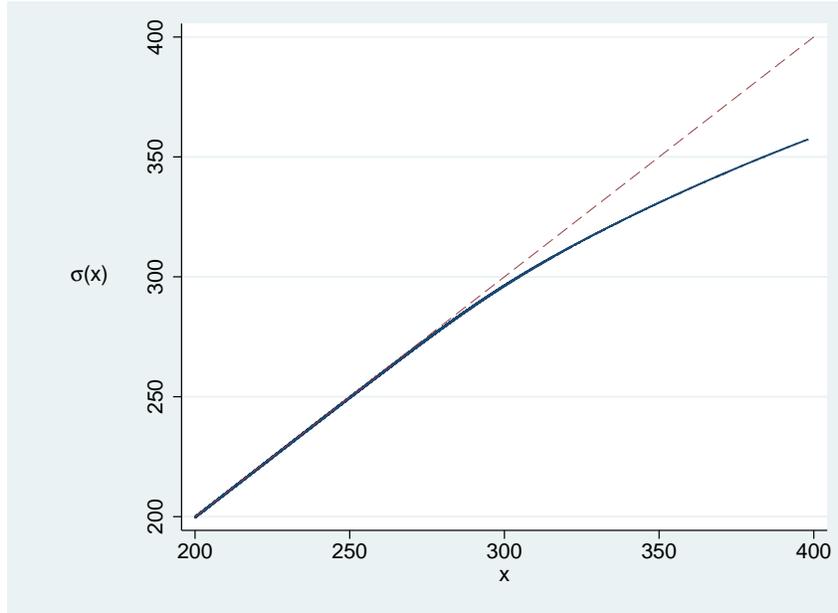
the method proposed by Guerre, Perrigne, and Vuong (2000). Applying our inverse bid function to the observed bids, we recover a set of pseudo-values; we then estimate the distributions of these pseudo-values nonparametrically with a kernel density estimator.

An important question in implementing this approach is which bids to include in the estimation. As noted above, incremental bidding is fairly common in our data, and this makes it difficult to draw inferences about the true intended bids of losing bidders. For instance, a bidder whose maximum intended bid is \$150 might initially bid \$50, but then lose when another bidder submits a bid of \$200. This bidder’s observed bid would then lead to a large underestimate of her true valuation. Since this censoring problem is most severe at low bids (because bid increments tend to be larger when the posted bid is low), and because incremental bidding appears to be most common among low-value bidders, we address this problem by simply excluding bids below \$150 when estimating f_E and f_L . Since \$150 is well below the lowest winning bid we observe in the data, the logic is that such bids were not serious bids: either they were initial bids submitted by incremental bidders, or they were submitted by bidders whose valuations were too low to have any chance of ever winning an auction.

Figure 4 shows kernel density estimates of f_E and f_L .²² The difference between the estimated densities is consistent with the model: the distribution of returning losers’ valuations looks like a

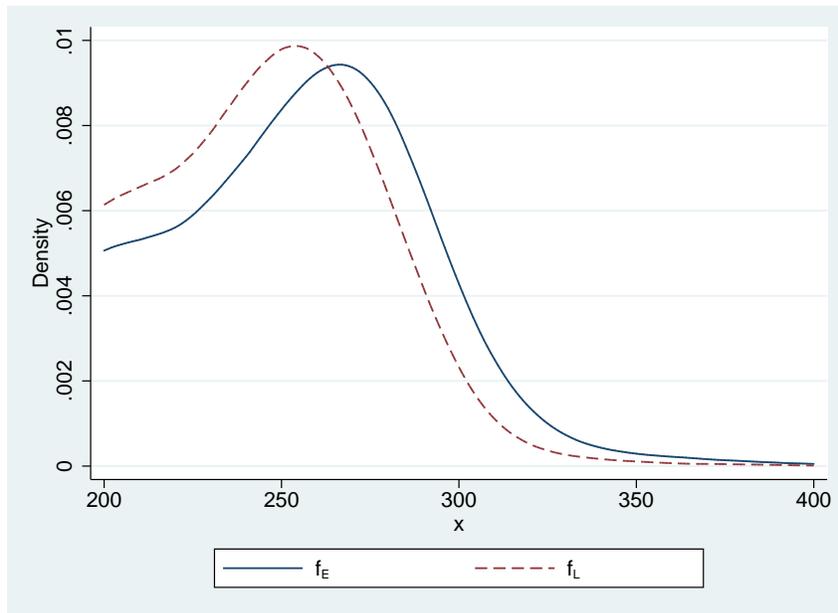
²²Even though there is positive density on very low valuations, we plot the estimates for values above \$200, since low-value bidders have virtually zero probability of winning and are essentially irrelevant.

Figure 3: Estimated bid function



resampling of new bidders' valuations, with less density in the upper tail. It is important to note that this difference is in no way imposed by our estimation procedure: since we can distinguish between new and returning bidders in the data, we simply estimate separate distributions for the two groups.

Figure 4: Estimated distributions of valuations, using all bids



Before moving on to tests of the model and counterfactual analyses, we note that our estimates imply that 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. A static model of bidding would especially mis-estimate the upper tail of the distribution of bidder values. For our sample, we estimate that the winning bidder’s true value (x) is on average roughly \$7.56 higher than the bid she submitted, and in some cases over \$25 higher.

Tests of over-identifying restrictions

Our model implies specific relationships between arrival rates of new and returning bidders as well as the distributions of the two groups’ values. Because we observe in the data whether a bidder is new or returning, we can estimate these arrival rates and distributions separately for each group. That is, we do not need to impose the restrictions implied by the model; we can instead treat them as testable implications.

Equation (13) implies that the number of returning buyers per auction should be given by

$$\beta \bar{k} = \frac{(1 - \alpha)(\lambda - q)}{\alpha}$$

where q represents the probability that an auction ends with a sale.²³ The size of the loser pool, \bar{k} , is not observable; but our estimates from the data of the exit rate α (0.50), the arrival rate of new bidders λ (5.47), and the probability of sale q predict an average of 4.58 returning bidders per auction, which is not far from the average of 4.86 we observe in the data.

