Efficiency in Dynamic Auctions
with Endogenous Matching

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

October 2025

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, largely due to inefficient matching of buyers to sellers. We show that a posted price mechanism can implement the efficient allocation and does so in our application.

1

1 Introduction

Many goods and services are bought and sold in decentralized, dynamic auction markets, where buyers who fail to purchase—and sellers who fail to sell—can return to the market and try again. Despite information and matching frictions, this opportunity to re-attempt trade makes the market for each trader large over time; and as the number of trading opportunities gets large, Satterthwaite and Shneyerov (2007, 2008) have shown that prices and allocations can converge to the efficient Walrasian equilibrium.

The goal of this paper is to evaluate whether this theoretical result is attained in a real-world dynamic auction like eBay, and to examine what a platform can do to increase market efficiency and/or revenues. Like other recent empirical studies,[1] we develop and estimate a model based on the main features of the eBay marketplace. However, our analysis incorporates an important aspect of the marketplace that has not been addressed in the literature: auction choice. Previous studies have assumed that each buyer who arrives at the marketplace bids in the soonest-to-close auction[2] and, as a result, the matching of buyers to sellers is randomly determined by their arrival times. In our model, as in reality, each buyer selects an auction from the set of available auctions, and this choice depends on the state of play in those auctions - namely, the highest losing bid (or start price if there are no bids) and time remaining. Thus, the matching process in our model is endogenous, and this feature needs to be taken into account in characterizing equilibrium bidding behavior and in estimating the model.

In principle, auction choice poses a serious modeling challenge, because it means that a forward-looking bidder needs to consider how his auction choices (and bid) can affect the decisions of subsequent buyers and, as a result, his re-entry payoff if and when she returns to try again. However, we show that these considerations become unimportant if the market is large enough. We prove that, as the market thickens, the equilibrium response of buyers to differences in posted bids and closing times across auctions causes many states to have nearly the same value. As a result, the buyer's expected re-entry payoff depends only on his type — not on his 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

---

[1] For example, Adachi (2016), Backus and Lewis (2023), and Bodoh-Creed et al (2023).

[2] Bodoh-Creed et al (2021) 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 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.

past actions on his expected re-entry payoff have largely washed away. We then prove that, in our second-price auction setting, it is nearly optimal (in the ϵ-equilibrium sense) for each buyer to always bid his value less his expected re-entry payoff. This bid is invariant to the choice of auction or state of play, and is strictly increasing in a bidder's value.

These properties of the equilibrium enable a surprisingly straightforward empirical strategy. Monotonicity implies that a buyer's value can be equivalently represented by his equilibrium bid, or what Backus and Lewis (2023) call her "pseudo-type". Invariance allows us to aggregate across states and express the bidder's expected re-entry payoff as a function of his bid and the distribution of the highest rival bid in the set of auctions in which he chooses to bid. We can account for this selection by estimating the probability that a buyer with pseudo-type b wins conditional on the set of auctions chosen (in different states) by buyers who bid b. Thus, we do not have to solve for the equilibrium auction choice rule to estimate a buyer's re-entry payoff. Bidder values can be identified and estimated directly from bids based on the (unknown) choice rule that they are actually using.

We estimate our model on data from the eBay mark for used iPads and obtain two important findings. One is that accounting for auction selection is empirically important. We find that we would have substantially over-estimated buyer values if we had assumed (like the previous literature) that buyers are randomly matched to sellers. We also find that random matching is not consistent with the data because the failure to account for auction selection leads to violations of the model's steady state conditions. The second is that, while the eBay mechanism significantly increases efficiency relative to a static benchmark (where buyers have only one chance to bid), it falls well short of the efficient equilibrium hypothesized by Satterthwaite and Shneyerov (2007, 2008).

One might expect eBay's real-time disclosure of the highest losing bids to improve market efficiency, since it allows buyers to avoid bidding in auctions where they have already been outbid and to respond to the differences in expected payoffs across auctions. We examine this issue by considering a counterfactual in which the platform does not post the highest losing bid during the auction, effectively making each auction a sealed bid auction.[3] In this case, we show that there is an equilibrium in our model where buyers always choose the soonest-to-close auction. We simulate this auction choice rule for the sample period of our data and find, contrary

---

[3]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.

to expectations, that posting the highest losing bid actually decreases efficiency. The intuition is that posting the highest losing bids in real time disproportionately generates auctions with two high-value bidders, which in turn makes it more likely that high value buyers lose and exit without winning a unit.

