Dynamics and Efficiency in Decentralized Online Auction Markets

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

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

Thomas Wiseman
University of Texas
wiseman@austin.utexas.edu

April 2024

Abstract

Economic theory suggests that decentralized markets can achieve efficiency if agents have many opportunities to trade. We examine this idea by using detailed eBay data to estimate a dynamic model of bidding in overlapping, sequential auctions. The matching of buyers to sellers in our model is in real time and endogenous: buyers use posted information about the state of play in each auction when choosing where and how much to bid. We prove that, in thick markets, a bidder's choices have vanishing influence on her re-entry payoff (conditional on losing), so optimal bids are invariant to the state of play and increasing in type. Given this result, our model is identified from bidding data. Our estimator accounts for the selection arising from auction choice and avoids the substantial underestimate of bidders' valuations that would occur under random matching. We find that outcomes fall well short of the efficient allocation. Sealed bid auctions can improve efficiency, but we show that a posted price mechanism can implement the efficient allocation in real time and does so in our data.

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. The 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. Trade in these types of markets may be inhibited by several frictions. The matching of buyers to sellers may be inefficient because it is based on arrival times instead of values, and trade within a matched set of traders may be inefficient due to private information and strategic behavior. However, the opportunity to trade many times can make the market for each trader large over time, and this feature can mitigate the effects of the matching and trading frictions. For example, Satterthwaite and Shenyerov (2007, 2008) show that equilibrium prices in their model of a decentralized, dynamic auction market converge to the Walrasian price and allocations converge to the efficient allocation as the number of opportunities to trade gets large. This theoretical result raises important empirical questions about real-world markets like eBay: How efficient are these markets? And what can the platform do to increase market efficiency and/or prices?

We address the first question by developing and estimating a structural model of the eBay market. Our model captures key features of the eBay marketplace. 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 have private values, arrive randomly over time and, upon arrival, observe the information that the platform provides about the state of play in each auction, namely, the highest losing bid (or start price if there are no bids) and time remaining. Buyers use this information to decide in which auction to bid and how much to bid. Winners exit, and buyers who lose (and sellers who fail to sell) either exit or return at some time in the future to try again.

The distinguishing feature of our model is that the matching of buyers to sellers is determined by the state of play in the auctions rather than by random assignment. It also presents the main challenge in taking the model to data. In principle, a forward-looking buyer needs to consider how his auction choice and bid can affect the decisions of subsequent bidders and, as a result, his re-entry payoff, if and when he returns to try again. But, for the model to be empirically tractable and useful, these considerations must be unimportant. We provide conditions under which this is the case. We prove that, as market thickens, arbitrage causes many states to have nearly the same value and, as a result, the buyer's continuation value depends only on her type, and not on her previous actions or on the losing state. Intuitively, in thick markets, the state is likely to undergo many transitions before a losing buyer returns, so the effects of the losing state and of the buyer's past actions on his expected re-entry payoff have largely washed away. We then prove that, in our setting, it is nearly optimal (in the ϵ-equilibrium sense) for each buyer to always bid her type less 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 invariance and monotonicity properties of the equilibrium bid function are crucial to our identification and estimation strategy. Invariance means that the distribution of buyer values can be identified from bids without solving for the equilibrium auction choice rule or imposing assumptions about that choice rule. Monotonicity implies that a bidder's type can be equivalently be represented by her equilibrium bid, or what Backus and Lewis (2023) 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 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 maximum rival bid arises primarily from the fact that buyers choose auctions based on the observed state of play. 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.

We use estimates of the model primitives to quantify the efficiency of the eBay trading mechanism and compare it to two hypothetical benchmarks. One is the Walraisan equilibrium, 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 static market in which sellers hold separate, second-price auctions, buyers are randomly assigned to those auctions, and each buyer gets only one chance to win a unit.[1] We find that the eBay mechanism meaningfully decreases price dispersion and increases efficiency relative to this second benchmark but falls well short of the efficient outcome. Prices in our data are quite dispersed, and the average price is significantly lower than the market-clearing price.

Should platforms like eBay provide information about the state of play? One might expect real-time disclosure of the highest losing bids in auctions to improve market efficiency since it allows buyers to avoid bidding in auctions where they have already been outbid and to arbitrage differences in expected payoffs across auctions. We examine this issue by considering a counterfactual in which eBay posts the closing time of each auction but not the highest losing bid. In other words, each auction is a sealed bid auction.[2] In this case, there is an equilibrium in our model where buyers always choose the soonest-to-close auction.[3] We simulate this auction choice rule on our data, taking as fixed the arrival times of sellers and new buyers. Contrary to expectations, we find that *not* posting the highest losing bid meaningfully increases efficiency, although it still falls well short of the efficient allocation. The intuition here is that real-time disclosure of the highest losing bid makes the auctions more competitive by disproportionately matching two high-value bidders, which in turn makes it more likely that high value buyers lose and exit without winning a unit.

Our analysis makes clear that the matching process in real-time, dynamic auction markets like eBay is an important source of inefficiency. Because it is based on arrival times, too many low value buyers win units at the expense of high value buyers. The uniform price auction solves this problem by pooling across time and then matching buyers to sellers based on their values. However, buyers may not want to wait until the platform conducts the auction to get their units This raises the question of alternative mechanisms, and the obvious one to consider is posted prices. Sellers on eBay have the option to sell their goods by auction or by posted price, and increasingly over time they have chosen the latter. Einav, Farronato, Levin and Sundaresan (2018) document that the share of posted price listings on eBay have grown from roughly 10%

[1]This benchmark is motivated by the matching that occurs in brick and mortar markets where buyers have to buy from a local seller, and sellers have to sell to local buyers. Online platforms create thicker markets 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.

[2]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.

[3]Note that, under this choice rule, buyers are randomly assigned to sellers based on their arrival times. As we discuss in the next section, this is the equilibrium that has been assumed in the structural literature.

in 2003 to more than 90% in 2015.

We consider a centralized mechanism in which the platform sets a price P and requires each seller to post that price. We show that, if P is chosen appropriately, then seller revenue, average price, and total surplus all increase, and the average time till sale decreases. In fact, for large markets, the optimal dynamic posted price mechanism achieves approximately the maximum possible level of both revenue and total surplus. We simulate this price mechanism on our data, taking as given the arrival times of sellers and new buyers. We set P equal to the market-clearing price of the uniform price auction and the posting duration at seven days, the auction duration chosen by most eBay sellers in our data. We find that this mechanism successively implements the efficient allocation. Supply never falls to zero, so every buyer with a value above P purchases a unit. There are sellers who fail to sell before their posting expires, but the number is very small. In addition, transaction costs falls, with the average time till sale falling from seven days to a little more than 3 days.

To implement the optimal posted price mechanism, the platform would need to know the market-clearing price. However, this should not be too difficult in a stationary environment. We show that, at the market-clearing price, the number of seller postings follow a random walk but, at any lower price, it converges to zero and, at any higher price, to an upper bound determined by the posting duration and the arrival rate of sellers. Thus, the platform could learn the market-clearing price by monitoring the number of available postings and adjusting prices accordingly.[4]

Our environment abstracts from product differentiation, but the results provide a potential explanation for the shift away from auctions towards posted prices. The growth of commerce on eBay have thickened markets, reduced the value of auctions as a mechanism for price discovery, and made it easier for sellers to learn the market-clearing prices and post them.

2 Related Literature

The theoretical literature on large, decentralized, dynamic auctions focuses on a setting in which a continuum of sellers and a continuum of buyers simultaneously arrive each period to trade units of a homogeneous good, each buyer is randomly matched to a seller, and traders who

[4]Many markets have both a posted price and an auction segment. When this is the case, it could use data from the auction segment to learn about the market-clearing price, much as we have done in this paper.

fail to trade either exit or return the following period.[5] Using this framework, McAfee (1993) studies competing mechanisms and shows that, in steady state, there is an equilibrium in which all sellers choose to sell via second-price sealed bid auction; Satterthwaite and Shneyerov (2007, 2008) study what happens to equilibrium prices and allocations when the number of trading opportunities for each trader is large; and Bodoh-Creed, Boehnke, and Hickman (2021) show that a steady state equilibrium of this model is an ϵ-equilibrium of the analogous model with a large but finite number of buyers and sellers.

An important implication of this framework is that a bidder's option value of losing depends only on his type. This is because buyers are randomly matched to sellers and, in steady state, the actions of any single buyer or seller today has negligible impact on the state of the market tomorrow. We contribute to this literature by providing an analogous result for real-time matching markets like eBay where the matching is endogenous and losing buyers return randomly over time. We 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. Thus, 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. As a result, the expected re-entry payoff can be independent of the losing state even though the expected return state is not.[6] This is important because in our application—the eBay market for iPads—losing buyers sometimes return 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.

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.[7]

[5] 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), Nekipelov (2007).

[6] 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 random matching model.

[7] 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.