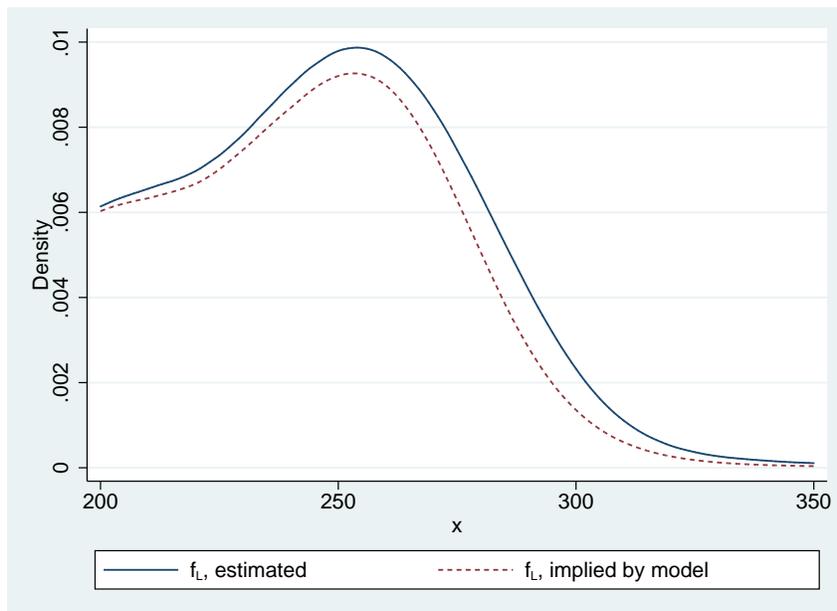
Equation (16) describes a more stringent test of the model’s underlying stationarity assumption: not only should the numbers of bidders flowing into and out of the loser pool be equal on average, but the flows should be equal at every type x . This puts a restriction on the relationship between the densities f_E and f_L :

²³For simplicity, section 3 develops the model under the assumption that $q = 1$, which means that for each auction there is one bidder who exits because she won. To fairly test the model’s implications, we need to modify equation (13) to account for the true probability of sale in the data.

$$f_L(x) = \frac{\lambda\alpha(1 - G_{M|B}(\sigma(x)|\sigma(x)))}{(\lambda - q) [1 - (1 - \alpha)(1 - G_{M|B}(\sigma(x)|\sigma(x)))]} f_E(x) \quad (17)$$

Figure 5 shows a comparison between the f_L we estimate directly from the data and the f_L implied by the model (as a function of the estimated f_E). The two densities are clearly not identical, but they are remarkably similar given that nothing in the test forces them to look the same. In principle, the rescaling of f_E in equation (17) could distort the shape of the resulting f_L and even cause it to not integrate to one. The test should fail if the model is simply incorrect, or if the estimates of λ , α , and/or $G_{M|B}$ are inaccurate or invalid. Indeed, if we run the same test using the empirical CDF of winning bids as our estimate of G , the results, shown in Figure 9 in Appendix C, are noticeably worse.

Figure 5: Test of restriction on f_L



Taken together, we view the results of the various tests in this section as reassuring evidence that the simplifying assumptions of our model are reasonably consistent with the true data-generating process.

7 Counterfactual analyses

The theoretical models developed by [Satterthwaite and Shneyerov \(2007, 2008\)](#) predict that a decentralized, dynamic market converges to the Walrasian equilibrium in the limit as the market dynamically thickens—i.e., as the period length shrinks to zero so that traders have infinitely many

opportunities to trade. The natural question to ask about a real-world decentralized, dynamic market like eBay is how close its stationary state comes to delivering the Walrasian equilibrium. In this market, the dynamics result from losing bidders’ ability to return and bid again in later auctions, which has two effects. First, it increases the level of competition in each auction because the number of bidders includes both new and returning bidders. We refer to this as the *dynamic participation effect*. Second, buyers bid less. Anticipating the possibility of returning to a later auction, they shade their bids to reflect the option value of losing. We refer to this as the *dynamic bidding effect*.

These two effects are related but distinct. The dynamic bidding effect requires buyers to be forward-looking. If buyers bid myopically—i.e., in a way that ignores the option value of losing—then there is no dynamic bidding effect, but the dynamic participation effect is still present. Also, the dynamic participation effect clearly enhances efficiency: it lowers the fraction of inefficient trades, since high-value bidders who lose can still win an item, and their ability to return also makes it more difficult for low-value buyers to win. By contrast, the dynamic bidding effect has no impact on efficiency. If all buyers are forward-looking, they shade their bids in a way that amounts to a monotone transformation of their types. Thus, the dynamic bidding effect only affects transfers between the buyer and seller (i.e., prices).

To evaluate the extent to which dynamics in the eBay marketplace yield convergence toward the efficient market outcome—and to measure the relative importance of the two distinct effects described above—we begin by using our estimates to calculate the market-clearing price that would prevail if the units in our data were sold in a uniform price auction. This is the price that would clear the market if eBay were to pool all buyers and pool all sellers and conduct a single uniform auction. Specifically, we calculate the total number of sellers $N_s = 5,002$ and total number of buyers $N_b = 27,380$ in our data, and then compute the market-clearing price as the $\left(1 - \frac{N_s}{N_b}\right) = 81.7^{th}$ percentile of the estimated distribution F_E . Since this is the competitive equilibrium price and allocation, it serves as the main benchmark against which to compare the prices and efficiency of other mechanisms.