The counterfactuals make clear that the primary source of inefficiency in a decentralized auction market like eBay is the matching process, which results in too many low value buyers winning units at the expense of high value buyers (who lost an auction to another high-value bidder, and decided not to return), a misallocation that also leads to lower prices. As noted by Bodoh-Creed et al (2023), a uniform price batch auction could solve this problem by pooling buyers and sellers who arrive in a period and matching them based on their values. However, buyers may not want to wait until the platform conducts the batch auction to get their units. As an alternative, we consider a mechanism in which the platform sets a single price P at which each seller must sell and a posting period. We show that, if the price is chosen appropriately and the posting period is sufficiently long, then seller revenue, average price, and total surplus all increase, and the average time until sale decreases. In fact, for large markets, the optimal posted price mechanism achieves approximately the maximum possible level of both revenue and total surplus.

We simulate this mechanism for the sample period of our data, taking the arrival times of sellers and new buyers as given by the data, and find that it successfully implements the efficient allocation. Supply never falls to zero, so there are no stockouts. As a result, every buyer with a value above P is able to purchase a unit when they arrive. Some sellers fail to sell before their posting expires, but the number is quite small. Finally, trading costs fall, since buyers visit the platform only once and sellers sell more quickly. The average time to sale falls from seven days to a little more than three days.

Of course, the platform needs to know the market-clearing price to implement the optimal posted price mechanism, but this should not be too difficult in a stationary environment.[4] We show that, at the market-clearing price, the number of seller postings follows a random walk, but at any lower price it converges to zero, and at any higher price it converges to an upper bound determined by the posting duration and the arrival rate of sellers. Thus, movements in the number of available postings would be enough for standard reinforcement learning algorithms

---

[4] By market-clearing price we mean the price at which all sellers would sell, and the successful buyers would be those with the highest willingness to pay.

4

to determine the market-clearing price.

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 single seller, and traders who fail to trade either exit or return the following period.[5] Papers in this literature include McAfee (1993), Satterthwaite and Shneyerov (2007, 2008), and Bodoh-Creed et al (2021). McAfee 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 examine what happens to prices and allocations as the number of trading opportunities for each trader gets 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 feature of this setting is that a bidder's option value of losing depends only on his type. This is because, in each period, buyers are randomly matched to sellers and, in steady state, the actions of any single buyer or seller today have negligible impact on the state of the market tomorrow. We contribute to this literature by providing an analogous result in a dynamic auction market where the matching is in real time and endogenous. An important empirical implication of our 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.[6]

There is a nascent literature on structural estimation of dynamic auction markets. The main papers in this literature are Adachi (2016), Backus and Lewis (2023), and Bodoh-Creed et al (2021), all of whom estimate their models using data from eBay. Backus and Lewis model the eBay auction marketplace as a sequence of sealed bid auctions and use data on buyers' product choices and bids to estimate substitution patterns in a differentiated good market. Adachi and Bodoh-Creed et al focus on homogenous good markets and estimate dynamic models of bidding to evaluate the efficiency of the eBay mechanism. These papers assume that the matching of

---

[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 (2023), 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]This is a property that Bodoh-Creed et al (2021) refer to as the "plug-and-play" property.

buyers to sellers is randomly determined by their arrival times. Our contribution to this literature consists of identifying and estimating a more general (and realistic) empirical model in which posted bids and auction closing times are informative and influence auction choices.

Our counterfactual on posted prices relates to work on posted prices versus auctions. Wang (1993) considers a setting in which buyers with private values arrive randomly over time to purchase an item from a monopoly seller, and shows that sellers prefer auctions when buyer values are more dispersed. Einav, Farronato, Levin, and Sundaresan (2018) study how sellers on eBay choose between auctions and posted prices, and how their incentives to use auctions have changed over time.[7] They propose a parsimonious choice model in which each seller is matched with at least two buyers who have private values and a common outside option. In their model, posted prices are favored when there is less dispersion in buyer values, better outside options, and/or greater demand for convenience.[8] In our counterfactual, it is the platform, not individual sellers, who chooses between the two trading mechanisms, but the conditions that favor posted price are similar: the product needs to be standardized, so the benefits of price discovery are low; the market needs to be competitive, so buyers have good alternatives (including the option to bid again); and the market needs to be thick and stationary, so the market-clearing price is stable.[9] Furthermore, when this is the case—as in our application—the posted price mechanism lowers time costs for both buyers and sellers.[10]

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)). These static models ignore the effect of the option value of losing on bids, and make no distinction between new and returning buyers, these models implicitly treating the steady state distribution of buyer values as the primitive rather than the distribution of new buyer values. This matters for counterfactuals, since changes in the auction mechanism likely result in different stationary distributions of buyer values.

