

# Dynamics and Efficiency in Decentralized Online Auction Markets

Kenneth Hendricks  
University of Wisconsin & NBER  
[hendrick@ssc.wisc.edu](mailto:hendrick@ssc.wisc.edu)

Alan Sorensen  
University of Wisconsin & NBER  
[sorensen@ssc.wisc.edu](mailto:sorensen@ssc.wisc.edu)

Thomas Wiseman  
University of Texas  
[wiseman@austin.utexas.edu](mailto:wiseman@austin.utexas.edu)

March 2021

## 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. The key features of the model are that the platform posts information about the state of play in each auction and bidders use that information when choosing in which auction to bid and how much to bid. We prove that when markets are thick enough, considerations about how a bidder's choices influence her re-entry payoff (conditional on losing) become unimportant, so that optimal bid strategies are invariant to the state of play and monotonically increasing in type. Given this result, we show that our model is identified from bidding data. Our estimator accounts for the selection effect that arises from the endogeneity of auction choice and avoids what would otherwise be a substantial underestimate of bidders' valuations. We use our estimated model to show 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.

# 1 Introduction

Many goods and services are sold or acquired through 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 trade through auctions. 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 the sellers sell their goods in separate auctions, and buyers choose in which auctions to bid. The online platform facilitates the matching by providing information about the state of play in each auction, but frictions resulting from private information and strategic behavior can still 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 several questions. The positive question is: How does the option to trade again affect prices and efficiency? The normative questions are: Should the platform provide information about the state of play in individual auctions? Can a centralized mechanism, like a double auction for all buyers and sellers in a period, achieve significantly better outcomes?

The goal of this paper is to address these questions for a real-world auction market like eBay. To achieve this goal, we develop a novel theoretical model that captures the main features of the eBay marketplace: namely, that buyers select in real time among a set of overlapping auctions, with the platform posting information about the state of bidding in each of those auctions. Sellers arrive sequentially over time, and each seller sells one unit of a homogenous good in a second-price, ascending price auction of fixed duration. The arrival rate is sufficiently high, and duration sufficiently long, that at any moment in time the market consists of a large number of overlapping auctions with different closing times. Buyers arrive randomly over time and upon arrival observe the highest losing bid (or start price if there are no bids) and time remaining in each available auction. Buyers use this information to decide in which auction to bid and how much to bid. Winners exit, and buyers who lose either exit or return at some random time in the future to bid again.

We discretize time, values, and bids with arbitrarily fine grids and use the theory of Markov chains to prove existence of a stationary equilibrium. We then provide a large market approximation result that makes the model empirically tractable and useful. When a buyer has the option to bid again after losing, her optimal bid depends on her beliefs about her re-entry payoff if and when she returns to try again. In our setting, those beliefs depend on the losing state since it is correlated with the return state. Her beliefs also depend on her choice of auction and the bid she submits in that auction, since these actions affect the decisions of subsequent buyers. We first use a law of

large numbers argument to prove that, as markets thicken, many buyers arbitraging across many auctions mean that these considerations become unimportant. In the limit, a buyer’s continuation value conditional on losing depends only on her type, not on her previous actions or on the losing state. We then prove that it is nearly optimal (in the  $\epsilon$ -equilibrium sense) for each buyer to always bid her type minus her continuation value. This bid is invariant to the state of play or choice of auction and is strictly increasing in a bidder’s value.

The intuition for our limit result is that in thick markets the state is likely to undergo many transitions before a losing buyer returns, so the effects of the losing state and her past actions on her expected re-entry payoff have largely washed away. However, our limit result does not require buyers to believe that they will face the steady-state distribution of states if they lose and bid again; instead, it implies that they expect to face the steady-state distribution of re-entry payoff values. Arbitrage causes many states to have nearly the same value, so the expected re-entry payoff can be independent of the losing state even though the expected return state is not. This is important because in our application—the eBay market for iPads—losing buyers sometimes return quickly before many of the auctions that were open when they lost have closed. They would thus expect the state when they return to be very similar to the state when they exited, rather than expecting a fresh draw from the steady state distribution.

The invariance and monotonicity properties of the equilibrium bid function are crucial to our identification and estimation strategy. Monotonicity implies that a bidder’s type can be equivalently be represented by her equilibrium bid, or what [Backus and Lewis \(2019\)](#) call her “pseudo-type.” 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, in our setting, the unobserved value of a buyer can be expressed as a function of her bid and the distribution of the maximum rival bid *conditional* on her pseudo-type. The dependence between a buyer’s pseudo-type and the maximum rival bid arises mainly from the fact that she chooses the auction based on the observed state of play—that is, the highest losing bids and closing times of the available auctions. Since we observe all bids, including the winning bid, and the identities of the bidders, we can account for this selection effect by estimating the probability that a buyer with pseudo-type  $b$  wins conditional on the set of auctions chosen (in the different states) by buyers who bid  $b$ . Thus, the distribution of values of new buyers (or the distribution of values of returning 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. 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 these buyers 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. We test this restriction. More generally, the flow of losing buyer types who return to the market must on average be equal to the flow of buyer types (new and returning) who 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 these restrictions to test the model. Finally, the estimated inverse bid function needs to be increasing in a buyer's pseudo-type. We find that the data are consistent with each of these implications.

Given estimates of the model primitives, we simulate a number of counterfactuals. The first set focuses on the efficiency of the eBay trading mechanism, which we compare to two hypothetical benchmarks. One is the fully efficient benchmark, which we compute by finding the price that would clear the market if the platform were to pool all buyers and all sellers and conduct a single uniform-price auction. The other is a fully decentralized, *static* mechanism in which the sellers hold separate, simultaneous second-price auctions and buyers are randomly assigned to those auctions. In this mechanism, each buyer gets only one chance to win a unit. We find that the actual eBay mechanism meaningfully increases efficiency relative to this second benchmark, but falls well short of the fully efficient outcome. Prices in the eBay mechanism are higher and much less dispersed than in the static, decentralized mechanism, but the average eBay price is significantly lower than the market-clearing price of the centralized mechanism.

A second set of simulations evaluates the impact of eBay's real-time disclosure of the highest losing bid in each auction on efficiency and prices. Specifically, we run simulations in which eBay does *not* post the highest losing bid in an auction before it closes, so the auctions are effectively sealed bid auctions.<sup>1</sup> In this case, there is an equilibrium in which buyers always bid in the soonest-to-close auction. 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. Our simulation solves for equilibrium bids and computes various auction outcomes. We find that not posting the highest losing bid lowers the average price but significantly increases efficiency relative to the outcomes we observe in the data. The intuition here is that disclosure tends to make the auctions more competitive by disproportionately matching two high-value bidders, but it also makes it more likely that high value buyers exit without winning a unit. This sorting of high-value buyers under the open auction mechanism also generates significantly more price dispersion than would be observed under random matching.

In the third set of simulations, we quantify the effects of dynamic competition on prices and

---

<sup>1</sup>This situation can also arise in open auctions if buyers wait until the last minute to submit their bids. Several papers (e.g., Ockenfels and Roth (2006), Bajari and Hortascu (2003)) have argued for this model of eBay auctions.

efficiency by letting the exit rate go to zero. The option to bid again leads buyers to shade their bids below their values, which we refer to as the dynamic bidding effect. It also increases the level of competition in the market, which we refer to as the dynamic participation effect. In theory, these two effects should cause the market to 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 below the market-clearing price should bid their value and lose almost surely. Thus, as the exit rate goes to zero, 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 remains.

In Section 2 below we review the related literature, both theoretical and empirical. Section 3 describes the model and presents our main results about equilibrium bidding, and also explains how the model can be empirically estimated. We describe the data and estimation results in Sections 5 and 6, respectively. Section 7 presents the results from our counterfactual simulations, and Section 8 concludes.

## 2 Related Literature

Our paper is related to the theoretical literature that studies decentralized, dynamic auction markets with a large number of buyers and sellers. This literature focuses on settings in which a continuum of sellers and a continuum of buyers arrive each period to trade units of a homogenous good, each buyer is randomly matched to one seller, each seller is matched to a random number of buyers, and traders who fail to trade either exit or return the following period. McAfee (1993) uses this framework to study competing mechanisms and shows that, in steady state, there is an equilibrium in which all sellers choose to sell via second-price sealed-bid auctions. Satterthwaite and Shneyerov (2007, 2008) consider a model with two-sided private information and examine what happens to prices and allocations when the number of trading opportunities for each trader is large. They show that in all steady state, Bayesian equilibria, prices converge to the Walrasian price and allocations converge to the efficient allocation. They 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 in markets with a large number of buyers and sellers. Bodoh-Creed, Boehnke, and Hickman (2020) show that a steady state equilibrium of this model is

an  $\epsilon$ -equilibrium of the analogous model with a finite number of buyers and sellers.<sup>2</sup>

Our paper contributes to this literature by characterizing equilibrium behavior in settings where buyers and sellers are matched and trade in real time based on the observed state of play in the auctions. The main challenge in analyzing this setting, however, is that in principle forward-looking buyers need to condition their decisions on the current state of play and consider how their actions affect subsequent play. This issue does not arise in the random matching (RM) models because of the random assignment and because the actions of any single buyer today have negligible impact on the state of the market tomorrow: in steady state, tomorrow’s state is the same as today’s state. We provide an analogous result for real-time matching and trading models like eBay with a large (but finite) number of buyers and sellers: arbitrage across auctions by buyers will equalize a buyer’s continuation value across possible return states. The key assumption is that losing buyers return randomly over time, not immediately as in the RM model. We then thicken the market by letting the time between seller arrivals go to zero, holding fixed the expected return time of a buyer and the expected number of new buyers per auction. In the limit, buyers in our model behave as they do in the RM models in the sense that they bid their value less a continuation value that depends only on their type, not on the current observed state of play or their actions.<sup>3</sup> That continuation value, however, differs from the one in the RM models, because buyers choose which auctions to bid in. This dependence of continuation values on the matching process has fundamental implications for empirical work.

On the empirical side, there is a nascent literature on structural estimation for dynamic auction markets in which buyers know their values and those values are perfectly persistent over time.<sup>4</sup> [Backus and Lewis \(2019\)](#) model eBay as a sequence of sealed bid auctions and use data on buyers’ product choices and bids to estimate substitution patterns in a differentiated good market. [Adachi \(2016\)](#) and [Bodoh-Creed et al \(2020\)](#) estimate models of eBay auctions for homogenous goods in order to evaluate the efficiency and prices of the eBay mechanism. In each of these three papers, the authors argue that because most of the bids that matter in an auction are submitted near the end of the auction, posted prices are not very informative about closing prices or winning bids, so buyers ignore this information when they make their bidding decisions. Specifically, the authors

---

<sup>2</sup>A second strand of this literature characterizes equilibrium bidding in settings where buyers arrive randomly over time and compete in an infinite sequence of single unit, sealed bid, second-price auctions (e.g., [Said \(2011\)](#), [Backus and Lewis \(2019\)](#), and [Zeithammer \(2006\)](#)). These two strands focus on inter-auction dynamics. There is another strand that studies intra-auction dynamics of equilibrium bidding in open second-price auctions (e.g., [Hopenhayn and Saeedi \(2017\)](#)).

<sup>3</sup>An important empirical implication of this result is that our dynamic model is identified under *any* auction format in which the static, one-shot auction is identified, not just the second-price auction. This is a property that [Bodoh-Creed et al \(2020\)](#) refer to as the “plug-and-play” property, and they show that it holds for the RM model.

<sup>4</sup>This setting is very different from the one studied by [Jofre-Bonet and Pesendorfer \(2003\)](#), [Groeger \(2014\)](#), [Balat \(2017\)](#), and [Raisingh \(2020\)](#). They study repeated auction environments in which each bidder gets an independent draw for each auction and does not learn that value until the auction is held.

assume that buyers participate in the soonest-to-close auction (or alternatively the last hour of the auction) and bid in that auction as if it were a sealed bid auction. Since entry times of buyers and sellers are random, this rule generates a random assignment of buyers to sellers. The authors also assume that, when buyers compute their re-entry payoffs, their beliefs about the probability of winning are given by the stationary distribution of the highest rival bid, regardless of how quickly they return. These assumptions may not be consistent with equilibrium play in general, but in our counterfactual analysis we show that they can be rationalized if eBay does not post the highest losing bid and if losing bidders do not return immediately. Thus, in making these assumptions, the literature is essentially assuming that the data-generating process can be approximated by the equilibrium of a model in which the eBay auctions are sealed bid auctions and losing buyers return randomly over time.

Our innovation relative to this literature is to develop a more general empirical model in which posted prices are informative and matching is endogenous. Buyers in our model arbitrage differences among the auctions by choosing the one that is the best match for them based on the observed state of play.<sup>5</sup> Furthermore, each buyer anticipates how her choice of auction and her bid can influence the choices and bids of subsequent buyers in ways that can change her payoff in that auction. The arbitrage activity is especially important when the market experiences a run of high value buyers and soon-to-close auctions fill up. As described above, we provide conditions under which a buyer's optimal bid depends only on her type, and not on which auction she chooses to bid in.

Given that constant-bidding result, we show that our model can be identified from bidding data without solving for the equilibrium auction choice rule. Our strategy involves computing the continuation value of a type- $b$  buyer using the distribution of the maximum rival bid in the set of auctions (in the different states) chosen by type- $b$  buyers. This value can be estimated directly from the data, since the expected re-entry payoff of a type- $b$  buyer is simply her win rate times the average price paid in the auctions that she wins. The continuation value function identifies the equilibrium (inverse) bid function, which can then be applied to the bids of new buyers to obtain the distribution of new buyer values. This approach nests random matching as a special case. If the set of auctions chosen by type- $b$  buyers is a random sample, then the distribution of the highest rival bid in the set of auctions chosen by type- $b$  buyers should be the same as the distribution of the highest rival bid in the set of all auctions.<sup>6</sup> We compute the continuation value functions associated with these two distributions and find that ignoring auction selection leads to a substantial overestimate of buyers' continuation values (and underestimate of their valuations),

---

<sup>5</sup>We assume that all buyers find a match, since there is always an auction that closes within a day that has no bids and a low start price.

<sup>6</sup>In our model, the numbers of new and returning buyers entering per period are assumed to be Poisson random variables, so the distribution of the highest rival bid is equivalent to the distribution of the winning bid under the null hypothesis of no selection.

especially for high-value buyers. We provide further evidence against the random matching model in our empirical analysis.

However, our identification strategy is more data-intensive than the ones used previously in the literature. Their strategies are based on the assumption that the price and winning bid of an auction are the second and first order statistics from an exogenous and random set of buyers. Given this assumption, the authors show that the distribution of values of new buyers can be identified from data on prices or winning bids. The basic idea is to obtain the parent distribution of bids from the distribution of the order statistic and then apply the inverse bid function to bids from the parent distribution to obtain the associated distribution of values.<sup>7</sup> By contrast, we require researchers to observe the identities of 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 obtained by [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 between a buyer's type and the maximum rival bid comes from the buyer's choice of auctions rather than from their values.