The opposite benchmark is a fully decentralized static mechanism, in which the N_s sellers hold separate simultaneous second-price auctions, and the N_b buyers are randomly allocated to those auctions, each buyer getting only one chance to win an auction.²⁴ This counterfactual serves as the inefficient benchmark: it tells us what the price distribution would be and how inefficient the allocation would be in the absence of any dynamic effects. We simulate outcomes under this

²⁴Even if sellers post reserve prices, [Peters and Severinov \(1997\)](#) prove in this setting that if the number of buyers and sellers is sufficiently large, there is an equilibrium in which sellers choose second-price auctions with zero reserve prices. The buyers would thus be indifferent between the auctions and would choose one at random.

benchmark by taking N_b buyers, with valuations drawn randomly from our estimated F_E , and randomly assigning them to N_s auctions, repeating the process 10,000 times in order to minimize any noise introduced by the simulation draws.

Table 3 shows prices and efficiency measures for the actual bidding we observe in the data compared to the two counterfactual benchmarks. We calculate the market-clearing price to be \$279.45. In the market-clearing (efficient) equilibrium, all buyers with valuations above the market-clearing price successfully purchase, and the average gross surplus of these buyers is \$307.73. At the other extreme, under simultaneous auctions with no dynamics, the price distribution exhibits considerable dispersion, and only 31% of the buyers who should get the item (i.e., buyers with valuations above the market-clearing price) actually do. The outcome we observe in the data is naturally in between these two extremes. It falls well short of complete convergence to the competitive equilibrium: price dispersion is still substantial, and we calculate that only 59% of the highest-value buyers successfully win an auction.

Table 3: Prices and efficiency compared to counterfactual benchmarks

	Simultaneous auctions, static bidding	Sequential auctions, dynamic bidding (i.e., data)	Market clearing
Average price	231.22	275.39	279.45
SD of prices	70.88	26.85	0.00
Average gross surplus	283.39	293.84	307.73
Prob(win $x > p^*$)	.305	.594	1.000

Notes: Average gross surplus is the average valuation (x) of the winning bidders. Prob(win | $x > p^*$) is the probability that a buyer whose x is greater than the market-clearing price wins an auction before exiting.

Since the dynamic bidding effect has no impact on efficiency, all of the efficiency differences between the data and the inefficient benchmark (simultaneous auctions with no dynamics) can be attributed to the dynamic participation effect. Interestingly, virtually all of the price differences also result from the dynamic participation effect. The dynamic bidding effect implies that the bids we observe in the data, and resulting prices, are shaded down. For the highest-value buyers, the difference between the bid and the true value can be substantial—but for the buyers who end up setting the prices, the differences are apparently small. If we calculate the prices we would have observed if buyers had bid their values directly, we find that the average price would have increased by only \$0.80, and the standard deviation would only have increased by \$0.65.

One potential explanation for the relative inefficiency of the outcome we observe in the data is

the buyers’ auction selection rules. The observed selection patterns are broadly consistent with buyers trying to arbitrage differences in expected payoffs across auctions. Since rational buyers do not bid in auctions with posted bids that exceed their optimal bid, lower value buyers should and do choose auctions with lower posted bids. The invariance property of the optimal bid implies a monotone relationship between the posted bid and time-to-close: if two auctions have the same posted bid, then the auction with the earlier closing time should dominate since the probability of being outbid by a subsequent buyer is lower in that auction. This tradeoff suggests that buyers with high values are more likely to choose auctions with higher posted bids that close sooner, causing buyers with low values to choose auctions that close later. The data are consistent with this prediction. We find that the posted bid in an auction declines roughly exponentially with its time-to-close. The consequence is a form of assortative matching in which, at any point in time, high-value buyers are matched to sellers with early closing times and low-value buyers are matched with sellers with later closing times. However, there is evidence in the data that high-value buyers are more likely to choose soon-to-close auctions than low-value buyers and, as a result, the selection process disproportionately places high-value buyers in the same auctions. Interestingly, in simulations we found that matching rules having this feature did generate separation between G_M and $G_{M|B}$ as is shown in Figure 2.²⁵

To explore how much this matters for prices and efficiency, we simulate outcomes with an alternative auction selection rule. We simulate a sequence of 10,000 auctions,²⁶ with the number of new buyers arriving before each auction being a draw from the Poisson distribution with a mean that matches the data ($\hat{\lambda} = 5.47$). New buyers’ valuations are drawn randomly from the estimated F_E . Losing bidders exit with probability $\hat{\alpha} = 0.502$, and otherwise enter a pool of losers. After entering the loser pool, a bidder has a probability $\hat{\beta} = 0.008$ of returning to bid each period,²⁷ where each period has one auction. When a new buyer arrives, we simply assign her to the soonest-to-close auction. This selection rule is a useful benchmark because it is analogous to the inefficient benchmark from Table 3, in the sense that buyers’ random arrivals lead them to be randomly matched to auctions, except that in this case they are able to return to try again if they lose. Note that while the simulations can be conducted in type space—i.e., assignment of bidders to auctions and determination of who wins can be done based on their actual valuations—in order to compute bids

²⁵We estimated a matching rule using a discrete-choice model in which each buyer’s choice set consisted of the thirty soonest-to-close auctions in which the posted bid was less than her eventual bid. The results indicated that the choices of low value buyers were approximately random but the choices of high value buyers were significantly more skewed toward the soonest-to-close auctions. We explain this in more detail in Appendix C.

²⁶To get 10,000 auctions, we simulate 30,000 and then drop the first and last 10,000. We drop the first 10,000 to ensure that we are sampling from auctions in steady state; we drop the last 10,000 auctions because for late-arriving buyers we cannot observe their eventual outcomes (e.g., whether they eventually succeed in winning an auction). At the start of the simulated sequence, we seed the loser pool with $\bar{k} = \hat{\lambda}(1 - \alpha)/(\alpha\hat{\beta})$ buyers whose valuations are drawn from F_E .

²⁷We estimate β as the inverse of the mean number of auctions before a losing bidder returns in our data, since if re-arrivals are a Poisson process then return times should be exponentially distributed.

(and prices) we need to find the new equilibrium continuation value function V induced by the counterfactual. The details of how we do this are explained in Appendix D.