Backus and Lewis (2023) use eBay data on product choices and bids of buyers to estimate substitution patterns in a differentiated good market. Adachi (2016) and Bodoh-Creed et al (2021) estimate models of eBay auctions for homogeneous goods in order to evaluate the efficiency and prices of the eBay mechanism. These papers assume that buyers are randomly assigned to sellers based on their arrival times so that each bidder's continuation value depends only on his type.[8] In Section 9, we show that, in making these assumptions, this literature is essentially assuming that the data-generating mechanism can be approximated by an equilibrium of a model in which eBay auctions are second-price, sealed bid auctions, and losing buyers who do not return immediately.[9]

We contribute to this literature by considering a more general empirical model in which posted prices are informative and influence auction choices. We control for endogenous auction selection by estimating the probability of a type-b bidder winning an auction when he returns on the set of auctions actually chosen by type-b buyers.in the different states. This 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 of all auctions. We compute the continuation value functions associated with these two distributions and find that ignoring auction selection leads to a substantial overestimate of buyers' values, especially for high-value buyers. More formally, in applying the steady state restrictions, we find that the data rejects the random matching model. We provide further (reduced form) evidence against random matching in the next section and in the online appendix.

There is an earlier structural literature that models eBay auctions as independent, *static* games (e.g., Bajari and Hortascu (2003), Gonzalez, Hasker, and Sickles (2009), Canals-Cerda and Pearcy (2006), Ackerberg, Hirano, and Shahriar (2017) and Lewis (2011)). Our results suggest that static models may not be useful for studying issues related to auction design because they ignore the effects of endogenous matching and the option value of losing on markups. Perhaps more importantly, by ignoring the distinction between new and returning buyers, these models implicitly treat the steady state distribution of buyer values as the primitive rather than the

[8] Bodoh-Creed et al (2020) assume that bidders are assigned to the soonest-to-close auction and bid sequentially in a random order, so some bidders are censored because they are outbid before their turn. By contrast, we assume that these bidders would bid in a later-to-close auction, since there is always an auction that closes within a day that has no bids and a low start price.

[9] See also Backus and Lewis (2023).

distribution of new buyer values. This matters for counterfactuals, since changes in the auction mechanism likely result in different stationary distributions of buyer values.

Our paper is also related to the empirical literature on dynamic search-and-bargaining models of trade, such as Gavazza (2011, 2016), Brancaccio et al (2021), Buchholz (2022), and Coey, Larsen, and Platt (2020). These papers use a continuum of agents to approximate markets with finite numbers of buyers and sellers *and* steady-state restrictions on entry and exit flows as the basis for estimation. 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 Background

In this section, we briefly describe some key patterns in our data and present some evidence on how bidders are choosing auctions. The objective is to motivate several of our modeling choices and our assumption that buyers are choosing auctions endogenously based on the state of play. We provide a more thorough description of the data, including details relevant to estimation, in Section 6.

Our data consist of all eBay listings for a specific iPad model during an 8-month period in 2013. 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. end date, start price, secret reserve price, shipping options, etc.). We also observe all bids (includng the high bid) submitted for each listed item, the times at which they are submitted, and the identities of the bidders submitting them. Importantly, the latter information allows us to track repeat bidding by the same bidder.

In posting an item for auction, each seller sets a start price, which serves as a public reserve price. She also has the option of setting a secret reserve for a small additional fee, but this option is rarely used—in our data only 9% of listed items had secret reserve prices. Most sellers choose very low start prices at which they are certain to sell. Of the sellers who set binding start or secret reserve prices, only a small fraction fail to sell, and an even smaller fraction return to sell again.[10]. With this in mind, the model we present in the next section ignores the role of reserve prices and treats sellers as non-strategic players.

Our model also assumes for simplicity that the interval between seller arrivals (or equivalently auction closings) is constant, and the arrival rate of buyers is measured relative to this period.

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

In reality the arrival rates of sellers and buyers fluctuate by time of day, but they tend to vary proportionately: in the data the number of bidders per auction closing is approximately constant, as shown in Table H.1 of the online appendix. Assuming constant arrival rates is thus an useful normalization.

Buyers in our model are assumed to have unit demands, and this is largely consistent with the data: over 94% of buyers (auction winners) in the data purchased only one unit. Buyers who lose an auction can return at some future time to bid again in a later auction, and roughly half of losing bidders do this. Some return quickly—one third of returning losers come back within an hour—but the median return time is nearly 12 hours. The median number of intervening auctions that have closed in the time it takes a losing bidder to return is 15.

Our model ignores intra-auction dynamics by assuming that buyers can bid only once in whatever auction they choose. This is true for many bidders, but bidding more than once in an auction is not uncommon and, when a bidder does this in our data, we consider her bid to be the maximum bid that she submitted in the auction. We refer to repeat bidding within an auction as incremental bidding, and discuss the prevalence of this kind of behavior and its implications for our empirical analysis in more detail in Section 6.

The market we study is a thick one: an average of 23.2 items are posted for auction each day (meaning sellers arrive at an average rate of one per 62 minutes), and an average of 5.47 new bidders and 4.86 return bidders participate in each auction. At any point in time, an arriving bidder can choose from over 100 auctions that are open for bidding.

Table 1 offers some information about how bidders make this choice. The table shows summary statistics by rank, where an auction's rank is its position in the sequence of soonest-to-close auctions (so the auction with rank 1 is the next to close, rank 2 is the soonest auction to close after that one, and so on). The first column shows how posted bids decline with rank. The posted bid of an auction with rank j is the highest losing bid at the time it becomes the j^{th} auction to close. Thus, the average posted bid of an auction when it becomes next to close — which typically happens when slightly more than one hour remains in the auction — is $259.54. By contrast, the average posted bid of an auction at rank 30 (typically more than 24 hours before its closing) is below $100.

The third column gives the distribution of ranks chosen by bidders. The typical auction sees an increase in bidding near its closing time, but most of the bids are submitted well before this final phase. The fourth column gives the distribution of ranks for winning bids. Most of the winning bids are submitted toward the end of the auction — 60 percent of winning bids are submitted when the auction is next-to-close (roughly the last hour) — but a lot of serious bidding occurs at lower ranks. For example, over 13 percent of the auctions received their winning bids when they were ranked 10 or above.

Table 1: Bids and prices by auction rank

Auction rank	Average posted bid	Std. dev. posted bid	Fraction of all bids submitted	Frac. of winning bids submitted
1	259.54	59.11	.152	.601
2	228.69	81.43	.054	.099
3	215.40	87.86	.039	.056
4	204.88	92.11	.032	.036
5	196.42	94.57	.029	.024
6	188.08	96.81	.025	.019
7	181.55	98.08	.023	.012
8	175.21	99.14	.020	.011
9	169.39	100.01	.020	.008
10+	72.96	74.14	.607	.134

The declining price pattern shown in Table 1 means that a low-value buyer who arrives to the platform cannot participate in the soonest-to-close auctions, as he has effectively been outbid already. But, even beyond this mechanical effect, high-value bidders appear to have stronger preferences for soon-to-close auctions.[11] The large standard deviations reported in the second column of Table 1 indicates a lot of dispersion in posted bids. This means that when choosing in which auction to bid, buyers may engage in a form of arbitrage by favoring auctions whose current posted bids are unusually low given their position in the queue.

In summary, our data make it clear that arriving bidders do not simply bid in the auction that is next to close. In fact, most of the bidders who arrive in any given period submit their bids in auctions that are not the soonest-to-close since they have effectively been outbid already. 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 soonest-to-close auction they chose to pass up, and 9.6% submit bids that were even higher than the eventual

[11] In section J of the online appendix, we show that even after conditioning on the set of auctions available to a buyer—meaning the set of open auctions whose posted bids were below the bid the buyer submitted—high bidders are significantly more likely to choose sooner-to-close auctions.

price of the auction they passed up.

4 A Dynamic Model of Trade

Our model is a discrete approximation to a continuous-time eBay auction market. Discretizing time, values, and bids with arbitrarily defined grids means we can analyze dynamics using the theory of Markov chains.

In our model, sellers arrive exogenously at fixed intervals to sell a homogenous good. We define a unit of time to be the length 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 exit the market when their auction closes. The open auctions, starting with the next-to-close, are indexed by $j = 1, .., J$. 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 $d(t) \in \{1, .., T\}$ 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.

New buyers with unit demand arrive randomly over time. The number of new buyers arriving in a period is a random variable drawn from a Poisson distribution with mean $\lambda\Delta$. Arrivals are independent over periods. Each new buyer's value for the good is drawn independently from distribution F_E with density f_E. The distribution has finite support $\mathcal{X} \subseteq (0, \overline{x}]$ where $f_E(\overline{x}) > 0$. 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 given by $\mathcal{B} = \{0, \underline{b}, .., \overline{b}\}$ 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.[12] If his bid is the winning bid, then he gets the good, pays the second highest bid, and exits. If he loses, then he exits with probability α and gets a payoff of zero;[13] 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 who continue do not return immediately.

[12] Proxy bidding rules out intra-auction bidding dynamics such as incremental bidding.