There is an earlier structural literature that models eBay auctions as independent, *static* games (e.g., [Bajari and Hortascu \(2003\)](#), [Gonzalez, Hasker, and Sickles \(2004\)](#), [Canals-Cerda and Percy \(2006\)](#), [Akerberg, Hirano, and Shahriar \(2006\)](#) and [Lewis \(2011\)](#)). Our results suggest that static models may not be a good approximation when studying issues related to auction design. First, by ignoring the option to bid in other, concurrent auctions, static models fail to account for buyers choosing an auction based on the state of play. We find that the selection effects from endogenous matching have a significant impact on our estimates of the distribution of buyer values. Second, by ignoring the option to bid again in a future auction, they tend to overestimate the values of the bidder (and underestimate markups), although our results suggests that this effect may be small on average. Third, and perhaps most importantly, by ignoring the distinction between new and returning buyers, these papers implicitly treat the steady state distribution of buyer values as the primitive rather than the distribution of new buyer values. This matters for counterfactuals, since changes in the auction mechanism are likely to lead to a different stationary distribution of buyer values.

Finally, our paper contributes to the empirical literature on dynamic search-and-bargaining models of trade, such as [Gavazza's \(2011, 2016\)](#) studies of the market for used aircraft, [Brancaccio et al's \(2018\)](#) study of global shipping, [Buchholz's \(2017\)](#) study of the New York City taxi market, and

---

<sup>7</sup>[Bodoh-Creed et al \(2020\)](#) cannot distinguish new and returning buyers, so they identify the distribution of new buyer values from the steady state condition that the flow of types entering each period must be equal to the flow that is exiting. [Adachi \(2016\)](#) observes bidder identities, so she simulates the model to identify the mapping between the distribution of bids by new buyers to the stationary distribution of the order statistic, and then finds the bid distribution that minimizes the difference between the simulated and observed distributions of the order statistic.



Coe, Larsen, and Platt’s (2020) study of the effect of buyer deadlines on bidding in eBay auctions. These papers approximate markets with finite numbers of buyers and sellers with a continuum of agents. They focus on steady states and use the restrictions on entry and exit flows as the basis for estimating the models’ primitives. By contrast, we work with the stationary state of a finite market, and use the restrictions on flows as over-identifying tests of our model.

### 3 A Dynamic Model of Trade

Our model is a discrete approximation to an eBay auction market in which buyers arrive and bid in continuous time. Discretizing time, values, and bids with arbitrarily defined grids means that we can analyze the dynamics of the game using the theory of Markov chains. Our model focuses on buyers and treats sellers as non-strategic players. 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 start prices at which they are certain to sell. Of the sellers who set binding start prices, only a small fraction fail to sell, and an even smaller fraction return to sell again. By contrast, most buyers in our application lose, and half of them return to bid again.

In our model, sellers arrive exogenously and deterministically to sell a homogenous good, with the same length of time between arrivals. We define a unit of time to be the length of time between arrivals. Upon arrival, each seller contracts with the platform to sell her good in an ascending, second price auction and sets a start price equal to zero. Each auction lasts for  $J$  (an integer) units, so in every unit of time one auction closes and another opens. Sellers who sell their goods exit the market. If an auction fails to attract any bids, then the seller exits. The open auctions, starting with the next-to-close, are indexed by  $j = 1, \dots, J$ . Time is discrete and indexed by  $t$ . We divide each unit of time into  $T$  periods of equal length, so  $\Delta \equiv 1/T$  is the length of a period. Thus, each auction lasts for  $J \cdot T$  periods. Let  $d(t) \in \{1, \dots, T\}$  denote the number of periods remaining in the next-to-close auction in period  $t$ . The remaining time in auction  $j$  is  $d_j(t) = d(t) + T \cdot (j - 1)$ . Thus, at any time  $t$ , the supply side of the market consists of  $J$  overlapping auctions.

On the demand side, new buyers arrive randomly over time to buy a single unit of the good. The number of new buyers arriving in a period is a random variable distributed according to a Poisson distribution with mean  $\lambda\Delta$ . Arrivals are independent over periods. Each new buyer’s value for the good is drawn independently according to distribution  $F_E$  with density  $f_E$ . The distribution has finite support  $\mathcal{X}$  contained in  $(0, \bar{x}]$  where  $f_E(\bar{x}) > 0$ . A buyer’s value is fixed and does not change over time. Upon arrival, a new buyer selects an auction in which to bid and which bid to submit.

The set of bids is finite and given by  $\mathcal{B} = \{0, \underline{b}, \dots, \bar{b}\}$  where 0 denotes no bid. A bid specifies the “maximum” amount that the bidder is willing to pay, and the platform bids on his behalf up to that level. These are known as *proxy* bids.<sup>8</sup> If his bid is the winning bid, then he gets the good, pays the second highest bid, and exits. If his bid is a losing bid, then he exits with probability  $\alpha$  and gets a payoff of zero;<sup>9</sup> otherwise he goes to the pool of losing buyers and returns at some future time to bid again.

An important feature of our model is that losing buyers that continue do not return immediately. The return process is a discrete version of a continuous time process in which a bidder’s return time is distributed exponentially with rate  $\gamma$ . In each period, the probability of returning is  $\gamma\Delta$ . This arrival rate is independent across time and buyers, and does not depend on when the buyer entered the pool, on how long she has been in the pool, or on her value. Thus, if the number of buyers in the losers’ pool in a period is  $n$ , then the number of returning buyers in that period is distributed Binomial with parameters  $(n, \gamma\Delta)$ .<sup>10</sup>

The platform runs the auction market as follows. At the beginning of each period, the platform lists the closing times of each open auction and posts the current highest *losing* bid in each auction if it has received at least two bids, or the start price of zero otherwise.<sup>11</sup> It does not disclose the highest bids. We will sometimes refer to the highest losing bid (or start price if there are no bids) in an auction as the *posted bid* in that auction. If multiple buyers (new or returning) arrive in the same period, then they are ordered randomly, and their (simultaneously placed) bids are processed in that order. When an auction with at least one bid closes, the platform awards the unit to the highest bidder at the second-highest bid or, if there is only one bid, the start price. Let  $w_j(t) \in \mathcal{B}$  denote the highest bid in auction  $j$  in period  $t$  and let  $r_j(t) \in \mathcal{B}$  denote the highest losing bid in auction  $j$  in period  $t$ . The vectors of highest bids and highest losing bids in period  $t$  are  $w(t)$  and  $r(t)$  respectively.

The payoff-relevant information in any period  $t$  consists of the closing times of the open auctions  $d(t)$ , the highest losing bids  $r(t)$  and highest bids  $w(t)$  in these auctions, the values of the highest bidders, and the size and composition of the losers’ pool. Let  $a_j(t) \in \{0\} \cup \mathcal{X}$  denote the value of the high bidder in auction  $j$  (where  $a_j(t) = 0$  means that no one has bid in that auction) and let  $a(t)$  denote the vector of  $a_j(t)$ ’s. The state of the pool at the beginning of period  $t$  is represented

---

<sup>8</sup>Proxy bidding rules out intra-auction bidding dynamics such as incremental bidding. In our application, we do observe some buyers submitting multiple bids in the same auction. We examine the prevalence of this kind of bidding behavior and discuss how we address it in our empirical analysis below.

<sup>9</sup>If 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.

<sup>10</sup>As period length  $\Delta$  shrinks, the Binomial distribution converges to Poisson distribution with mean  $n\gamma\Delta$ .

<sup>11</sup>The platform also discloses the number of bids so a buyer can distinguish between an auction with no bids and an auction with one bid.

by the distribution  $\mathcal{N}_L(t) \in (\mathbb{Z}_+)^{|\mathcal{X}|}$ , which gives the number of losers of each type in the pool. Then the (countable) set of states that can occur is given by

$$\Omega \equiv \{1, \dots, T\} \times \mathcal{B}^J \times \{\{0\} \cup \mathcal{X}\}^J \times \mathcal{B}^J \times (\mathbb{Z}_+)^{|\mathcal{X}|}.$$

A buyer bids in the period of his arrival. When he arrives in period  $t$ , he observes  $d(t)$  and  $r(t)$  in the open auctions.<sup>12</sup> We call  $\tilde{\omega}(t) \equiv (d(t), r(t))$  the *observed* state;  $\tilde{\Omega}$  is the set of observed states that can occur. We restrict buyers to stationary strategies that condition only on their value and the observable state. That is, given a value  $x$ , a (pure) strategy  $\sigma_x$  is a function from  $\tilde{\Omega}$  to  $\{1, \dots, J\} \times \mathcal{B}$ . Let  $\Sigma$  denote the set of such mixed strategies. In what follows, when we refer to “strategies,” we mean “stationary strategies” unless otherwise noted.<sup>13</sup>

### 3.1 State Transitions

Any profile  $\sigma = (\sigma_x)_{x \in \mathcal{X}}$  of mixed strategies, together with an initial condition  $\omega_0$ , defines a Markov process  $\Phi(\sigma)$  on  $\Omega$ . The number of arrivals in a period of new and returning buyers of each type is determined by the probabilities given above. These arrivals are randomly ordered, and they choose in which auction to bid and what bid to submit according to  $\sigma$ . Returning buyers leave the losers’ pool at the beginning of each period, and losing buyers (new and returning) that fail to exit enter the pool at the end of the period. When an auction closes, the winner exits.

The platform only accepts bids that are at least an increment above the posted bid, so any bids submitted to auction  $j$  in period  $t$  must be strictly greater than  $r_j(t)$ . If more than one bidder submits the same bid, then the tie goes to the one who submitted first. We describe these state transitions more precisely in the appendix.

Those transitions, and the probabilities associated with them, define the one-step transition matrix  $P(\sigma)$  generated by  $\sigma$ . We denote the  $n$ -step-ahead transition function as  $P^n(\sigma)$  and define the probability of reaching state  $\omega$  from an initial state  $\omega_0$  in  $n$ -steps by  $P^n([\omega_0, \omega]; \sigma)$ .

---

<sup>12</sup>In our application, the platform reports the history of highest losing bids in an auction and partially masked identities of losing buyers for each auction that potential buyers can access at some small cost (of time). The assumption here is that buyers do not bother to use this information in forming beliefs about the high bid or the pool of losers. The value of this information is likely to be quite small in thick markets where buyers have the option to bid again.

<sup>13</sup>The restriction to stationary strategies means that a returning buyer cannot condition on any private information about his previous bidding experiences. That is, a returning buyer behaves the same way as a new buyer of the same type.

### 3.2 Ergodicity

Given the observed state, buyers have to form beliefs about the high bids in the open auctions and the state of the losers' pool. We need these conditional beliefs to be well-defined. This requires showing that a stationary strategy profile induces an ergodic distribution.

The first point to note is that  $\Phi(\sigma)$  is not ergodic, because the  $d(t)$  component that tracks the number of periods until the next auction closes is obviously periodic. We aim instead for a result like the following. For each  $d \in \{1, \dots, T\}$ , let  $\{\omega_d, \omega_{T+d}, \dots, \omega_{nT+d}\dots\}$  track the state every time there are  $d$  periods left in the next-to-close auction, and let  $\Phi(\sigma, d)$  denote that Markov process. Given state  $\omega$ , let  $d(\omega)$  denote the component that specifies the number of periods remaining in the next-to-close auction. Thus, the state space of  $\Phi(\sigma, d)$  is  $\Omega(d) \equiv \{\omega \in \Omega \mid d(\omega) = d\}$  and the  $n$ -step-ahead transition function for  $\Phi(\sigma, d)$  is  $P^{nT}(\sigma)$ . Proposition 1 establishes that the Markov process  $\Phi(\sigma, d)$  is ergodic—that is, it converges to a unique invariant distribution,  $\pi(\sigma, d)$ , regardless of the initial state.

**Proposition 1** *For any  $d$  and any initial state  $\omega_0 \in \Omega(d)$ , there exists a unique invariant distribution  $\pi(\sigma, d)$  such that the Markov process  $\Phi(\sigma, d)$  satisfies*

$$\max_{\omega \in \Omega(d)} |P^{nT}([\omega_0, \omega]; \sigma) - \pi(\omega; \sigma, d)| \rightarrow_{n \rightarrow \infty} 0.$$

**Proof.** See Appendix. ■

It would be sufficient to show that  $\Phi(\sigma, d)$  is an irreducible, recurrent, and aperiodic<sup>14</sup> process. In general,  $\Phi(\sigma, d)$  may not be irreducible. For example, suppose that there is an equilibrium where arriving buyers bid only in the soonest-to-close auction. Under such strategies, states in which later-to-close auctions have already received bids never occur. However, we can show that  $\Phi(\sigma, d)$  has a single absorbing communicating class  $\Omega^C(\sigma, d)$ , because every state leads to the empty state in which there are no bidders in any auction and the losers' pool has no bidders.<sup>15</sup> Thus, there exists a Markov process  $\Phi^C(\sigma, d)$  with state space  $\Omega^C(\sigma, d)$  and the same transition probabilities as above, restricted to  $\Omega^C(\sigma, d)$ , that is irreducible and recurrent. It therefore has a unique invariant measure, and because  $\Phi(\sigma, d)$  eventually leads to  $\Omega^C(\sigma, d)$ ,  $\Phi(\sigma, d)$  has the same unique invariant measure. The process  $\Phi(\sigma, d)$  is also aperiodic since the empty state transitions with positive probability to itself. This establishes the proposition because an aperiodic process on a countable

<sup>14</sup>A Markov process is irreducible if every state can be reached from every other state; it is recurrent if in expectation each state is visited infinitely often; and it is aperiodic if there is a state that transitions in one step to itself with positive probability.

<sup>15</sup>A state  $\omega$  leads to state  $\omega'$  if the probability of reaching  $\omega'$  from  $\omega$  is strictly positive. Two states communicate if each leads to the other.

state space with a unique invariant probability measure is ergodic.

Given a state  $\omega$ , let  $\tilde{\omega}(\omega)$  denote the observable component. We say that an observable state  $\tilde{\omega}$  is “on the long-run path of  $\sigma$ ” if there exists a state  $\omega \in \Omega^C(\sigma, d(\omega))$  such that  $\tilde{\omega}(\omega) = \tilde{\omega}$ . That is, on-long-run path observable states are those that occur in the absorbing communicating class of  $\sigma$ . Proposition 1 implies that there are well-defined long-run conditional beliefs  $\pi(\sigma, \tilde{\omega}) \in \Delta(\Omega)$ , given by Bayes’ rule, for every on-long-run-path observable state  $\tilde{\omega}$ . ( $\Delta(\Omega)$  denotes the set of probability distributions over  $\Omega$ .) Given a strategy profile  $\sigma$ , buyers can use these beliefs to compute their expected payoffs from choosing an auction and submitting a bid.

### 3.3 Stationary Equilibrium

We need to specify what it means for a strategy to be a best response. Let  $\sigma$  be the strategy profile used by other players, and let  $p : \tilde{\Omega} \rightarrow \Delta(\Omega)$  specify buyer  $i$ ’s beliefs about the state conditional on the observable state. Suppose buyer  $i$  with value  $x$  submits a bid  $b$  in auction  $j$  in observable state  $\tilde{\omega}$ . Then  $\sigma$  and  $p(\tilde{\omega})$  determine the buyer’s beliefs over future states, and specifically over other bids in auction  $j$ . Let  $M_j$  denote the highest rival bid in auction  $j$ . If buyer  $i$  wins the auction, then only the value of  $M_j$  affects his payoff. For each weakly lower bid  $m \in \{0, \dots, b\}$ , let  $g_{\sigma,p}(m; \tilde{\omega}, j, b)$  denote the probability of the event that buyer  $i$  wins and that the highest rival bid submitted before the auction closes (including bids submitted before or at the same time as  $b$ ) is  $m$ . Buyer  $i$  wins for sure when  $m < b$ , but he also wins when  $m = b$  and  $m$  is submitted after  $b$ . If buyer  $i$  loses the auction and enters the losers’ pool, then what matters for his expected continuation value is the state of the market in the period immediately following his loss. Let  $\omega^l$  denote this state and, in what follows, we will refer to it as the *losing* state. For each  $\omega^l \in \Omega$ , let  $h_{\sigma,p}(\omega^l; \tilde{\omega}, j, b)$  denote the probability of the event that buyer  $i$  loses the auction and that the losing state is  $\omega^l$ . These winning and losing probabilities depend not only on the observable state, but also upon the buyer’s auction choice and bid since they can affect the distribution over future states, and specifically over other bids in auction  $j$ .