Table 4 shows the comparison of outcomes. The alternative auction selection rule significantly reduces price dispersion and increases efficiency relative to the outcome in the data. With buyers randomly matched to auctions, but participating dynamically, 72 percent of the highest-value buyers successfully win an auction, as opposed to the 59 percent from the data. This happens because high-value buyers are less likely to end up in auctions where they are bidding against other high-value buyers. In a sense, high-value buyers are spread more evenly across auctions, which reduces price dispersion and improves allocative efficiency.

Table 4: Prices and efficiency under alternative auction selection rule

	Actual selection rule (data)	Next-to-close selection rule (simulation)
Average price	275.39	274.02
SD of prices	26.85	16.65
Average gross surplus	293.84	301.86
Prob(win $x > p^*$)	0.594	0.717

Notes: Results in column 2 are based on a simulation in which bidders enter the next-to-close auction when they arrive. Average gross surplus is the average valuation (x) of the winning bidders. Prob(win | $x > p^*$) is the probability that a buyer whose x is greater than the market-clearing price wins an auction before exiting.

We have explored other counterfactual selection rules—for example, assigning bidders to the soonest-to-close auction in which their optimal bid exceeds the posted bid, or having bidders randomize over auctions in which their optimal bid exceeds the posted bid—and found that they result in outcomes that are quantitatively similar to those described in Table 4. The relevant aspect of these alternative rules seems to be that they are homogeneous across types, meaning that even though the buyer’s type determines the set of auctions in which she is mechanically eligible to bid, the rule chooses an auction from that set in a way that does not depend on type. This distinguishes them from the selection rule that actual bidders in the data appear to be using (see footnote 25).

Another possible explanation for the relative inefficiency of the outcome in the data is that the exit rate may be too high: losing bidders come back, but not often enough for dynamic effects to be significant. Using the same simulation procedure described above, we can examine whether the market would converge to the efficient outcome as the exit rate (α) goes to zero, and whether the dynamic bidding effect would be large if the exit rate were near zero. Note that to conduct these

counterfactuals, we must impose an auction selection rule: we cannot simply use the arrivals and re-arrivals observed in the data, because changing α fundamentally changes the re-arrival process. We use the same auction selection rule described above, assigning bidders to the soonest-to-close auction when they arrive. If we use alternative homogeneous rules, the results we describe are essentially unchanged. Table 5 shows that as α declines, average prices increase, and dispersion decreases. But even with an exit rate of 0.10, prices do not come close to complete convergence. By contrast, efficiency does come reasonably close to the market-clearing benchmark. When the exit rate is 0.10, the average gross surplus (average valuation of winning bidders) is almost as high as under market-clearing, and 85 percent of the highest-value bidders succeed in winning an auction. Note that as the exit rate gets small, the number of bidders per auction increases, because any given auction will have many returning bidders. (In the simulations with $\alpha = 0.10$, the average number of bidders per auction is 46.) This is the dynamic participation effect, and its main impact on prices is to eliminate low prices, because it makes it difficult for low-value bidders to be the price-setters. The table shows that this effect is significant: the lowest prices when $\alpha = 0.10$ are close to the median price when $\alpha = 0.50$.

Table 5: Price distributions for different exit rates

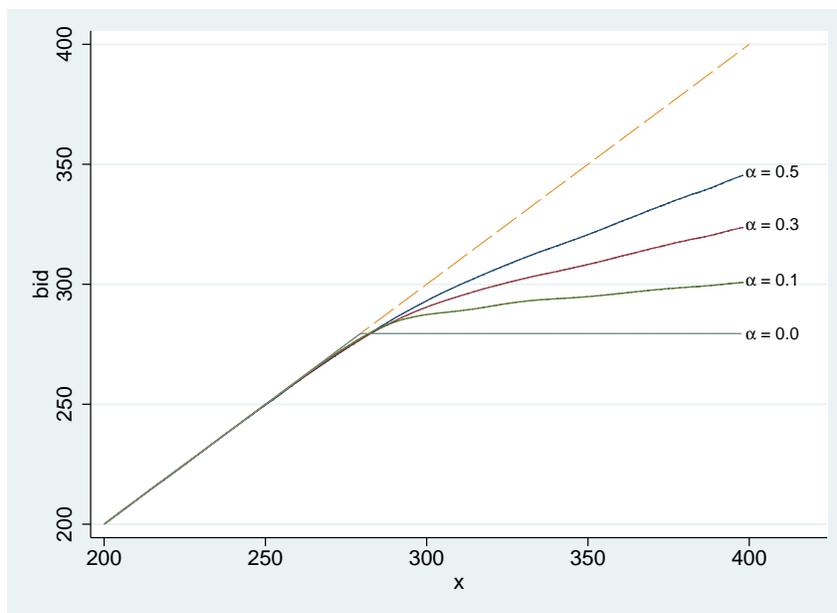
	$\alpha = 0.50$		$\alpha = 0.10$	
	Dynamic bidding	Myopic bidding	Dynamic bidding	Myopic bidding
Price percentiles:				
.01	219.73	219.77	270.89	271.41
.10	254.28	254.65	276.00	277.09
.50	275.78	277.74	282.33	284.84
.90	291.43	297.17	287.96	298.98
.99	308.16	324.58	292.46	325.35
Average gross surplus	301.86	301.86	306.55	306.55
Prob(win $x > p^*$)	0.717	0.717	0.868	0.868

Notes: These results are based on simulations in which bidders choose the soonest-to-close auction in which their valuations exceed the posted bid. Dynamic bidding means buyers shade their bids to reflect their continuation values; myopic bidding means they simply bid their values.