[13] 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.

The return process is a discrete version of an exponential return rate γ. 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 $(n, \gamma\Delta)$.[14]

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 a start price of zero otherwise.[15] It does not disclose the highest bids. We call the highest losing bid (or start price if there are no bids) in an auction the *posted bid*. 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 a price equal to the posted bid. 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 posted bid in auction j in period t. The vectors of highest bids and posted bids in period t are $w(t)$ and $r(t)$ respectively. The platform only accepts bids above the posted bid, so any bids submitted to auction j in period t must be strictly greater than $r_j(t)$.

The payoff-relevant information in period t consists of the closing times of the open auctions $d(t)$, the posted bids $r(t)$ and highest bids $w(t)$ in these auctions, the values of the highest bidders, and the composition of the losers' pool. Let $a_j(t) \in \{0\} \cup \mathcal{X}$ denote the value of the high bidder in auction j ($a_j(t) = 0$ means that no one has bid) 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 by the distribution $\mathcal{N}_L(t) \in (\mathbb{Z}_+)^{|\mathcal{X}|}$, which gives the number of losers of each type in the pool. The (countable) set of states is

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

A buyer bids in the period of his arrival. When he arrives in period t, he observes $d(t)$ and $r(t)$.[16] We call $\widetilde{\omega}(t) \equiv (d(t), r(t))$ the *observed* state; $\widetilde{\Omega}$ is the set of observed states. We restrict

[14] As period length Δ shrinks, the Binomial distribution converges to the Poisson distribution with mean $n\gamma\Delta$.

[15] 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.

[16] 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.

buyers to stationary strategies that condition only on their value and the observable state. Let Σ denote the set of such mixed strategies. In what follows, "strategy" means "stationary strategy" unless otherwise noted.[17]

A profile $\sigma = (\sigma_x)_{x \in \mathcal{X}}$ of mixed strategies, together with an initial state ω_0, defines a Markov process $\Phi(\sigma)$ on Ω, with one-step transition matrix $P(\sigma)$. (We describe the state transitions more precisely in section A of the online appendix.) We denote the n-step-ahead transition function as $P^n(\sigma)$ and define the probability of reaching state ω from an initial state ω_0 in n-steps by $P^n([\omega_0, \omega]; \sigma)$.

Given the observed state $\widetilde{\omega}$, buyers have to form beliefs about the high bids in the open auctions and the state of the losers' pool. In section B of the online appendix we prove that the Markov process induced by σ has a unique limiting distribution, implying that buyers' conditional beliefs are given by Bayes' rule at every $\widetilde{\omega} \in \widetilde{\Omega}$ that is on the long-run path of σ; i.e., that has positive probability under the invariant measure on $\widetilde{\Omega}$ induced by σ. At off-path observable states, Bayes' rule does not pin down beliefs. We deal with this issue by requiring that conditional beliefs be the limit of some sequence of full-support strategies that converges to σ, as in sequential equilibrium.

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

Given a strategy profile σ, buyers can use these beliefs to compute their expected payoffs from choosing an auction and submitting a bid.

4.1 Stationary Equilibrium

We next specify what it means for a strategy to be a best response, given the strategy profile σ used by other players and conditional belief system p. Suppose buyer i with value x submits a bid b in auction j in observable state $\widetilde{\omega}$. Then σ and $p(\widetilde{\omega})$ 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

[17] 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.

weakly lower bid $m \in \{0, .., b\}$, let $g_{\sigma,p}(m; \widetilde{\omega}, j, b)$ denote the probability that buyer i wins and that the highest rival bid (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 and enters the losers' pool, then what matters for his expected continuation value is the state of the market in the period after his loss. Let ω^l denote this *losing state*. For each $\omega^l \in \Omega$, let $h_{\sigma,p}(\omega^l; \widetilde{\omega}, j, b)$ denote the probability that buyer i loses the auction and that the losing state is ω^l. These winning and losing probabilities depend not only on the observable state, but also on the buyer's auction choice and bid since they can affect the distribution over future bids in auction j.

We now define payoffs. Given (σ, p), let $v(x, \omega; \sigma, p)$ be the expected payoff to a buyer of type x who arrives at state ω and plays his optimal strategy. Define

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

as his maximized payoff when he arrives and observes $\widetilde{\omega}$, given conditional beliefs $p(\widetilde{\omega})$. His expected re-entry payoff if he loses and the losing state is ω^l is given by

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

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 ω^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 each observable state $\widetilde{\omega} \in \widetilde{\Omega}$,

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

A strategy is a best response to (σ, p) if it achieves these optimal values for every observable state $\widetilde{\omega} \in \widetilde{\Omega}$ and 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 its probability. The second term is the payoff if he loses: probability of losing, times probability of not exiting, times expected re-entry payoff. The summation is over the set of possible losing states ω^l.[18]

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 ω^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. Losing after a low b is less discouraging. The observed state $\widetilde{\omega}$ when the buyer arrives also is informative about the losing state ω^l.

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

Definition 2 *An equilibrium is a strategy profile σ^* and a conditional belief system p^* such that (i) σ^* is a best response to (σ^*, p^*), and (ii) p^* is consistent with σ^*.*

Proposition 1 *An equilibrium exists.*

The proof is in the online appendix. In general, there may be multiple equilibria. In estimation, we will use a somewhat broader equilibrium concept, ϵ-equilibrium. The existence of an ϵ-equilibrium is implied by Proposition 1.

4.2 An Approximation Result

Here we characterize equilibrium bidding in thick markets. We leverage the fact that losing buyers do not return immediately—it takes time for them to learn and respond to the news that they have lost. As a result, if the arrival rates of buyers and sellers are high, then many buyers may have bid and a large number of auctions may have closed before a buyer returns. A buyer's re-entry payoff may then be largely independent of the losing state since, by the time he returns, the market has undergone many transitions. In that case, his continuation value after losing

[18]Note that, because the buyer either wins or loses, $\sum_{\omega^l \in \Omega} h_{\sigma,p}(\omega^l; \widetilde{\omega}, j, b) = 1 - \sum_{m \in \{0,..,b\}} g_{\sigma,p}(m; \widetilde{\omega}, j, b)$.

depends only on his type, not on his previous actions or 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. Then the following bid is weakly dominant:

$$b(x) = x - (1 - \alpha)V(x; \sigma, p) \tag{3}$$

That result (Proposition D.1 in the online appendix) states that each buyer should bid his value less his expected re-entry payoff. When 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. Call such a strategy a *constant* bidding strategy. Adachi (2016), Backus and Lewis (2019), and Bodoh-Creed, Boehnke, and Hickman (2020) get a similar result, but in their models the standard weak dominance argument for second-price auctions applies. Our proof extends this argument to settings where a buyer's bid can influence the bidding decisions of subsequent buyers.

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 γ to 0 while holding the number of open auctions J fixed. The expected return time is then so far in the future that the buyer might expect the re-entry state to come from the steady-state distribution.[19] However, there is a subtle problem with that limit: the state of the losers' pool still matters. All buyers in the losers' pool return at the same rate, so they are likely to compete against each other when they return. That limit is also inconsistent with the data. It implies that all the open auctions will have closed long before the buyer returns. By contrast, in our application, many of the auctions 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. The number of open auctions is over a hundred, so a returning bidder sees mostly the same auctions he saw last time he bid.

An alternative limit is more plausible and does make the expected re-entry payoff independent of the losing state: fix the expected return time and number of new buyers per auction, but thicken the market by increasing the arrival rates for sellers and new buyers, letting the interval between seller arrivals go to zero. Because we normalized that interval as one unit of time, we

[19] This is the approach that the literature has taken and that we took in an earlier version of this paper.

take that limit by sending the return rate γ to 0 while increasing the number of auctions, J, so that $\gamma \cdot J$ is constant.

In this limit, the expected number of auctions that close before a loser returns $(1/\gamma)$ gets large, but the fraction of currently open auctions that close $((1/\gamma)/J)$ 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 that many states have nearly the same value, so 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).

That is, this limit does not imply that a returning loser expects to face the steady-state distribution of states, but he does expect the steady-state distribution of re-entry payoff values. However, for any fixed γ a constant bidding strategy will not be exactly optimal. The effect of the current observable state on the expected re-entry payoff is small but not quite zero, and the losing bid and choice of auction do influence the actions of future bidders a little bit. A bidder would thus want to "fine tune" his bid.

We show, though, that constant bidding can be an approximate best response for any level of precision. In order to capture the idea that buyers ignore those details that vanish in the limit, we use a weaker solution concept, ϵ-equilibrium. For $\epsilon > 0$, an ϵ-equilibrium is a strategy profile where no buyer has a deviation that can improve his steady-state expected payoff by more than ϵ.[20] Our main result is the following.

