

Dynamics and Efficiency in Decentralized Online Auction Markets

Online Appendix

A State transitions

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

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

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

- If auction j receives no bids in period t , then the high bidder with value $a_j(t)$ exits.

- If auction j receives at least one bid, then there are two possible transitions. As above, let b_j be maximum of the bids, submitted by bidder i with value x , and, if there are multiple bids, let b'_j denote the second-highest.
 - If $b'_j > w_j(t)$, then bidder i with value x exits; the displaced bidder with value $a_j(t)$ enters the losers' pool with probability $1 - \alpha$ and otherwise exits, as do the buyers other than i .
 - If $b_j \leq w_j(t)$, then the high bidder with value $a_j(t)$ exits; all other bidders enter the losers' pool with probability $1 - \alpha$ and otherwise exit.

B Conditional beliefs and ergodicity

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

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

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

Our proof relies on standard results about Markov chains on a countable state space. (See, for example, Meyn and Tweedie (1993).) It would be sufficient to show that $\Phi(\sigma, d)$ is an irreducible, recurrent, and aperiodic¹ process. In general, $\Phi(\sigma, d)$ may not be irreducible. For example, suppose that arriving buyers bid only in the soonest-to-close auction. Under such strategies, states in which later-to-close auctions have already received bids never occur. However, we can first show that

¹A Markov process is irreducible if every state can be reached from every other state; it is recurrent if in expectation each state is visited infinitely often; and it is aperiodic if there is a state that transitions in one step to itself with positive probability.

$\Phi(\sigma, d)$ has a single absorbing communicating class $\Omega^C(\sigma, d)$.² Second, we show that the Markov chain confined to that class, $\Phi^C(\sigma, d)$, is ergodic. The proposition follows.

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

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

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

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

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

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

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

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

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

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

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

²A state ω leads to state ω' if the probability of reaching ω' from ω is strictly positive. Two states communicate if each leads to the other.

d again.

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

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

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

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

C Proof of Proposition 1

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

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

Then take a sequence $\{\epsilon_n\}$ of $\epsilon_n > 0$ converging to 0, and a sequence $\{\sigma_n^*\}$ of Nash equilibria of Γ^{ϵ_n} . The set of strategy profiles is compact, so without loss of generality assume that $\{\sigma_n^*\}$

converges to a limit σ^* . As noted above, conditional beliefs are continuous in σ , so the sequence of conditional beliefs $\{\pi(\sigma_n^*, \tilde{\omega})\}$ also has a limit; call it p^* . We want to show that (σ^*, p^*) is an equilibrium. First, the upper hemicontinuity of the best response correspondence ensures that σ^* is a best response to (σ^*, p^*) . Similarly, to establish that p^* is consistent with σ^* , it is enough to show that the set of conditional belief systems consistent with a strategy profile σ is upper hemicontinuous in σ .

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

Thus, (σ^*, p^*) is an equilibrium.

D Weakly dominant bidding

Here we prove the claim in Section 3.2 that if a type- x buyer's expected re-entry payoff $V(x; \sigma, p)$ is independent of the losing state, then the bid $b(x) = x - (1 - \alpha)V(x; \sigma, p)$ is weakly dominant.

Proposition D.1 *Suppose that the expected re-entry payoff for a type- x buyer is $V(x; \sigma, p)$, regardless of the losing state ω^l . Then the following bid is weakly dominant for type- x buyer given any strategy profile and conditional beliefs (σ, p) :*

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

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

Suppose that the buyer submits a bid in period t in an auction that will close after d more periods. For $s \in \{t, \dots, t + d\}$, let X_s denote the highest competitor's bid in the auction up through period s . Let $\{x_t, \dots, x_{t+d}\}$ denote the realized increasing sequence of highest competing bids if the buyer

submits a bid of $b(x)$. If the buyer submits a bid of $b(x)$ and $x_{t+d} > b(x)$, then the buyer loses the auction and gets expected continuation payoff $(1 - \alpha)V(x; \sigma, p)$. If $x_{t+d} < b(x)$, then the buyer wins the auction, pays x_{t+d} , and gets payoff $x - x_{t+d} > x - b(x) = (1 - \alpha)V(x; \sigma, p)$. And if $x_{t+d} = b(x)$, then depending on timing and tie-breaking, the buyer either loses the auction and gets continuation payoff $(1 - \alpha)V(x; \sigma, p)$, or wins the auction, pays x_{s+d} , and gets the same payoff: $x - x_{t+d} = x - b(x) = (1 - \alpha)V(x; \sigma, p)$.

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

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

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

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

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

above $b(x)$ are weakly dominated.