By contrast, the main impact of the dynamic bidding effect should be to eliminate high prices. Low-value buyers have continuation values near zero, since they are so unlikely to ever win, so for them the dynamic bidding effect is negligible. But high-value buyers' higher continuation values lead them to shade their bids toward the market-clearing price. The table shows prices that would result under myopic bidding—i.e., if buyers directly bid their values—to indicate the magnitude

of the dynamic bidding effect. When $\alpha = 0.10$, the effect is significant—for example, the highest prices are \$33 lower than they would be in its absence—but there is still substantial price dispersion above the market-clearing price. This can also be seen by looking at the bid functions, which are shown in Figure 6. In the limit as $\alpha \rightarrow 0$, buyers with values above the market-clearing price submit bids equal to that price. When $\alpha = 0.10$, we find that the highest-value buyers substantially reduce their bids, but still submit bids well above the market-clearing price. Thus, while a lower exit rate would lead the dynamic bidding effect to be meaningfully larger than what we observe in the data, it still would not deliver complete price convergence.

Figure 6: Counterfactual bid functions



Taken together, these findings suggest that in real-world markets, convergence to the Walrasian price may occur more quickly from below than from above. Even a moderate amount of dynamic participation can mostly eliminate prices below the Walrasian price, but high prices aren't entirely eliminated even when buyers anticipate having a large number of opportunities to trade. Interestingly, this means that the average price in a decentralized market may be *above* the Walrasian price, as is the case in our simulations when $\alpha = 0.10$.

8 Discussion and Conclusions

In contrast to the early literature on online auction marketplaces, recent papers have explicitly incorporated dynamics into models of bidding behavior. We view our study as making three contributions to this nascent literature. First, the model we propose is simple and empirically

tractable, while still capturing the important dynamic aspects of the bidding environment. The key simplifying assumption is that bidders believe the relevant future state will be a random draw from the stationary distribution of states. This approach is analogous to the oblivious equilibrium concept proposed by [Weintraub et al \(2008\)](#), which simplifies the analysis of dynamic games in markets with a large number of firms. Our central assumption is reasonable in thick markets like the one we study, since the large number of auctions and bidders leads to a high rate of churn in the state—so that bidders’ own actions do little to influence the state, and they can learn little about future states by observing the current state. While we believe this approach is likely to be quite useful in many markets, we caution that it is less suitable in thin markets.

The second main contribution of our analysis is to highlight the importance of accounting for buyers’ endogenous selection of which auction to bid in. We recover the primitive distribution of buyers’ valuations using a dynamic version of the technique proposed by [Guerre et al \(2000\)](#), in which inverting the bids requires an estimate of the distribution of maximum rival bids. Our results show that when estimating this distribution, it is critical to condition on the auctions in which buyers of a given type choose to bid. In other words, one cannot simply use an estimate of the unconditional distribution of maximum rival bids; it is necessary to estimate the distribution of rival bids faced by a buyer who bids a particular amount. Interestingly, while our empirical analysis shows this distinction is important, our modeling approach allows us to identify the model’s primitives without actually solving for equilibrium auction selection rules. Hence, it is not necessary to explicitly model how buyers are choosing auctions, but it *is* necessary to condition on their actual choices when estimating key quantities from the data.

The third main contribution of the paper is to show the quantitative impact of dynamic competition on prices and efficiency. Relative to an environment in which buyers can only bid once, the option to return and try again after a losing bid leads to two main effects. The dynamic participation effect comes from a mechanical increase in competition, as the presence of returning buyers inflates the number of buyers per auction. The dynamic bidding effect comes from buyers strategically shading their bids to reflect the option value of losing and potentially trying again. Our counterfactual simulations indicate that both effects are quantitatively meaningful, but that the dynamic participation effect appears to have a more substantial impact—not just on allocative efficiency, which is unaltered by the dynamic bidding effect, but also on prices. This finding is reminiscent of the famous result of [Bulow and Klemperer \(1996\)](#) that adding a bidder has more impact on revenues than changes to auction design. In our case, the mere presence and participation of returning bidders is more impactful than the strategic changes in bids that result from buyers’ ability to return.

References

- Akerberg, 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. (2016). “Competition in a Dynamic Auction Market: Identification, Structural Estimation, and Market Efficiency,” *Journal of Industrial Economics*, 64(4), pp. 621-655.
- Ashenfelter, O. (1989). “How Auctions Work for Wine and Art,” *Journal of Economic Perspectives*, 3, pp. 23-36.
- Ashenfelter, O. and K. Graddy (2003). “Auctions and the Price of Art,” *Journal of Economic Literature*, 41(3), pp. 763-87.
- Athey, S. and P. Haile (2002). “Identification of standard auction models,” *Econometrica* 70(6), pp. 2107-2140.
- Backus, M. and G. Lewis (2017). “Dynamic Demand Estimation in Auction Markets,” 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.
- Bodoh-Creed, A., B. Hickman and J. Boehnke (2017). “How Efficient are Decentralized Auction Platforms?” working paper.
- Brancaccio, G., Kalousptsi, M. and T. Papageorgiou (2018). “Geography, Search Frictions and Endogenous Trade Costs,” NBER Working Paper 23581.
- Buchholz, N. (2017). “Spatial Equilibrium, Search Frictions, and Efficient Regulation in the Taxi Industry,” working paper.
- Budish, E. and R. Zeithammer (2016). “An Efficiency Ranking of Markets Aggregated from Single-Object Auctions,” working paper.
- Bulow, J. and P. Klemperer (1996). “Auction versus Negotiations,” *American Economic Review* 86(1), pp. 180-194.
- Canals-Cera, J.J. and Percy, J. (2006). “Econometric Analysis of English Auctions: Applications to Art Auctions on eBay,” working paper.
- Coey, D., Larsen, B., and B. C. Platt (2015). “A Theory of Bidding Dynamics and Deadlines in Online Retail,” 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.
- Engelbrecht-Wiggans, R. (1994). “Sequential Auctions of Stochastically Equivalent Objects,” *Economics Letters*, 44, pp. 87-90.
- Freyberger, J. and B. Larsen (2017). “Identification in Ascending Auctions, with an Application to Digital Rights Management,” NBER Working Paper #23569.
- 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.
- Haile, P. and E. Tamer (2003). “Inference with an Incomplete Model of English Auctions,” *Journal of Political Economy*, 111(1).
- Hopenhayn, H. and M. Saeedi (2015). “Dynamic Bidding in Second Price Auction,” working paper.
- 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. 1535-1546.
- Loertscher, S., Muir E., and P. G. Taylor (2018). “Optimal Market Thickness and Clearing,” working paper.
- McAfee, R.P. (1993). “Mechanism Design by Competing Sellers,” *Econometrica*, 61(6), pp.1281-1312.
- McAfee, R. and Vincent, D. (1993). “The Declining Price Anomaly,” *Journal of Economic Theory*,

60, pp. 191-212.

