A Theory of Bidding Dynamics and Deadlines in Online Retail

Dominic Coey,* Bradley Larsen,† and Brennan C. Platt‡

April 1, 2015

Abstract

We present an equilibrium search model that parsimoniously rationalizes several defining patterns of online markets for new goods. We model buyers in these markets as having a deadline by which the good must be purchased and we model sellers' choice between auctions and fixed-price mechanisms. As the deadline approaches, buyers increase their bids and are more likely to buy through fixed-price listings. The model predicts equilibrium price dispersion even for brand new, homogeneous goods. Using data on over one million auction and posted-price listings for new-in-box items on eBay.com, we find robust evidence consistent with our model: bidders increase their bids from one auction to the next, equilibrium price dispersion exists, and auctions and fixed-price listings coexist, with auction listings selling faster and at a discount relative to fixed-price listings. The bulk of the evidence points toward deadlines playing a significant role in online retail markets. Numerical simulations from our model match realized market outcomes surprisingly well. We highlight insights from comparative statics in our model and distinctive features that would be missed by overlooking the existence of deadlines and dynamic behavior in studying these marketplaces.

JEL Classifications: C73, D44, D83

Keywords: Sequential auctions, mechanism choice, search, deadlines, endogenous valuations

1 Introduction

Standard static auction models depict bidders as having a single opportunity to acquire the good in question, receiving a payoff of zero upon losing. While this seems appropriate for unique one-of-a-kind treasures, online auction sites frequently offer multiple listings of new products that are readily available at retail outlets. With popular items, one can encounter dozens of identical auction offerings over the course of a week. Thus, losers of the current auction still have ample opportunity to obtain the same good, whether through subsequent

*eBay Research labs, dcoey@ebay.com

†Stanford University, Department of Economics and NBER, bjlarsen@stanford.edu

‡Brigham Young University, Department of Economics, brennan_platt@byu.edu

auctions or outright purchase (at local stores, other internet vendors, or on eBay itself through a posted-price "buy-it-now" listing).

Our study is motivated by three striking regularities in the sale new goods via auctions and posted-price listings which have previously not been noted in the literature. First, losers in prior auctions tend to bid more for identical items in subsequent auctions, raising their bids by about 2.1%. Why are bidders increasingly willing to pay more, even though identical auctions keep coming at a steady pace? Second, auctions produce considerable dispersion in their final price, with an interquartile range of 31% of the retail price. Why do bidders differ in their willingness to pay when they share the common outside option of purchasing the posted price good? Finally, the average auction closing price provides a modest discount of 17% below the average posted-price sale. Why then are 48% of products sold by auctions, when sellers could earn more using posted prices?

We show that buyer deadlines can parsimoniously explain all three of these phenomena. Deadlines can be interpreted literally — for instance, if the item is needed as a present for a child's birthday or to be used at a given time — but can also represent any other limit on how long a consumer is willing to spend procuring a good.[1] We model buyers as having unit demand for a homogeneous item, which they must buy within a fixed time of entering the market. Buyers encounter second-price sealed-bid auctions for the item at random intervals, and may purchase the item at any time at a posted price.

Deadlines introduce non-stationarity into the search problem and lead to bids increasing with the duration of search. If the consumer reaches his deadline, he will purchase the item at the posted price, but before his deadline, he might win it in an auction at a lower price. Thus, his bid is shaded down from the true valuation, with more time remaining providing more opportunity and hence greater shading. Even when anticipating a steady flow of auctions, bidders will offer higher prices in each subsequent auction. This also explains why auction prices are necessarily lower than posted prices in the model, which consumers only use as a last resort.

Deadlines also generate price dispersion. To demonstrate this, we consider the stark case in which all buyers are identical in their valuation and the amount of time they have to buy the item. Despite this ex-ante homogeneity, differences arise among them ex-post because some customers take longer than others to win an auction, leading to an endogenous, non-degenerate distribution of bids at any time, even though all buyers enjoy the same eventual utility from the good. This distribution, along with variation in the number of participants

[1] A time limit could arise because the cost of auction participation increases, perhaps because customers cannot sustain the same level of attention to the auctions, or become increasingly frustrated with repeatedly losing auctions. Alternatively, the consequence of not winning could deteriorate with time; for instance, customers could be shopping for a replacement part (such as an engine timing belt) that hasn't yet failed but is increasingly likely to do so.

in each auction, allows the model to replicate the observed dispersion in auction closing prices.

Our model also allows sellers to choose whether to list their product in an auction or with a posted price. If auctions sell faster than posted price listings, as indeed we observe in our data, then auctions and posted price listings can coexist, and even be offered by the same individual. A mixed strategy equilibrium allows identical sellers to be indifferent between the two selling mechanisms: auctions yield less revenue than posted price sales, but yield it sooner.

We then use our eBay sample to calibrate our model parameters. Despite the model's simplifying assumptions, it is quantitatively consistent with the key patterns in the data, which are not used as calibration targets. Beyond rationalizing various key features in our data, our model highlights some limitations of standard static empirical auction techniques. Ignoring bidding dynamics and instead assuming that bidders are bidding their true valuations, as in a static second-price auction, would understate consumer welfare by 47.5%. Market design also takes on more nuance, as the distribution of valuations responds endogenously to changes in other parameters. We illustrate this with a simple increase in listing fees. This initially discourages seller entry into either sector; but with fewer auctions offered, more buyers will have to purchase in the posted price sector. Thus, even with fees increased equally on both selling mechanisms, sellers shift from auctions to posted prices. Moreover, auctions are less plentiful and thus less valuable to buyers, making them willing to bid more. Despite only being a direct concern for sellers, increased fees will compress the distribution of bids.

We model a dynamic market of optimizing buyers and sellers, with each choosing between auctions and posted-price sales. Despite this complexity our model remains tractable, due to two techniques common in the search literature but largely unexploited in the auction literature: analysis in continuous time, and steady state distributions.[2] Specified in a continuous time framework, the model's equilibrium conditions can be translated into a solvable system of differential equations, which was the methodological innovation of Akın and Platt (2012).[3] The restriction to steady state behavior implies that the distribution of buyers with respect to their deadline remains stable. Population dynamics are consistent

[2]The benefits of continuous-time modeling extend beyond analytic tractability. Since buyers' bid functions depends on their state variable (time until the deadline), having a continuous state variable allows for a continuous distribution of buyer willingness to pay, as is typically assumed in static auction models. If our model were depicted in discrete time, there would be a mass of buyers at each state/bid level, and ties between bidders would be a real concern.

[3]Relative to the sequential auctions literature, only Ingster (2009) uses a continuous time model, but as bidder heterogeneity is exogenous there, our approach is unnecessary for that solution. Rather, the solution to the discrete time model in Said (2011) has more parallels to our method; there, the relevant state variable affecting bid shading is the (discrete) number of active bidders. The equilibrium conditions are then depicted as a system of first order difference equations.

with steady offsetting flows of incoming buyers and outgoing winners. This can best be interpreted as focusing on long-run behavior in a market fairly thick with buyers, which seems reasonable for the auctioning of standardized retail items. Using these search-theoretic techniques makes evaluating comparative statics relatively straightforward, in contrast to much of the literature focusing on the structural estimation of dynamic auction models, which lacks comparative static analysis.[4]

While we see deadlines as a natural explanation for the key features of our data, we recognize that they are not the only possible explanation for any one feature in isolation. Increasing bids over time, for example, can be also rationalized by a rather subtle story of bidders' learning about the competition they face, which affects the option value of waiting for future auctions and motivates bidders to revise each subsequent bid. We present empirical evidence showing that bidder learning is unlikely to fully explain the pattern of increasing bids over time. Even the most experienced bidders show the same tendency to increase their bids, and the degree of competition that bidders face in one auction explains only a small portion of their bid increases in subsequent auctions. Moreover, while separate explanations might be given for each of our motivating facts individually, buyer deadlines provide a single, unified explanation of all of these facts together.

Our paper contributes to the competing mechanisms literature, which considers a seller's choice between auctions and posted price mechanisms. Julien et al. (2001), Einav et al. (2013), and Wang (1993) provide models in which one mechanism is strictly preferred over the other except in "knife-edge" or limiting cases. Etzion et al. (2006), Caldentey and Vulcano (2007), Hammond (2013), and Bauner (2015) present models in which both mechanisms coexist, but rely on buyer or seller heterogeneity. In contrast, we show that both mechanisms may be employed in equilibrium under generic parameters, even though buyers and sellers are homogeneous ex ante.[5]

We also contribute to the nascent literature on infinite sequential auctions (Zeithammer, 2006; Ingster, 2009; Said, 2011, 2012; Backus and Lewis, 2012; Bodoh-Creed, 2012; Hendricks et al., 2012), in which bidders shade their bids below their valuations, taking into account the continuation value of participation in future auctions.[6] These papers, as well

[4]Notable exceptions are Zeithammer (2006) and Said (2011), which feature comparative statics with respect to the discount rate, the time between auctions, and (in the latter) the expected number of competitors. Reservation prices respond to these in the same direction as in our model.