Theorem 1 *Pick any $\epsilon > 0$. Fix a sequence $\{\gamma_k, J_k\}_{k=1}^{\infty}$ such that $\lim_{k\to\infty} \gamma_k = 0$ and $\gamma_k J_k$ is constant. Then there exists a sequence of ϵ-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.*

In the proof (details in the appendix), we construct for each k a constant bidding strategy σ_k^* in which each buyer bids in one of the M next-to-close auctions, $M < J_k$. We treat the choice

[20] ϵ-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, ϵ-equilibrium is appropriate if players find it costly to compute the optimal strategy, or if they believe that other players may make small mistakes.

of auction as a sequence of binary participation decisions in static auctions with a fixed outside option. We then show that, fixing a large M and taking k to infinity, there is a high probability of arriving at a state where σ_k^* is an approximate best response. The idea is to make M large enough that arbitrage across auctions equalizes expected payoffs, while making γ_k small enough that a losing buyer expects all of those M auctions to have closed before he returns.

The theorem says that we can make the cost 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 make small mistakes most of the time. At some unlikely extreme observable states, such as when all M auctions have received multiple high bids or when none of the auctions have received bids, bidders may make large mistakes. Both kinds of mistakes have a small effect on the ex ante expected payoff: 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 and is increasing in type. As we shall see in the next section, these two properties are critical to making the model useful for empirical work.

5 Empirical model

In what follows, we assume the data are generated by an ϵ-equilibrium (σ^*, 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 model's primitives. 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 $m \in \{0, .., b^*\}$, define

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

as the probability that pseudotype b^* pays m in the set of auctions in which he chooses to bid and wins, where $\tilde{\pi}^*(\widetilde{\omega})$ denotes the steady state probability of observable state $\widetilde{\omega}$ under σ^*, and

$j^*(\widetilde{\omega}; b^*)$ denotes the auction chosen by pseudotype b^* at $\widetilde{\omega}$. 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 $G_{\sigma^*,p^*}(m|b^*)$ 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^*, applying Theorem 1, and taking expectations over $\widetilde{\omega}$ we obtain

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

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 η. Substituting V into the constant bid function from equation (3) and solving for x yields

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

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

5.1 Identification and Estimation

Our data on buyers consist of their identities, their bids (including winning bids), the times at which the bids are submitted, and the auctions in which the bids are submitted. Buyer identities are crucial because they allow us to distinguish between new and returning buyers and to observe who exited. We assume that 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 practice 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 : $\widehat{\lambda}$ is the average number of new buyers arriving per period, $\widehat{\gamma}$ is the mean return time of a loser who does not exit, and $\widehat{\alpha}$ is the fraction of losing buyers who exit. The distribution F_E is identified from Expression 4 which we can rewrite as

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

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-paramtric 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 density 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 she chooses or what the observable state is. This invariance property allows the econometrician to use each buyer's bid to directly infer her type, effectively conditioning on the set of auctions she chooses in the data. However, this convenience comes at a cost: the econometrician needs to observe 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 below 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.

5.2 Tests

Our model generates testable implications about bidding behavior. First, bid functions need to be strictly increasing. Since this is the case if and only if η is increasing, we can test for monotonicity by checking whether our estimates of $G_{\sigma^*,p^*}(b^*|b^*)$ and $E[M|M < b^*, b^*]$ imply that the expression on the RHS of equation (4) is increasing. Second, if buyers are using a constant bid strategy, then returning bidders should bid approximately the same amount as they bid in the auction they previously lost. 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.

A second set of testable restrictions are 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

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

where $\overline{n}$ 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 $\gamma\overline{n}f_L(x)$, 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

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

Equating these two flows yields

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

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 section G of the online appendix we also show that the model can be extended to allow for endogenous exit.

6 Data details

As noted above, our primary data consist of all eBay listings for iPads posted during an 8-month period (February-September 2013), obtained from eBay's internal data warehouse. 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.[21]

When a seller posts an item for sale on eBay, she has the option to create a fixed price listing,

[21] In our analysis we only include used items that were fully functioning—i.e., we exclude items identified as "For parts or not working."

in which case the price is fixed and the listing remains active on the site for up to 30 days until the item is sold. For the specific product and time period we are studying, auctions are the most common form of sale: 65% of successfully sold items were sold by auction. An additional option is to offer a "Buy it Now" (BIN) price, which allows buyers to purchase the item at the specified price without waiting for the auction to end. However, this option disappears when a bid below the BIN price is submitted. In the analyses below we focus on auction listings only, meaning we drop fixed-price listings and any BIN listings that were sold at the BIN price. We also drop auctions that received no bids, which in most cases happened because the start price was unrealistically high.

Table 2 shows summary statistics for the 5,002 auction listings in our sample. The majority of these listings ended successfully with a sale, and the average sale price (conditional on sale) was \$288.86 with an average shipping fee of \$7.26. 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 10.42, but this number varies substantially across auctions.

Table 2: Summary statistics for auction listings (N=5,002)

			Percentiles		
	Mean	Std. Dev.	0.10	0.50	0.90
Start price	122.42	111.64	0.99	100.00	275.00
Positive reserve price (0/1)	0.10	0.29	0.00	0.00	0.00
Reserve price (if positive)	274.40	42.84	216.00	280.00	325.00
Offered BIN	0.37	0.48	0.00	0.00	1.00
Sale price (if sold)	288.86	31.32	255.60	290.00	325.00
Shipping fee	7.26	5.62	0.00	6.60	15.00
Number of bids	23.74	17.06	4.00	21.00	47.00
Number of unique bidders	10.42	5.84	3.00	10.00	18.00
Minutes since last auction	69.63	112.97	4.63	32.07	155.10
Cover included (0/1)	0.19	0.39	0.00	0.00	1.00
Seller feedback (#)	7515.94	45020.92	13.00	131.00	3327.00
Seller feedback (% positive)	99.03	5.81	98.40	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.

Our model abstracts away from intra-auction dynamics, since buyers are assumed to bid when they arrive, and bid exactly once in whichever auction they choose. Of the various simplifying assumptions we make, this one is perhaps the most at odds with the data, since in reality "incremental bidding" (submitting multiple, increasing bids within a single auction) is relatively common. Roughly 44% of the bidders in our data submit multiple bids for the same item. However, 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, in the interest of simplicity 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 7 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 more than half 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 H.1 of the online appendix.

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 also 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.

7 Estimation

In this section we first explain how we obtain estimates of new buyersâ€ arrival rates (λ), returning buyers' re-arrival rates (γ), and their exit rates (α). 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.

7.1 Estimating bidder arrival and exit rates

As noted above, sellers' and new 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 λ, 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, α, is thus 0.502.[22] Return times are fairly short: conditional on returning to bid again, 21% of bidders return within an hour, and 10% return within 5

[22] 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 α, since most bidders return relatively quickly if they are going to return at all.

minutes. The full distribution of return times is highly skewed, however, since there is a long right tail reflecting bidders who take 24 hours or more to come back. We approximate this distribution using an exponential distribution and estimate its mean (γ) as the inverse of the mean number of auctions before a losing buyer returns, which is .008.

Although our model can be extended to allow for endogenous exit, as shown in the online appendix, 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. We find that exit rates are relatively flat with respect to bidders' types (see Figure H.2 in the online appendix), and since allowing for endogenous exit makes computing counterfactuals meaningfully more difficult, we use the inverse bid function from the simpler model with exogenous exit when we estimate the distributions of bidders' valuations below.[23]

7.2 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 by equation (4). This inversion requires estimates of the exit rate, α; 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. Intuitively, competition forces every bidder to adjust its bid by $Z\hat{\gamma}$ since this value is common to all bidders in the auction. The normalizing regression includes indicators for color (white vs.

[23] 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.

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.

One issue of concern is that bidders may observe auction characteristics that we do not. As is well known, unobserved heterogeneity can lead to misleading estimates of bidder values. We examined this issue for a different eBay sample (the product was iPads, not iPad Minis) by contracting with MTurk workers to read product descriptions for 2,000 listings. They reported a list of additional product characteristics, including subtleties like payment restrictions and whether the iPad had been unlocked. These hard-to-observe characteristics explained very little of the price variation (R^2 went from 0.65 to 0.68). Another reason for thinking that the effects of unobserved heterogeneity are likely to be small is the fact that losing bidders who return bid on average approximately the same amount as they did in previous auctions.