Milgrom, P. and R. Weber (1999). "A Theory of Auctions and Competitive Bidding, II," in *The Economic Theory of Auctions*, P. Klemperer (ed.), Edward Elgar Publishing.

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.

Peters, M. and S. Severinov (1997). "Competition Among Sellers who offer Auctions instead of Prices," *Journal of Economic Theory*, 75, pp. 141-179.

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.

Said, M. (2011). "Sequential Auctions with Randomly Arriving Bidders," *Games and Economic Behavior*, 73, pp. 236-243.

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

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.

Song, U. (2004). "Nonparametric Estimation of an eBay Auction Model with an Unknown Number of Bidders," working paper, University of British Columbia.

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

Weintraub, G. Y., C. L. Benkard, and B. Van Roy (2008). "Markov perfect industry dynamics with many firms," *Econometrica* 76, no. 6: 1375-1411.

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

Appendix A: Proof of Proposition 1

A buyer is willing to bid in the auction as long as the surplus she earns from winning the auction exceeds the expected value of her outside option. Define

$$R = x - (1 - \alpha)V(x)$$

as the difference between her willingness-to-pay for the good and the expected value of her outside option. Then the buyer bids in the auction if $R < P$, the price she pays if she wins the auction.

Let B denote a buyer's bid and let $\{W_0, \dots, W_t\}$ denote a sequence of high bids. Here W_0 is the high bid at the time the buyer submits his bid B and W_t is the highest losing bid in the event that B is the winning bid and the first rival bid that exceeds B in the event that B is a losing bid. Note that, in the former case, the sequence of high bids generate the posted price sequence $\{W_0, \dots, W_t\}$ and, in the latter case, the posted price sequence is $\{W_0, \dots, B\}$.

Suppose $B > R$. There are three cases to consider.

1. $B > R > W_t$: buyer wins and his surplus is $R - W_t$.
2. $B > W_t > R$: buyer wins and loses $W_t - R$.
3. $W_t > B$: buyer loses and earns zero.

In Case 1, lowering B to R has no impact on the buyer's payoff; $B = R$ is still the highest bid and the sequence of posted prices determining the decisions of the buyers who arrive before or after buyer does not change.

In Cases 2 and 3, lowering B to R has no impact on the sequence of posted prices if $t = 0$ or $t > 0$ and $W_{t-1} < R$. However, if $t > 0$ and $W_{t-1} > R$, then lowering B below W_{t-1} means that B becomes a posted price and could influence the decisions of subsequent buyers (e.g., the buyer submitting W_t). However, it can only do so as a *losing* bid so there is always at least one bidder who bids more than R and does so based on a posted price less than R . As a result, the buyer's payoff remains zero in Case 3 and he eliminates losses in Case 2. Thus, bidding more than R is dominated by bidding R .

Suppose $B < R$. Once again there are three cases.

1. $B > W_t$: buyer wins and his payoff is $R - W_t$.
2. $B < W_t < R$: buyer loses and his payoff is zero.
3. $B < R < W_t$: buyer loses and his payoff is zero.

In Cases 1 and 3, increasing B to R clearly has no impact on the buyer's payoff since, by definition, $W_{t-1} < B$ if $t > 0$. However, in case 2, the buyer could gain be outbidding W_t . Thus, bidding less than your valuation is dominated by bidding R .

Q.E.D.

Appendix B: Endogenous exit

Suppose losing buyers find it costly to stay in the market and bid again. The cost is denoted by c , and it is randomly drawn from a distribution F_C with support $[0, \bar{c}]$. The buyer draws the cost after she bids and loses, and it is independently distributed across a buyer's losses. The probability that a buyer with type x exits is then given by

$$\Pr\{c > V(x; \rho)\} \equiv 1 - F_C(V(x; \rho)), \quad (18)$$

and the optimal bid function is

$$\sigma(x) = x - F_C(V(x; \rho))V(x; \rho). \quad (19)$$

The ex ante value function is given by the function

$$V(x) = \frac{\int_0^{\sigma(x)} (x - m) dG_{M|B}(m|\sigma(x))}{[1 - F_C(V(x; \rho))(1 - G_{M|B}(\sigma(x)|\sigma(x)))]}. \quad (20)$$

Therefore, given $G_{M|B}$, F_C and x , we have three equations to solve for three unknowns: the bid $b = \sigma(x)$, the continuation value $v = V(x; \rho)$, and the exit probability $\alpha = 1 - F_C(v)$. F_C is not known, but it can be identified from the data. To see why, note that we can use the transformation $x = \eta(b)$ and express the above three equations in bid space. The probability of exit becomes

$$\alpha(b) = 1 - F_C(V(\eta(b); \rho)). \quad (21)$$

The inverse bid function is

$$\eta(b) = b + (1 - \alpha(b))V(\eta(b); \rho), \quad (22)$$

and the value equation becomes

$$V(\eta(b); \rho) = \frac{\int_0^b (\eta(b) - m) dG_{M|B}(m|b)}{[1 - (1 - \alpha(b))(1 - G_{M|B}(b|b))]} \quad (23)$$

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

$$\eta(b) = b + \frac{(1 - \alpha(b))}{\alpha(b)} G_{M|B}(b|b) [b - E(M|M < b, b)]. \quad (24)$$