We now define a buyer’s payoffs. Given  $(\sigma, p)$ , let  $v(x, \omega; \sigma, p) : \Omega \rightarrow [0, x]$  be the expected payoff to a buyer of type  $x$  who arrives at state  $\omega$  and plays his optimal strategy (which depends only on the observable component  $\tilde{\omega}$ ). Then define

$$\tilde{v}(x, \tilde{\omega}; \sigma, p) \equiv \sum_{\omega \in \Omega} v(x, \omega; \sigma, p) \cdot p(\omega; \tilde{\omega})$$

as his maximized payoff when he arrives and observes  $\tilde{\omega}$ , given conditional beliefs  $p(\tilde{\omega})$ . His expected

re-entry payoff if he loses and the losing state is  $\omega^l$  is given by

$$V(x, \omega^l; \sigma, p) \equiv \sum_{t=1}^{\infty} \gamma \Delta (1 - \gamma \Delta)^{t-1} \left( \sum_{\omega' \in \Omega} P^{t-1}([\omega^l, \omega']; \sigma) v(x, \omega'; \sigma, p) \right). \quad (1)$$

The term  $P^{t-1}([\omega^l, \cdot])$  gives the distribution over the state when the buyer returns after  $t$  periods given the losing state  $\omega^l$ . The number of periods until he returns is itself random, and the term  $\gamma \Delta (1 - \gamma \Delta)^{t-1}$  represents its distribution.

We can now write down the Bellman equation for the type- $x$  buyer. For each observable state  $\tilde{\omega} \in \tilde{\Omega}$ ,

$$\tilde{v}(x, \tilde{\omega}, ; \sigma, p) = \max_{j \in \{1, \dots, J\}, b \in \mathcal{B}} \left[ \begin{array}{l} \sum_{m \in \{0, \dots, b\}} (x - m) \cdot g_{\sigma, p}(m; \tilde{\omega}, j, b) \\ + (1 - \alpha) \sum_{\omega^l \in \Omega} V(x, \omega^l; \sigma, p) h_{\sigma, p}(\omega^l; \tilde{\omega}, j, b) \end{array} \right] \quad (2)$$

A bidder's best response to  $(\sigma, p)$  is a stationary strategy that achieves these optimal values for every observable state  $\tilde{\omega} \in \tilde{\Omega}$ . A strategy is a best response to  $(\sigma, p)$  if it specifies a best response for each value  $x \in \mathcal{X}$ .

The first term in Expression 2 represents the payoff to a type- $x$  buyer who wins auction  $j$ . The summation is over the highest losing bid  $m$ , weighted by the probability of each such  $m$ . The second term is the payoff to the buyer if he loses. Since the buyer does not actually observe  $\omega^l$ , the summation is over the set of possible losing states, weighted by the probabilities of those states. Note that, because the buyer either wins or loses,

$$\sum_{\omega^l \in \Omega} h_{\sigma, p}(\omega^l; \tilde{\omega}, j, b) = 1 - \sum_{m \in \{0, \dots, b\}} g_{\sigma, p}(m; \tilde{\omega}, j, b)$$

Overall, the second term in Expression 2 is the bidder's continuation value (probability of losing, times probability of not exiting, times expected re-entry payoff) given the observed state and the buyer's auction choice and bid.

Both the probability of losing and the expected re-entry payoff depend on  $b$  and  $j$ . The first dependence is straightforward—how much the buyer bids in which auction affects the probability that he wins. The second dependence is less obvious. It operates through two channels. First,  $b$  and  $j$  may directly influence continuation play (and thus the re-entry payoff) by changing the actions of future buyers who observe them. Second, the buyer's expectation of his re-entry payoff depends on the losing state, and different  $j$ 's and  $b$ 's lead to different distributions over  $\omega^l$ . For

example, if the buyer submits a very high  $b$  and loses, then he may conclude that the losers' pool is likely to have lots of high types, and so his expected re-entry payoff is low. In contrast, if he submits a low  $b$  and loses, then that event is not very informative about the underlying state, so he becomes relatively more optimistic about his re-entry payoff. The observed state  $\tilde{\omega}$  also affects the buyer's beliefs about  $\omega^l$ . For example, if the buyer submits a bid and it is not the highest bid submitted in that period, then he enters the losers' pool in the following period, so  $\tilde{\omega}$  will be highly informative of the losing state  $\omega^l$ .

Before defining an equilibrium, we need to define a conditional belief system. Recall that a strategy profile  $\sigma$  induces an ergodic distribution  $\pi(\sigma, d)$  over states given  $d$  periods until the next auction closes, which induces conditional beliefs  $\pi(\sigma, \tilde{\omega})$  at every observable state  $\tilde{\omega} \in \tilde{\Omega}$  that is on the long-run path of  $\sigma$ . The challenge here is that  $\sigma$  does not pin down conditional beliefs at observable states that are off-path, that is, not in the support of the ergodic distribution  $\pi(\sigma, d)$ . Our approach for dealing with this issue is to require that conditional beliefs be the limit of some sequence of full-support strategies that converges to  $\sigma$ , as in sequential equilibrium.

**Definition 1** *A conditional belief system  $p$  is consistent with strategy profile  $\sigma$  if there exists a sequence of full-support strategies  $\{\sigma_k\}$  such that (i)  $\sigma_k \rightarrow \sigma$ , and (ii)  $\pi(\sigma_k, \tilde{\omega}) \rightarrow p(\tilde{\omega})$  for every observable state  $\tilde{\omega} \in \tilde{\Omega}$ .*

Having defined best responses and consistent beliefs, we can now define an equilibrium.

**Definition 2** *An equilibrium is a (stationary) strategy profile  $\sigma^* \in \Sigma$  and a conditional belief system  $p^* : \tilde{\Omega} \rightarrow \Delta(\Omega)$  such that (i)  $\sigma^*$  is a best response to  $(\sigma^*, p^*)$ , and (ii)  $p^*$  is consistent with  $\sigma^*$ .*

**Proposition 2** *An equilibrium exists.*

**Proof.** See Appendix. ■

The proof involves showing that the set of conditional belief systems consistent with a strategy profile  $\sigma$  is upper hemicontinuous in  $\sigma$ . We then use a standard fixed point theorem to establish existence. In general, there may be multiple equilibria.

In estimation, we will use a somewhat broader equilibrium concept,  $\epsilon$ -equilibrium. The existence of an  $\epsilon$ -equilibrium is implied by Proposition 2.

### 3.4 An Approximation Result

The goal of this section is to provide a characterization of equilibrium bidding in thick markets. The main idea is to leverage the fact that buyers do not return immediately after losing—it takes some time for them to learn and respond to the news that they have lost and to bid again. As a result, if the arrival rates of buyers and sellers are sufficiently high, then many buyers may have bid and a large number of auctions may have closed between the time that a buyer loses and the time that he returns. In these kinds of markets, a buyer’s re-entry payoff may be largely independent of the losing state since, by the time he returns, the market has undergone so many transitions that the effects of the losing state have dissipated. If this is the case, then the buyer’s continuation value after losing depends only on his type, and not on his previous actions or on the observable state. We look for an equilibrium that satisfies that condition, which greatly simplifies the buyer’s bidding decision.

To see why, suppose that a type- $x$  buyer’s expected re-entry payoff  $V(x; \sigma, p)$  does not depend on the losing state. It would then follow from Expression 2 that the Bellman equation for a type- $x$  buyer becomes

$$\tilde{v}(x, \tilde{\omega}; \sigma, p) = \max_{j \in \{1, \dots, J\}, b \in \mathcal{B}} \left[ \sum_{m \in \{0, \dots, b\}} (x - m) \cdot g_{\sigma, p}(m; \tilde{\omega}, j, b) + (1 - \alpha) \left( 1 - \sum_{m \in \{0, \dots, b\}} g_{\sigma, p}(m; \tilde{\omega}, j, b) \right) V(x; \sigma, p) \right] \quad (3)$$

**Proposition 3** *Suppose that the expected re-entry payoff for a type- $x$  buyer is  $V(x; \sigma, p)$ , regardless of the losing state  $\omega^l$ . Then the following bid is weakly dominant for type- $x$  buyer given any strategy profile and conditional beliefs  $(\sigma, p)$ :*

$$b(x) = x - (1 - \alpha)V(x; \sigma, p).$$

**Proof.** See Appendix. ■

The proposition states that each buyer should bid his value less his expected re-entry payoff. Since the latter depends only on the buyer’s type, each type- $x$  buyer submits the same bid regardless of which auction he chooses and what the observed state is. We will refer to such a strategy as a *constant* bidding strategy. Adachi (2016), Backus and Lewis (2019), and Bodoh-Creed, Boehnke, and Hickman (2020) get the same optimal bid function, but in their models the standard weak dominance argument for second-price auction applies. Our proof extends this argument to settings in which a buyer’s bid can influence the bidding decisions of subsequent buyers in ways that can



cause the probability distribution of the highest rival bid to change.

The challenge is to provide plausible conditions under which the expected re-entry payoff does not depend on the losing state. One possibility is to take the return rate  $\gamma$  to 0 while holding the number of open auctions  $J$  fixed. In that case, the expected return time is so far in the future that the buyer might expect the re-entry state to be drawn from the steady state distribution.<sup>16</sup> However, there is a subtle problem with that limit: all the buyers in the losers' pool return at the same rate, so they are likely to compete against each other when they return. Thus, the state of the losers' pool at the time that a buyer loses affects the expectation of his re-entry state even in the limit as  $\gamma$  shrinks. That limit is also inconsistent with the data. Specifically, it implies that all the open auctions will have closed (in fact, multiple cycles of open auctions will have opened and closed) before the buyer returns. By contrast, in our application, many of the auctions that were open when the buyer loses are still open when he returns. By the time a losing bidder returns to bid again, on average over fifty other buyers have bid, but only six auctions have closed. Given that the number of open auctions is over a hundred, this means that a returning bidder will see mostly the same auctions he saw last time he bid.

An alternative limit, which is more plausible and which does deliver an expected re-entry payoff that is independent of the losing state, is the following: fix the buyer's expected return time and the expected number of new buyers per auction, but thicken the market by increasing the arrival rates for sellers and new buyers, letting the time between seller arrivals go to zero. Recall that we normalized the time between seller arrivals as one unit of time. Therefore, in our model, we take the above limit by letting the return rate  $\gamma$  shrink to zero while increasing the number of auctions,  $J$ , so that  $\gamma \cdot J$  is constant.<sup>17</sup>

In this limit, the expected number of auctions that will close before a loser returns ( $1/\gamma$ ) gets large, but the fraction of currently open auctions that will close before a loser returns ( $(1/\gamma)/J$ ) is constant. As a result, the expected return state is not independent of the losing state, because most of the auctions open when the bidder loses will still be open when he returns. However, the large numbers of auctions that close in the meantime imply that the effect of the losing state on the bidder's re-entry payoff will have largely washed away. Many buyers arbitraging across many auctions mean, by a law of large numbers argument, that the losing state will not matter much (except in extreme cases, as when a run of high value buyers have filled all open auctions with high bids).

---

<sup>16</sup>This is the approach that the literature has taken and that we took in an earlier version of this paper.

<sup>17</sup>Suppose time is measured in terms of hours rather than time between seller arrivals. Fix the duration of each auction at  $J^*$  hours and the return rate of a buyer at  $\gamma^*$  per hour. Let  $\delta^*$  denote the number of hours between seller arrivals and let  $\lambda^*$  be the arrival rate of a new buyer per hour. Then  $\gamma = \gamma^* \delta^*$ ,  $\lambda = \lambda^* \delta^*$ , and  $J = J^* / \delta^*$ . Letting  $\delta^*$  shrink to zero implies that  $\gamma \rightarrow 0$ ,  $J \rightarrow \infty$ , and  $\gamma J = \gamma^* J^*$ .

That is, this limit does not imply that a losing buyer expects to face the steady-state distribution of states when he returns, but it does imply that he expects to face the steady-state distribution of re-entry payoff values. Arbitrage with lots of buyers and sellers implies that many states have nearly the same value, so that the expected re-entry payoff can be independent of the losing state even though the expected return state is not. However, for any fixed  $\gamma$  a constant bidding strategy will not be exactly optimal, for at least two reasons. First, the expectation of what the value of the state will be when a losing buyer returns does depend a little on the current observable state. Second, the value of the losing bid and choice of auction do influence the actions of future bidders a little bit. For both reasons, a bidder would want to “fine tune” his bid. As a result, an exact equilibrium will not feature constant bids.

Instead, we consider a weaker solution concept,  $\epsilon$ -equilibrium, in order to capture the idea that buyers will ignore those details that vanish in the limit. We show that a constant bidding strategy can be an approximate best response for any precision of approximation. For  $\epsilon > 0$ , an  $\epsilon$ -equilibrium is a strategy profile where no buyer has a deviation that can improve his expected payoff by more than  $\epsilon$ . In our model, “expected payoff” means the steady-state expected payoff.<sup>18</sup> Our main result is stated in the following Theorem.

**Theorem 1** *Pick any  $\epsilon > 0$ . Fix a sequence  $\{\gamma_k, J_k\}_{k=1}^\infty$  such that  $\lim_{k \rightarrow \infty} \gamma_k = 0$  and  $\gamma_k J_k$  is constant. Then there exists a sequence of  $\epsilon$ -equilibria  $\{(\sigma_k^*, p_k^*)\}_{k=1}^\infty$  such that for high enough  $k$ , (i) each type of bidder  $x$  submits a single bid  $b_k^*(x)$  on the equilibrium path, and (ii)  $b_k^*(\cdot)$  is increasing.*

**Proof.** See Appendix. ■

The idea behind the proof of part (i) is to treat the buyer’s choice of auction as a sequence of binary participation decisions in static auctions with a fixed outside option. We then construct a constant bidding strategy profile for each  $\gamma_k$  and show that that strategy profile is an  $\epsilon$ -equilibrium for  $\gamma_k$  sufficiently small. More precisely, define an arbitrary function  $v_0 : \mathcal{X} \rightarrow (0, \bar{x}]$ , representing the value of the outside option for each type. Given  $v_0$ , define

$$b_k(x, v_0) = x - (1 - \alpha)v_0(x)$$

as the bid of each type- $x$  buyer. Define arbitrary beliefs  $p_0$  that specify, for each observable state, beliefs about the current highest bids  $w$  in the  $M$  next-to-close auctions, where  $M < J_k$ . Given  $v_0$  and  $p_0$ , we recursively construct a constant bidding strategy profile  $\sigma_k$  in which each buyer bids

---

<sup>18</sup> $\epsilon$ -equilibrium is widely used in game theory. As Mailath, Postlewaite, and Samuelson (2005) explain, the solution concept captures the idea that a slight mis-specification of the underlying game should not cause the modeler to rule out reasonable predicted outcomes. Similarly,  $\epsilon$ -equilibrium is appropriate if players find it costly to compute the optimal strategy, or if they believe that other players may make small mistakes.

in one of the  $M$  next-to-close auctions (details in the appendix). Given such a strategy, on-path steady state beliefs are pinned down by Bayes' rule, and we can fill in off-path beliefs in a consistent way; call the result  $p(\sigma_k)$ . Similarly,  $\sigma_k$  determines the expected payoff in steady state to a bidder of each type  $x$ ; call that function  $v(\sigma_k)$ . We look for a fixed point: a strategy profile that determines steady-state payoffs and beliefs that in turn generates the same strategy profile. Let  $\sigma_k^*$  be this fixed point and call the corresponding payoffs  $v_k^*(x)$ .