[5]The above papers also differ from ours in that each assumes a static setting. Kultti (1999) examines a dynamic setting, obtaining a continuum of payoff-equivalent equilibria in which wait times are equivalent in auctions vs. posted prices. In contrast with the Kultti (1999) model, we find a unique equilibrium predicting that buyer wait times are shorter in posted-price listings than in auctions (and vice-versa for sellers). Bauner (2015) is only dynamic on the sellers' side, whose unsold product can be sold in the next period until a common deadline; but each buyer only participates in the market for one period.

[6]This literature also focuses on online auctions. Earlier work considered a finite sequence of auctions (Milgrom and Weber, 2000; Engelbrecht-Wiggans, 1994; Jeitschko, 1999), which induces a common deadline for all potential buyers, beyond which the good cannot be obtained.

as ours, focus on dynamics between auctions rather than within an auction, which occur instantaneously via a second-price sealed bidding.[7] We contribute to this literature by providing an alternative source of heterogeneity among bidders. Rather than assuming buyers fundamentally differ in their valuations, which are drawn from an exogenous distribution, we assume all buyers are identical at the time they enter the market. This seems plausible in the context of auctioning standardized products.[8] Thus, differences arise endogenously because some unlucky bidders search longer than others, and therefore are nearer their deadline and willing to pay more. A distinguishing prediction of our model, and one for which we find strong empirical evidence, is that a bidder's bid will increase the longer the bidder has participated in auctions for a given item.[9]

Finally, our paper contributes to the literature in industrial organization seeking to explain equilibrium price dispersion, where search frictions allow sellers to charge different prices for a homogeneous good. This behavior is rational for sellers when buyers differ in their search costs (Salop and Stiglitz, 1977; Stahl, 1989; Sorensen, 2000; Schneider, forthcoming) or are inattentive (Malmendier and Lee, 2011a); or sellers obfuscate (Ellison and Wolitzky, 2012) or do not honor previous quotes (Akın and Platt, 2014). In our model, buyers value the good identically, but their differing deadlines create a continuum of dispersed prices (as it did for labor markets in Akın and Platt, 2012). We also provide the first analysis of price dispersion in an auction environment.[10]

We proceed by first discussing several motivating facts from eBay data in Section 2. We then develop the model for buyers and characterize its solution in Section 3, and then close the model in Section 4 by introducing the seller's problem and describing its equilibrium behavior. Section 5 assesses the evidence for bidder learning, an alternative explanation for increasing bids over time. Section 6 concludes. All proofs are found in the Appendix.

[7]The exceptions are Hendricks et al. (2012), where new bidders move last to avoid disclosing their arrival, and Said (2012), where each period features a multi-unit ascending auction. Within-auction dynamics have mostly been studied in the context of a single auction (Nekipelov, 2007; Ambrus et al., 2014) or concurrent auctions (Peters and Severinov, 2006; Ely and Hossain, 2009). In those environments, bidders can benefit from incremental bidding or waiting till the last minute (*sniping*) rather than submitting their true valuation as their only bid (Backus et al., 2014a).

[8]An identical valuation among buyers is also assumed in Hendricks et al. (2012) and Einav et al. (2013).

[9]In Said (2011) and Bodoh-Creed (2012), a bidder's valuation is redrawn between each auction; thus, those participating in their first auction bid the same on average as those participating in their n^{th} auction. In Ingster (2009), a bidder holds the same valuation across all auctions and always submits the same bid. Bidder valuations are also constant in Zeithammer (2006), Backus and Lewis (2012) and Hendricks et al. (2012), but bids adjust in response to other state variables (number of upcoming auctions, information about other past bidders, and number of past bidders, respectively). These states evolve in a Markov process, and would likely dampen any drift in bids over the participation spell.

[10]The model used in both Kultti (1999) and Julien et al. (2001) provides a first step toward price dispersion in an auction environment. A bidder wins for free if he is alone, but competes to his valuation (which is common to all) if anyone else participates.

2 Motivating Facts on Buyer and Seller Behavior

2.1 Data and Descriptive Statistics

Our sample consists of auctions and posted price sales on eBay.com for the year from October 1st 2013 to September 30th 2014.[11] As our model describes the sale of homogeneous goods, we restrict attention to new items which have been matched by the seller to an item in one of several commercially available product catalogs, called a *catalogued product category*. These are narrowly defined, matching a product available at retail stores, such as: "Microsoft Xbox One, 500 GB Black Console", "Chanel No.5 3.4oz, Women's Eau de Parfum", and "The Sopranos - The Complete Series (DVD, 2009)". We remove listings in which multiple quantities were offered for sale; listings which have any outlier bids or posted prices (defined as bids in the top or bottom 1% of bids for auctioned items in that product category, and similarly for posted price sales); and product categories with under 25 auction or posted price sales. All of our bid and price data are calculated inclusive of shipping fees.

Table 1 presents descriptive statistics for this sample. There are over 1.1 million sales in over 6,700 product categories, split roughly evenly between auctions and posted prices. Conditional on sale, the average selling price is higher with posted prices than auctions ($100 versus $83), although some of this difference is attributable to differences in the composition of goods sold in the two selling mechanisms. The average number of bidders per sold auction is 5.54.

To adjust for differences in product values across categories, we follow Einav et al. (2013) and rescale all bids, dividing by the mean price of posted price sales in that product category. In the context of our model below, we demonstrate that bids scale multiplicatively with product values, and thus our rescaling is equivalent to "homogenizing" bids (Haile et al., 2003), as is now standard in the empirical auctions literature.

2.2 Bids Over Time

Our first observation requires that we follow the same bidder across multiple auctions in the same product category. Such bidders constitute an economically significant fraction of all bidding activity — 20% of bidders participate in more than one auction of some product, and collectively, these repeat bidders place 45% of all bids. We find that bidders tend to increase their bids in a given product category from one auction to the next.

[11] For a discussion of the eBay auction mechanism, see Lucking-Reiley et al. (2007). Not included in our analysis are hybrid formats, namely the *buy-it-now auction*, in which the seller simultaneously auctions the item and offers a posted price (Budish and Takeyama, 2001; Kirkegaard and Overgaard, 2008; Bauner, 2015), and posted price sales which allow for buyer-seller bargaining (Backus et al., 2014b).

To compute this, we keep each bidder's highest bid in each auction he participates in. We then arrange these bids in a chronological sequence for each bidder and product category pair, ending when the bidder either wins an auction or does not participate in any more auctions in our sample. We then compute the average and standard deviation of the bids, separately for each sequence length and each step within the sequence.

Figure 1 displays the resulting trend across repeated auction participation. Each line corresponds to a different sequence length, and each point to the mean normalized bid for the corresponding auction in the sequence. Due to our normalization, the bids can be read as a percentage of the item's retail price. For each sequence length, the average bid steadily rises over time from the first to the last auction in the sequence.[12] The slope reflects the average increase in the bid, which varies by individual. Indeed, 12% of bidders hold their bid constant between auctions, 48% increase it, and 40% decrease it. Averaging across all sequence lengths and auction numbers (or from the results of a regression), the bid increases by 2.1% in each successive auction.

To get a sense of whether this increase in bids over time is statistically significant, we regress the normalized bid on fixed effects for the auction number (where the auction appears in the sequence) and the length of the auction sequence, after removing outliers in the auction number variable (defined as the largest 1% of observations). Figure 2 plots the estimated coefficients on the auction number variables, together with 95% confidence intervals, and shows a steady upward trend in bids over time.

2.3 Price Dispersion

Across the repeated auctions of a given product category, there is still substantial dispersion in transaction prices. Figure 3 shows the distribution of the normalized auction closing price (the second highest bid in each auction). This is suggestive that even similar goods may vary considerably in price; the interquartile range of the normalized second-highest bid is 31 percentage points.[13] Some of this dispersion is due to low price items, which sometimes sell for a relatively small or large percentage of their fixed price. Figure 4 displays the same histogram for product categories with a mean fixed price of over $100. Large percentage deviations of auction prices from fixed prices are rarer for these goods, but there remains a good deal of price dispersion: the interquartile range is 14 percentage points. This dispersion remains even after controlling for seller and product fixed-effects: the residuals from the

[12] This increase in bids is not driven by the fact that the final bid in a sequence may be a winning bid, while by construction previous bids are not. Figure 1 shows that even before the final bid in a sequence, bids tend to increase. Nor is it due to selection in the product category mix across the auction number variable, as the sequences are constructed at the bidder by product category level, so conditional on sequence length the product category mix is constant across auction number.