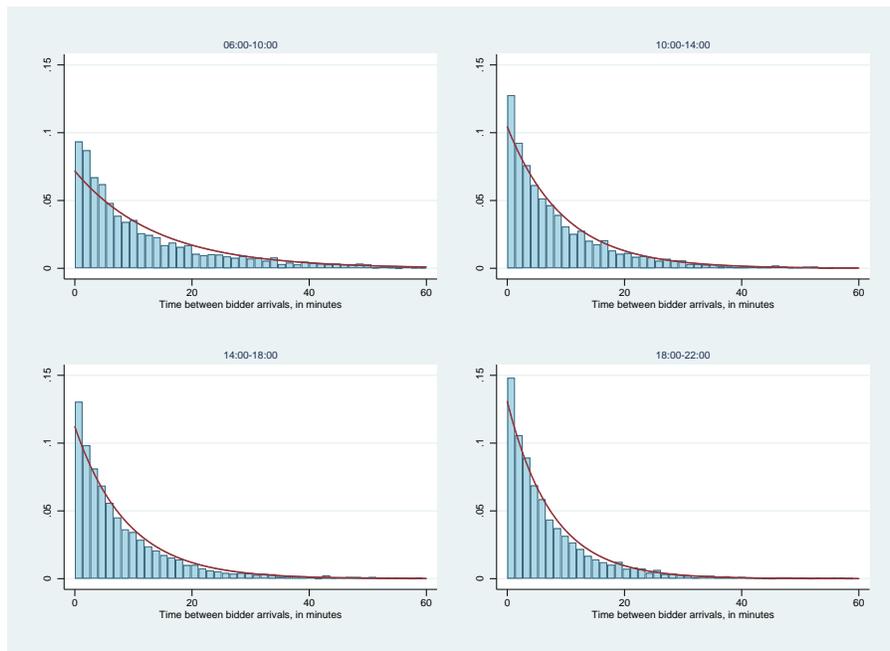
Once again, estimates of the private values can be obtained directly from data on bids and exits.

Thus, F_E (and F_L) are identified. To identify F_C , we solve $v(b) = V(\eta(b); \rho)$ for each bid b and then plot $\alpha(b)$ against $v(b)$ to determine the distribution F_C .

Appendix C: Additional figures

Figure 7 shows the distributions of times between bids, across all bidders and auctions, compared to the exponential distribution. Figure 8 shows the distributions of new bidder arrivals per hour, compared to the Poisson distribution.

Figure 7: Time between bids



In Section 6 we tested our model by comparing the f_L estimated directly from the data to the f_L implied by our model—i.e., the f_L that satisfies the flow restrictions in equation (17). To show that the test has power to reject the model, we ran it using an incorrect estimate of G . Instead of estimating $G_{M|B}$, we simply used the empirical CDF of the winning bid as our estimate of G , applying it both when estimating f_E and when imposing the rescaling implied by equation (17). The results, shown in Figure 9, are clearly worse than those in Figure 5, which is based on the valid estimate of $G_{M|B}$.

Footnote 25 briefly describes an auction selection rule that we estimated. For each buyer in our data, we constructed a choice set consisting of the thirty soonest-to-close auctions in which the posted bid was less than the buyer’s eventual bid—i.e., the thirty soonest-to-close auctions in which the bid she submitted would have been an allowable bid. Using these choice sets, we then estimated a multinomial logit model in which the only explanatory variable was the rank of the auction (by soonest to close), allowing for 10 different coefficients corresponding to the 10 deciles of the submitted bid. In other words, high bidders were allowed to have different preferences than low bidders for bidding in soon-to-close auctions. The magnitudes of the coefficients were monotonically decreasing in the decile: the highest-value bidders were significantly more likely to choose soon-to-