We first prove that a fixed point  $\sigma_k^*$  exists. We then prove that, fixing a large  $M$  and taking  $k$  to infinity, for any type  $x$ , there is a high probability of arriving at a state where playing according to  $\sigma_k^*(x)$  gives at least (close to)  $v_k^*(x)$ , and submitting a bid in auction  $M + 1$  or later gives payoff no higher than (a value close to)  $v_k^*(x)$ . The idea is to make  $M$  large enough that arbitrage across auctions (as in [Burguet and Sákovic \(1999\)](#)) equalizes expected payoffs, while making  $\gamma_k$  small enough that a losing buyer expects that all of those  $M$  auctions will have closed before he returns. The proof of part (ii) is a standard mechanism design argument.

The theorem says that we can pick  $\epsilon > 0$  as close to zero as we want, and make the value of the “mistakes” that buyers make in best responding with a constant bidding strategy arbitrarily small. At most observable states, the constant bid is almost but not quite optimal, so bidders may be making small mistakes most of the time. At some unlikely extreme observable states, such as those in which all  $M$  auctions have received multiple high bids or those in which none of the  $M$  auctions have received bids, bidders may make large mistakes. But both kinds of mistakes have a small effect on the expected payoff evaluated ex ante: neither a high probability of a small mistake or a small probability of a big mistake affects the expectation much. Furthermore, these unlikely extreme states are not states that we see in the data.

To summarize, we have shown that the constant bidding strategy is nearly optimal nearly all the time, and increasing in type. As we shall see in the next section, these two properties are critical to making the model useful for empirical work. Furthermore, our results do not depend upon the length of a period or the bid increment. We can make these intervals arbitrarily small so that the probability of ties shrinks to zero.

## 4 Empirical model

In what follows, we assume that the data are generated by an  $\epsilon$ -equilibrium  $(\sigma^*, p^*)$  in which buyers use a constant bidding strategy. Under this assumption, we obtain closed form solutions for the value function and the (inverse) bid function and show that the latter can be expressed in terms of bid distributions. We then show that our model is identified and outline a strategy for estimating

the primitives of the model. Finally, we develop and discuss several tests of the model.

Given Theorem 1, a buyer's type  $x$  can equivalently be represented by his bid,  $b^*(x)$ . We follow [Backus and Lewis \(2019\)](#) and refer to  $b^*(x)$  as type  $x$ 's "pseudotype." The constant bidding result allows us to aggregate across states. For each  $m \in \{0, \dots, b^*\}$ , define

$$g_{\sigma^*, p^*}(m|b^*) = \sum_{\tilde{\omega} \in \tilde{\Omega}} g_{\sigma^*, p^*}(m|\tilde{\omega}, j^*(\tilde{\omega}; b^*), b^*) \tilde{\pi}^*(\tilde{\omega}),$$

as the probability that pseudotype  $b^*$  pays  $m$  in the set of auctions in which he chooses to bid and wins, where  $\tilde{\pi}^*(\tilde{\omega})$  denotes the steady state probability of observable state  $\tilde{\omega}$  under  $\sigma^*$  and  $j^*(\tilde{\omega}; b^*)$  denotes the auction chosen by pseudotype  $b^*$  at  $\tilde{\omega}$ . In order to simplify notation, we assume here that the auction choice rule  $j^*$  is a pure strategy, but none of our identification results below depend on that assumption. Similarly, define

$$G_{\sigma^*, p^*}(m|b^*) = \sum_{m \in \{0, \dots, b^*\}} g_{\sigma^*, p^*}(m|b^*)$$

as the probability that he wins in those auctions. Note that  $b^*$  plays two roles here: it accounts for the set of auctions that type  $x$  selects and the bid he submits in those auctions.

Evaluating the Bellman equation (2) at  $j^*$  and  $b^*$  and taking expectations over  $\tilde{\omega}$  gives

$$\sum_{\tilde{\omega} \in \tilde{\Omega}} \tilde{v}(x, \tilde{\omega}; \sigma^*, p^*) \tilde{\pi}^*(\tilde{\omega}) = \sum_{\tilde{\omega} \in \tilde{\Omega}} \left( \begin{array}{c} \sum_{m \in \{0, \dots, b^*\}} (x - m) g_{\sigma^*, p^*}(m|\tilde{\omega}, j^*(\tilde{\omega}; b^*), b^*) + \\ (1 - \alpha) \left( 1 - \sum_{m=0}^{b^*} g_{\sigma^*, p^*}(m|\tilde{\omega}, j^*(\tilde{\omega}; b^*), b^*) \right) V(x; \sigma^*, p^*) \end{array} \right) \tilde{\pi}^*(\tilde{\omega})$$

Changing the order of summation on both sides, summing over  $\tilde{\omega}$ , and solving for  $V$ , we obtain

$$\begin{aligned} V(x; \sigma^*, p^*) &= \sum_{m \in \{0, \dots, b^*\}} (x - m) g_{\sigma^*, p^*}(m|b^*) + (1 - \alpha)(1 - G_{\sigma^*, p^*}(m|b^*))V(x; \sigma^*, p^*) \\ &= \frac{\sum_{m \in \{0, \dots, b^*\}} (x - m) g_{\sigma^*, p^*}(m|b^*)}{[1 - (1 - \alpha)(1 - G_{\sigma^*, p^*}(m|b^*))]} \end{aligned}$$

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

We use this expression for  $V$  to solve for the inverse bid function, which we denote by  $\eta$ . Substituting  $V$  into the constant bid function from Proposition 3 and solving for  $x$  yields

$$\eta(b^*) = b^* + \left(\frac{1-\alpha}{\alpha}\right) \sum_{m \in \{0, \dots, b^*\}} (b^* - m) g_{\sigma^*, p^*}(m|b^*) \quad (4)$$

Thus, the private values of bidders can be obtained directly from data on their 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.

#### 4.1 Identification and Estimation

Our data on buyers consists 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 the number of potential buyers (i.e, buyers who visit the platform) is equal to the number of actual buyers (i.e., buyers who submit a bid). 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 unobserved model primitives are the entry, return and exit parameters  $(\alpha, \gamma, \lambda)$  and the distribution of values,  $F_E$ . Given data on bidder identities and participation, identification of the parameters is straightforward :  $\hat{\lambda}$  is the average number of new buyers arriving per period,  $\hat{\gamma}$  is the mean return time of a loser who does not exit, and  $\hat{\alpha}$  is the fraction of losing buyers who exit. The distribution  $F_E$  is identified from Expression 4 which we can rewrite as

$$\eta(b^*) = b^* + \left(\frac{1-\alpha}{\alpha}\right) G_{\sigma^*, p^*}(b^*|b^*) [b - E(M|M < b^*, b^*)]$$

A non-parametric estimate of  $G_{\sigma^*, p^*}(b^*|b^*)$  can be obtained by computing the fraction of auctions in which pseudo-type  $b^*$  bids and wins. Similarly, a non-parametric estimate of  $E[M|M < b^*, b^*]$  is the average price that the pseudotype  $b^*$  pays when he wins these auctions. Given these estimates, we can use the sample of bids by new buyers to compute their values, and obtain a non-parametric

estimate of  $F_E$ . Note that we can also use Expression 4 to derive estimates of the private values of *returning* buyers, and use these estimates to obtain a non-parametric estimate of the stationary distribution of values in the loser’s pool. We denote this distribution by  $F_L$  and its probability distribution by  $f_L$ .

The remarkable aspect of our analysis is that  $F_E$  is identified without solving for the equilibrium selection rule. This result is due to the fact that each buyer type submits a single bid, regardless of which auction he chooses or what the observable state is. This invariance property allows the econometrician to use each buyer’s bid to directly infer his 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 the bids of every buyer and assume that they are realizations of pseudotypes. The latter is a strong assumption 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 this bid as her pseudotype. Nevertheless, bid censoring may still be a problem.

## 4.2 Tests

Our model generates several testable implications. These tests are important because they shed light on the validity of our key approximations.

The first set of predictions concern bidding behavior. Bid functions need to be strictly increasing. Since this is the case if and only if  $\eta$  is increasing, we can test for monotonicity by checking that, our estimates of  $G_{\sigma^*, p^*}(b^*|b^*)$  and  $E[M|M < b^*]$  imply that the expression on the RHS of equation (4) is increasing. Second, if buyers are using a constant bid strategy, then buyers who lose and return should bid approximately the same amount. The data on bidder identities allow 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. Since a buyer’s maximum bid may not be his pseudotype due to incremental bidding, this test also provides information on the extent to which bid censoring is a problem.

The second set of predictions concern the restrictions implied by steady state. 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 the time between auction closings is

$$\gamma \bar{n} = \frac{(1 - \alpha)(\lambda - q)}{\alpha}, \tag{5}$$

where  $\bar{n}$  denotes the steady state size of the losers' pool and  $q$  is the probability that an auction ends successfully with a sale. We test this condition using the data on bidder identities. Second, and relatedly, 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 the time between closings is  $\gamma\bar{n}f_L(x)$ , where  $f_L$  is the probability distribution of types in the loser's pool. The flow back into the pool over this time is on average

$$(1 - \alpha)[1 - G_{\sigma^*, p^*}(b^*|b^*)][\gamma\bar{n}f_L(x) + \lambda f_E(x)].$$

Equating these two flows yields

$$f_L(x) = \frac{\lambda\alpha(1 - G_{\sigma^*, p^*}(b^*|b^*))}{(\lambda - 1)[1 - (1 - \alpha)(1 - G_{\sigma^*, p^*}(b^*|b^*))]} f_E(x). \quad (6)$$

Equation (6) shows that, in steady state, the probability distribution of values in the losers' pool is a rescaling of the probability distribution of values of new buyers. The relationship reflects the censoring due to auction outcomes. The scaling factor approaches 0 for very high types since they are almost certain to win, and it approaches  $\lambda/(\lambda - 1) > 1$  for very low types who are almost certain to lose. As a result,  $f_L$  has more density than  $f_E$  at low values and less density at high values.

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 F 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  $G_{\sigma^*, p^*}$ , 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%). Also, 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.<sup>19</sup>

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 reserve prices that are clearly intended to be non-binding: 22% of listed items had 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

---

<sup>19</sup>In our analysis we only include used items that were fully functioning—i.e., we exclude items identified as “For parts or not working.”



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,<sup>20</sup> 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.

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

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.

Items in our data rarely fail to sell, but in that event sellers have the option to come back 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, we focus on dynamics among bidders and largely ignore the seller dynamics.

<sup>20</sup>eBay requested that we not report the exact conversion rate, but it is higher than 85%.

Our model also 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.

The presence of incremental bidding also raises the important question of which bids to take seriously when estimating the model, since it complicates inference 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  $f_L$  in Section 6 below. 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.

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 G. 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 use constant bidding strategies: 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 bidder fixed effects explain 87% of the variance in bids. This result also suggests that, for most buyers, the maximum bid is approximately equal to the pseudotype.

## 6 Estimation

In this section we first explain how we obtain estimates of bidders' arrival and exit rates ( $\lambda$  and  $\alpha$ ) from the data. We then turn to our method for estimating the distribution of bidders' valuations, which is an adaptation of the method proposed by [Guerre, Perrigne, and Vuong \(2000\)](#) to a dynamic setting.

### 6.1 Estimating bidder arrival and exit rates

Our theoretical model assumes for simplicity that the interval between seller arrivals (or equivalently auction closings) is constant, and the arrival rate of buyers is the measured relative to this period. In reality, the arrival rates of sellers and buyers differ by time of 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 unit of time equals one auction closing) is approximately correct. We therefore estimate  $\lambda$ , the arrival rate of new buyers, as the average number of new buyers per auction closing, which is 5.47.

Conditional on losing an auction, 49.8% of bidders come back to bid again in a subsequent auction. Our estimate of the exit rate,  $\alpha$ , is thus 0.502.<sup>21</sup> 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

---

<sup>21</sup>We 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  $\alpha$ , since most bidders return relatively quickly if they are going to return at all.

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

reflecting bidders who take 24 hours or more to come back.<sup>22</sup>

Although our model can be extended to allow for endogenous exit, as shown in Appendix F, 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 very low bids are 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.<sup>23</sup>


## 6.2 Estimating the distribution of bidders’ valuations

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  $b^*(x)$  (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

<sup>22</sup>Some of the losing bidders return to purchase an item on eBay at a fixed price, but this is rare: we find that only 1% of losing bidders do this.

<sup>23</sup>Since we can estimate the  $\alpha(b)$  function, estimating the model with endogenous exit is not much more difficult than with exogenous exit. However, when simulating counterfactuals in a model with endogenous exit, we must compute a new equilibrium in which value functions accurately reflect exit functions and exit functions are optimal given the value functions.

Figure 1: Exit rate as a function of bid



by equation (4). This inversion requires estimates of the exit rate,  $\alpha$ ; the probability of winning,  $G_{\sigma^*, p^*}(b^*|b^*)$ ; and the expected price conditional on winning,  $E(M|M < b^*, b^*)$ .

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

Estimating  $G_{\sigma^*, p^*}(b^*|b^*)$  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 6<sup>th</sup>-order Hermite polynomial.

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  $M$  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  $\alpha$  and the functions  $G_{\sigma^*, D^*}(b^*|b^*)$  and  $E(M|M < b^*, b^*)$ , we can compute the bid function and directly check monotonicity. Figure 2 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 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.

Figure 2: Estimated bid function





Figure 3 shows kernel density estimates of  $f_E$  and  $f_L$ .<sup>24</sup> The difference between the estimated densities is consistent with the model: the distribution of returning losers’ valuations looks like a 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.

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

<sup>24</sup>Even 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 distributions of valuations, using all bids



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.

### 6.3 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 (5) implies that the number of returning buyers per auction should be given by


$$\gamma \bar{n} = \frac{(1 - \alpha)(\lambda - q)}{\alpha}$$

where  $q$  represents the probability that an auction ends with a sale.<sup>25</sup> The size of the loser pool,  $\bar{n}$ , is not observable, but  $\gamma\bar{n}$  is observable: it is the average number of returning bidders per auction, which is 4.86 in the data. Our estimates of the exit rate  $\alpha$  (0.50), the arrival rate of new bidders  $\lambda$  (5.47), and the probability of sale  $q$  predict an average of 4.58 returning bidders per auction, which is not far off.<sup>26</sup>

Equation (6) 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  $f_L$ , as expressed in equation (6) above.