[13] Malmendier and Lee (2011b) also find price dispersion for equivalent goods.

regression of the normalized second-highest bid on seller and product fixed effects have a standard deviation and interquartile range of 8 and 4 percentage points.

2.4 Coexistence of Auctions and Posted Price Sales

Most product categories are sold using a mix of auctions and posted-price listings. This finding necessarily holds in our preferred sample (which only includes product categories with at least 25 transactions of each type), and therefore to document evidence of the mix of sales mechanisms, we turn now to a larger sample that nests our preferred sample. This broader sample includes all products sold at least 50 times in our sample period without regard to listing method, yielding 11,634 product categories. For each product category, we compute the fraction of listings sold via auction. The histogram in Figure 5 indicates the distribution of this fraction across categories. Under 2% of product categories are sold exclusively by auction or by posted price. It is far more common to see a nontrivial fraction of sales through both formats: over 90% of product categories have between 10% and 90% of sales by auction. Averaging across product categories in this broader sample, the average rate of auction use is 48%.

2.5 Auctions and Posted Price Sales Rate

Finally, it is noteworthy that auctions typically yield a completed transaction faster than a posted price listing. The seller explicitly chooses the auction length for either 1, 3, 5, 7 or 10 days. In contrast, posted-price listings are available until a buyer purchases it, and can be renewed if not purchased after 30 days. Figure 6 graphs the cumulative fraction of listings sold against the number of days after listing the item for sale, for auctions and posted price listings. Auctions are about as likely to sell as posted price listings within a day (20% vs 22%), and are over twice as likely to sell within 10 days (85% vs 39%). This is a reason why sellers might prefer to sell by auction: on average, auctions are considerably more likely than posted price listings to sell.

3 Model

We now develop a search-theoretic model of auction participation which provides a unified explanation of the facts discussed above. In this section, we focus exclusively on the bidding decisions of buyers; in the next section, we examine the incentives of sellers in listing their items.

3.1 Buyers

Consider a market for a homogeneous good in a continuous-time environment. Buyers randomly enter the market at Poisson rate δ, seeking one unit of the good which is needed for consumption in T units of time. Each buyer will enjoy the same utility x (measured in dollars) from the good at the time of consumption; if purchased early, that utility is discounted at rate ρ. Thus, if the good is purchased with s units of time remaining until the buyer's deadline, his realized utility is $xe^{-\rho s}$ minus the purchase price.

The good is offered in second-price sealed-bid auctions that occur at Poisson rate α. When such an auction occurs, each buyer in the market participates with exogenous independent probability τ, reflecting that buyers can be distracted by other commitments. Each participant submits a bid and immediately learns the auction outcome, with the highest bidder winning and paying the second highest bid. Alternatively, at any time, a buyer can obtain the good directly at a posted price of z, either as posted-price listings on eBay or elsewhere. We assume throughout that $x \geq z$, so that buyers weakly benefit from purchasing via the posted-price option.

Every buyer shares the same utility x and deadline T on entering the market; but because they enter the market at random times, they differ in their remaining time s. In any given auction, bidders do not know the number of other competing bidders and each bidder's state s is private information. However, the distribution of bidder types in the market (represented by cumulative distribution $F(s)$) is commonly known, as is the stock of buyers in the market, which is Poisson distributed with mean H. Both $F(s)$ and H are endogenously determined, as described below. Note that the Poisson processes governing the entry of buyers (rate δ) and the arrival of auctions (rate α) are independent, and both are independent of the stock of buyers in the market, but all will be related to one another in equilibrium by the steady state conditions in the next subsection.

The strategic question for buyers is what bid to submit and when to purchase from the posted-price listings. This dynamic problem can be expressed recursively, letting $V(s)$ denote the discounted expected utility of a buyer with s units of time remaining until his deadline. Each participant submits a bid $b(s)$ that depends on his remaining time until deadline.[14]

[14] One abstraction in our model is that we assume that bidders do not infer any information about their rivals from prior rounds. Such information is unimportant if valuations are redrawn between each auction, but if valuations are persistent, this *leakage* of private information can hinder future success and leads to further shading of bids (Backus and Lewis, 2012; Said, 2012). In a large market, though, the cost to a bidder of recording hundreds of opposing bidders' past actions seems impractical; also, since not every bidder participates in every auction, the expected value of such information diminishes quickly. Zeithammer (2006) uses the same assumption and justification.

The optimal behavior is to bid one's reservation value, setting:

$$b(s) = xe^{-\rho s} - V(s); \tag{1}$$

that is, the present value of the item minus the opportunity cost of skipping all future auctions. As in the standard second-price sealed-bid auction, this strategy is weakly dominant.[15] We assume that $b(s)$ is decreasing in s (that is, willingness to pay increases as the deadline approaches), and later verify that this holds in equilibrium. In this section, the auctioneer is assumed to open the bidding at $b(T)$, which is relevant in the case that only one buyer participates.[16]

As a fraction τ of buyers participate in a given auction, the number of participants per auction is Poisson distributed with mean $\lambda = \tau H$; that is, the probability that n bidders participate is $e^{-\lambda}\lambda^n/n!$. While this literally governs the *total* number of participants per auction, the same density function also describes (from the perspective of a bidder who has just entered) the probability that n *other* bidders will participate. This convenient parallel between the aggregate distribution (which enters expected revenue and the steady state conditions) and the distribution faced by the individual (which enters his expected utility) is crucial to the tractability of the model but is not merely abuse of notation. Myerson (1998) demonstrates that in Poisson games, the individual player would assess the distribution of other players the same as the external game theorist would assess the distribution for the whole game.

In light of this, a buyer's expected utility in state s can be expressed in the following Bellman equation:

$$\begin{aligned}\rho V(s) = -V'(s) + \tau\alpha\Bigg(&\sum_{n=0}^{\infty}\frac{e^{-\lambda}\lambda^n}{n!}(1-F(s))^n\left(xe^{-\rho s} - V(s)\right)\\ &- e^{-\lambda}b(T) - \sum_{n=1}^{\infty}\frac{e^{-\lambda}\lambda^n}{n!}\int_s^T b(t)n(1-F(t))^{n-1}F'(t)dt\Bigg).\end{aligned} \tag{2}$$

We now explain the continuous-time Bellman equation in (2) piece by piece.[17] The left-hand side represents the flow of expected utility that a buyer with s units of time remaining receives each instant. On the right hand side, we depict any potential changes in

[15] Suppose he instead bids price $p > b(s)$, and the second highest bid is q. This results in the same payoff whenever $q \leq b(s)$, but yields negative surplus when $q \in (b(s), p]$. Similarly, if he bids $p < b(s)$, he has the same payoff whenever $q \leq p$, but misses out on positive surplus when $q \in (p, b(s))$.

[16] We consider seller optimization of the reserve price in the Online Appendix. Provided that the search length T is reasonably short, $b(T)$ is in fact the optimal reservation price due to competition among potential sellers.

[17] A discrete-time formulation of this Bellman equation is provided in the Appendix.

(net) utility times the rate at which those changes occur.[18] For example, the term $-V'(s)$ accounts for the steady passage of time: just by remaining in the market for another unit of time, the buyer's state s decreases by 1 unit (hence the negative sign) and his utility changes by $-V'(s)$.

When an auction occurs and the individual participates in it — which occurs at a rate of $\tau\alpha$ auctions per unit of time — the expected payoff depends on the number n of other participants, which is Poisson distributed with mean λ. The buyer in state s will only win (have the highest bid $b(s)$) if all n other participants have more than s units of time remaining; this occurs with probability $(1-F(s))^n$. When this occurs, the term $xe^{-\rho s}-V(s)$ depicts the change in utility due to winning.

The terms on the second line compute the expected cost of winning (*i.e.* the average second-highest bid times the probability of winning and thus paying it). If there are no other participants (which occurs with probability $e^{-\lambda}$), the bidder pays the starting price $b(T)$. Otherwise, inside the sum we find the probability of facing n opponents, while the integral computes the expected highest bid among those n opponents, which has a probability density of $n(1-F(t))^{n-1}F'(t)$.[19]

Buyers also have the option to purchase from the posted-price listings at any time, receiving utility $xe^{-\rho s}-z$. However, a buyer in state s can obtain a discounted expected utility of $(x-z)e^{-\rho s}$ by waiting until $s=0$ to make the purchase, which is strictly preferred. This delay strategy has even greater payoff due to the possibility of winning an auction in the meantime. Hence, the posted-price option is exercised if and only if $s=0$, and the expected utility of a buyer who reaches his deadline is simply the consumer surplus from making the purchase:

$$V(0) = x - z. \tag{3}$$

3.2 Steady State Conditions

In most auction models, the distribution of willingness to pay is exogenously given as a primitive of the model (Milgrom and Weber, 1982; Athey and Haile, 2002). Here, the buyers' reservation prices $b(s)$ are endogenous, affected by the value of further search ($V(s)$ in Equation 1) in addition to the underlying utility x. The distribution $F(s)$ of buyer states is also endogenously determined by how likely a bidder is to win and thus exit the market at each state, which itself depends on the distribution of competitors he faces.