E Proof of part (ii) of Theorem 1

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

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

Using that notation, we can write

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

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

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

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

F Equilibrium with sealed-bid auctions

Here we prove that when the platform provides no information about the state of bidding, there is an epsilon equilibrium in which buyers always bid in the soonest-to-close auction.³

³Under the assumption that buyers use stationary strategies, the outcome is in fact an exact equilibrium. That assumption, though, is very restrictive in a sealed bid environment where the only public information is the auction closing

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

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

F.1 Sealed Bid Counterfactual

We simulate equilibrium outcomes for a sequence of 10,000 auctions,⁴ Each period is an auction, and the number of new buyers arriving before each auction being a draw from the Poisson distribution with estimated mean $\hat{\lambda} = 5.47$. Their valuations are drawn randomly from the estimated F_E . Similarly, the number of returning buyers in each auction is a random draw from a Poisson distribution with estimated mean of $\hat{\gamma}\bar{n} = 4.68$. Their valuations are drawn random from the loser pool. When a buyer arrives, she is assigned to the soonest-to-close auction. Losing bidders in each auction exit with probability $\hat{\alpha} = 0.502$, and otherwise enter a pool of losers, which evolves stochastically over the sequence.⁵ Table F.1 compares the simulation outcomes to the data outcomes. The sealed bid auction significantly reduces price dispersion and increases efficiency relative to the outcome in the data. With buyers randomly matched to auctions, but participating dynamically, 72 percent of the highest-value buyers successfully win an auction, as opposed to the 59 percent from the data.

Table F.1: Prices and efficiency compared to counterfactual benchmark

	Endogenous matching (i.e., data)	Random matching (simulation)
Avg. price	275.39	274.02
SD of prices	26.85	16.65
Pr(win $x > P^*$)	.594	0.717

The simulation provides further evidence on how the matching in the data differs from random times. The proposition holds if we drop that assumption and allow buyers to condition on their private history, and it holds regardless of what information about outcomes the platform releases when an auction closes.

⁴To get 10,000 auctions, we simulate 30,000 and then drop the first and last 10,000. We drop the first 10,000 to ensure that we are sampling from auctions in steady state; we drop the last 10,000 auctions because for late-arriving buyers we cannot observe their eventual outcomes (e.g., whether they eventually succeed in winning an auction). At the start of the simulated sequence, we seed the loser pool with $\bar{k} = \hat{\lambda}(1 - \alpha)/(\alpha\hat{\beta})$ buyers whose valuations are drawn from F_E .

⁵At the start of the simulated sequence, we seed the loser pool with $\bar{n} = \hat{\lambda}(1 - \hat{\alpha})/(\hat{\alpha}\hat{\gamma}) \approx 600$ buyers whose valuations are drawn from F_E .

Table F.2: Price distributions: data vs. random matching

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

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

matching. Table F.2 compares the variation of the price distributions in the number of bidders. In the random matching simulation, the average price increases with the number of observed bidders, and the standard deviation decreases. By contrast, in the data, average prices are fairly flat with respect to the number of bidders, and the standard deviation declines much less sharply that it would under random matching.

G Endogenous exit

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

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

and the optimal bid function is

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

The ex ante value function is given by the function

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

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

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

The inverse bid function is

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

and the value equation becomes

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

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

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

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

H Additional tables and figures

Table H.1 shows how the number of bidders per auction varies across times of day.

Table H.1: Bidders per auction closing, by time of day

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

Figure H.1 shows the distributions of times between bids, across all bidders and auctions, compared to the exponential distribution.

Figure H.1: Time between bids

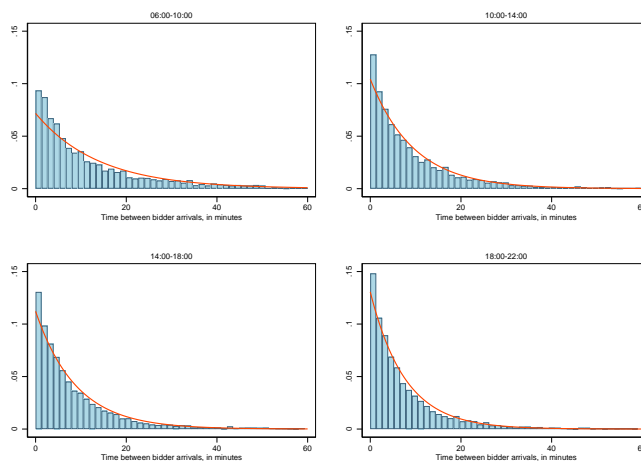


Figure H.2 shows binned exit frequencies as a function of the submitted bid b , along with a semi-nonparametric estimate of the exit rate $\alpha(b)$.

Figure H.3 shows the distribution of the number of bidder arrivals per hour, which in our model is assumed to be Poisson.

Figure H.4 plots the evolution of the number of available items in our posted-price simulation.