Figure 4 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 (6) 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  $\lambda$ ,  $\alpha$ , and/or  $G_{\sigma^*, D^*}(b|b)$  are inaccurate or invalid.

Figure 4: 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.

<sup>25</sup>This probability is very high in our data, but eBay preferred that we not publish its exact value.

<sup>26</sup>Re-arrival times in the data imply a value for  $\gamma$  of approximately 0.01, which would imply the average size of the loser pool is between 450-500.

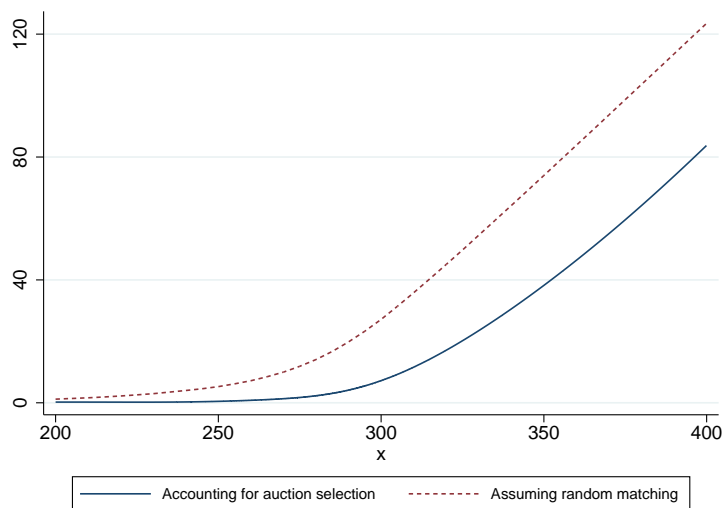


## 6.4 Auction selection

As noted in Section 2, a key distinction of our model compared to the prior literature is that it allows buyers to endogenously choose in which auction to bid: buyers with different valuations will not only submit different bids, they will choose different auctions in which to bid. This means that the distribution of the maximum rival bid depends on the buyer’s type, so we are careful in our estimation procedure to condition on the set of auctions chosen by bidders of type  $b^*$  when computing  $G_{\sigma^*, p^*}(b^*|b^*)$  and the expected price conditional on winning,  $E(M|M < b^*, b^*)$ . By contrast, in a model with random matching of buyers to auctions one could simply use the *unconditional* distribution of the highest rival bid, which is equivalent to the distribution of the winning bid under the assumption of Poisson arrivals. This is the approach taken by Adachi (2016) and Bodoh-Creed et al (2020), for example.

To check whether the distinction is quantitatively important, we can compute and plot bidders’ estimated continuation values under both assumptions, as shown in Figure 5. Ignoring auction selection leads to a substantial overestimate of bidders’ continuation values, especially for high-value bidders. Our estimation strategy does not require us to model buyers’ actual auction selection rules, but the discrepancy shown in Figure 5 indicates the importance of using a method that allows for auction selection instead of simple random matching.

Figure 5: Estimated continuation values: selection vs. random matching



Random matching is also inconsistent with some other basic patterns in the data. For instance, price distributions in the data are roughly invariant to the number of bidders in the auction, as shown in Table 3. If assignment of bidders to auctions were random, we would expect the mean

Table 3: Price distributions: data vs. random matching

# of bidders	% of auctions		Average prices		Std. deviations	
	Data	RM*	Data	RM	Data	RM
0-3	12.61	18.56	269.71	268.15	32.11	20.96
4-6	18.17	70.52	274.92	275.03	28.02	15.43
7-9	16.79	10.74	275.13	280.31	27.09	13.91
10-12	16.75	0.18	276.37	283.20	26.15	6.71
13+	35.67	0.00	277.29	–	24.00	–

\* RM represents a simulation with random matching of bidders to auctions, with censoring of bidders who are outbid before their turn (so the table reports the number of bidders whose bids would have been observed).

price to increase with the number of bidders in an auction, and the standard deviation of prices to decrease. To highlight this, the table compares averages and standard deviations from the data to those from a simulation in which bidders (with valuations drawn from our estimated distribution) arrive randomly and are assigned to bid in the next-to-close auction when they arrive. We discuss this simulation in more detail in Section 7 below, but one important feature for the present purposes is that bidders assigned to a given auction bid sequentially in a random order, so some bidders are outbid before their turn. Consistent with the [Bodoh-Creed et al \(2020\)](#) model, we assume these bidders are censored: the table shows the number of bidders who would have been observed in the data. With this mechanism we would expect average prices to increase with the number of observed bidders, and the standard deviation to decrease substantially as well. Instead what we observe in the data is that average prices are fairly flat with respect to the number of bidders, and the standard deviation of prices declines much less sharply with the number of bidders than it would under random matching.

Another implication of the random matching models is that the winning bid and price in an auction are first and second order statistics from the set of buyers assigned to bid in that auction. However, while it is reasonable to assume that bidders arrive to the platform randomly, our data make it clear that bidders do not simply bid in the auction that is next to close when they arrive. In fact, we find that 79.9% of bidders submit their bids in auctions that are *not* the next to close. Furthermore, these are not merely low-value bidders submitting irrelevant bids: 19.8% of them submit bids in their chosen auctions that are higher than the posted bid in the next-to-close auction they chose to pass up, and 9.6% submit bids that were even higher than the eventual price of the auction they passed up. Thus, the order statistics from the sample of bidders who arrive in any given period are distributed across auctions.

## 7 Counterfactual analyses

In this section we present counterfactual analyses that examine the effect of dynamic competition on efficiency and prices. 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).

### 7.1 Efficiency

Our first set of counterfactuals measure the efficiency of the eBay mechanism relative to two benchmarks. One benchmark is the efficient allocation. 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.

To evaluate the extent to which dynamics in the eBay marketplace yield convergence toward the efficient market outcome, we begin by using our estimates to calculate the market-clearing price  $P^*$  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 second benchmark is motivated by the matching process that occurs in brick and mortar markets. In these markets, the matching is determined by the physical locations of buyers and sellers: buyers have to buy from a local seller, and a seller has to sell to local buyers. Online platforms create thicker markets at any moment by eliminating location as a factor in the matching process. But they also make the market large over time by allowing buyers who fail to purchase to return to the market in a later period and try again with a different seller. Our aim is to measure the efficiency gains that result from this pooling of buyers and sellers over locations and time.

We address this by considering a counterfactual in which the  $N_s$  sellers hold separate second-price auctions with no reserve price, and the  $N_b$  buyers are randomly allocated to those auctions, each buyer getting only one chance to win an auction. In other words, the sellers are local monopolists, and demand is stochastically the same in each local market. This counterfactual 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 benchmark by taking  $N_b$  buyers, with valuations drawn randomly from our estimated  $F_E$ , and randomly assigning them to  $N_s$  auctions, taking the averages from 10,000 repetitions in order to minimize any noise introduced by the simulation draws.

Table 4 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.

These results resemble those of [Bodoh-Creed et al \(2020\)](#), who conduct a counterfactual welfare exercise very similar to ours. They find a larger welfare loss relative to the efficient market-clearing benchmark (14%), but also point out that the eBay mechanism achieves three quarters of the potential welfare gain relative to a lottery that randomly allocates items to bidders.<sup>27</sup> They also show that meaningful increases in efficiency can be achieved by selling items in uniform auctions

---

<sup>27</sup>If we make the same calculation for our data, we find that the eBay mechanism achieves 88% of the potential welfare gain of market-clearing over a random lottery.

Table 4: 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  $P^*$  wins an auction before exiting.

of small batches—e.g., auctioning four or eight units at a time, instead of one at a time—without going all the way to a single uniform auction.

The efficiency differences shown in Table 4 between the data and the inefficient benchmark (simultaneous auctions with no dynamics) can be attributed entirely to the dynamic participation effect, since the dynamic bidding effect has no impact on efficiency. 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.

## 7.2 Posted Bid vs. Sealed Bid Auctions

One potential explanation for the relative inefficiency of the outcome we observe in the data is endogenous matching. As explained previously, buyers allocate themselves across auctions in a way that is clearly nonrandom, instead choosing auctions to arbitrage differences in expected payoffs. To explore how much this matters for prices and efficiency, we consider a counterfactual in which the platform posts the closing times of the auctions, but does not provide any information about the state of bidding in any of the auctions. In other words, the auctions are sealed bid auctions. Each buyer observes the closing times of the available auctions, chooses one, and then waits until the end of the auction to learn whether she is a winner or a loser and, if she has won, the price she has to pay. In that setting, there is an epsilon equilibrium where buyers always bid in the

soonest-to-close auction.<sup>28</sup>

**Proposition 4** *Suppose the auctions are sealed bid auctions. Pick any  $\epsilon > 0$ . Fix a sequence  $\{\gamma_k, J_k\}_{k=1}^{\infty}$  such that  $\lim_{k \rightarrow \infty} \gamma_k = 0$  and  $\gamma_k J_k$  is constant. Then there exists a sequence of  $\epsilon$ -equilibria  $\{(\sigma_k^*, p_k^*)\}_{k=1}^{\infty}$  such that for high enough  $k$ , each type of bidder  $x$  always chooses the soonest-to-close auction upon arrival and submits a bid equal to*

$$b^*(x) = x - (1 - \alpha)V(x; \sigma^*, p^*)$$

where

$$V(x; \sigma^*, p^*) = \frac{\sum_{m \in \{0, \dots, b^*\}} (x - m)g_{\sigma^*, p^*}(m)}{[1 - (1 - \alpha)(1 - G_{\sigma^*, p^*}(m))]}.$$

In this equilibrium, the random arrival times of buyers implement a random allocation of buyers to sellers. The argument is straightforward: a buyer cannot gain by deviating and choosing a later auction because her rivals do not observe her deviation and, given their soonest-to-close choice strategy, the level of competition is the same in every auction. As a result, each buyer is indifferent as to which auction to join, so there is no selection effect. The buyers' continuation values (and bids) are determined by the distribution of the highest rival bid among the new and returning buyers assigned to the auction which, in our model, consist of the buyers who arrive during the period in which that auction is the soonest to close. In steady state, this distribution does not vary across auctions. Thus, this equilibrium is a dynamic version of the equilibrium in the RM models, and the one that the previous structural literature has estimated.

We simulate equilibrium outcomes for a sequence of 10,000 auctions,<sup>29</sup> 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

---

<sup>28</sup>Under the assumption that buyers use stationary strategies, the outcome in Proposition 4 is in fact an exact equilibrium. That assumption, though, is very restrictive in a sealed bid environment where the only public information is the auction closing times. The proposition holds if we drop that assumption and allow buyers to condition on their private history, and it holds regardless of what information about outcomes the platform releases when an auction closes.

<sup>29</sup>To 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  $F_E$ .

the loser pool, a bidder has a probability  $\hat{\gamma} = 0.008$  of returning to bid each period,<sup>30</sup> where each period has one auction. When a buyer arrives, she is assigned to the soonest-to-close auction. This selection rule is analogous to the inefficient benchmark from Table 4, 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 (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 H.

Table 5 shows the comparison of outcomes. The sealed bid auction 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, and therefore they are also less likely to exit. In steady state, there are more high-value buyers and they are spread more evenly across auctions, which reduces price dispersion and improves allocative efficiency.

Table 5: 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, which means they are randomly assigned to auctions. 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.

<sup>30</sup>We estimate  $\gamma$  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.

### 7.3 Convergence

In our third counterfactual, we examine whether the equilibrium of our dynamic model would converge to the efficient outcome as the exit rate ( $\alpha$ ) 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 need to specify an auction selection rule: we cannot simply use the arrivals and re-arrivals observed in the data, because changing  $\alpha$  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.<sup>31</sup>

Table 6 shows that as  $\alpha$  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  $\alpha = 0.50$ .

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.

---

<sup>31</sup>If we use alternative rules that are homogeneous—meaning the rule that assigns bidders to auctions does not depend on the bidder's type—the results are essentially unchanged.




Table 6: 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.

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. Inter-

estingly, 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 result is an approximation result: in thick enough markets, it is approximately optimal for a buyer's bid to be invariant to the choice of auction and the observed state. This is what makes the model empirically tractable, and it is in some ways 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. Relying on this approximation result 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. We believe this approach will likely be useful in many markets, but 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. On the one hand, our modeling approach allows us to identify the model's primitives without actually solving for equilibrium auction selection rules. On the other hand, there is an important sense in which we must control for auction selection. To recover the primitive distribution of buyers' valuations we use 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. 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 that a buyer faces in the auctions in which he chooses to bid. Hence, while it is not necessary to explicitly model how buyers are choosing auctions, 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 (2019). “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.
- Balat, J. (2017). “Highway Procurement and the Stimulus Package: Identification and Estimation of Dynamic Auctions with Unobserved Heterogeneity,” working paper.
- 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 (2020). “How Efficient are Decentralized Auction Platforms?” Forthcoming, *Review of Economic Studies*.
- 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 (2017). “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.
- Burguet, R. and J. Sákovic (1999). “Imperfect Competition in Auction Designs,” *International Economic Review* Vol. 40, No. 1, pp. 231-247.
- 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 (2020). “Discounts and Deadlines in Consumer Search,” *American Economic Review* 110(12), pp. 3748-85.

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.

Groeger, J. (2014). “A Study of Participation in Dynamic Auctions,” *International Economic Review*, 55(4).

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 (2017). “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.

Kreps, D. and R. Wilson (1982). "Sequential Equilibria," *Econometrica* 50:863-894.

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.

Mailath, G. J., A. Postlewaite, and L. Samuelson (2005). "Contemporaneous perfect epsilon-equilibria," *Games and Economic Behavior* 53, 126-140.

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.

Meyn, S.P. and R.L. Tweedie (1993). *Markov Chains and Stochastic Stability*, Springer Verlag.

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. (1981). "Optimal Auction Design," *Mathematics of Operations Research* 6, 58-73.

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.

Raisingh, D. (2020). "The Effect of Pre-announcements on Participation and Bidding in Dynamic Auctions," working paper.

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.

Selten, R. (1975). "A Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory* 4 (1): 25-55.

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.

## A State transitions

The transitions for an auction  $j$  that does not close at the end of the current period (i.e.,  $d_j(t) > 1$ ) are as follows.

- If auction  $j$  receives no bids in period  $t$ , then  $w_j(t+1) = w_j(t)$ ,  $r_j(t+1) = r_j(t)$ , and  $a_j(t+1) = a_j(t)$ .
- If auction  $j$  receives exactly one bid  $b_j$  from a buyer  $i$  with value  $x$  in period  $t$ , then there are two possible transitions:
  - If  $b_j > w_j(t)$ , then  $w_j(t+1) = b_j$ ,  $r_j(t+1) = w_j(t)$ , and  $a_j(t+1) = x$ ; the displaced bidder with value  $a_j(t)$  enters the losers' pool with probability  $1 - \alpha$  and otherwise exits.
  - If  $b_j \leq w_j(t)$ , then  $w_j(t+1) = w_j(t)$ ,  $r_j(t+1) = b_j$ , and  $a_j(t+1) = a_j(t)$ ; buyer  $i$  with value  $x$  enters the losers' pool with probability  $1 - \alpha$  and otherwise exits.
- If auction  $j$  receives bids from multiple buyers in period  $t$ , then there are three possible transitions. Let  $b_j$  be the maximum of the bids, submitted by bidder  $i$  with value  $x$ , and let  $b'_j$  denote the second-highest.
  - If  $b'_j > w_j(t)$ , then  $w_j(t+1) = b_j$ ,  $r_j(t+1) = b'_j$ , and  $a_j(t+1) = x$ ; the displaced bidder with value  $a_j(t)$  enters the losers' pool with probability  $1 - \alpha$  and otherwise exits, as do the buyers other than  $i$ .
  - If  $b_j > w_j(t) \geq b'_j$ , then  $w_j(t+1) = b_j$ ,  $r_j(t+1) = w_j(t)$ , and  $a_j(t+1) = x$ ; the displaced bidder with value  $a_j(t)$  enters the losers' pool with probability  $1 - \alpha$  and otherwise exits, as do the buyers other than  $i$ .
  - If  $b_j \leq w_j(t)$ , then  $w_j(t+1) = w_j(t)$ ,  $r_j(t+1) = b_j$ , and  $a_j(t+1) = a_j(t)$ ; each of the arriving buyers enters the losers' pool with probability  $1 - \alpha$  and otherwise exits.