[18] If buyers received any enjoyment from the search process itself, it would appear as a positive constant on the left-hand side. This would have negligible effect on the solution other than lowering willingness to pay, since making a purchase would cut off this flow of enjoyment.

[19] Note that since the range of integration is only from s to T, $n(1-F(t))^{n-1}F'(t)$ would only sum to $(1-F(s))^n$ rather than 1. This coincides with the probability that this buyer's bid is higher than all opponents, as required.

We require that the distribution of buyers remains constant over time. As buyers exit the market, they are exactly replaced by new customers; as one group of buyers get closer to their deadline, a proportional group follows behind. These steady state requirements are commonly used in equilibrium search theory to make models tractable. Rather than tracking the exact number of current buyers and sellers, which would change with each entry or exit and require a large state space, the aggregate state is always held at its average. This does not eliminate all uncertainty — for instance, the number of bidders in a given auction need not equal the average λ — but these shocks are transitory, as the number of bidders in the next auction is independently drawn from a constant (though endogenous) Poisson distribution. Thus, steady state conditions smooth out the short-run fluctuations around the average, and are interpreted as capturing the long-run average behavior in a market.

To begin the steady state analysis, consider the relative density of bidders over their time until deadline. For instance, consider a cohort in state $s > 0$. In each of the next Δ units of time, suppose on average that a fraction w of these buyers win an auction and exit. Then steady state requires that $F'(s - \Delta) = F'(s) - w\Delta F'(s)$. After rearranging and letting $\Delta \to 0$, we obtain $F''(s) = wF'(s)$. The steady-state law of motion is therefore:

$$F''(s) = F'(s)\tau\alpha \sum_{n=0}^{\infty} \frac{e^{-\lambda}\lambda^n}{n!}(1 - F(s))^n. \tag{4}$$

Here, $\tau\alpha$ is the rate at which buyers participate in an auction, and the summation indicates how likely they are to have the highest bid of n participants and thus win. Recall that all bidders enter at $s = T$; at all other states, some bidders exit. Thus the bidder density $F'(s)$ must decrease as s falls, which holds because the right-hand side of Equation 4 is always positive.

Equation 4 defines the law of motion for the interior of the state space $s \in (0, T)$. The end points are defined by requiring $F(s)$ to be a continuous distribution:

$$\lim_{s \to 0} F(s) = F(0) = 0 \tag{5}$$

$$\lim_{s \to T} F(s) = F(T) = 1. \tag{6}$$

These two conditions prevent a discontinuous jump at either end of the distribution — that is, a positive mass (an *atom*) of buyers who share the same state, $s = 0$ or T. Atoms cannot occur in our environment because all buyers who reach state $s = 0$ immediately purchase from a posted-price listing and exit the market; hence, no stock of state 0 buyers can accumulate. Similarly, no stock of state T buyers can accumulate because as soon as they enter the market, their clock begins steadily counting down. Conveniently, a continuous

distribution ensures that no two bids will tie with positive probability.

Finally, we ensure that the total population of buyers remains steady. Since H is the average number of buyers in the market, $HF(s)$ depicts the average number of buyers with less than s units of time remaining, and $HF'(s)$ denotes the average flow of state s buyers over a unit of time. Thus, we can express the steady state requirement as:

$$\delta = H \cdot F'(T). \tag{7}$$

Recall that buyers exogenously enter the market, averaging δ new buyers in one unit of time. Since all buyers enter the market in state T, this must equal $HF'(T)$, the average number of state T buyers in one unit of time.

3.3 Buyer Equilibrium Definition

The preceding optimization by buyers constitutes a dynamic game. We define a *buyer steady-state equilibrium* of this game as a bid function $b^* : [0, T] \to \mathbb{R}$, a distribution of buyers $F^* : [0, T] \to [0, 1]$, an average number of buyers $H^* \in \mathbb{R}^+$, and an average number of participants per auction $\lambda^* \in \mathbb{R}^+$, such that:

1. Bids b^* satisfy the Bellman equations 1 through 3, taking F^* and λ^* as given.
2. The distribution F^* satisfies the steady state equations 4 through 6.
3. The average number of active buyers H^* satisfies Steady State Equation 7.
4. The average number of participants per auction satisfies $\lambda^* = \tau H^*$.

The first requirement requires buyers to bid optimally; the last three require buyers' beliefs regarding the population of competitors to be consistent with steady state.

3.4 Buyer Equilibrium Characterization

We now present the unique equilibrium of this sequential auctions game for a homogeneous retail good. Our equilibrium requirements can be translated into two second-order differential equations regarding $F(s)$ and $b(s)$. The equations themselves have a closed-form analytic solution, but one boundary condition does not. We solve for the equilibrium λ^* which implicitly solves the boundary condition. Defining $\phi(\lambda)$ as:

$$\phi(\lambda) \equiv \alpha \left(1 - e^{-\lambda}\right) - \delta + \delta e^{\lambda - \tau T\left(\delta + \alpha e^{-\lambda}\right)}, \tag{8}$$

the boundary condition is equivalent to $\phi(\lambda^*) = 0$, which ensures that no atom of buyers occurs at state $s = 0$. The rest of the equilibrium solution is expressed in terms of λ^*.

First, the distribution of buyers over time remaining until deadline is:

$$F^*(s) = 1 - \frac{1}{\lambda^*} \ln \left(\frac{\delta + \alpha e^{-\lambda^*}}{\delta e^{\tau(s-T)\left(\delta + \alpha e^{-\lambda^*}\right)} + \alpha e^{-\lambda^*}} \right). \tag{9}$$

The density function F' is strictly increasing in s. Intuitively, this occurs because some buyers win auctions and thus exit as the clock s counts down. The rate of increase depends on s, with F' typically changing from convex to concave as s increases. This is because buyers rarely win at the beginning of their search, but increasingly do so as time passes and they increase their bids. Those near their deadline win quite frequently, but few of them remain in the population, so their rate of exit decelerates.

The average number of buyers in the market is simply:

$$H^* = \frac{\lambda^*}{\tau}. \tag{10}$$

Equilibrium bids are expressed as a function of the buyer's state, s, as follows:

$$b^*(s) = z e^{-\rho s} \cdot \frac{\tau \left(\delta + \alpha e^{-\lambda^*}\right) \left(\delta e^{\lambda^*} + \alpha e^{\rho(s-T)}\right) + \rho \left(\delta e^{\lambda^*} + \alpha e^{\tau(T-s)\left(\delta + \alpha e^{-\lambda^*}\right)}\right)}{\tau \left(\delta + \alpha e^{-\lambda^*}\right) \left(\delta e^{\lambda^*} + \alpha e^{-\rho T}\right) + \rho \left(\delta e^{\lambda^*} + \alpha e^{\tau T\left(\delta + \alpha e^{-\lambda^*}\right)}\right)}. \tag{11}$$

The following result establishes that this proposed solution is both necessary and sufficient to satisfy the equilibrium requirements. In the proof (provided in the Appendix), we translate the equilibrium conditions into first-order differential equations of $F(s)$ and $b(s)$. Our proposed solution not only satisfies these differential equations, but is the unique solution to them.

Proposition 1. *Equations 9 through 11 satisfy equilibrium conditions 1 through 4, and this equilibrium solution is unique.*

As previously conjectured, one can readily show that $b'(s) < 0$; that is, bids increase as buyers approach their deadline. Moreover, this increase accelerates as the deadline approaches, since $b''(s) > 0$.

Proposition 2. *In equilibrium, $b'(s) < 0$ and $b''(s) > 0$.*

Discounting plays a critical role in creating dispersion among the bidder valuations. For instance, if buyers become extremely patient ($\rho \to 0$), the bidding function approaches $b(s) = z$ regardless of time until deadline.[20] In other words, prices become less dispersed as patience increases, all else equal.

[20] The fractional term of Equation 11 approaches zero. Note that the equilibrium λ^* is unaffected (Eq. 8), as is the distribution of bidders (Eq. 9).