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^{th}-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 α and the functions $G_{\sigma^*,p^*}(b|b)$ and $E(M|M < b, b)$, we can compute the bid function and directly check monotonicity. The left panel of Figure 1 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 1: Estimated bid function, and estimated distributions of valuations

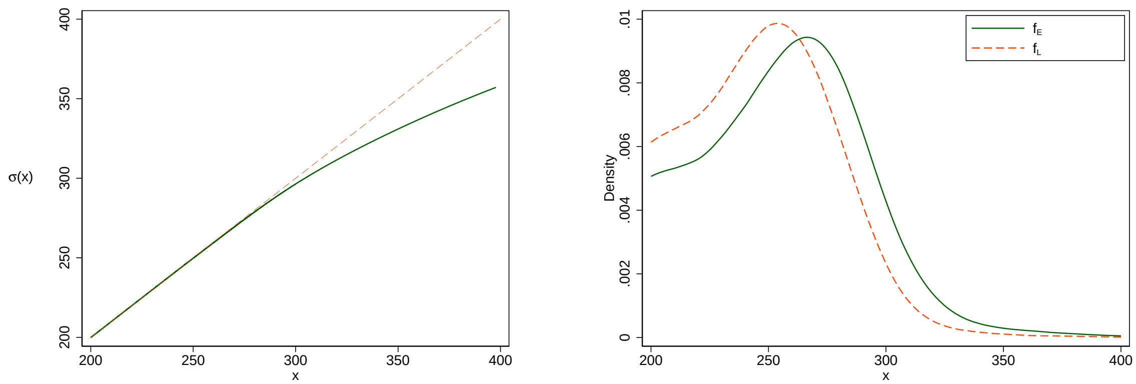

The right panel of Figure 1 shows kernel density estimates of f_E and f_L.[24] The difference between the estimated densities is consistent with the model: the distribution of returning losers' valuations looks like a resampling of new bidders' valuations, with less density in the upper tail. It is important 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 bidders. 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.

7.3 Tests of over-identifying restrictions

Equation (5) describes how the number of returning buyers per auction ($\gamma\bar{n}$) should depend on bidders' arrival rates (λ) and exit rates (α) and the probability of a successful sale (q). The average number of returning bidders per auction in the data is 4.86. Our estimates of the exit rate α (0.50), the arrival rate of new bidders λ (5.47), and the probability of sale q predict an average of 4.58 returning bidders per auction, which is not far off.[25]

[24] We used the first bids of new bidders to estimate f_E and the bids of each returning bidder for f_L. We plot the estimates for values above $200, since low value bidders have virtually zero probability of winning and are essentially irrelevant.

[25] Our estimates of λ, α, and γ would imply the average size of the loser pool is between 550 and 600.

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.

The left panel of Figure 2 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 λ, α, and/or $G_{\sigma^*,p^*}(b|b)$ are inaccurate or invalid.

Figure 2: Test of restriction on f_L

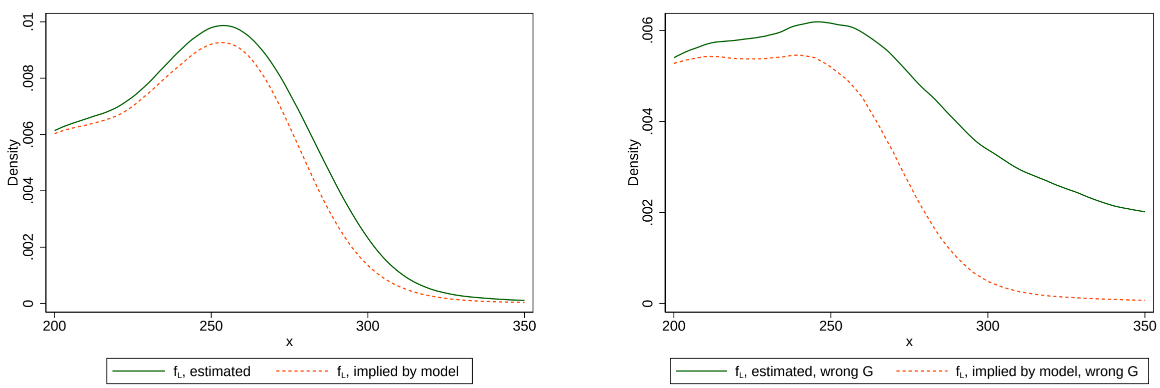

Indeed, if we estimate the model without allowing for endogenous auction selection, this test clearly fails. The fact that buyers self-sort into auctions means that the distribution of the maximum rival bid depends on the buyer's type, so we are careful in our estimation procedure to condition on the 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 where buyers are randomly matched 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. When we use the distribution of the winning bid as our estimate of G, however, we find that it overestimates high-value bidders' probability of winning, which would lead to an overestimate of their values. Using this invalid estimate of G leads to an obvious violation of the model's steady-state conditions, as shown in the right panel of Figure 2.

We interpret the results of these tests as both validating the simplifying assumptions of our model and also confirming the importance of accounting for auction selection.[26]

8 Efficiency

In this section, we 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 does it come to delivering the Walrasian equilibrium.

To answer this question, we use the data and our estimate of F_E to calculate the market-clearing price P^* that would prevail if eBay were to conduct a uniform price auction for all buyers and all sellers who arrive over a period of time. Specifically, we use the entire sample period, calculate the total number of sellers $N_s = 5,002$ and 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 an efficient benchmark against which to compare the prices and efficiency of other mechanisms.

The second benchmark is a counterfactual in which the sellers hold separate second-price auctions with no reserve price, and buyers are randomly allocated to those auctions, with each buyer getting only one chance to win an auction. This benchmark tells us what the price distribution would be and how inefficient the allocation would be if buyers do not have the option to return and bid again in a future auction. 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 3 shows prices and efficiency measures for the bidding we observe in the data compared to the two counterfactual benchmarks. In the uniform price auction, we calculate the

[26] A model that assumes random matching of buyers to auctions (instead of endogenous matching) would also be inconsistent with some basic price patterns in the data, which we describe in the Appendix.

market-clearing price to be $\$279.45$. All buyers with valuations above the market-clearing price successfully purchase, and the average gross surplus of these buyers is $\$307.73$. In the other benchmark, the average price is much lower at $\$231.22$ and efficiency is also much lower with only 31% of the buyers who should get the item (i.e., buyers with valuations above the market-clearing price) actually do. The outcomes that we observe in the data are naturally in between these two extremes but fall well short of the efficient benchmark: price dispersion is substantial and only 59% of the highest-value buyers successfully win an auction. The eBay mechanism achieves only 43.0% of the potential welfare gain relative to the benchmark in which bidders can only bid once.

Table 3: Prices and efficiency compared to counterfactual benchmarks

	Simultaneous auctions, static bidding	Sequential auctions, dynamic bidding (i.e., data)	Uniform price auction
Avg. price	231.22	275.39	279.45
SD of prices	70.88	26.85	0.00
Avg. gross surplus	283.39	293.84	307.73
Pr(win $\mid x > P^*$)	.305	.594	1.000

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

These results resemble those of Bodoh-Creed et al (2021), 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.[27]

The option to return and bid again after a losing bid has two effects. First, it makes the auctions more competitive, as the presence of returning buyers increases the number of buyers per auction. Second, buyers bid less since they shade their bids to reflect the option value of losing. Our counterfactual simulation indicates that both effects are quantitatively meaningful, but the former appears to have a more substantial impact — not just on allocative efficiency, 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,

[27] 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.

the mere presence and participation of returning bidders is more impactful than the strategic changes in bids that result from buyers' ability to return.

9 Endogenous vs. Random Matching

Should platforms like eBay provide buyers with information about the state of play in the auctions? More specifically, does posting the highest losing bid in an auction deliver a more efficient allocation? There is no question that it helps individual buyers choose better matches since it allows them to avoid selecting auctions where they have already been outbid.and to arbitrage expected payoffs across auctions. But the equilibrium effect of more efficient sorting on outcomes is not clear. We examine this issue by considering 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.

To run this counterfactual, we need to specify an auction choice rule. Our theoretical model provides a useful benchmark. We show that, when auctions are sealed-bid, there exists an equilibrium in which buyers always bid in the soonest-to-close auction (see section F of the online appendix for a formal statement of the Proposition). 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 expected level of competition is the same in every subsequent auction. As a result, each buyer is indifferent as to which auction to join, so there is no selection effect.[28] Furthermore, in this equilibrium, buyers are randomly matched to sellers based on their arrival times. Thus, the soonest-to-close choice rule is a dynamic version of the equilibrium of the random matching models studied in the literature.

One way to run the counterfactual is to simulate the equilibrium outcomes of our model under the soonest-to-close choice rule.[29] However, our model makes strong assumptions about the arrival process of buyers and sellers, and we did not want the comparison between the two choice rules to depend on these assumptions. Therefore, we implement the counterfactual by simulating the entire sequence of auctions in the data under the soonest-to-close choice rule, taking as given the observed closing times of the auctions and the arrival times of new and returning

[28] A buyer's continuation value (and bids) in this equilibrium is determined by the (marginal) distribution of the highest rival bid among the new and returning buyers who arrive during the period in which the auction is soonest-to-close. This distribution does not depend upon the bidder's choice and, in steady state, it does not vary across auctions.