The transitions for an auction  $j$  that closes at the end of the period (i.e.,  $d_j(t) = 1$ ) are as follows.

- If auction  $j$  receives no bids in period  $t$ , then the high bidder with value  $a_j(t)$  exits.
- If auction  $j$  receives at least one bid, then there are two possible transitions. As above, let  $b_j$  be maximum of the bids, submitted by bidder  $i$  with value  $x$ , and, if there are multiple bids, let  $b'_j$  denote the second-highest.
  - If  $b'_j > w_j(t)$ , then bidder  $i$  with value  $x$  exits; the displaced bidder with value  $a_j(t)$  enters the losers' pool with probability  $1 - \alpha$  and otherwise exits, as do the buyers other than  $i$ .
  - If  $b_j \leq w_j(t)$ , then the high bidder with value  $a_j(t)$  exits; all other bidders enter the losers' pool with probability  $1 - \alpha$  and otherwise exit.



## B Proof of Proposition 1

Our proof relies on standard results about Markov chains on a countable state space. (See, for example, Meyn and Tweedie (1993).) First, we show that  $\Phi(\sigma, d)$  has a unique absorbing communicating class; call it  $\Omega^C(\sigma, d)$ . Second, we show that the Markov chain confined to that class,  $\Phi^C(\sigma, d)$ , is ergodic. The proposition follows.

The following lemma immediately implies that  $\Phi(\sigma, d)$  has a unique absorbing communicating class.

**Lemma 1** *State  $(\mathbf{0}, \emptyset, \mathbf{0}, \mathbf{0}, d)$  is recurrent under  $\Phi(\sigma, d)$ .*

**Proof.** We will show that starting from any state, the process reaches the empty state  $(\mathbf{0}, \emptyset, \mathbf{0}, \mathbf{0}, d)$ , where there are no active buyers and no bids in any open auction, with probability 1. Consider the total number of buyers in the losers' pool,  $n(t)$ . The probability that such a buyer reenters over the next  $T$  periods (the length of an auction) is  $1 - (1 - \gamma\Delta)^T$ . The expected number of buyers who leave the losers' pool over  $T$  periods, then, is at least

$$\alpha \left( n \left[ 1 - (1 - \gamma\Delta)^T \right] - J - 1 \right) :$$

the returning losers, minus the  $J + 1$  spots available as high bidders in open auctions ( $J$  auctions are open at a time, and at most one new auction can open up over  $T$  periods), times the probability of exit  $\alpha$ .

The expected number of buyers entering the losers' pool over  $T$  periods is at most  $(\lambda + J)(1 - \alpha)$ : the expected number of new bidders arriving, plus the  $J$  high bidders at period  $t$ , times the probability  $1 - \alpha$  that a losing bidder enters the losers' pool rather than exiting. Thus, whenever

$$n(t) > \frac{(\lambda)(1 - \alpha) + J}{\alpha \left[ 1 - (1 - \gamma\Delta)^T \right]},$$

$n$  is falling on average over the next  $T$  periods. Pick an

$$n^* > \frac{(\lambda)(1 - \alpha) + J}{\alpha \left[ 1 - (1 - \gamma\Delta)^T \right]},$$

and it follows from the law of large numbers that any  $n > n^*$  will reach a state less than or equal to  $n^*$  with probability 1.

Starting from any state  $\omega_0 \in \Omega(d)$  such that  $n \leq n^*$ , the probability of reaching  $(\mathbf{0}, \emptyset, \mathbf{0}, \mathbf{0}, d)$  is bounded below by  $L(n^*)$ , defined as follows: the probability  $(\gamma\Delta\alpha)^{n^*}$  that  $n^*$  losers enter in the current period and exit if they do not win an auction, times the probability  $(e^{-\lambda\Delta})^{JT} = e^{-\lambda J}$  that no new bidders enter over the next  $JT$  periods until all the current auctions close and the state hits  $d$  again.

Thus, the process reaches  $(\mathbf{0}, \emptyset, \mathbf{0}, \mathbf{0}, d)$  with probability 1: the set of states satisfying  $n \leq n^*$  is reached infinitely often, and the probability of reaching  $(\mathbf{0}, \emptyset, \mathbf{0}, \mathbf{0}, d)$  from any state in that set is bounded below by  $L(n^*) > 0$ . ■

Lemma 1 implies that the unique absorbing communicating class of  $\Phi(\sigma, d)$ ,  $\Omega^C(\sigma, d)$ , is the set of states that communicate with  $(\mathbf{0}, \emptyset, \mathbf{0}, \mathbf{0}, d)$ .

Define the Markov process  $\Phi^C(\sigma, d)$  with state space  $\Omega^C(\sigma, d)$  as having the same transition probabilities as  $\Phi(\sigma, d)$ , restricted to  $\Omega^C(\sigma, d)$ . As constructed,  $\Phi^C(\sigma, d)$  is irreducible and recurrent. It therefore has a unique invariant distribution  $\pi(\sigma, d)$ . And because the empty state  $(\mathbf{0}, \emptyset, \mathbf{0}, \mathbf{0}, d)$  follows itself under  $P^T(\sigma)$  with probability at least  $(e^{-\lambda\Delta})^T = e^{-\lambda}$  (the probability that no new buyers enter over the next  $T$  periods until the state hits  $d$  again),  $\Phi^C(\sigma, d)$  is aperiodic.

Because  $\Phi^C(\sigma, d)$  is aperiodic with a unique invariant distribution, it is ergodic. Lemma 1 then immediately implies that  $\Phi(\sigma, d)$  is ergodic as well, with the same invariant distribution.

## C Proof of Proposition 2

The proof mirrors Kreps and Wilson's (1982) existence result for sequential equilibrium, which in turn relies on Selten's (1975) result for extensive form trembling hand perfect equilibrium. The idea is that the limit of Nash equilibria of a sequence of perturbed games where each action must be played with positive probability is an equilibrium in our setting.

For any small  $\epsilon > 0$ , define the  $\epsilon$ -perturbed game  $\Gamma^\epsilon$  as our model with the restriction that each type of buyer must choose each possible action with probability at least  $\epsilon$  at every observable state. It is straightforward to show that a Nash equilibrium of  $\Gamma^\epsilon$  exists using Kakutani's fixed point theorem: a pure strategy is a function from the finite set  $\mathcal{X} \times \bar{\Omega}$  to the finite set  $\{1, \dots, J\} \times \mathcal{B}$ , so the set of mixed strategies satisfying the  $\epsilon$  restriction is a compact, convex subset of a finite dimensional simplex. Expression 2 is continuous in the strategies of other players  $\sigma$  and conditional beliefs  $p$ , so the best response correspondence is upper hemicontinuous in  $\sigma$  and  $p$ . Given a full-support strategy  $\sigma$ , every observable state  $\tilde{\omega}$  is on the long-run path, so all conditional beliefs  $\pi(\sigma, \tilde{\omega})$  are pinned down by Bayes' rule. Those conditional beliefs are continuous in  $\sigma$ , because for each  $d$  the stationary distribution  $\pi(\sigma, d)$  is continuous in  $\sigma$ . Thus, the mapping from  $\sigma$  to best responses is upper hemicontinuous, and Kakutani's fixed point theorem applies.

Then take a sequence  $\{\epsilon_n\}$  of  $\epsilon_n > 0$  converging to 0, and a sequence  $\{\sigma_n^*\}$  of Nash equilibria of  $\Gamma^{\epsilon_n}$ . The set of strategy profiles is compact, so without loss of generality assume that  $\{\sigma_n^*\}$  converges to a limit  $\sigma^*$ . As noted above, conditional beliefs are continuous in  $\sigma$ , so the sequence of conditional beliefs  $\{\pi(\sigma_n^*, \tilde{\omega})\}$  also has a limit; call it  $p^*$ . We want to show that  $(\sigma^*, p^*)$  is an equilibrium. First, the upper hemicontinuity of the best response correspondence ensures that  $\sigma^*$  is a best response to  $(\sigma^*, p^*)$ . Similarly, to establish that  $p^*$  is consistent with  $\sigma^*$ , it is enough to show that the set of conditional belief systems consistent with a strategy profile  $\sigma$  is upper hemicontinuous in  $\sigma$ .

That argument is straightforward: for any strategy profile  $\sigma$  and conditional belief system  $p$ , take a sequence  $\{\sigma_n, p_n\}$  such that (i)  $\sigma_n \rightarrow \sigma$ , (ii)  $p_n \rightarrow p$ , and (iii) for each  $n$ ,  $p_n$  is consistent with  $\sigma_n$ . We want to show that  $p$  is consistent with  $\sigma$ . By definition, for each  $n$  there exists a sequence of full-support strategies  $\{\sigma_{n,k}\}_k$  such that as  $k \rightarrow \infty$ ,  $\sigma_{n,k} \rightarrow \sigma_n$  and  $\pi(\sigma_{n,k}, \tilde{\omega}) \rightarrow p_n(\tilde{\omega})$  for every observable state  $\tilde{\omega} \in \tilde{\Omega}$ . Define the sequence  $\{\sigma'_k, p'_k\}$  by  $\sigma'_k = \sigma_{k,k}$  and  $p'_k = p_{k,k}$ . By construction,  $\sigma'_k \rightarrow \sigma$  and  $\pi(\sigma'_k, \tilde{\omega}) \rightarrow p(\tilde{\omega})$  for every observable state  $\tilde{\omega} \in \tilde{\Omega}$ , so we conclude that  $p$  is consistent with  $\sigma$ .

Thus,  $(\sigma^*, p^*)$  is an equilibrium.

## D Proof of Proposition 3

For now, suppose that  $b(x)$  is a feasible bid; that is, that  $b(x) \in \mathcal{B}$ . Because the bid that a buyer submits in an auction may influence the actions of future bidders who arrive before the auction closes, the argument that  $b(x)$  is weakly dominant is slightly more complicated than in the case of a static second price auction. The key observation is that a buyer's bid  $b$  can affect future bidders' behavior only through the observable state. Because only the second highest current bid  $r$  is visible,  $b$  is observed only when the highest competitor's bid exceeds  $b$ .

Suppose that the buyer submits a bid in period  $t$  in an auction that will close after  $d$  more periods. For  $s \in \{t, \dots, t+d\}$ , let  $X_s$  denote the highest competitor's bid in the auction up through period  $s$ . Let  $\{x_t, \dots, x_{t+d}\}$  denote the realized increasing sequence of highest competing bids if the buyer submits a bid of  $b(x)$ . If the buyer submits a bid of  $b(x)$  and  $x_{t+d} > b(x)$ , then the buyer loses the auction and gets expected continuation payoff  $(1 - \alpha)V(x; \sigma, p)$ . If  $x_{t+d} < b(x)$ , then the buyer wins the auction, pays  $x_{t+d}$ , and gets payoff  $x - x_{t+d} > x - b(x) = (1 - \alpha)V(x; \sigma, p)$ . And if  $x_{t+d} = b(x)$ , then depending on timing and tie-breaking, the buyer either loses the auction and gets continuation payoff  $(1 - \alpha)V(x; \sigma, p)$ , or wins the auction, pays  $x_{s+d}$ , and gets the same payoff:  $x - x_{t+d} = x - b(x) = (1 - \alpha)V(x; \sigma, p)$ .

Next consider a bid  $b > b(x)$ . There are three cases. If  $x_{t+d} < b(x)$ , then the outcome is the same as with a bid of  $b(x)$ : the buyer wins the auction, pays  $x_{t+d}$ , and gets payoff  $x - x_{t+d}$ . If  $x_{t+d} = b(x)$ , then again both bids give the same payoff: with a bid of  $b$ , the buyer wins and gets payoff  $x - b(x) = (1 - \alpha)V(x; \sigma, p)$ . A bid of  $b(x)$  may win or lose, but the payoff is  $x - b(x) = (1 - \alpha)V(x; \sigma, p)$  either way. Otherwise (if  $x_{t+d} > b(x)$ ), let

$$\underline{s} \equiv \min \{s \in \{t, \dots, t+d\} \mid x_s > b(x)\},$$

and let  $\{x_t, \dots, x_{\underline{s}}, x'_{\underline{s}+1}, \dots, x'_{t+d}\}$  denote the realized increasing sequence of highest competing bids if the buyer submits a bid of  $b$ . (Note that the sequence is the same as under  $b(x)$  up until the first period that a competing bid strictly exceeds  $b(x)$ ; up until then the observable second highest bid is the same.) In this case, a bid of  $b(x)$  loses, and the buyer gets continuation payoff  $(1 - \alpha)V(x; \sigma, p)$ . A bid of  $b$  gives a weakly lower payoff: if  $x'_{t+d} > b$ , the buyer loses and gets  $(1 - \alpha)V(x; \sigma, p)$ . If  $x'_{t+d} \in (b(x), b)$ , then the buyer wins and gets payoff  $x - x'_{t+d} < x - b(x) = (1 - \alpha)V(x; \sigma, p)$ . Finally, if  $x'_{t+d} = b$ , then the buyer may either lose and get payoff  $(1 - \alpha)V(x; \sigma, p)$  or win and get payoff  $x - x'_{t+d} = x - b < (1 - \alpha)V(x; \sigma, p)$ . Thus, bidding  $b(x)$  always gives a weakly higher payoff than bidding  $b > b(x)$  and sometimes a strictly higher payoff.

Finally, consider a bid  $b < b(x)$ . If  $x_{t+d} < b$ , then the outcome is the same as with a bid of  $b(x)$ : the buyer wins the auction, pays  $x_{t+d}$ , and gets payoff  $x - x_{t+d}$ . Otherwise, bidding  $b(x)$  gives a weakly higher payoff than bidding  $b$ . If  $x_{t+d} = b$ , then by submitting  $b(x)$  the buyer wins, pays  $b$ , and gets payoff  $x - b > x - b(x) = (1 - \alpha)V(x; \sigma, p)$ . With a bid of  $b$ , the buyer may win and get payoff  $x - b$ , but also may lose and get only the continuation payoff  $(1 - \alpha)V(x; \sigma, p)$ . Finally, if  $x_{t+d} > b$ , then a bid of  $b$  loses, and the buyer gets continuation payoff  $(1 - \alpha)V(x; \sigma, p)$ . A bid of  $b(x)$  gives a weakly higher payoff: if  $x_{t+d} \in (b, b(x))$ , then the buyer wins and gets payoff

$x - x_{t+d} > x - b(x) = (1 - \alpha)V(x; \sigma, p)$ . If  $x_{t+d} \geq b(x)$ , then win or lose a bid of  $b(x)$  gives the buyer a payoff of  $(1 - \alpha)V(x; \sigma, p)$ . Thus, bidding  $b(x)$  always gives a weakly higher payoff than bidding  $b < b(x)$  and sometimes a strictly higher payoff.