The average time between auctions ($1/\alpha$) is of similar importance. In effect, this is the search friction that buyers face, as it prevents them from making unlimited attempts at winning an auction. In the extreme, if auctions almost never occurred ($\alpha \to 0$), the value of search $V(s)$ drops to zero,[21] so a bidder's reservation price would simply equal his present value of the good: $b(s) = xe^{-\rho s}$. In contrast, for larger values of α, a bidder would optimally reduce his bid well below this, since he is likely to have several opportunities to win a deal before his deadline.

3.5 Calibration

We next calibrate the model to match stylized facts regarding the average product auctioned in our eBay data, using this parameterization to illustrate the equilibrium behavior and to compare the predicted outcomes to the facts reported in Section 2. We begin by normalizing the posted price $z = 1$. This has no effect on the distribution $F(s)$, and $b(s)$ will scale proportionally. We similarly transform bidding data: for each item, we compute the average price among all sold posted-price listings (the analog of z), then divide all bids for that item by the average. We consider one unit of time to be a month, which simply affects the interpretation of T and ρ.

Other parameters are then computed as depicted in Table 2. Targets are computed for each catalogued product category, and then averaged across all categories. We use completed auctions (*i.e.* auctions in which at least one bidder arrived) for our calibration targets rather than listed auctions because transactions in the latter require rational choices by both buyer and seller, while auctions that are listed but not sold could have failed due to bad luck (as in the model) or bad seller decisions (like setting the opening price too high).

The first four items in Table 2 are either directly observed or require minimal adjustment to extract from the data.[22] Note that bidders per auction, λ, is endogenous but observable. We use this to back out the number of new buyers per month, δ, from the equilibrium condition that $\phi(\lambda) = 0$ (Eq. 8).

The final item in Table 2 uses the mean of the price distribution (whose formula is reported in the Appendix with the proof of Proposition 3) to pin down the discount rate, and which must be solved with numeric methods.

The right panel of Figure 7 depicts equilibrium bids as a function of time remaining. Since $z = 1$, these can be read as the factor by which bidders shade their bids below the posted price. Note that prices are dispersed across a range equal to 24% of the posted

[21] In this limit, the equilibrium $\lambda^* = \tau\delta T$, while the is the distribution of bidders is $F(s) = s/T$.

[22] For the fourth row of Table 2, we observe the distribution $d(k)$ of the number k of auctions that a bidder participates in. The model predicts this will follow the Poisson distribution $d(k) = ke^{-\tau\alpha T}(\tau\alpha T)^k/k!$; thus, $k \cdot d(k+1)/d(k) = \tau\alpha T$. While this computation should be the same for all k, in the data it rises from 0.11 for $k = 1$ to 7.9 for $k = 9$, after which it stabilizes. We take the average of these estimates across $k = 5$ to 9.

price; this dispersion becomes larger with higher discounting or longer deadlines. Initially (for near $s = T$), the price path is more or less linear; but as the deadline approaches, greater curvature is introduced.

On average, an individual will increase his bid at a rate of 3.5 percentage points per month. Since the average bidder participates in 1.2 auctions per month, that translates to increase of 2.9 percentage points between each auction of a given item. In the data, we see an increase of 2.1% between each auction attempt, which is the average slope of each line in Figure 1. Bear in mind that the calibration process did not exploit any details about the bids of an individual over time; even the moments of the distribution were across all auctions and all bidders.

Our model also explains why bidders with longer sequences start with a lower initial bid, also seen in Figure 1. Recall that auction participation is stochastic, so some unlucky bidders will wait longer than others to place their first bid. These unlucky bidders will place a higher first bid (because their s is smaller), but they will also have less time until their deadline, and thus expect to participate in fewer additional auctions. Indeed, if we simulate the model under the calibrated parameters and then summarize the data similarly, the resulting Figure 8 is qualitatively similar.

Next, consider the distribution of bids. The left panel of Figure 7 illustrates the equilibrium density of bidders. Note that $F'(s)$ is nearly constant from $s = 3$ to 4.5. With an average of 5 bidders per auction, those with lower valuations (hence longer time remaining) are highly unlikely to win. Yet the relative density cuts in half between $s = 3$ and $s = 1$, and then does so again before $s = 0$. Those closest to their deadline are far more likely to win and exit.

Figure 9 provides another perspective on the realized bids in the auction. The dot-dashed line plots the density of a randomly selected bid in the typical auction. That is, for any price p on the x-axis, the y value indicates the relative likelihood of that price being placed as a bid. Effectively, this is $F'\left(b^{-1}(p)\right) \cdot (b^{-1})'(p)$, obtained via a parametric plot since $b^{-1}(p)$ cannot be analytically derived. This plot shares much in common with the plot of $F'(s)$ in the left panel of Figure 7, reversed in direction since the highest bids come from the bidders closest to their deadline. This is contrasted with the dotted line, which plots the density of the highest bid in each auction. Note that this is concentrated more towards higher prices, even though bidders with those valuations are somewhat scarce. This is because the highest of 5 bids tends to be closer to the top of available bids. Of course, only the second-highest price is actually paid; this density is depicted with the solid line. Despite the uniform nature of the auctioned goods, closing prices are significantly dispersed, with an interquartile range of 4.4 percentage points.

The notable departure of the calibrated model from the data is in the distribution of

bid sequences. In our data, 85% of bidders only bid once on a product, while the model predicts that only 24% should place a single bid. These single-auction participants have a low success rate at winning auctions, so that only 42% of bidders win an auction; in contrast, our model predicts that 95% will eventually win an item.

3.6 Comparative Statics

We next examine how the equilibrium behavior reacts to changes in the underlying parameters. Although our equilibrium has no closed form solution, these comparative statics can be obtained by implicit differentiation of $\phi(k)$, which allows for analytic derivations reported in the Online Appendix.

Table 3 reports the sign of the derivatives of three key statistics. First is the average number of participants per auction, λ^*, which reflects how competitive the auction is among buyers. Note that the average number of buyers in the market, H^* is always proportional to λ^*. Second is the flow of buyers who never win an auction and thus resort to the posted-price listings; in the next section we will see that this crucially affects the profitability of the posted-price market. Third is the bid of new buyers in the market, indicating the effect on buyers' willingness to pay. This comparative static can be derived at any s and has a consistent effect, but the computation is easiest to report at $s = T$.

Changes in α have a intuitive impact. If auctions arrive more frequently, this reduces search frictions; that is, the value of continued search is greater as there are more opportunities to bid. Moreover, the increase in auctions will reduce the stock of bidders and hence the number of competitors per auction. Both of these effects lead bidders to lower reservation prices.

Changes in τ have nearly the reverse effect, though here there are opposing forces at work. Having a higher likelihood of participating also reduces the search friction of a given bidder, as he will participate in more of the existing auctions. On the other hand, all other bidders are more likely to participate as well. This greater number of competitors dominates the increased auction participation to reduce the value of search, increasing bidders' bids.

The discount rate, the value of the good, and the fraction of immediate consumption utility all have no impact on the number or distribution of bidders, which can be seen mathematically in the fact that these parameters do not enter into Equations 8 through 10. Intuitively, this is because the rate at which bidders exit is a matter of how often auctions occur, which is exogenous here. Also, who exits is a matter of the ordinal ranking of their valuations, which does not change even if the cardinal values are altered. Indeed, the bids react as one would expect: buyers offer more when their utility from consumption is higher or more immediate.

We can also consider the effect of the parameter change on the expected revenue gen-

erated in an auction, which we formally derive in the next section. In all the preceding comparative statics, revenue moves in the same direction as bids because the number of participants per auction is either constant or moves in the same direction. For instance, more auctions will reduce the bids and reduce the number of bidders; thus expected revenue must be lower. The intriguing exception is when the deadline changes. Intuitively, one would expect that an increase in T would work to buyers' advantage. Indeed, bids are lower, but only because bidders now enter with more time and hence a lower starting bid, $b(T)$. At the same time, the longer period of search allows for more bidders to accumulate (H^* increases). Thus, each auction includes more bidders, driving up expected revenue.

3.7 Consumer Surplus

A distinctive aspect of our model is the endogenous distribution of bidders' willingness to pay. This feature has important consequences for interpreting empirical data from bidding, a fact that is easily illustrated in terms of the consumer surplus generated in the market.

In a static second-price auction with private valuations independently drawn from an exogenous distribution, the consumer surplus from the auction is the difference between the first price (which truthfully reveals the winner's valuation) and the second price (which the winner actually pays). In our setting, however, all bids have been shaded down by each bidder's continuation value, $V(s)$; thus, the true consumer surplus is the first price plus $V(s)$ of the winner minus the second price. Thus, if the data generating process came from our deadline model but were interpreted as in the static approach, the econometrician would underestimate the true valuation of each bidder and hence the consumer surplus in the market.