---

[7]They document that the share of posted price listings by sellers on eBay grew from roughly 10% in 2003 to more than 90% in 2015.

[8]The model estimates suggest that the main driver of the increase in posted price listings over time is an increase in the buyer's "hassle" costs of auctions.

[9]To quote Milgrom (1989): "When the good is not standardized or when the market clearing prices are highly unstable, posted prices work poorly, and auctions are usually preferred."

[10]In a dynamic auction market, bidders often have to bid more than once to win. In our application, winners bid on average 2.33 auctions.

Finally, 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 choose auctions. The objective is to motivate several of our modeling choices and our assumption that buyers choose 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 (including the winning 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. 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.[11]. 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. 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 bidder arrivals per auction closing is approximately constant, as shown in Table H.1 of the online appendix. Assuming constant arrival rates is thus

---

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

7

a 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 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 returning 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.[12] The large standard deviations reported in the second column of Table 1 indicate a lot of dispersion in posted bids. This means that when choosing in which auction to bid, buyers may favor auctions whose current posted bids are unusually low given their position in the queue.[13]

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 in the sooner-to-close auctions they have already been outbid. Furthermore, these are not merely low-value bidders submitting irrelevant bids: 19.8% of bidders 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 price (i.e., the second-highest bid) of the auction they passed up.

---

[12]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.

[13]Section J of the online appendix shows that final prices become more predictable over time, in the sense that the variance of the final price conditional on posted information—including the posted bid and the number of bidders so far—declines as the auction nears its end.

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 $\mathcal{X} \subseteq (0, \bar{x}]$ where $f_E(\bar{x}) > 0$. A buyer's value is fixed and does not change over time. Upon arrival, a new buyer selects an auction in which to bid and which bid to submit. The set of bids is given by $\mathcal{B} = \{0, \underline{b}, .., \bar{b}\}$ where 0 denotes no bid. A bid specifies the maximum amount the bidder is willing to pay, and the platform bids on his behalf up to that level. These are known as *proxy* bids.[14] 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;[15] 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. 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

---

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

[15]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.

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 $(n, \gamma\Delta)$.[16]

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.[17] 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. Because we focus on small period lengths Δ, such simultaneous arrivals are very unlikely. 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. Those are the variables that affect an arriving buyer's payoff from participating in an auction and his expectations of future competition. 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 $r(t)$.[18] We call $\widetilde{\omega}(t) \equiv (d(t), r(t))$ the *observed* state; $\widetilde{\Omega}$ is the set of observed states. We restrict

---

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

[17]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. In order to avoid further complicating our notation, we omit from our analysis the number of bids and other information available to bidders, such as the timing of those bids. It would be straightforward to include those variables in our model, and none of our results would change.

[18]In our application, the platform reports the history of highest losing bids in an auction and partially masked

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.[19]

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. In our equilibrium definition in the next section, we deal with this issue by requiring that conditional beliefs be the limit of the beliefs $\pi(\sigma_k, \widetilde{\omega})$ induced by some sequence of full-support strategies $\{\sigma_k\}$ 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 about the unobserved current high bids in open auctions and about the amount of future competition from the losers' pool 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

---

identities of losing buyers for each auction. The assumption here is that buyers do not bother to use this information in forming beliefs about the 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, and there is a time cost involved in accessing the information because it is not prominently displayed.

[19]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.