[29] The results are reported in Section F.1 of the online appendix.

bidders.[30] That is, a bidder who bids in auction k at time t in the data is reassigned to whichever auction j is soonest-to-close at that time. These reassignments change who participates, who wins, and who loses in each auction, and, in steady state, the composition of the pool of losers. We treat arrival times of new bidders and returning bidders differently. If t is the arrival time of a new bidder, then the identity of the bidder is the same as in the data. In that case, if she loses the soonest-to-close auction, we first check what she did in the data. If she lost in the data and returned to bid again, then in the simulation we assume she will do the same and add her to the loser pool. If she won in the data but lost in the simulation, then we have her exit with probability equal to the exit rate $\widehat{\alpha}$. However, if t is the arrival time of a returning bidder, then the bidder's identity is randomly chosen from the simulated loser pool.[31] The winner of each auction is the bidder with the highest value.[32]

Comment: *we need to include a statement or footnote decribing how we truncated the sample to ensure that we can identify new bidders and whether they returned.*

The top panel of Table 4 compares the efficiency of these simulated auctions (with random matching) to what we observe in the data (with endogenous matching). We find that outcomes are more efficient under random matching than under endogenous matching: the percentage of highest value buyers who successfully win an auction increases from 56% to 64%. The reason for these results can be seen in the bottom panel, which shows the frequency with which high-value bidders[33] face competition from 0, 1, or 2 or more other high value bidders. Disclosure increases the likelihood of two high-value bidders choosing the same seller, which in turn makes it more likely that a high value buyer exits without winning a unit. This leads to lower overall efficiency due to the higher rate at which high types lose (and exit).

The surprising implication of our analysis is that providing information about the second-highest bid does not enhance efficiency. Instead, it appears to induce too much competition

[30] In this setting, however, the soonest-to-close matching rule may not be an equilibrium because the expected level of competition may vary across auctions due to stochastic variation in the time between auction closings.

[31] To initiate the loser pool, we first seed it with $\overline{n} = \widehat{\lambda}(1-\widehat{\alpha})/(\widehat{\alpha}\widehat{\gamma})$ buyers whose valuations are drawn from F_E, then simulate the sequence of auctions 50 times to let the distribution of valuations in the loser pool "burn in." We use this as the initial pool for the main simulations.

[32] The assignment of bidders to auctions and determination of who wins can be done based on their actual values, so we conduct the simulations in type space. To compute bids (and prices), we would need to find the new equilibrium continuation value function V induced by the counterfactual. The details of how this can be done are explained in section I of the online appendix.

[33] For purposes of the table we define high bidders to be those whose values exceed the market clearing price. In calculating the number of rival high bidders, we condition on auctions won by high bidders and count each auction only once.

Table 4: Prices and efficiency compared to counterfactual benchmark

	Endogenous matching (i.e., data)	Random matching (simulation)
Pr(win $\mid x > P^*$)	.557	0.638
# of rival high bidders		
0	40.87	45.24
1	36.23	27.99
2+	22.90	26.77

Notes: Average gross surplus is the average valuation (x) of the winning bidders. Pr(win $\mid x > P^*$) is the probability that a buyer whose x is greater than the market-clearing price P^* wins an auction before exiting. The second column is based on a simulation in which bidders enter the next-to-close auction when they arrive, which means they are randomly assigned to auctions.

between high value bidders, which — combined with the relatively high exit rate — leads to a less efficient outcome than would be achieved if bidders were simply randomly assigned to auctions.

10 A Posted Price Mechanism

The main source of inefficiency in large, decentralized dynamic markets like eBay is the matching process. Because buyers are matched to sellers based on their (random) arrival times instead of their values, low value buyers win units at the expense of high value buyers. The uniform price auction solves this problem by pooling across time and then matching buyers to sellers based on their values.[34] However, buyers may not want to wait until the platform conducts the auction to get their units In this section, we consider how the platform can implement the efficient allocation in real time using a posted price mechanism. In this mechanism, the platform sets a price P and requires each seller to post that price. We show that, if P is chosen appropriately, then seller revenue, average price, and total surplus all increase, and the average time till sale decreases. In fact, for large markets, the dynamic posted price mechanism achieves approximately the maximum possible level of both revenue and total surplus.

Formally, we adjust our theoretical model to consider the following mechanism. When she

[34] Bodoh-Creed et al (2020) conduct a counterfactural in which they show that that meaningful increases in efficiency can be achieved by selling items in uniform auctions 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.

arrives, each seller posts price P. The first buyer to accept the seller's price offer immediately gets the good and pays P. An available posting is one that has not yet been accepted and is open (that is, its closing time has not been reached). A buyer who arrives when there are no available postings is treated like a losing bidder: he exits with probability α and gets a payoff of zero; otherwise he goes to the losers' pool. Note that buyers with values below P are irrelevant in this setting, so we will ignore them.

The market clearing price P^* that would prevail in a posted price mechanism is given by $\lambda\left[1-F_E\left(P^*\right)\right]=1$, so that the expected number of new arriving buyers with values above P^* per auction is 1.

Proposition 2 *Suppose that the market clearing price $P^*$is above the static optimal reserve price.*[35] *Fix a sequence $\{J_k\}_{k=1}^{\infty}$ that grows without bound. Then there is a sequence of prices $\{P_k\}_{k=1}^{\infty}$ converging to P^*such that for large enough k the posted price mechanism using price P_k yields higher average revenue, higher average surplus, and lower average time to sell than the posted bid auction environment.*

The observable information for a buyer arriving in a period t consists of the closing times of the available postings $d_A(t)$. In this setting, it is an equilibrium for each arriving buyer (with valuation above P) to buy from the next-to-close seller. For any fixed parameter values, the size and composition of the losers' pool and number of positings converges to a unique invariant distribution, by arguments analogous to those from our baseline model.

That invariant distribution is continuous in the posted price P. Because buyers may also return from the losers' pool, for any fixed number J of available postings, there is a price $P^{**}\left(J\right)>P^*$ such that in steady state the average total number of buyers arriving per posting equals 1. As J grows, $P^{**}\left(J\right)$ converges to P^*, for the following reason: for any $P>P^*$, the steady state size of the losers' pool shrinks to 0 as J grows. If $P>P^*$, then $\lambda\left[1-F_E\left(P\right)\right]<1$, meaning that new high-value buyers arrive less frequently than new postings, so the probability that a buyer cannot find an available posting is very low. As a consequence, the size of the losers' pool is very small.

Similarly, the probability that a seller does not find a buyer is very small. At $P=P^{**}\left(J\right)$, the number of available postings $J_A(t)$ follows a random walk with no drift, bounded below by 0

[35]The optimal static reserve price is the value at which the virtual valuation, $\psi(x)\equiv x-(1-F_E(x))/f_E(x)$, is 0.

and above by J. A seller who arrives at a time when there are $J_A(t) \leq J-1$ other available postings fails to sell only if fewer than $J_A(t)+1$ high-value buyers arrive over the J units of time that the posting is available. With a mean of 1 buyer arriving per unit of time, the likelihood of that event is very low when $J_A(t)$ is far from its upper bound, as will be the case with very high probability when J is large.

Thus, as J grows, a posted price close to P^* achieves the maximum possible surplus per posting, $E\left[x|x \geq P^*\right]$. All units are sold, and they go to the buyers who value them most. The mechanism also maximizes revenue. We can derive an upper bound on average revenue by considering a hypothetical single seller who controls all items. That seller would maximize revenue by using a uniform price auction with the optimal static reserve price. When the market clearing price P^* is above that optimal reserve, the expected price in that auction is very close to P^* and all items are sold. For large J, the posted price mechanism achieves nearly that same revenue of P^* per posting.

Finally, by the same reasoning as above, the expected time to sell for a seller who arrives when there are $J_A(t) \leq J$ other available posting will be close to the expected time it takes for $J_A(t)+1$ high-value buyers to arrive, which is

$$\frac{J_A(t)+1}{\lambda\left[1-F_E\left(P\right)\right]}.$$

At $P = P^{**}\left(J\right)$, $J_A(t)$ follows a random walk with expected value $J/2$. The average time to sell, then, is close to

$$\frac{J}{2\lambda\left[1-F_E\left(P^*\right)\right]} = \frac{J}{2}.$$

That expected selling time is half as long as in our baseline, posted bid model, where a seller sells only when the auction closes after J units of time.

We summarize those results in the following theorem.

Theorem 2 *In the limit as $J \to \infty$, a posted price mechanism using the market clearing price P^* has the following properties:*

1. *both the probability that a seller makes a sale and the probability that a buyer with valuation at least P^* purchases approach 1;*
2. *average seller revenue approaches its upper bound P^*;*

3. *average surplus approaches its upper bound $E[x|x \geq P^*]$;*
4. *the expected time for a seller to sell is half as long as under the posted bid mechanism.*

To implement the optimal posted price mechansim, the platform would need to know the market clearing price. This should not be too difficult in a stationary environment. One way is use the auction data to estimate F_E as we have done in this paper. The other, more ad hoc approach would be to learn about P^* by monitoring the number of available postings. If the number is stochastically increasing over time at P, then P is likely to be higher than P^* and if it is stochastically decreasing at P, then P is likely to be lower than P^*. In either case, the platform can respond by adjusting the posted price accordingly.