To provide precise comparison, consider the expected consumer surplus from a random auction, conditional on that auction having any participants. In our model, this would be:

$$\begin{aligned} CS_{\text{model}} = \frac{1}{1-e^{-\lambda}} \Bigg(& \int_0^T x e^{-\rho s} \lambda e^{-\lambda F(s)} F'(s) ds \\ & - \int_0^T b(s) \lambda^2 e^{-\lambda F(s)} F(s) F'(s) ds - e^{-\lambda} b(T) \Bigg). \end{aligned} \tag{12}$$

In the first integral, the remaining time until deadline, s, of the highest bidder is distributed according to $\lambda e^{-\lambda F(s)} F'(s)$ (after evaluating the sum over the Poisson distribution of bidders in the auction), and the utility enjoyed by the highest bidder is $xe^{-\rho s}$. In the second integral, we determine the average of the second highest bid, which is distributed according to $\lambda^2 e^{-\lambda F(t)} F(t) F'(t)$. The final term accounts for when only one bidder arrives and thus wins at price $b(T)$.

In contrast, if the bids were mistakenly interpreted as being truthful, as in a static model, we would replace $xe^{-\rho s}$ in Equation 12 with $b(s)$, which simplifies to:

$$CS_{\text{static}} = \frac{\lambda}{1-e^{-\lambda}} \left(\int_0^T b(s)e^{-\lambda F(s)}(1-\lambda F(s))F'(s)ds - e^{-\lambda}b(T) \right). \tag{13}$$

The difference between these calculations in smallest when $x = z$. Evaluating the remaining parameters at their values from the calibration, we find $CS_{\text{model}} = 0.098$. In other words, the true consumer surplus is on average 9.8% of the retail value of the good, or 11.7% of the average revenue raised in the auction (0.84). However, if the same data were interpreted in a static context, we would conclude that $CS_{\text{static}} = 0.052$; that is, 5.2% of the retail value of the good, or 6.1% of the auction revenue. Thus, the static calculation underestimates consumer surplus by 47.5%.

The same miscalculation would be relevant when the empirical distribution of bids is used to estimate the demand curve; a static interpretation would report demand that is artificially lower (by 4.6 percentage points, under our calibration). To be fair, the observed distributions of bids reflects what buyers are willing to pay at that moment in that market format. However, it falls short of measuring the full value to the consumers that one expects from a demand curve. Using the incorrect demand curve could easily distort calculations needed for profit maximization, price discrimination, regulation, and many other applications.

4 Seller Incentives

We next examine optimization by sellers in this environment, allowing them to decide whether to enter the market and whether to sell their product via auctions or the posted-price listing. We consider a continuum of sellers producing an identical good. Each has negligible effect on the market, taking the behavior of other sellers and the distribution and bidding strategy of buyers as given; yet collectively, their decisions determine the frequency with which auctions occur. In other words, by modeling seller choices we endogenously determine α from the preceding section.

Each seller can produce one unit of the good at a marginal cost of $c < z$, with fraction γ of this cost incurred at the time the good is sold (the *completion cost*), and $1-\gamma$ incurred when seller first enters the market (the *initial production cost*).[23] For either selling format,

[23] In the extreme, $\gamma = 1$ would indicate the ability to build-to-order or just-in-time inventories, while $\gamma = 0$ indicates a need to build in advance (like a spec home built without a committed buyer). Intermediate values could be taken literally as partial production, or as full initial production followed by additional expenses (such as shipping costs) at the time of sale. It could also reflect producing in advance but delaying full payment of the cost through the use of credit.

sellers also pay a listing fee of ℓ each unit of time from when the seller enters the market to when the good is sold. We assume that there are no barriers to entry for sellers. Upon entry, each seller must decide whether to join the auction or the posted-price market.

4.1 Auction Sellers

The advantage of the auction sector is that the sale occurs more quickly. Let η denote the exogenous Poisson rate of auction closing, so $1/\eta$ is the average time delay between the listing and closing of an auction.[24] At its conclusion, the auction's expected revenue (conditional on at least one bidder participating) is denoted θ and computed as follows:

$$\theta \equiv \frac{1}{1-e^{-\lambda}}\left(\lambda e^{-\lambda} b(T) + \sum_{n=2}^{\infty} \frac{e^{-\lambda}\lambda^n}{n!} \int_0^T b(s) n(n-1) F(s)(1-F(s))^{n-2} F'(s) ds\right). \quad (14)$$

Inside the parentheses, the first term applies when only one bidder participates and therefore wins at the opening price of $b(T)$; this happens with probability $\lambda e^{-\lambda}$. The sum handles cases when there are $n \geq 2$ simultaneous bidders, with the integral computing the expected bid $b(s)$ of the second-highest bidder. All of this must be divided by the probability that at least one bidder arrives, $1 - e^{-\lambda}$.

To determine the expected profit of the auction seller, we must also account for the costs of production and the expected time delay. Let Π_a denote the expected profit from the vantage of someone who has already incurred the initial production cost $(1-\gamma)c$ and has just posted the listing. This profit can be computed in the following Bellman equation:

$$\rho \Pi_a = -\ell + \eta\left(1 - e^{-\lambda}\right)(\theta - \gamma c - \Pi_a). \quad (15)$$

The right-hand side of (15) indicates that the seller incurs the listing fee per unit of time. The listing closes at Poisson rate η, but if no bidders have arrived (which occurs with probability $e^{-\lambda}$) then the seller re-lists the item and continues waiting for the new auction to close. If at least one bidder participates, the seller's net gain is the realized benefit (revenue minus the completion cost) relative to the expected profit. Of course, from the perspective of a potential entrant as an auction seller, the expected profits from entry are net of the initial production cost: $\Pi_a - (1-\gamma)c$.

4.2 Posted-Price Sellers

The posted-price listing will always sell the good at a higher price, since for all bids, $b(s) < z$; the disadvantage of this format is that sellers may wait a considerable time before being

[24] Under the default setting at eBay, an auction concludes one week after creating the listing, which provides time for bidders to examine the listing. In our sample, 45% of listings use this default setting.

chosen by a buyer. Let ζ denote the rate of encountering a customer, so $1/\zeta$ is the average wait of a posted-price seller. Sellers take ζ as given, but it will be endogenously determined as described in the next subsection.

The discounted expected profit of posted-price sellers already in the market, denoted Π_p, is computed in the following Bellman equation:

$$\rho\Pi_p = -\ell + \zeta\left(z - \gamma c - \Pi_p\right). \tag{16}$$

Like auction sellers, posted-price sellers incur the listing fee ℓ per unit of time that they await a buyer.[25] When they encounter a buyer (which they do at rate ζ), the purchase always occurs, with a net gain of $z - \gamma c$ relative to Π_p. For sellers contemplating entry into the posted-price market, their expected profit is $\Pi_p - (1-\gamma)c$.

4.3 Steady State Conditions

As with the population of buyers, the stock and flow of sellers are also assumed to remain stable over time. In the aggregate, recall that δ buyers enter (and exit) the market over a unit of time; thus, we need an identical flow of δ sellers entering per unit of time so as to replenish the δ units sold.[26]

In addition, the number of sellers in each market must remain steady. Let σ be the fraction of newly-entered sellers joining the auction market, so that $\sigma\delta$ choose to list an auction over a unit of time. This must equal the number of auctions that close with at least one bidder over the same unit of time:

$$\sigma\delta = \alpha\left(1 - e^{-\lambda}\right). \tag{17}$$

The remaining $(1-\sigma)\delta$ sellers flow into the posted-price market over a unit of time. This must equal the flow of purchases made by buyers that hit their deadline:

$$(1-\sigma)\delta = HF'(0). \tag{18}$$

At any moment, both markets will have a stock of active listings — sellers who are waiting for a buyer to make a purchase or for their auction to close. Let A denote the measure of auction sellers with active listings, and P denote the same for posted-price sellers. From the perspective of the individual auction seller, his auction will close at rate η; but with A sellers in the market at any instant, there will be ηA auctions that close over

[25] Very little changes if we allow the listing fee to differ across the two selling formats.