The arguments above generalize to show that for any bids  $b'', b' \in \mathcal{B}$  such that either  $b'' > b' \geq b(x)$  or  $b(x) \geq b' > b''$ , bidding  $b'$  weakly dominates bidding  $b''$ . Thus, if bidding exactly  $b(x)$  is not feasible – that is, if  $b(x) \notin \mathcal{B}$  – then any bids other than the closest feasible bids just below and above  $b(x)$  are weakly dominated.

## E Proof of Theorem 1

We first construct the strategy profile and then show that it is an  $\epsilon$ -equilibrium. By construction, the strategies satisfy condition (i) of the theorem (each type uses a constant bid). The last step is to show condition (ii), that the bids are increasing in the bidder's type.

### E.1 Constructing the strategy $\sigma_{M,J}^*$

For any integer  $M \geq 1$ , define an  $M$ -horizon strategy profile as one that specifies that each type of bidder at each state 1) chooses to submit a bid in one of the  $M$  next-to-close auctions, and 2) bases his choice of auction and bid only on the observable state of those  $M$  auctions. That is, an  $M$ -horizon bidding strategy profile is a mapping from  $\mathcal{X} \times \mathcal{B}^M \times \{1, \dots, T\}$  to  $\mathcal{B} \times \{1, \dots, M\}$ .

Pick an  $M$ , and let  $J$  be greater than  $M$ . We will construct an  $M$ -horizon constant bidding strategy profile  $\sigma_{M,J}^*$  by looking for a fixed point: a payoff function together with a strategy profile for other types and beliefs about bids will determine the strategy for each type, which will in turn determine steady-state payoffs and beliefs.

First, define an arbitrary function  $v_0 : \mathcal{X} \rightarrow (0, \bar{x}]$ , representing the payoff for each type. Given  $v_0$ , define the net value  $x^{net}(x, v_0) \equiv x - (1 - \alpha)v_0(x)$  for each type; the bid that a buyer of type  $x$  will submit,  $b(x, v_0)$ , is the closest feasible bid to  $x^{net}(x, v_0)$ . Next, define an arbitrary conditional belief system  $p_0$  with the property that beliefs about the vector of highest bids  $w$  in the  $M$  next-to-close auctions depend only on the vector of highest losing bids  $r$  in those auctions. (When players use  $M$ -horizon strategies, later auctions will not have any bids yet.) Finally, define a one-period-ahead belief function  $r_0 : \mathcal{B}^M \times \{1, \dots, T\} \rightarrow \Delta\mathcal{B}^M$  that specifies beliefs about what next period's vector of  $M$  highest losing bids will be as a function of this period's vector.

Given  $v_0$ ,  $p_0$ , and  $r_0$ , we construct the corresponding consistent  $M$ -horizon strategy profile  $\sigma_{M,J}(v_0, p_0, r_0)$  recursively as follows. Bidders arrive and see 1) the number of periods remaining in the next to close auction  $d$ , and 2) the vector  $r$  of highest losing bids for each of the next  $M$  auctions. We start with behavior in the next-to-close auction ( $j = 1$ ), first when it has one period remaining ( $d = 1$ ) and then for higher values of  $d$ . For each  $d$  we specify behavior as the equilibrium of a particular simultaneous-move game. Then we do the same for the next auction ( $j = 2$ ), and so on up to  $j = M$ .

- Step  $j = 1$ ,  $d = 1$ : Start with auction 1, the next to close, with one period remaining.

Arriving buyers have beliefs given by  $p_0(r, d = 1)$  about the highest standing bid in that auction,  $w_1$ . Define a hypothetical static game where the (random) set of players equals the buyers who arrive in that period, and a player of type  $x$  can either 1) submit a bid of  $b(x, v_0)$  in auction 1, or 2) not bid. The payoffs of the hypothetical game to a type- $x$  player are as follows: the payoff to not bidding is  $v_0(x)$ . If he submits a bid, and  $b(x, v_0)$  is higher than any competing bid (from among the bids by other players arriving that period, plus the realized standing high bid  $w_1$  drawn from  $p_0(r, d = 1)$ ), then he gets the usual auction payoff:  $x$  minus the highest competing bid. If he submits a bid that is not the highest, then he gets a payoff of  $(1 - \alpha)v_0(x)$ . Ties are broken as in the Model section. A Nash equilibrium of that hypothetical game exists – use it to define for each type of buyer the probability of bidding in auction 1 when  $d = 1$  under  $\sigma_{M,J}(v_0, p_0, r_0)$ . (If the equilibrium is not unique, select one arbitrarily.)

- Step  $j = 1, d = 2$ : Next we define the probability of bidding in the next-to-close auction for a buyer who arrives when that auction has two periods remaining. Arriving buyers have beliefs given by  $p_0(r, d = 2)$  about the highest standing bid in auction 1,  $w_1$ . Define another hypothetical static game where, again, the set of players are the buyers who arrive this period, and a type- $x$  player can either 1) submit a bid of  $b(x, v_0)$  in auction 1, or 2) not bid and get a sure payoff of  $v_0(x)$  instead. The difference from the previous step is the function mapping bids to payoffs: because auction 1 will still be open in the next period, the behavior of bidders arriving then will affect payoffs. In defining the expected payoffs of this step’s hypothetical static game, we treat the behavior of next period’s arriving buyers as exogenous. Specifically, next period’s buyers will submit bids in this auction according to the strategies determined in Step  $j = 1, d = 1$ , so the players in this period need to have beliefs about what the vector of highest losing bids,  $r'$ , will be next period. The new second highest bid in auction 1 will be determined by the actions of the current bidders, together with the standing high bid  $w_1$ . Beliefs about the new second highest bids in the other auctions are given by the one-period-ahead belief function  $r_0(r, d = 2)$ . After next period’s bids are submitted, auction 1 closes, and this period’s bidders get the corresponding realized payoffs: the winning bidder gets a payoff equal to his type  $x$  minus the highest competing bid, and a losing bidder of type  $x$  gets payoff  $(1 - \alpha)v_0(x)$ . An equilibrium of this hypothetical game exists – use it to define the probabilities of bidding in auction 1 when  $d = 2$  under  $\sigma_{M,J}(v_0, p_0, r_0)$ .
- Steps  $j = 1, d = 3$  through  $d = T$ : We iterate the process above in order to specify the probability of bidding in the next-to-close auction for a buyer who arrives when that auction has three or more periods remaining. Arriving bidders have beliefs given by  $p_0(r, d)$  about the highest standing bid in auction 1,  $w_1$ . Define a hypothetical static game where, again, the set of players are the buyers who arrive this period, and a type- $x$  player can either 1) submit a bid of  $b(x, v_0)$  in auction 1, or 2) not bid and get a sure payoff of  $v_0(x)$  instead. Buyers arriving in the remaining periods before auction 1 closes will submit bids in this auction according to the strategies determined in the previous steps. Next period’s second highest bid in auction 1 will be determined by the actions of the current bidders, together with the standing high bid  $w_1$ . Beliefs about next period’s second highest bids in the other auctions are given by the one-period-ahead belief function  $r_0(r, d)$ . The actions of next period’s buyers will then determine the new second highest bid in auction 1 in the period after that, and beliefs about that period’s new second highest bids in the other auctions are given by applying the one-period-ahead belief function  $r_0(r', d - 1)$  to next period’s vector of highest losing bids,  $r'$ . Continue that process to predict future bids in auction 1 until the auction closes, and

this period's bidders get the corresponding realized payoffs: the winning bidder gets a payoff equal to his type  $x$  minus the highest competing bid, and a losing bidder of type  $x$  gets payoff  $(1 - \alpha) v_0(x)$ . An equilibrium of this hypothetical game exists – use it to define the probabilities of bidding in auction 1 for each  $d$  between 3 and  $T$  under  $\sigma_{M,J}(v_0, p_0, r_0)$ .

Under the strategies described above, some buyers choose not to submit a bid in auction 1. For those buyers, we next specify in a similar way their decision whether or not to participate in auction 2, the second in line to close.

- Step  $j = 2, d = 1$ : Start with the case where auction 1 has one period remaining. Arriving bidders have beliefs given by  $p_0(r, d = 1)$  about the highest standing bid in that auction,  $w_2$ . Think of the hypothetical static game where the set of players are those who 1) arrive this period and 2) in the equilibrium of Step  $j = 1, d = 1$  chose not to bid. (This set of players is doubly random. The initial arrival of buyers is exogenously random, and then the strategies of whether or not to bid in auction 1 are potentially mixed.) In this hypothetical game too, each type- $x$  player can either submit a bid of  $b(x, v_0)$  in auction 2, or not bid and get a payoff of  $v_0(x)$ . In the next period, this auction will become the next-to-close with  $d = T$ , so we can determine the expected payoffs of the hypothetical game as we did for Steps  $j = 1, d = 3$  through  $d = T$ . An equilibrium of this hypothetical game exists – use it to define the probabilities of bidding in auction 2 when  $d = 1$  under  $\sigma_{M,J}(v_0, p_0, r_0)$ .
- Steps  $j = 2, d = 2$  through  $d = T$ : Next we define the probability of bidding in auction 2 for a buyer who arrives when auction 1 has two periods remaining. Arriving bidders have beliefs given by  $p_0(r, d = 2)$  about the highest standing bid in that auction,  $w_2$ . Define a hypothetical static game where, again, the set of players are those who 1) arrive this period and 2) in the equilibrium of Step  $j = 1, d = 2$  chose not to bid. Again, each type- $x$  player can either submit a bid of  $b(x, v_0)$  in auction 2, or not bid and get a payoff of  $v_0(x)$ . Buyers arriving in the remaining periods before this auction closes will submit bids in this auction according to the strategies determined in the previous steps. Next period's new second highest bids in auctions 1 and 2 will be determined by the actions of the current bidders, together with the standing high bids  $w_1$  and  $w_2$ . Beliefs about next period's second highest bids in the other auctions are given by the one-period-ahead belief function  $r_0(r, d = 2)$ . The actions of next period's buyers will then determine the new second highest bids in auction 1 (until it closes) and auction 2 in the period after that, and beliefs about that period's new second highest bids in the other auctions are given by applying the one-period-ahead belief function  $r_0(r', d - 1)$  to next period's vector of highest losing bids,  $r'$ . Continue that process to predict future bids in auction 2 until the auction closes, and this period's bidders get the corresponding realized payoffs. An equilibrium of this hypothetical game exists – use it to define the probabilities of bidding in auction 2 for each  $d$  between 2 and  $T$  under  $\sigma_{M,J}(v_0, p_0, r_0)$ .

Finally, we repeat that process to construct the probabilities of bidding in each auction 3 through  $M - 1$  for each  $d$  between 1 and  $T$  under  $\sigma_{M,J}(v_0, p_0, r_0)$ . For each  $d$ , let  $\sigma_{M,J}(v_0, p_0, r_0)$  specify that any buyer who does not submit a bid in one of the first  $M - 1$  bids in auction  $M$ .

Those  $T \times (M - 1)$  steps tell us how to construct  $\sigma_{M,J}(v_0, p_0, r_0)$  given values  $v_0$ , conditional beliefs  $p_0$ , and one-period-ahead belief function  $r_0$ . We then look for a fixed point. Given a strategy  $\sigma_{M,J}$ , on-path steady-state conditional beliefs are pinned down by Bayes' rule, and off-path beliefs can

be specified in a consistent way so that 1) beliefs about the vector of highest bids  $w$  in the  $M$  next-to-close auctions depend only on the vector of highest losing bids  $r$  in those auctions, and 2) beliefs assign probability one to an off-path bid being the minimal feasible bid compatible with the observable state. Given a strategy  $\sigma_{M,J}$  and conditional beliefs  $p_{M,J}$ , the value of the one-period-ahead belief function is pinned down at each observable state; call the result  $r(\sigma_{M,J}, p_{M,J})$ . The strategy  $\sigma_{M,J}$  also determines the expected payoff in steady state to a bidder of each type  $x$ ; call that function  $v(\sigma_{M,J})$ . Then the  $M$ -horizon constant bidding strategy profile  $\sigma_{M,J}^*$  that we want is one that satisfies the following property: we can find a payoff function  $v_{M,J}^*$ , conditional beliefs  $p_{M,J}^*$ , and a one-step-ahead function  $r_{M,J}^*$  such that

1.  $\sigma_{M,J}^* = \sigma_{M,J}(v_{M,J}^*, p_{M,J}^*, r_0)$ ;
2.  $v_{M,J}^* = v(\sigma_{M,J}^*)$ ;
3.  $p_{M,J}^*$  is consistent with  $\sigma_{M,J}^*$ ; and
4.  $r_{M,J}^* = r(\sigma_{M,J}^*, p_{M,J}^*)$ .

A fixed point argument paralleling the proof of Proposition 2 establishes that such a  $\sigma_{M,J}^*$  exists for any  $M$  and  $J$ .

## E.2 Showing that $\sigma_{M,J}^*$ is an $\epsilon$ -equilibrium

We next establish that if  $M$  is large enough given any  $\epsilon > 0$ , and  $k$  is large enough given  $\epsilon$  and  $M$  (so that  $\gamma_k$  is close to 0 and  $J_k$  is large), then the  $\sigma_{M,J_k}^*$  defined above is an  $\epsilon$ -equilibrium. Let  $v_{M,J_k}^*$  and  $p_{M,J_k}^*$  be the corresponding payoff function and conditional belief system. In Lemma 2, we show that conditional on choosing one of the next  $M$  auctions, a bid of  $b(x, v_{M,J_k}^*)$  is nearly optimal for a buyer of type  $x$ . Lemma 3 completes the proof by showing that choosing one of the  $M$  next-to-close auctions is in fact nearly optimal with very high probability under the steady state distribution.

As a preliminary, we note that it is without loss of generality to assume that the sequences

$$\{\sigma_{M,J_k}^*, p_{M,J_k}^*, v_{M,J_k}^*\}_{k=1}^{\infty}$$

are such that the transition probabilities over the state of the next  $M$  auctions (bids and high bidders' types) converge. The set of  $M$ -horizon constant bidding strategy profiles is compact, as is the set of beliefs over the state of the next  $M$  auctions; any sequence of those strategy profiles and beliefs thus has a convergent subsequence. Along that subsequence the behavior of each type of buyer converges. To show that the transition probabilities converge, we also need to show that the arrival rates of each type of buyer converge. The arrival rates of new buyers are fixed with respect to  $k$  by assumption. Finally, although the distribution over the size of the losers' pool  $n$  does not converge, we do get convergence (at least along a subsequence) if we normalize the numbers of each type in the losers' pool,  $n(x)$ , to  $\gamma_k n(x)$ . Let  $\bar{n}_k$  denote the steady-state size of the losers' pool. Then the steady-state expected number of returning losers each period  $\bar{n}_k \gamma_k \Delta$  is bounded above

by  $\lambda\Delta(1-\alpha)/\alpha$ , because over time the inflow into the losers' pool comes from new buyers who lose an auction and do not exit (rate bounded above by  $\lambda(1-\alpha)$ ), while the outflow comes from those losers who return and either win an auction or lose and exit (rate at least  $n\gamma_k\alpha$ ). Given that normalization, conditional beliefs at on-path observable states converge as well.