We simulated the optimal posted price mechanism for the sample period of our data to see if it successfully implements the efficient allocation. Given our estimate of F_E, the market-clearing price is $P^* = \$279.45$.[36] We take sellers' and bidders' arrival times as given by the data, and assume that every bidder whose estimated valuation exceeds 279.45 purchases an item upon arrival if one is available. We set the initial stock of available items to 70, equal to the expected value of $J_A(t)$ if $J = 140$.[37] Figure H.4 in the online appendix plots the evolution of the number of available items in the simulation. It spends almost all of the sample period below the upper bound of 140 and never reaches the lower bound of 0. Thus, a very small number of sellers fail to sell before their postings expire, but every bidder with a value above 279.45 successfully purchases an item. The sellers' average time to sell is 3.2 days. We conclude that the posted price mechanism would have achieved the Walrasian equilibrium in our data.

11 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 several 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

[36]Given our estimate of F_E, the static optimal reserve price is approximately \$190, which is well below the market clearing price, so our data satisfy the condition for optimality.

[37]We choose $J = 140$ based on 20 seller arrivals per day (as in our data) and postings lasting 7 days.

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 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 contribution is showing that a simple posted price mechanism can eliminate the matching frictions of a dynamic auction market, implement the efficeint allocation, and maximize seller revenue. This result extends to differentiated good markets, although it would likely be harder for a platform (or sellers) to learn the market-clearing prices. The key would be to obtain estimates of the joint distribution of buyer values for the products and, of course, the arrival rates of buyers relative to sellers. Backus & Lewis (2023) show how this can be done. One can then compute the market-clearing prices for the different products, just as we do for the one good case.

References

Ackerberg, D., Hirano, K., and Q. Shahriar (2017). "Identification of Time and Risk Preferences in Buy Price Auctions", *Quantitative Economics*, 8(3), pp. 809-849.

Adachi, A. (2016). “Competition in a Dynamic Auction Market: Identification, Structural Estimation, and Market Efficiency,” *Journal of Industrial Economics*, 64(4), pp. 621-655.

Athey, S. and P. Haile (2002). “Identification of standard auction models,” *Econometrica* 70(6), pp. 2107-2140.

Backus, M. and G. Lewis (2023). “Dynamic Demand Estimation in Auction Markets,” *Review of Economic Studies,* forthcoming.

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.

Bodoh-Creed, A., Hickman B.,and J. Boehnke (2021). “How Efficient are Decentralized Auction Platforms?” *Review of Economic Studies,* 88(1), pp. 91–125,

Brancaccio, G., Kalousptsidi, M. and T. Papageorgiou (2020). “Geography, Search Frictions and Endogenous Trade Costs,” *Econometrica*, 88(2), pp. 657-691.

Buchholz, N. (2022). “Spatial Equilibrium, Search Frictions, and Dynamic Efficiency in the Taxi Industry,” *Review of Economic Studies*, 89(2), pp. 556–591,

Bulow, J. and P. Klemperer (1996). “Auction versus Negotiations,” *American Economic Review* 86(1), pp. 180-194.

Canals-Cera, J.J. and Pearcy, 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.

Einav, L., Farronato, C., Einav L., Levin J. and N. Sundaresan (2018) ”Auction versus Posted Prices in Online Markets,” *Journal of Polictical Economy*, 126(1), pp. 178-215

Gallant, A. R. and D. W. Nychka (1987). “Semi-nonparametric maximum likelihood estimation.” *Econometrica,* 55(2), pp. 363-390.

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

Gavazza, A. (2016). “An Empirical Equilibrium Model of a Decentralized Asset Market,” *Econometrica*, 84(5), pp. 1755-1798

Gonzalez, R., Hasker, K., and Sickles, R. (2009). “An Analysis of Strategic Behavior in eBay Auctions,” T*he Singapore Economic Review* 54(3), pp. 441-472.

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.

Hopenhayn, H. and M. Saeedi (2017). "Dynamic Bidding in Second Price Auction," working paper.

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

Lewis, G. (2011). "Asymmetric Information, Adverse Selection and Seller Disclosure: The Case of eBay Motors," *American Economic Review*, 101(4), pp. 1535-1546.

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.

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

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.

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.

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 Proof of Theorem 1

We construct a constant bidding strategy and show that it is an ϵ-equilibrium. The proof (in section E of the online appendix) that bids are increasing in type is standard.

A.1 Constructing the strategy $\sigma^*_{M,J}$

For any integer M, define an *M-horizon strategy* as one where each type of bidder at each state 1) submits 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.

Pick an M, and let J be greater than M. We construct an M-horizon constant bidding strategy $\sigma^*_{M,J}$ 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.

Fix an arbitrary payoff function $v_0 : \mathcal{X} \to (0, \bar{x}]$, and define the net value $x^{net}(x, v_0) \equiv x - (1-\alpha)\, v_0(x)$; 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, fix 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 posted bids r in those auctions. (When players use M-horizon strategies, later auctions will not have any bids yet.) Finally, fix a one-period-ahead belief function $r_0 : \mathcal{B}^M \times \{1, \ldots, T\} \to \triangle\mathcal{B}^M$, which specifies beliefs about next period's vector of M posted bids as a function of this period's vector.

Given v_0, p_0, and r_0, we construct the corresponding M-horizon strategy $\sigma_{M,J}(v_0, p_0, r_0)$ recursively. We start with the next-to-close auction ($j = 1$), first when it has one period remaining ($d = 1$) and then for higher d's. For each d we specify behavior as the equilibrium of 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 and get a payoff of $v_0(x)$ instead.

- Step $j = 1$, $d = 1$: Arriving buyers have beliefs given by $p_0(r, d = 1)$ about the highest standing bid, w_1. The payoff of the hypothetical game to a type-x player who submits a bid is as follows: if $b(x, v_0)$ is higher than any competing bid (the bids by other players arriving that period, plus the realized w_1 drawn from $p_0(r, d = 1)$), then he gets x minus the highest competing bid m. If $b(x, v_0) < m$, then he gets payoff $(1-\alpha)\, v_0(x)$. 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 $j = 1$, $d = 2$: Beliefs about w_1 are given by $p_0(r, d = 2)$. The difference from the previous step is the function mapping bids to payoffs in the hypothetical static game:

because auction 1 will still be open next period, bidders arriving then will affect payoffs. Specify that next period's buyers will submit bids according to the strategies in Step $j = 1$, $d = 1$. Players in this period then need beliefs about next period's posted bids. The new posted bid in auction 1 will be determined by the actions of the current bidders, together with w_1. Beliefs about the new posted bids in 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 resulting payoffs: the winning bidder's payoff equals his type x minus the highest competing bid, and a losing bidder of type x gets $(1 - \alpha) v_0 (x)$. Use an equilibrium of this game 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. Beliefs about w_1 are given by $p_0 (r, d)$. Buyers arriving in the remaining periods of auction 1 will submit bids according to the strategies in the previous steps. Next period's posted bid in auction 1 will be determined by w_1 and the actions of the current bidders. Beliefs about next period's posted bids in other auctions are given by $r_0 (r, d)$. The actions of next period's buyers then determine the new posted bid in auction 1 in the period after that, and beliefs about that period's new posted bids in other auctions are given by applying $r_0 (r', d - 1)$ to next period's vector of posted bids r'. Continue that process to predict future bids in auction 1 until the auction closes, and this period's bidders get the resulting payoffs. Use an equilibrium of this game to define the probabilities of bidding in auction 1 for each d between 3 and T under $\sigma_{M,J} (v_0, p_0, r_0)$.

For auction $j = 2$, the second in line to close, the set of players in the hypothetical static game for each d are those who 1) arrive this period and 2) in the equilibrium of Step 1, d choose not to bid in auction 1.

- Step $j = 2$, $d = 1$: When auction 1 has one period remaining, beliefs about the highest standing bid in auction 2, w_2, are given by $p_0 (r, d = 1)$. 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 in Steps $j = 1$, $d = 3$ through $d = T$. Use an equilibrium of this game 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 $d = T$: When auction 1 has d periods remaining, beliefs about w_2 are given by $p_0 (r, d = 2)$. Buyers arriving in the remaining periods of this auction will submit bids in this auction according to the strategies in the previous steps. Next period's new posted bids in auctions 1 and 2 will be determined by w_1, w_2, and the actions of the current bidders. Beliefs about next period's posted bids in the other auctions are given by $r_0 (r, d)$. The actions of next period's buyers then determine the new posted bids in auction 1 (until it closes) and auction 2 in the period after that, and beliefs about that period's new posted bids in the other auctions are given by applying $r_0 (r', d - 1)$

to next period's vector of posted bids, r'. Continue that process to predict future bids in auction 2 until the auction closes, and this period's bidders get the resulting payoffs. Use an equilibrium of this game to define the probabilities of bidding in auction 2 for each d between 2 and T under $\sigma_{M,J}(v_0, p_0, r_0)$.