[26] While we refer to each seller as producing a single unit, one could also think of a seller offering multiple units so long as the production and listing costs scale up proportionately.

a unit of time. From the buyer's perspective, α auctions close over a unit of time; thus, these must equate in equilibrium:

$$\eta A = \alpha. \tag{19}$$

A similar condition applies to posted-price sellers. In aggregate, $HF'(0)$ purchases occur over a unit of time (sold to buyers who reach their deadline). From the individual seller's perspective, he can sell ζ units over one unit of time; collectively, these sellers expect to sell ζP units. In equilibrium, the expected sales must equal the expected purchases:

$$\zeta P \quad = \quad HF'(0). \tag{20}$$

4.4 Market Equilibrium Definition

With the addition of the seller's problem, we augment the equilibrium definition with three conditions. A *market steady-state equilibrium* consists of a buyer equilibrium as well as expected revenue $\theta^* \in \mathbb{R}^+$, expected profits $\Pi_a^* \in \mathbb{R}^+$ and $\Pi_p^* \in \mathbb{R}^+$, arrival rates $\alpha^* \in \mathbb{R}^+$ and $\zeta^* \in \mathbb{R}^+$, seller stocks $A^* \in \mathbb{R}^+$ and $P^* \in \mathbb{R}^+$, and fraction of sellers who enter the auction sector, $\sigma^* \in [0, 1]$, such that:

1. Expected revenue θ^* is computed using the bidding function $b^*(s)$ and distribution $F^*(s)$ derived from the buyer equilibrium, given α^*.

2. Prospective posted-price entrants earn zero expected profits: $\Pi_p^* = (1-\gamma)c$, given ζ^*.

3. Prospective auction entrants earn zero expected profits: $\Pi_a^* = (1-\gamma)c$ if $\alpha^* > 0$, or $\Pi_a^* \leq (1-\gamma)c$ if $\alpha^* = 0$.

4. α^*, ζ^*, σ^*, A^*, and P^* satisfy the steady state equations 17 through 20.

The first requirement simply imposes that buyers behave optimally as developed in Section 3, given the endogenously-determined auction arrival rate. The fourth imposes the steady state conditions.

The second and third requirements impose zero profits for both types of sellers, which is necessary because of the large, unrestricted pool of potential entrants. If either market offered positive profits, additional sellers would be attracted to that market. This in turn would reduce profits: more posted-price sellers P would reduce the rate of selling ζ, and more auction sellers A would increase the auction arrival rate and thus decrease expected revenue θ. Together, these two requirements also ensure that sellers are indifferent about which market they enter, thereby allowing them to randomize according to mixed strategy σ.

In the third requirement, we allow for the possibility that no auctions are offered, but this can only occur if the expected revenue from an auction would be weakly less than that of a posted-price listing. A similar possibility could be added to the posted-price market, but that market would never shut down in equilibrium. Due to the search friction, a fraction of buyers will inevitably reach their deadline; as a consequence, the posted-price market can always break even by reducing the stock of sellers waiting to serve these desperate buyers.

Since the posted price always exceeds the realized auction price, expected profits can only be equated if auction listings are sold more quickly. In other words, if both types of listings are offered in equilibrium, then $\zeta^* < \eta$.

4.5 Market Equilibrium Characterization

While the market equilibrium conditions simplify considerably, they do not admit an analytic solution and we must numerically solve for α^* and λ^* simultaneously. All other equilibrium objects can be expressed in terms of these. The key additional equation comes from the third equilibrium requirement:

$$\theta^* = c + \frac{\ell + \rho c(1-\gamma)}{\eta\,(1 - e^{-\lambda^*})}. \tag{21}$$

This ensures that the expected revenue from each auction precisely covers the expected cost of listing and producing the good. The costs (on the right-hand side) are affected by λ because of the (small) chance that no bidders arrive, while expected revenue (on the left-hand side) is affected by both λ and α because of their influence on the bidding function and distribution of buyers. To compute θ^*, Equation 14 must be evaluated using $b(s)$ and $F(s)$ from the buyer equilibrium; the resulting equation is cumbersome and is reported in the proof of Proposition 3 in the Appendix. At the same time, a buyer equilibrium requires that $\phi(\lambda^*) = 0$ from Equation 8; here, we note that this equation involves both λ and α. Equilibrium is attained when both Equation 8 and 21 simultaneously hold, which can only be solved numerically.

Once α^* and λ^* are found, the remaining equilibrium objects are easily solved as follows:

$$\Pi_a^* = (1-\gamma)c \tag{22}$$

$$\Pi_p^* = (1-\gamma)c \tag{23}$$

$$A^* = \frac{\alpha^*}{\eta} \tag{24}$$

$$P^* = \frac{(z-c)\left(\delta - \alpha^*\left(1-e^{-\lambda^*}\right)\right)}{\ell + \rho(1-\gamma)c} \tag{25}$$

$$\zeta^* = \frac{\ell + \rho(1-\gamma)c}{z-c} \tag{26}$$

$$\sigma^* = \frac{\alpha^*\left(1-e^{-\lambda^*}\right)}{\delta}. \tag{27}$$

It is readily apparent that $\sigma^* \geq 0$. To see that $\sigma^* < 1$, note that the equilibrium condition $\phi(\lambda^*) = 0$ requires that $\alpha\left(1-e^{-\lambda}\right) < \delta$. This also ensures that $P^* > 0$.

The following proposition demonstrates that these solutions are necessary for any equilibrium in which auctions actually take place.

Proposition 3. *A market equilibrium with active auctions ($\alpha^* > 0$) must satisfy $\phi(\lambda^*) = 0$, Equations 9 through 11, and Equations 21 through 27.*

The solution described in Proposition 3 can appropriately be called a *dispersed equilibrium*, to use the language of equilibrium search theory, because we observe the homogeneous good being sold at a variety of prices. Contrast this with a *degenerate equilibrium*, in which the good is always sold at the same price. In our context, this only happens if all goods are purchased via posted-price listings and no auctions are offered ($\alpha^* = \sigma^* = 0$). We can analytically solve for this degenerate equilibrium and for the conditions under which it exists, as described in the following proposition.

Proposition 4. *The degenerate market equilibrium, described by Equations 9 through 11 and Equations 23 through 27 with $\alpha^* = 0$ and $\lambda^* = \tau\delta T$, exists if and only if*

$$z \cdot \frac{\delta\tau\left(\delta\tau + e^{-(\delta\tau+\rho)T}\left(-\delta\tau + \delta\rho\tau T + \rho^2 T\right)\right)}{(\delta\tau+\rho)^2\left(1-e^{-\tau\delta T}\right)} \leq c + \frac{\ell + \rho c(1-\gamma)}{\eta\left(1-e^{-\tau\delta T}\right)}. \tag{28}$$

The left-hand side of Equation 28 is the expression for θ when $\alpha = 0$; it calculates the expected revenue that a seller would earn by deviating from $\alpha = 0$, offering an auction when no one else does. For this equilibrium to exist, the expected revenue must be lower than the expected cost of entering the market (the right-hand side of Eq. 28), so that not selling in the auction market is a best response. We can consider such a deviation because buyers still wait until their deadline before purchasing via the posted-price listing. Thus,in

this equilibrium there would be $H^* = \delta T$ buyers in the market, uniformly distributed on $s \in [0, T]$, who would be available as bidders in the measure-zero event that an auction occurs and would bid their reservation price $b(s) = ze^{-\rho s}$.

Equation 28 provides insight on when a degenerate solution will occur. For instance, on the right-hand side, one can see that a high production cost or listing cost can make the auction market unprofitable. The posted-price market can compensate for these costs by keeping a low stock of sellers so that the item is sold very quickly. Long delays before closing the auction (a small η) also increase the cost of the auction. On the left-hand side, δ and τ have the largest impact of any of the parameters on expected revenue. With either a small flow of new buyers entering the market or a tiny fraction of them paying attention to a given auction, the number of participants per auction will be low. Without much competition in the second-price auction, expected revenue will be too low to cover expected costs.

In equilibrium search models, a degenerate equilibrium often exist regardless of parameter values, essentially as a self fulfilling prophecy. Buyers won't search if there is only one price offered, and sellers won't compete with differing prices if buyers don't search. Yet in our auction environment, the degenerate equilibrium does not always exist. This is because our buyers do not incur any cost to watch for auctions; even if no auctions are expected, buyers are still passively available should one occur.[27] In that sense, they are always searching, giving sellers motivation to offer auctions (if Equation 28 does not hold).

In fact, it appears that the degenerate equilibrium is mutually exclusive of the dispersed equilibrium, such that one exists if and only if the other does not. The complicated expression for θ^* in the dispersed equilibrium precludes an analytic proof of this conjecture, but we have observed this outcome in numerous calculations across a wide variety of parameters.

The underlying cause is quite intuitive: all else equal, more auctions lead to lower expected revenue per auction. Indeed, the comparative statics in Section 3.6 revealed that more auctions lead to fewer bidders per auction and cause buyers to lower their bids in anticipation of extra opportunities to win future auctions — both of which reduce expected revenue. That is to say, we conjecture that θ is a decreasing function of α (including the indirect effect of α on λ), and if so, uniqueness of equilibrium is assured. For example, if θ is less than the expected cost when $\alpha = 0$, then the degenerate equilibrium exists but the dispersed equilibrium cannot, since increasing α will further lower θ. On the other hand, if θ equals the expected cost for some $\alpha^* > 0$, then the dispersed equilibrium exists but the degenerate equilibrium cannot, since decreasing α to 0 will increase θ. Thus, the expected revenue from deviating from the no-auction equilibrium would strictly exceed the expected cost.