Figure H.2: Exit rate as a function of bid

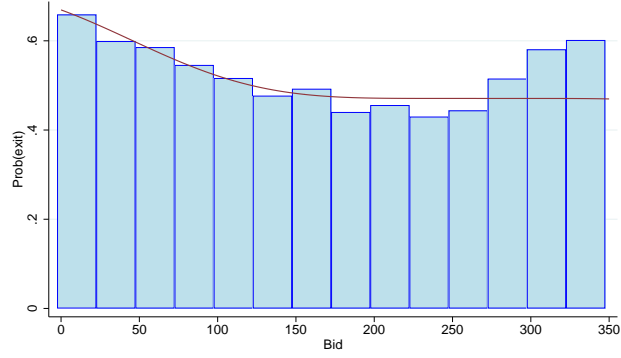
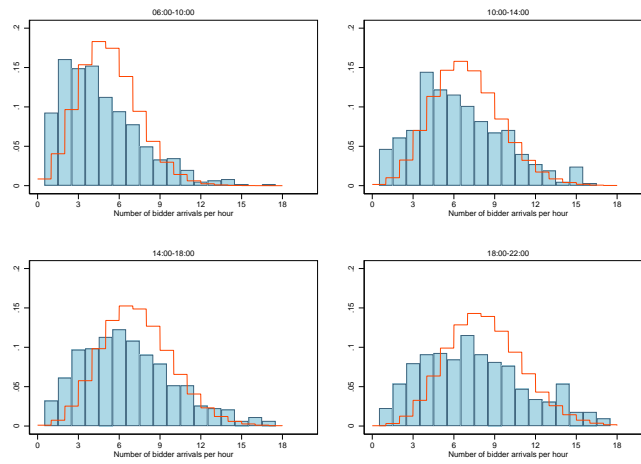


Figure H.3: Bidder arrivals per hour



I Computing bids in counterfactual simulations

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

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

Figure H.4: Evolution of the number of available items



and

$$V(x) = \frac{\int_0^{\sigma(x)} (x - p) dG_{M|B}(p|\sigma(x))}{[1 - (1 - \alpha)(1 - G_{M|B}(\sigma(x)|\sigma(x)))]} \quad (\text{I.2})$$

When the state of the market is a stationary process, $G_{M|B}(\sigma(x)|\sigma(x))$ can be computed as the average probability that a buyer of type x wins. As long as the bid function is monotone, this probability does not depend on the bids, so we can find an equilibrium by first simulating a large number of auctions to compute $G_{M|B}(\sigma(x)|\sigma(x))$, and then numerically solving for the value function $V(x)$ that satisfies conditions (I.1) and (I.2). The latter step is a search for a fixed point in function space, and can be accomplished with a simple iterative procedure. We set $V(x)$ equal to zero initially, so that $\sigma(x) = x$, and then compute the surplus that the simulated bidders would have earned in that case. This computed surplus becomes the new estimate of $V(x)$, and the bids are updated according to (I.1). Surplus is then recomputed for all bidders, and the process is iterated until the newest estimate of $V(x)$ is unchanged relative to the previous one.

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

simulated auctions, we estimate $G_{M|B}(\sigma(x)|\sigma(x))$ with a local polynomial regression of the win dummy on x .

J Auction selection

For each buyer in our data, we constructed a choice set consisting of the thirty soonest-to-close auctions in which the posted bid was less than the buyer’s eventual bid—i.e., the thirty soonest-to-close auctions in which the bid she submitted would have been an allowable bid. Using these choice sets, we then estimated a multinomial logit model in which the only explanatory variable was the rank of the auction (by soonest to close), allowing for 10 different coefficients corresponding to the 10 deciles of the submitted bid. In other words, high bidders were allowed to have different preferences than low bidders for bidding in soon-to-close auctions. The magnitudes of the coefficients were monotonically decreasing in the decile: the highest-value bidders were significantly more likely to choose soon-to-close auctions. Figure J.1 shows the predicted probabilities for the highest, lowest, and middle-decile bidders. Interestingly, the selection probabilities are much less skewed for low value buyers — they are essentially randomizing over the set of auctions in their choice set. This may reflect the fact that posted bids in auctions with later closing times are not very informative of expected payoffs.⁶

References

- Kreps, D. and R. Wilson (1982). “Sequential Equilibria,” *Econometrica* 50:863-894.
- Meyn, S.P. and R.L. Tweedie (1993). *Markov Chains and Stochastic Stability*, Springer Verlag.
- Selten, R. (1975). “A Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games,” *International Journal of Game Theory* 4 (1): 25–55.

⁶Since the default setting in eBay is to list the auctions by their closing times from earliest to latest, it could also reflect the declining saliency of the listing.

Figure J.1: Auction selection probabilities for high vs. low bidders