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 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 (which depends on the strategies of other buyers σ, the observable component $\widetilde{\omega}$, and his beliefs p about the unobserved component). Then, taking expectations over the unobserved component, 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

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

The term $P^{t-1}([\omega^l, \cdot])$ gives the distribution over the state when the buyer returns after t periods given the losing state 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 the Bellman equation. For each observable state $\widetilde{\omega} \in \widetilde{\Omega}$,

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

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.[20]

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.[21]

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.*

---

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

[21]In the thick market limit that we consider below, learning about the values of the bidders in a single auction has only a tiny effect on expectations of future competition. Thus, a bidder has vanishing incentive either to shade his bid down in order to lower future bids from other buyers who expect weaker competition, or to shade his bid up in order to learn more about others' values.

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 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.[22] However, there is a subtle problem with that limit: the state of the losers' pool still matters. All buyers in

---

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

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 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 bidding up low prices 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.[23] 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

---

[23]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.~~

16

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.[24] 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 L next-to-close auctions, $L < J_k$. We treat the choice of auction as a sequence of binary participation decisions in static auctions with an exogenous, type-dependent outside option set equal (by a fixed point argument) to the buyer's ex ante expected payoff. Buyers submit a bid in the first auction where their expected payoff conditional on the observable state exceeds their outside option. We then show that, fixing a large L 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 L large enough that buyer competition across auctions equalizes expected payoffs, while making γ_k small enough that a losing buyer expects all of those L 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 L 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

---

[24]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.

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 develop and discuss several tests of the model.

5.1 Inverse bid function

Given Theorem 1, a buyer's type x can equivalently be represented by his bid, $b^*(x)$. We follow Backus and Lewis (2023) and refer to $b^*(x)$ as type x's pseudotype.[25] The constant bidding result allows us to aggregate across states. For each $m \in \{0, .., b^*\}$, define

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

as the probability that pseudotype b^* pays m in the set of auctions in which he chooses to bid and wins, where $\widetilde{\pi}^*(\widetilde{\omega})$ 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 results below depend on that assumption. Similarly, define $G_{\sigma^*,p^*}(m|b^*)$ as the probability that pseudotype b^* 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

---

[25]In what follows we use the notation b^* in lieu of $b^*(x)$ in order to keep the equations uncluttered.

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.

5.2 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 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. Let F_L denote the stationary distribution of values in the losers' pool and let f_L denote its density. Then, on average, the flow of x types that leave the pool during the time between

19

closings is $\gamma \overline{n} f_L(x)$ and 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

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

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 $\lambda/(\lambda - q) > 1$ 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.

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 new vs. used items and between models (e.g. 16GB vs. 32GB), one might be concerned this definition of the market is too narrow. However, 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. Substitution between used models is indeed evident in the bidding data: when buyers return to bid on an 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 used 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%). Thus, while there is obviously some substitutability between item conditions and models, 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 characteristics 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.[26]

When a seller posts an item for sale on eBay, she has the option to create a fixed price listing, in which case the price is fixed and the listing remains active on the site for up to 30 days until the item is sold. For the specific product 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.

Even though we are looking only at auctions for a specific model, 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 "in-

---

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

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

| | Mean | Std. Dev. | Percentiles | | |
			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

cremental 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.[27]

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 do not observe users' browsing behavior prior to their bid submissions, we can at least check for irregular bunching in the timing of bids. Unlike some other studies using eBay data,[28] 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 him 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   Identification and Estimation

In this section we first explain how we obtain estimates of new buyer arrival rates (λ), returning buyer re-arrival rates (γ), and their exit rates (α). Next turn to our method for estimating the distribution of bidders' valuations, which adapts the method of Guerre, Perrigne, and Vuong (2000) to a dynamic setting. We then use the steady state restrictions to test the validity of our modeling assumptions and of our estimates of the model primitives.

---

[27]Note that the censoring problem described here is distinct from the one addressed by Platt (2017) and Hickman et al (2017), which is that we may not observe what a low-value bidder would have bid in a given auction if he was outbid before he arrived. In our setting, while we may not observe what an arriving bidder would have bid in some auction in which he has already been outbid, we *do* observe the bid he submits in whichever auction he chooses.

[28]See, for example, Ockenfels and Roth (2006), Bajari and Hortacsu (2003), and Backus et al (2015).

7.1 Estimating bidder arrival and exit rates

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, has a zero start price, and closes within a day or so.

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.[29] Return times are fairly short: conditional on returning to bid again, 21% of bidders return within an hour, and 10% return within 5 minutes. The full distribution of return times is highly skewed, however, since there is a long right tail 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.[30]

---

[29]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.