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

Having constructed $\sigma_{M,J}(v_0, p_0, r_0)$ given v_0, p_0, and r_0, we next look for a fixed point. Given a strategy $\sigma_{M,J}$, Bayes' rule pins down on-path steady-state conditional beliefs, and off-path beliefs can be specified in a consistent way so that 1) beliefs about the vector of highest bids in the M next-to-close auctions depend only on the posted bids in those auctions, and 2) beliefs assign probability 1 to an off-path bid being the minimal feasible bid compatible with the observable state. Given a $\sigma_{M,J}$ and conditional beliefs $p_{M,J}$, the one-period-ahead belief function is pinned down at each observable state; call the result $r(\sigma_{M,J}, p_{M,J})$. Strategy $\sigma_{M,J}$ also determines the steady-state expected payoff to each type of bidder; call that function $v(\sigma_{M,J})$. Then a fixed point argument paralleling the proof of Proposition 1 (in section C of the online appendix) establishes that for any M and J there exists $\sigma^*_{M,J}$ that satisfies the following: we can find $v^*_{M,J}$, $p^*_{M,J}$, and $r^*_{M,J}$ such that $\sigma^*_{M,J} = \sigma_{M,J}\left(v^*_{M,J}, p^*_{M,J}, r_0\right)$; $v^*_{M,J} = v\left(\sigma^*_{M,J}\right)$; $p^*_{M,J}$ is consistent with $\sigma^*_{M,J}$; and $r^*_{M,J} = r\left(\sigma^*_{M,J}, p^*_{M,J}\right)$.

A.2 Showing that $\sigma^*_{M,J}$ is an ϵ-equilibrium

We now show that if M is large given $\epsilon > 0$, and k is large given ϵ and M (so that γ_k is close to 0 and J_k is large), then the σ^*_{M,J_k} above is an ϵ-equilibrium, with value function v^*_{M,J_k} and conditional belief system p^*_{M,J_k}. Lemma 1 shows that conditional on choosing one of the next M auctions, bid $b\left(x, v^*_{M,J_k}\right)$ is nearly optimal for a type-x buyer. Lemma 2 completes the proof by showing that under the steady state distribution, choosing one of those M auctions is nearly always nearly optimal.

We can assume without loss of generality that along the sequences

$$\left\{\sigma^*_{M,J_k}, p^*_{M,J_k}, v^*_{M,J_k}\right\}_{k=1}^{\infty}$$

the transition probabilities over the state of the next M auctions (bids and high bidders' types) converge. The set of M-horizon constant bidding strategies is compact, as is the set of beliefs over the state of the M auctions; any sequence of those strategies and beliefs thus has a convergent subsequence. Along that subsequence the behavior of buyers converges. To show that transition probabilities converge, we also need to show that the arrival rates of each buyer type converge. The arrival rates of new buyers are fixed with respect to k by assumption, and for returning buyers we get convergence (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 pool. The steady-state expected number of returning losers each period $\bar{n}_k \gamma_k \triangle$ is bounded above by $\lambda\triangle (1-\alpha)/\alpha$: flow into the pool comes from new buyers who lose and do not exit (rate bounded above by $\lambda(1-\alpha)$), while the outflow comes from losers who return and either win 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. Given (σ, p), for each type x define $\hat{v}(x, j, b; \widetilde{\omega}, \sigma, p)$ as the expected payoff from bid b in auction j at observable state $\widetilde{\omega}$.

Lemma 1 *Fix M, and pick any observable state $\widetilde{\omega}$ such that only the M next-to-close auctions have received bids. Then for any type x and auction $j \leq M$,*

$$\lim_{k\to\infty} \left| \max_{b\in\mathcal{B}} \hat{v}\left(x, j, b; \widetilde{\omega}, \sigma^*_{M,J_k}, p^*_{M,J_k}\right) - \hat{v}\left(x, j, b\left(x, v^*_{M,J_k}\right); \widetilde{\omega}, \sigma^*_{M,J_k}, p^*_{M,J_k}\right)\right| = 0.$$

Proof. We first show that at such an $\widetilde{\omega}$, in the limit the expected re-entry value is independent of b, j, and $\widetilde{\omega}$. 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

$$\lim_{k\to\infty} \left| EV\left(x, \widetilde{\omega}; \sigma^*_{M,J_k}, p^*_{M,J_k}, j, b\right) - v^*_{M,J_k}(x)\right| = 0,$$

where

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

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

As $\gamma_k \to 0$, 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 $j \leq M$ have no further effect on the observable state (no active bid has been placed by a bidder who saw b and j, or by a bidder who saw a bid placed by a bidder who saw b and j, and so on) and thus have no effect on the actions of other buyers.

As noted above, when k grows conditional beliefs at on-path observable states and transition probabilities over bids and high bidders' types in the next M auctions converge. Because that limit process is ergodic and the number of on-path observable states is finite, if the number of periods before the buyer re-enters is high, then beliefs over re-entry state are close to the stationary distribution, conditional on any observable state when the buyer enters the losers' pool. As $\gamma_k \to 0$, the re-entry time is very high with very high probability. Therefore, in the limit the expected re-entry payoff is independent of the observable state $\widetilde{\omega}$ when the buyer chooses b and j.

Because the expected continuation value after losing for a type-x buyer is close to $(1-\alpha)\, v^*_{M,J_k}(x)$, regardless of which bid or which of the next M auctions the buyer chooses, the arguments of Proposition 3 imply that $b\left(x, v^*_{M,J_k}\right) \approx x - (1-\alpha)\, v^*_{M,J_k}(x)$ is a nearly optimal bid in any of the next M auctions. ■

Lemma 1 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 2, we show that one of those M auctions is nearly always a nearly optimal choice. By construction of σ^*_{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 σ^*_{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.

Given scalar $\eta > 0$ and M-horizon constant bidding strategy σ 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 σ 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 $\pi\left(\Omega^\eta(\sigma, v)\,|\sigma\right)$ denote the steady-state probability of $\Omega^\eta(\sigma, v)$ under σ.

Lemma 2 *For any* $\eta > 0$*,* $\lim_{M\to\infty} \lim_{k\to\infty} \pi\left(\Omega^\eta\left(\sigma^*_{M,J_k}, v^*_{M,J_k}\right)|\sigma^*_{M,J_k}\right) = 1.$

Proof. The hypothetical static game in Step $1,1$, where at observable state $\widetilde{\omega}$ a type-x buyer decides between submitting a bid of $b\left(x, v^*_{M,J_k}\right) \approx x - (1-\alpha)\, v^*_{M,J_k}(x)$ in an auction that closes at the end of the period or taking payoff of $v^*_{M,J_k}(x)$, corresponds to the following second-price auction: the same random set of bidders, but a bidder of type x has value $x^{net}\left(x, v^*_{M,J_k}\right) = x - (1-\alpha)\, v^*_{M,J_k}(x)$ and an outside option of $\alpha v^*_{M,J_k}(x)$. (That is, the payoffs from winning and from not participating are both measured as surplus over continuation value $(1-\alpha)\, v^*_{M,J_k}(x)$.) There is a hidden reserve price equal to the standing high bid w_1, distributed according to $p^*_{M,J_k}(\widetilde{\omega})$.

Equilibrium in the hypothetical game is equivalent to the outcome of that auction: bidders choose whether or not to bid, and if they do they bid their value. From Myerson (1981), we know that the expected auction payoff to a bidder with value v equals the integral up to v of the probability of winning for each bid, $P\left(x^{net}\right)$. In the hypothetical game, then, the payoff to a type-x player who chooses to bid is

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

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 in the other steps of defining σ^*_{M,J_k} are similarly equivalent to auctions, even when auction j will not close until a later period, so that in principle the buyer's bid could affect the entry choices of future buyers. Because only the highest losing bid is observed, his bid can influence future behavior only when he has already lost. Thus, the set of competing bidders is effectively exogenous. The expected payoff from submitting a bid in any of the first $M-1$ auctions is an integral of the probability that a bid exceeds both the standing high bid and the net values of bidders who enter that auction.

We can now show that for each type x, $\sigma^*_{M,J_k}(x)$ gives an expected payoff, conditional on the observable state at arrival, of at least (close to) the average payoff $v^*_{M,J_k}(x)$ with high probability in the limit. Under $\sigma^*_{M,J_k}(x)$, a buyer bids with positive probability in the earliest auction j that gives him expected payoff at least $v^*_{M,J_k}(x)$ (if there is one among the next M auctions). Bidding in auction j lowers j's expected value for future buyers, all else equal, and so lowers their equilibrium probability of participating in auction j and pushes them toward later auctions. The type-x buyer can get less only if all M auctions yield an expected payoff below $v^*_{M,J_k}(x)$ – that is, if the expected distribution of the numbers and types of bidders is worse than average in all M auctions. When M is large, the probability of so many deviations from long-run averages of arrival rates and mixed strategy auction choices is low.

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 $v^*_{M,J_k}(x)$ by submitting a bid in an auction later than specified by $\sigma^*_{M,J_k}(x)$. If buyer i participates in an auction that should not under σ^*_{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. The conditional belief system p^*_{M,J_k} assigns probability 1 to that unknown bid being the lowest amount feasible, and so the off-path bid does not deter subsequent buyers from participating in that auction. ■