Figure 8: Bidder arrivals per hour

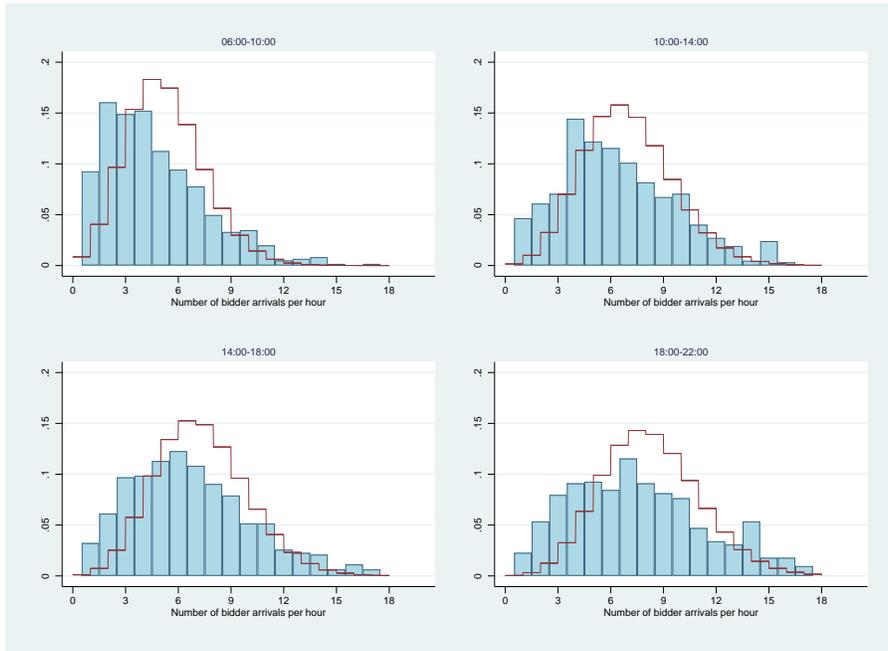
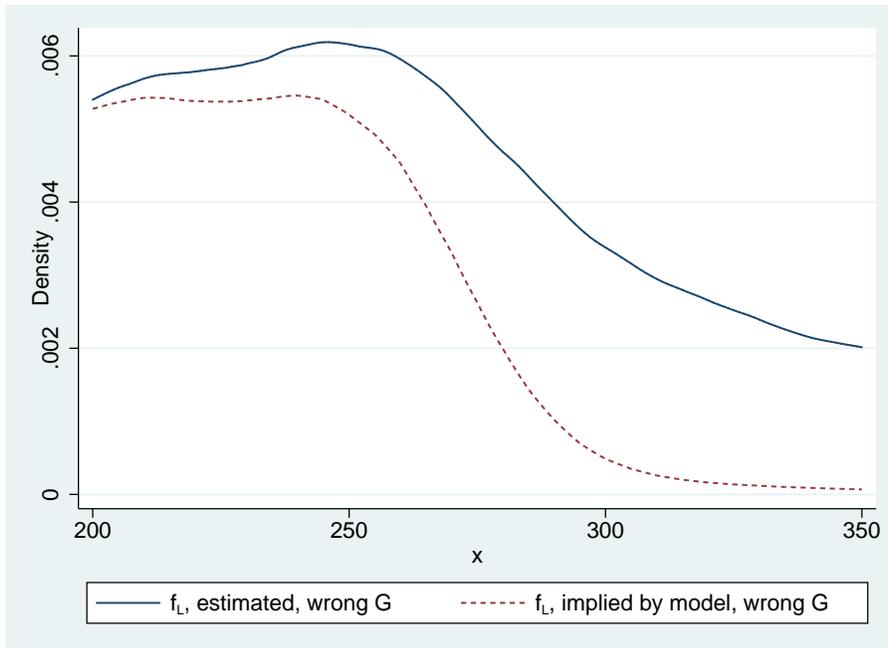
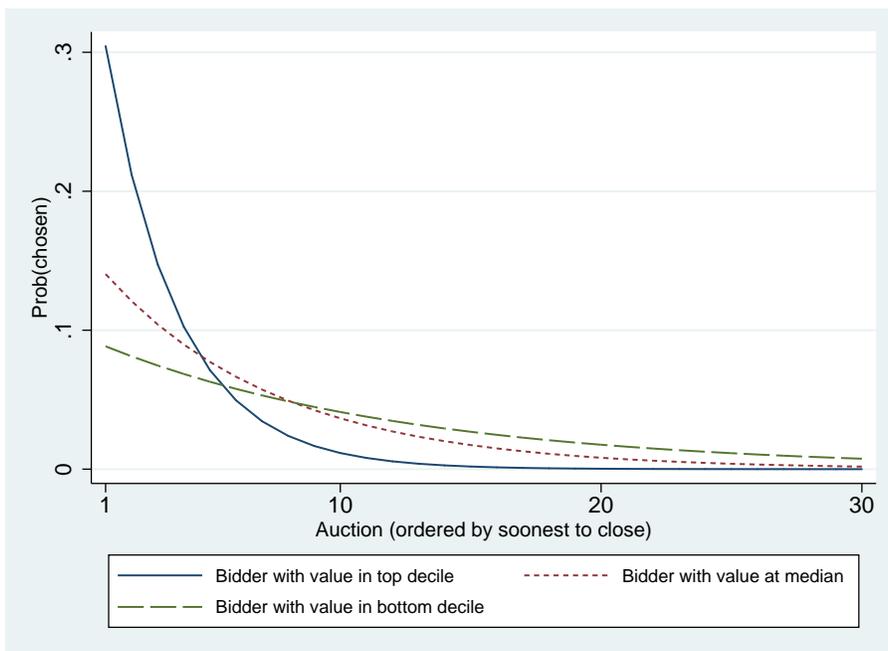


Figure 9: Test of restriction on f_L using wrong estimate of G



close auctions, and the lowest-value bidders' choices were closer to being random across auctions. Figure 10 shows the predicted probabilities for the highest, lowest, and middle-decile bidders.

Figure 10: Auction selection probabilities for high vs. low bidders



Appendix D: Computing bids in counterfactual simulations

An equilibrium of our model consists of a bid function $\sigma(x)$ and a distribution of the maximum rival bid $G_{M|B}$ such that $\sigma(x)$ is optimal given bidders' beliefs, and $G_{M|B}$ is the stationary distribution generated when bidders bid according to $\sigma(x)$. Formally, an equilibrium must satisfy

$$\sigma(x) = x - (1 - \alpha)V(x) \quad (25)$$

and

$$V(x) = \frac{\int_0^{\sigma(x)} (x - p) dG_{M|B}(p|\sigma(x))}{[1 - (1 - \alpha)(1 - G_{M|B}(\sigma(x)|\sigma(x)))]} \quad (26)$$

When the state of the market is a stationary process, $G_{M|B}(\sigma(x)|\sigma(x))$ can be computed as the average probability that a buyer of type x wins. As long as the bid function is monotone, this probability does not depend on the bids, so we can find an equilibrium by first simulating a large number of auctions to compute $G_{M|B}(\sigma(x)|\sigma(x))$, and then numerically solving for the value function $V(x)$ that satisfies conditions (25) and (26). The latter step is a search for a fixed point in function space, and can be accomplished with a simple iterative procedure. We set $V(x)$ equal to zero initially, so that $\sigma(x) = x$, and then compute the surplus that the simulated bidders would have earned in that case. This computed surplus becomes the new estimate of $V(x)$, and the bids are updated according to (25). Surplus is then recomputed for all bidders, and the process is iterated until the

newest estimate of $V(x)$ is unchanged relative to the previous one.

In each simulated auction, we compute the winner's surplus as $x - p$, setting $p = y - (1 - \alpha)V(y)$ where y is the type of the second-highest bidder. To get lifetime surplus (the full continuation value), we scale this result by $1/[1 - (1 - \alpha)(1 - G_{M|B}(\sigma(x)|\sigma(x)))]$. Using the data from the simulated auctions, we estimate $G_{M|B}(\sigma(x)|\sigma(x))$ with a local polynomial regression of the win dummy on x .