[30]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.

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 his pseudotype, and recover his true type with the inverse bid function given by Expression (4), which we can rewrite as

$$\eta(b^*) = b^* + \left(\frac{1-\alpha}{\alpha}\right) G_{\sigma^*, p^*}(b^*|b^*) \left[b - E(M|M < b^*, b^*)\right] \tag{7}$$

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, $E(M|M < b^*, b^*)$. Given these estimates, we can use the sample of bids by new buyers to compute their values and obtain a non-parametric estimate of F_E. Note that we can also use Expression (7) to derive estimates of the private values of *returning* buyers and use these estimates to obtain a non-parametric estimate of 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 constant bidding strategy: in equilibrium, each buyer type submits the same bid regardless of which auction he chooses or what the observable state is. This invariance property allows the econometrician to use each buyer's bid to directly infer his type, effectively conditioning on the set of auctions he chooses in the data. However, this convenience comes at a cost: the researcher needs to observe the bids of every buyer and assume that they are realizations of pseudotypes.

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 $\widehat{\phi}$ from this regression to normalize bids as $\hat{b} = \bar{b} + (b - Z\widehat{\phi})$, where $\bar{b}$ is the baseline bid for an item with average characteristics. 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 her bid by $Z\widehat{\phi}$ since this value is common to all bidders in the auction. The normalizing regression includes indicators for color (white vs. black); indicators for whether the auction included a cover, keyboard, screen protector, stylus, headphones,

and/or extra charger; seller feedback ratings; shipping fee; and month dummies (to control for a gradual downward trend in prices over time). In all that follows, when we refer to bids we mean normalized bids.

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).

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^*.

Given 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.

The right panel of Figure 1 shows kernel density estimates of f_E and f_L.[31] The difference between the estimated densities is consistent with the model: the distribution of returning losers'

---

[31]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.

Figure 1: Estimated bid function, and estimated distributions of valuations

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 (for winners with especially high values) over \$25 higher.[32]

7.3 Steady state restrictions

Equation (5) describes how the number of returning buyers per auction ($\gamma \bar{n}$) should depend on the new buyer arrival rate (λ), buyer exit rate (α), and the probability of a successful sale (q) in steady state. 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.[33] This test should fail if the model's underlying assumption of stationarity did not hold, or if our estimates of buyer arrival, return, and exit rates are inaccurate

Equation (6) describes a more stringent test of our modeling assumptions: not only should the

---

[32]In their sample of auctions for digital cameras, Backus and Lewis (2025) estimate the average continuation value of winning bidders to be \$15.08, which is around 6.4 percent of the average sale price.

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

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 parameter estimates and 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 invalid or if the estimate of $G_{\sigma^*,p^*}(b|b)$ is incorrect.

Figure 2: Test of restriction on f_L

Indeed, if we estimate the model without allowing for endogenous auction selection, this test clearly fails. If buyers are randomly matched to auctions, then the set of auctions chosen by type-b buyers is a random sample, and one can equivalently use the distribution of the highest rival bid in the set of *all* auctions to compute the bidder's expected re-entry payoffs. This is the approach taken by Adachi (2016) and Bodoh-Creed et al (2020), for example.[34] But, when we use this (unconditional) distribution as our estimate of G, we find that it overestimates the probability of winning for high value buyers, which in turn leads to an overestimate of their values. The right panel of Figure 2 shows a comparison between the f_L that we estimate directly from the data using this estimate of G and the f_L predicted by the model. The main takeaways from using an estimate of G that ignores selection are that it leads to substantial overestimates of high values and an obvious violation of the steady state restrictions.

---

[34]More precisely, Bodoh-Creed et al use the distribution of the winning bid, but this is equivalent to the distribution of the highest rival bid under the assumption of Poisson arrivals.

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.

8 Efficiency

Our model allows us to compare the efficiency of the eBay mechanism 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 is whether a real-world dynamic market like eBay comes close 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.[35] 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.

---

[35]This counterfactual is motivated by the matching that occurs in brick and mortar markets where buyers have to buy from a local (monopoly) 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 and try again with a different seller.~~ Although the counterfactual doesn't perfectly match the operation of such markets—for example, local markets aren't purely static—online markets are thicker by comparison in two important ways. They increase bidder participation by eliminating location as a factor in the matching process, and they make the market large over time by allowing losing bidders to easily return and try again with a different seller.