We now state and prove the two lemmas. Define

$$\hat{v}(j, b; \tilde{\omega}, \sigma, p) \equiv \left[ \begin{aligned} & \sum_{m \in \{0, \dots, b\}} (x - m) \cdot g_{\sigma, p}(m; \tilde{\omega}, j, b) \\ & + (1 - \alpha) \left( 1 - \sum_{m \in \{0, \dots, b\}} g_{\sigma, p}(m; \tilde{\omega}, j, b) \right) V(x; \sigma, p) \end{aligned} \right]$$

as the expected payoff to submitting bid  $b$  in auction  $j$  at observable state, given  $(\sigma, p)$ .

**Lemma 2** *Fix  $M$ , and pick any observable state  $\tilde{\omega}$  with the feature that only the  $M$  next-to-close auctions have received bids. Then for any type  $x$  and auction  $j \leq M$ , we have*

$$\lim_{k \rightarrow \infty} \left| \max_{b \in \mathcal{B}} \hat{v}(j, b; \tilde{\omega}, \sigma_{M, J_k}^*, p_{M, J_k}^*) - \hat{v}(j, b(x, v_{M, J_k}^*); \tilde{\omega}, \sigma_{M, J_k}^*, p_{M, J_k}^*) \right| = 0.$$

**Proof.** We first show that given such an observable state  $\tilde{\omega}$ , in the limit the expected re-entry value is independent of the buyer's choices of  $b$  and  $j$  and also independent of  $\tilde{\omega}$ . It therefore equals the unconditional expectation,  $v_{M, J_k}^*(x)$ . That is, we show that for any type  $x$ , bid  $b$ , and auction  $j \leq M$ , we have

$$\lim_{k \rightarrow \infty} \left| EV(x, \tilde{\omega}; \sigma_{M, J_k}^*, p_{M, J_k}^*, j, b) - v_{M, J_k}^*(x) \right| = 0,$$

where

$$EV(x, \tilde{\omega}; \sigma, p, j, b) \equiv \frac{\sum_{\omega^l \in \Omega} V(x, \omega^l; \sigma, p) h_{\sigma, p}(\omega^l; \tilde{\omega}, j, b)}{\sum_{\omega^l \in \Omega} h_{\sigma, p}(\omega^l; \tilde{\omega}, j, b)}$$

is the expectation of the re-entry payoff conditional on having submitted bid  $b$  in auction  $j$  at observable state  $\tilde{\omega}$ , losing, and entering the losers' pool.

As  $\gamma_k \rightarrow 0$ , the buyer's return time is arbitrarily far in the future with arbitrarily high probability, so the probability of the event that no buyers arrive as  $M$  consecutive auctions go by before the buyer returns goes to 1. That event implies that the choice of  $b$  and  $j \leq M$  can have no further effect on the observable state (no active bid has been placed by a bidder who saw the buyer's choices, or by a bidder who saw a bid placed by a bidder who saw the buyer's choices, and so on) and thus can have no effect on the actions of other buyers.

As noted above, when  $k$  grows the Markov chain over the state space with the normalized size of the losers' pool converges, and in particular conditional beliefs at on-path observable states and the transition probabilities over the state of the next  $M$  auctions (bids and high bidders' types) converge. Because that limit process is ergodic (Proposition 1) and the number of on-path observable states is finite, if the number of periods before the buyer re-enters is high enough, then beliefs over what the state at re-entry will be are close to the stationary distribution, conditional on any observable state when the buyer enters the losers' pool. As  $\gamma_k \rightarrow 0$ , the re-entry time is very high with very high probability. Therefore, for high enough  $k$  the distribution of the re-entry



state and thus the expected re-entry payoff are independent of the observable state  $\tilde{\omega}$  when the buyer chooses  $b$  and  $j$ .

Because the expected continuation value conditional on losing for a buyer of type  $x$  is very close to  $(1 - \alpha) v_{M, J_k}^*(x)$ , regardless of which of the next  $M$  auctions the buyer chooses or which bid he submits, the arguments of Proposition 3 imply that  $b(x, v_{M, J_k}^*) \approx x - (1 - \alpha) v_{M, J_k}^*(x)$  is an optimal (or nearly optimal) bid to submit in any of the next  $M$  auctions. ■

Lemma 2 establishes that because a buyer who has chosen one of the first  $M$  auctions faces an approximately constant continuation value after losing, an approximately optimal bid in any of those auctions is his value minus his continuation value. In Lemma 3, we show that one of those  $M$  auctions is nearly always a nearly optimal choice. By construction of the strategy profile  $\sigma_{M, J_k}^*$ , a type- $x$  buyer who chooses one of the first  $M - 1$  auctions expects to get at least  $v_{M, J_k}^*(x)$ . The only way to get less is if he turns down all of the first  $M - 1$  auctions; in that event, under  $\sigma_{M, J_k}^*$  he bids in auction  $M$  even if it yields a low expected payoff. That event, though, is very unlikely for large  $M$ . (Note that Lemma 2 still applies in that event: the bid  $b(x, v_{M, J_k}^*)$  is optimal in auction  $M$  even when auction  $M$  is a sub-optimal choice.)

We will use the following notation in the formal statement of Lemma 3: given a scalar  $\eta > 0$  and an  $M$ -horizon constant bidding strategy profile  $\sigma$  with corresponding payoff function  $v$ , let  $\Omega^\eta(\sigma, v)$  be the set of states at which for each buyer type  $x$ , 1) playing according to  $\sigma$  gives an expected payoff no lower than  $v(x) - \eta$ , and 2) submitting a bid in auction  $M + 1$  or later gives an expected payoff no higher than  $v(x) + \eta$ . Let  $\pi(\Omega^\eta(\sigma, v) | \sigma)$  denote the steady-state probability of state in  $\Omega^\eta(\sigma, v)$  under  $\sigma$ .

**Lemma 3** *For any  $\eta > 0$ ,*

$$\lim_{M \rightarrow \infty} \lim_{k \rightarrow \infty} \pi(\Omega^\eta(\sigma_{M, J_k}^*, v_{M, J_k}^*) | \sigma_{M, J_k}^*) = 1.$$

**Proof.** First consider the hypothetical static game in Step  $j = 1$ ,  $d = 1$ , where at observable state  $\tilde{\omega}$  a type- $x$  buyer decides between submitting a bid of  $b(x, v_{M, J_k}^*) \approx x - (1 - \alpha) v_{M, J_k}^*(x)$  in an auction that will close at the end of the period or getting a sure payoff of  $v_{M, J_k}^*(x)$ . That game corresponds to the following second-price auction environment: the same random set of bidders, but a bidder of type  $x$  gets gross payoff  $x^{net}(x, v_{M, J_k}^*) = x - (1 - \alpha) v_{M, J_k}^*(x)$  from winning the auction, and has an outside option of  $\alpha v_{M, J_k}^*(x)$  from not participating. (That is, both the payoff from winning and the value of the outside option are measured as the surplus over the continuation value  $(1 - \alpha) v_{M, J_k}^*(x)$ .) The auction has a hidden reserve price equal to the standing high bid  $w_1$ , distributed according to  $p_{M, J_k}^*(\tilde{\omega})$ .

Equilibrium behavior in the hypothetical game thus is equivalent to the outcome of that auction: bidders choose optimally whether or not to participate in the auction, and if they participate they bid their valuation. From Myerson (1981), we know that the expected payoff from an auction to a bidder with valuation  $v$  is given by the integral up to  $v$  of the probability of winning for each bid,  $P(x^{net})$ . In the hypothetical static game, then, the payoff to a type- $x$  player who chooses to bid

can be written as

$$(1 - \alpha) v_{M, J_k}^*(x) + \int_0^{x^{net}(x, v_{M, J_k}^*)} P(x^{net}) dx^{net}.$$

In the auction, the probability that a bidder wins equals the probability that his net value exceeds the standing high bid  $w_1$  and that no other bidder with a higher net value participates (with a small adjustment for the possibility of a tie).

The hypothetical static games constructed in the other steps are similarly equivalent to auctions. That equivalence is less obvious when a buyer decides whether or not to participate in an auction that will not close until a later period, because in principle the buyer's bid could influence the entry choices of buyers who arrive in the future. In fact, though, because only the highest losing bid is observed, his bid can influence future behavior only in the case where he has already lost. Thus, the set of competitors that the buyer will face is effectively exogenous. In each step of the definition of  $\sigma_{M, J_k}^*$ , the expected payoff from submitting a bid in any of the first  $M - 1$  auctions is a function of the probability that the bid exceeds the standing high bid and that no bidder with a higher net value enters that auction.

We can now establish that for each type  $x$ , playing according to  $\sigma_{M, J_k}^*(x)$  gives at least (close to)  $v_{M, J_k}^*(x)$  with high probability in the limit. The strategy profile  $\sigma_{M, J_k}^*$  specifies that a type- $x$  buyer bids with positive probability in the earliest auction  $j$  that gives him an expected payoff of at least  $v_{M, J_k}^*(x)$  (as long as there is such an auction among the next  $M - 1$  auctions). Submitting a bid in auction  $j$  lowers the expected value of that auction for future buyers, all else equal, and so lowers their equilibrium probability of participating in auction  $j$  and pushes them toward later auctions. Thus, under  $\sigma_{M, J_k}^*$ ,  $v_{M, J_k}^*(x)$  is the long-run average payoff to a type- $x$  buyer, and he expects to get at least  $v_{M, J_k}^*(x)$  if he participates in one of the first  $M - 1$  auction. He can get less only if all  $M$  auctions would yield an expected payoff below  $v_{M, J_k}^*(x)$  – that is, if the expected distribution of the numbers of bidders of each type is worse than average in all  $M - 1$  auctions. But when  $M$  is large, the probability of so many deviations from the long-run averages of arrival rates and mixed strategy auction choices is extremely low. Thus, with very high probability the state when the buyer arrives is such that playing according to  $\sigma_{M, J_k}^*$  gives an expected payoff of at least  $v_{M, J_k}^*(x)$ .

Similarly, as  $M$  grows and the number of potential bidders in each of the  $M$  auctions becomes large, the same type of arbitrage implies that a buyer  $i$  of type  $x$  would be unlikely to get a payoff much above the average payoff  $v_{M, J_k}^*(x)$  by submitting a bid in an auction later than specified by  $\sigma_{M, J_k}^*$ . (Participating in an earlier auction gives a payoff below  $v_{M, J_k}^*(x)$  by the construction of  $\sigma_{M, J_k}^*$ .) If buyer  $i$  participates in an auction that should not under  $\sigma_{M, J_k}^*$  receive a bid, then buyers in the next period see that a bid has been placed but do not observe the amount of the bid. Recall that the conditional belief system  $p_{M, J_k}^*$  assigns probability one to that unobserved bid being the lowest amount feasible, and so the off-path bid does not deter subsequent buyers from participating in that auction. Whether the deviation to a later auction is on-path or off-path, then, as future buyers follow  $\sigma_{M, J_k}^*$ , the probability that buyer  $i$  winds up with a payoff above  $v_{M, J_k}^*(x)$  shrinks to zero. ■

### E.3 Showing that bids are increasing in type

To complete the proof of Theorem 1, we need to show that condition (ii) is satisfied; that is, that  $x - (1 - \alpha) v_{M, J_k}^*(x)$  is increasing in  $x$ . It is sufficient to show that for high enough  $k$ , the derivative of  $v_{M, J_k}^*(x)$  is less than  $1/(1 - \alpha)$ ; more precisely, that for types  $y > x$ ,  $v_{M, J_k}^*(y) - v_{M, J_k}^*(x) < (y - x)/(1 - \alpha)$ .

The argument is standard. A buyer's expected payoff in the dynamic game equals his type times the probability that he eventually wins an auction, minus the expected price that he pays conditional on winning. Let  $q_k(x)$  denote the steady-state probability that a buyer who plays the strategy  $\sigma_{M, J_k}^*(x)$  (that is, the strategy of a type- $x$  buyer) eventually wins an auction, given that all other buyers play according to  $\sigma_{M, J_k}^*$ . Similarly, let  $t_k(x)$  denote the expected payment of such a buyer. Note that neither  $q_k(\cdot)$  nor  $t_k(\cdot)$  depends on the buyer's type – they depend only on his strategy.

Using that notation, we can write

$$v_{M, J_k}^*(x) = x \cdot q_k(x) - t_k(x).$$

Let  $\bar{\epsilon} \equiv \min\{|x'' - x'| : x', x'' \in \mathcal{X}\}$ , and pick an  $\epsilon' \in \left(0, \frac{\alpha}{1-\alpha} \bar{\epsilon}\right)$ . Because for high enough  $k$  playing according to  $\sigma_{M, J_k}^*$  is an  $\epsilon'$ -best response for all types, we have that

$$\begin{aligned} v_{M, J_k}^*(x) &= x \cdot q_k(x) - t_k(x) \\ &\geq x \cdot q_k(y) - t_k(y) - \epsilon' \\ &= [x - y] \cdot q_k(y) + y \cdot q_k(y) - t_k(y) - \epsilon' \\ &= [x - y] \cdot q_k(y) + v_{M, J_k}^*(y) - \epsilon'. \end{aligned}$$

Because  $q_k(\cdot) \leq 1$ , we get  $v_{M, J_k}^*(y) - v_{M, J_k}^*(x) \leq y - x + \epsilon'$ . Because  $\epsilon' < \frac{\alpha}{1-\alpha} \bar{\epsilon}$ , we conclude that  $v_{M, J_k}^*(y) - v_{M, J_k}^*(x) < (y - x)/(1 - \alpha)$ , as desired.

## F 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

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

and the optimal bid function is

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

The ex ante value function is given by the function

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

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

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

The inverse bid function is

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

and the value equation becomes

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

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

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

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  $F_C$ .

## G 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.

In Section 6.3 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 (6). 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 (6). The results, shown in Figure 9, are clearly worse than those in Figure 4, which is based on the valid

Figure 7: Time between bids


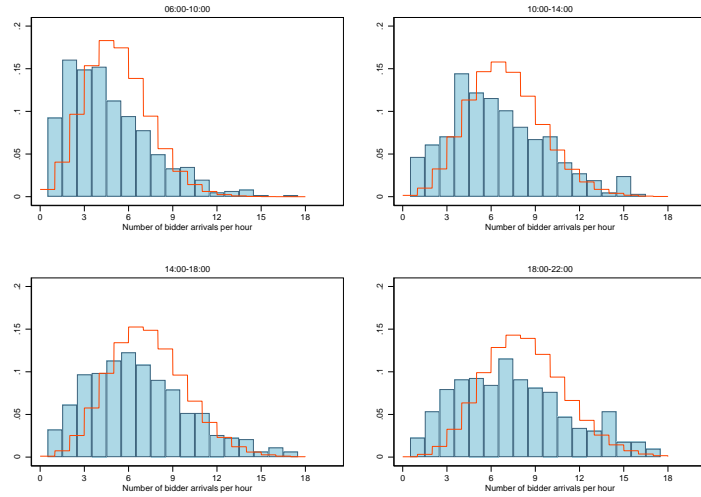


Figure 8: Bidder arrivals per hour




estimate of  $G_{M|B}$ .

## H 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

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



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

and

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

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 (H) and (H). 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 (H). 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  $x$ .