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 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: only 31% of the buyers who should get the item (i.e., buyers with valuations above the market-clearing price) actually do. The outcomes we observe in the data are naturally in between these two extremes, but they 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.[36]

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

---

[36]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.

on prices. This finding is reminiscent of the famous result of Bulow and Klemperer (1996) that adding a bidder has more impact on revenues than changes to auction design. In our case, the mere presence and participation of returning bidders is more impactful than the strategic changes in bids that result from buyers' ability to return.[37]

9 Posted Bid versus Sealed Bid Auctions

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 on outcomes is less clear. To examine this issue, we consider a counterfactual in which the platform posts the closing times of the auctions but does not provide any information about the state of bidding in any of the auctions. In other words, the auctions are sealed bid auctions.

To run this counterfactual, we have to specify an auction choice rule. Our theoretical model provides a useful benchmark. 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 his rivals do not observe his 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.[38] 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 sealed bid auction counterfactual is to simulate the equilibrium outcomes of our model under the soonest-to-close choice rule.[39] However, our model makes strong as-

---

[37]We discuss this further in Section J of the appendix, where we examine how the equilibrium of our dynamic model converges to the Walrasian outcome as the exit rate (α) goes to zero.

[38]A buyer's continuation value (and bid) 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 on the bidder's choice and, in steady state, it does not vary across auctions.

[39]The results of this simulation are reported in Section F.1 of the online appendix. They are qualitatively similar to the results reported below in Table 4.

sumptions about the arrival process of buyers and sellers, and we did not want the comparison between the soonest-to-close choice rule and the data-generating choice rule 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 bidders.[40] 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.[41] 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, adding 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.[42] The winner of each auction is the bidder with the highest value.[43]

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 Table 4, which shows the frequency with which high-value bidders[44] 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).

---

[40]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.

[41]The first three weeks of the sample are excluded from the simulation because they are needed to classify bidders as new vs. returning.

[42]To initiate the loser pool, we first seed it with $\overline{n} = (\widehat{\lambda} - q)(1 - \widehat{\alpha})/(\widehat{\alpha}\widehat{\gamma})$ 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.

[43]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.

[44]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)
$\Pr(\text{win} \mid x > P^*)$	.557	0.638
# of rival high bidders		
0	40.87	45.24
1	36.23	27.99
2+	22.90	26.77

Notes: $\Pr(\text{win} \mid x > P^*)$ 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.

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 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 results above show that 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 sometimes 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.[45] 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 to sale decreases. In fact, for large markets, the dynamic posted price mechanism achieves approximately the maximum possible level of

---

[45]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.

both revenue and total surplus.

Formally, we adjust our theoretical model to consider the following mechanism. When she 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.*[46] *Fix a sequence $\{J_k\}_{k=1}^{\infty}$ along which the number of open auctions J_k 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 available postings converge 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 open postings, there is a price $P^{**}(J) > P^*$ such that in steady state the average total number of buyers arriving per posting equals 1. As J grows, $P^{**}(J)$ 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.

---

[46]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.

At $P = P^{**}(J)$, the number of available postings $J_A(t)$ follows a random walk with no drift, bounded below by 0 and above by J. When J is large, therefore, J hits either of its bounds very infrequently. As a consequence, the probability of a stock out, when an arriving high-value buyer fails to find an available posting ($J_A(t) = 0$), is close to zero. Similarly, the probability that a seller does not find a buyer is very small. 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 open. 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, $E[x|x \geq P^*]$. 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 $J_A(t)$ other available postings will be close to the expected time it takes for $J_A(t) + 1$ high-value buyers to arrive, which is

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

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

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

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 the number of open postings $J \to \infty$, 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 $E\left[x|x \geq P^*\right]$;*

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 mechanism, the platform would need to know the market clearing price. This should not be too difficult in a stationary environment. One way is to use the auction data to estimate F_E, as we have done in this paper. An alternative pragmatic approach would be to use a reinforcement learning algorithm that monitors the number of available postings. If the number of postings is stochastically increasing over time, then the current price must be too high; if it is decreasing, then the price must be too low.

We simulated the optimal posted price mechanism for the sample period of our data to see if it successfully implements the efficient allocation. As noted above, given our estimate of F_E, the market-clearing price is $P^* = \$279.45$.[47] 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 $J = 140$.[48] Figure H.4 in the online appendix plots the evolution of the number of available items in the simulation. For almost all of the sample period it lies below the upper bound of 140, and it 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. 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

---

[47]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.

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

a buyer's bid to be invariant to the choice of auction and the observed state. This is what makes the model empirically tractable, and it is in some ways analogous to the oblivious equilibrium concept proposed by Weintraub et al (2008), which simplifies the analysis of dynamic games in markets with a large number of firms. Relying on this approximation result is reasonable in thick markets like the one we study, since the large number of auctions and bidders leads to a high rate of churn in the state. We believe this approach will likely be useful in other 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.

In addition to explaining how endogenous auction selection must be accounted for in estimation, we also show that it results in inefficient matching of buyers to sellers. Our third contribution is to show that a simple posted price mechanism can eliminate these matching frictions, implement the efficient 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., K. Hirano, 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., Blake, T., Masterov, D. V., and Tadelis, S. (2015). "Is sniping a problem for online auction markets?" In *Proceedings of the 24th International Conference on World Wide Web*, (pp. 88-96).

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

Bajari, P. and A. Hortacsu (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., B. Hickman, and J. Boehnke (2021). "How Efficient are Decentralized Auction Platforms?" *Review of Economic Studies,* 88(1), pp. 91–125,

Brancaccio, G., M. Kalouptsidi, 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 J. Pearcy, (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., C. Farronato, J. Levin, and N. Sundaresan (2018) "Auction versus Posted Prices in Online Markets," *Journal of Political 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., K. Hasker, and R. Sickles (2009). "An Analysis of Strategic Behavior in eBay Auctions," *The 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.

Hickman, B.R., Hubbard, T.P., Paarsch, H.J. (2017). "Identification and estimation of a bidding model for electronic auctions," *Quantitative Economics*, 8(2), 505-551.

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.

Milgrom, P. (1989). "Auctions and Bidding: A Primer," *Journal of Economic Perspectives*, 3: 3-22.

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

Ockenfels, A. and A. Roth (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.

Platt, Brennan C. (2017). "Inferring ascending auction participation from observed bidders," *International Journal of Industrial Organization* 54 (2017): 65-88.

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.

Wang, R. (1998). "Auctions versus Posted-Price Selling," *American Economic Review* 83 (4): 838-851.

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 the online appendix) that bids are increasing in type is standard.

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

For any integer L, define an *L-horizon strategy* as one where each type of bidder at each state 1) submits a bid in one of the L next-to-close auctions, and 2) bases his choice of auction and bid only on the observable state of those L auctions.

Pick an L, and let J be greater than L. We construct an L-horizon constant bidding strategy $\sigma^*_{L,J}$ 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 L next-to-close auctions depend only on the vector of posted bids r in those auctions. (When players use L-horizon strategies, later auctions will not have any bids yet.) Finally, fix a one-period-ahead belief function $r_0 : \mathcal{B}^L \times \{1, \ldots, T\} \to \triangle \mathcal{B}^L$, which specifies beliefs about next period's vector of L posted bids as a function of this period's vector.

Given v_0, p_0, and r_0, we construct the corresponding L-horizon strategy $\sigma_{L,J}(v_0, p_0, r_0)$ 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 $\sigma_{L,J}(v_0, p_0, r_0)$. (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.

41

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 $\sigma_{L,J}(v_0, p_0, r_0)$.

- 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 $\sigma_{L,J}(v_0, p_0, r_0)$.

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 $\sigma_{L,J}(v_0, p_0, r_0)$.

- 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 $\sigma_{L,J}(v_0, p_0, r_0)$.

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

Having constructed $\sigma_{L,J}(v_0, p_0, r_0)$ given v_0, p_0, and r_0, we next look for a fixed point. Given a strategy $\sigma_{L,J}$, 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 L 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 $\sigma_{L,J}$ and conditional beliefs $p_{L,J}$, the one-period-ahead belief function is pinned down at each observable state; call the result $r(\sigma_{L,J}, p_{L,J})$. Strategy $\sigma_{L,J}$ also determines the steady-state expected payoff to each type of bidder; call that function $v(\sigma_{L,J})$. Then a fixed point argument paralleling the proof of Proposition 1 (in the online appendix) establishes that for any L and J there exists $\sigma_{L,J}^*$ that satisfies the following: we can find $v_{L,J}^*$, $p_{L,J}^*$, and $r_{L,J}^*$ such that $\sigma_{L,J}^* = \sigma_{L,J}(v_{L,J}^*, p_{L,J}^*, r_0)$; $v_{L,J}^* = v(\sigma_{L,J}^*)$; $p_{L,J}^*$ is consistent with $\sigma_{L,J}^*$; and $r_{L,J}^* = r(\sigma_{L,J}^*, p_{L,J}^*)$.

A.2 Showing that $\sigma_{L,J}^*$ is an ϵ-equilibrium

We now show that if L is large given $\epsilon > 0$, and k is large given ϵ and L (so that γ_k is close to 0 and J_k is large), then the σ_{L,J_k}^* above is an ϵ-equilibrium, with value function v_{L,J_k}^* and conditional belief system p_{L,J_k}^*. Lemma 1 shows that conditional on choosing one of the next L auctions, bid $b(x, v_{L,J_k}^*)$ 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 L auctions is nearly always nearly optimal.

We can assume without loss of generality that along the sequences

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

the transition probabilities over the state of the next L auctions (bids and high bidders' types) converge. The set of L-horizon constant bidding strategies is compact, as is the set of beliefs over the state of the L 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 L, and pick any observable state $\widetilde{\omega}$ such that only the L next-to-close auctions have received bids. Then for any type x and auction $j \leq L$,*

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

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, $v^*_{L,J_k}(x)$. That is, we show that for any type x, bid b, and auction $j \leq L$, we have

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

where

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

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 L consecutive auctions go by before the buyer returns goes to 1. That event implies that the choice of b and $j \leq L$ 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 L 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 $(1 - \alpha) v_{L,J_k}^*(x)$, regardless of which bid or which of the next L auctions the buyer chooses, the arguments of Proposition 3 imply that $b\left(x, v_{L,J_k}^*\right) \approx x - (1 - \alpha) v_{L,J_k}^*(x)$ is a nearly optimal bid in any of the next L auctions. ∎

Lemma 1 establishes that because a buyer who has chosen one of the first L 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 L auctions is nearly always a nearly optimal choice. By construction of σ_{L,J_k}^*, a type-x buyer who chooses one of the first $L - 1$ auctions expects to get at least $v_{L,J_k}^*(x)$. The only way to get less is if he turns down all of the first $L - 1$ auctions; in that event, under σ_{L,J_k}^* he bids in auction L even if it yields a low expected payoff. That event, though, is very unlikely for large L.

Given scalar $\eta > 0$ and L-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 $L+1$ or later gives an expected payoff no higher than $v(x) + \eta$. Let $\pi (\Omega^\eta (\sigma, v) | \sigma)$ denote the steady-state probability of $\Omega^\eta (\sigma, v)$ under σ.

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

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

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, $P(x^{net})$. In the hypothetical game, then, the payoff to a type-x player who chooses to bid is

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

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^*_{L,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

46

competing bidders is effectively exogenous. The expected payoff from submitting a bid in any of the first $L-1$ 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, $\sigma^*_{L,J_k}(x)$ gives an expected payoff, conditional on the observable state at arrival, of at least (close to) the average payoff $v^*_{L,J_k}(x)$ with high probability in the limit. Under $\sigma^*_{L,J_k}(x)$, a buyer bids with positive probability in the earliest auction j that gives him expected payoff at least $v^*_{L,J_k}(x)$ (if there is one among the next L 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 L auctions yield an expected payoff below $v^*_{L,J_k}(x)$ – that is, if the expected distribution of the numbers and types of bidders is worse than average in all L auctions. When L is large, the probability of so many deviations from long-run averages of arrival rates and mixed strategy auction choices is low.

Similarly, as L grows and the number of potential bidders in each of the L auctions becomes large, the same type of competition implies that a buyer i of type x would be unlikely to get a payoff much above $v^*_{L,J_k}(x)$ by submitting a bid in an auction later than specified by $\sigma^*_{L,J_k}(x)$. If buyer i participates in an auction that should not under σ^*_{L,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^*_{L,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. ∎