[27] For example, eBay allows users to create "Searches you follow," in which the user enters search terms and is periodically informed of any newly listed items that match that search criteria.

To illustrate the outcome of this augmented model, we continue our calibration from Section 3.5. First, we can directly observe the average listing fee paid and the average time for which an auction is listed. We also observe the fraction of posted price listing that sell within the 30 day window of their listing, which is equivalent to ζ, the endogenous rate at which posted-price listings sell. We also note that α, which was calibrated in Table 2, is now endogenous. Here, we can use the equilibrium condition for α (Eq. 21) to determine the underlying cost of production. The last two conditions are solved jointly to recover γ and c.

Under these parameters, a sale in the posted-price market generates $z - \theta = 16.2\%$ more revenue, but this is offset in that the sale occurs after 35 periods on average. Note that we did not use the fraction of sellers in each format as a calibration target, providing us with another check for goodness of fit. Our model predicts that $\sigma^* = 84\%$ of new sellers will use the auction format, and that the stock of posted-price listings should be $P^*/A^* = 1.72$ times bigger than the stock of auction listings.

4.6 Comparative Statics

We now present comparative statics for the market equilibrium. Here, the computation of θ^* prevents analytic determination of the sign of the comparative statics, but numeric evaluation remains consistent over a large space of parameter values. We highlight a few of these, noting that even if these phenomena may not always occur, it is striking that they occur at all.

First, we note that for increases in τ under the market equilibrium, the bidding response is the opposite of in the buyer equilibrium. The difference is that when buyers are more attentive, sellers are willing to offer more auctions are offered in the market equilibrium. Additional auctions will improve the continuation value of buyers, and thus reduce their bids. There are still more participants per auction (which led to higher bids in the buyer equilibrium), but the effect of more auctions dominates to produce a net decline in bids and expected revenue.

In the buyer equilibrium, an increase in ρ reduced bids but had no effect on the distribution of buyers. In a market equilibrium, bids will still fall, but sellers offer fewer auctions. Surprisingly, this leads to higher revenue per auction, as it concentrates more buyers per auction. Changes in T behave similarly under either equilibrium definition.

For c, it is remarkable that even though increased production costs do not raise the retail price (by assumption), they still affect auctions in the distribution of buyers and their bids. Higher costs will shrink the margins in both markets, which the auction market responds to by reducing its flow of sellers. Fewer auctions necessarily mean that more buyers reach their deadline; and this increased demand for posted-price listings more than compensates

for the smaller margin. That is, a larger stock of posted-price sellers is needed to return to normal profits. Also, with fewer available auctions, buyers have a lower continuation value from waiting for future auctions. This drives up bidders' reservation prices, but not enough to prevent a smaller flow of auction sellers.

A higher listing fee, ℓ, has a similar effect, but this is surprising because the listing fee applies to both markets and will be paid more often by posted-price sellers (who have to relist their good more times). This drives sellers away from the auction market, and any reduction in the rate of auctions will increase bids.

This subtle response illustrates a potential hazard in auction design if buyer valuations are not fundamental but rather the endogenous results of deeper factors. A seemingly neutral change in the auction listing fee not only alters which market sellers use, but also warps the distribution of buyer valuations.

5 The Role of Bidder Learning

Deadlines provide an explanation for the robust pattern observed in our data of bidders increasing their bids over time. Another possible explanation might involve bidder learning. Consider the case where bidders are uncertain about the degree of competition they face, and form different estimates of its intensity.

A bidder who underestimates the number of competitors or the bids of competitors will overestimate his likelihood of winning in future auctions; this raises his continuation value and causes him to shade his bid lower. Such a bidder will gradually revise his estimates upwards as he fails to win auctions, and thus tend to bid more over time. On the other hand, bidders who overestimate the amount of competition will bid more aggressively than those who underestimate. However, their initial aggressive bidding will tend to result in their winning auctions early on; they may not remain in the market for long enough to learn their way to lower bids. Thus in principle, bidder learning could also explain the pattern of bidders increasing their bids over time. However, we find substantial evidence that bidder learning is not the sole driver of increasing bids.

First, consumers do not need to participate in auctions to learn about the prices at which they are closing. An eBay user can easily view current prices for numerous auctions of a given product through a simple search, including those that will end soon and are thus near their final price. Moreover, any user can review auctions that ended in the last three months by selecting the "Sold Listings" checkbox on the search results page. This readily indicates the final price and the bid history in each auction. This would provide much more information from which to learn than participation in a single auction.

If bidders were to see this research as too costly, they might rationally choose to learn

through participation instead. However, learning through participation carries its own cost, in that those who overestimate initially are more likely to win and thus overpay. We would expect this type of learning by doing to be more costly for expensive items as the potential loss from mistakenly overpaying would be higher. The cost of research through searching on eBay, however, is more or less fixed. Thus, if bids only increase due to learning, we would very little increase among expensive items. Yet Figure 10 shows the same pattern of bidders increasing their bids with time, even for those products which sell at a median price of over $100.

Alternatively, perhaps some bidders are unaware that they can search completed auctions. But even users who have a good deal of experience bidding on eBay, who are likely to be aware of these kinds of basic site features, also increase their bids over time. Figure 11 makes this point. It shows the same graphs as in Figure 1, for the subset of people who have bid in at least 50 other auctions in the 12 months preceding the month of the current auction. Average bids over time are also increasing for these subsets of heavy eBay users, and the magnitude of the increases are similar to the increases seen in the whole sample in Figure 1.

Another category of people for whom learning is less likely to be occurring at the time of bidding are those who have participated in auctions for similar goods in the past. If someone has participated in multiple auctions for DVDs in the past, they are likely to be aware of the kind of discount they can expect on eBay when bidding on DVDs, and thus have little to learn through additional participation. Figure 12 is again constructed analogously to Figure 1, but for the subsample of users who have participated in at least 10 auctions in the past year for products in the same product grouping as the product currently being bid on.[28] Again, there appears to be an upward trend in average bids over time.

Table 6 provides another metric to compare these several groups. For each subsample of bidders described above, we report the fraction of consecutive auctions in which bids increase, remain constant, or decrease. Note that the fraction who decrease their bid remains essentially constant near 40%, while bidders are 5 to 9 percentage points more likely to increase their bid. Indeed, the small amount of variation among these mostly comes from whether the bidders hold their bid constant or increase it. The fact that a robust trend of increasing bids over time exists for all groups indicates that bidder learning is unlikely to be the driving factor for this trend.

A separate piece of evidence comes from the relation between past and future bidding activity. If learning is the sole explanation for increasing bids over time, then a bidder's bid increase from one auction to the next should be well predicted by the amount by which the first auction's revenue exceeds his bid. If this gap is large, then the second highest

[28] Examples of product groupings are DVDs, video games, and cell phones.

bidder (who determines auction revenue) bid much more than he did, revealing that he is likely underestimating the degree of competition in the market. Consequently, he should revise downwards his beliefs about the option value of continuing to wait for auctions, and increase his bids in the next auction. Table 7 shows the results of regressing the bid increase from the n to $n+1^{th}$ auction in a bidder's bidding sequence, on the amounts by which the auction revenue exceeded that bidder's bids in the previous n auctions, for $n \in \{1, 2, 3, 4\}$. The estimated coefficients on revenue minus bid in the last auction are indeed positive and significant, consistent with the learning hypothesis. However quantitatively the explanatory power of the bidder learning story is rather weak. Bidding experiences in prior auctions only explain about 4-7% of the total variation in the bid increases from one auction to the next, suggesting that other factors — including deadlines — also have a substantial role to play.

6 Conclusion

This work reexamines the auction environment as a venue for selling retail goods. Our analysis leverages methods frequently used in search theory, which provide analytic tractability and plausibly match the auction setting. Standard auction theory ascribes all variation in bids as generated by exogenous differences in valuations; but this seems less compelling when there is a readily-available outside option for the same item. In our model, all buyers enjoy the same eventual utility from the good, but endogenously differ in how soon they must acquire it. This produces an increasing and accelerating path for an individual's bidding over time, and a rich continuous distribution of closing prices across auctions.

The model makes heavy use of homogeneity, with buyers who only differ in how long they have been in the market and sellers who only differ because of their mixed strategy of which market to enter. This homogeneity makes the solution more tractable, but also serves to set the model in sharp relief to those with exogenous valuations. Indeed, recovering the distribution of valuations is a key focus of empirical work in the auction literature, yet it is plausible that these distributions are themselves a product of more fundamental factors, such as impending deadlines. If so, one should not expect the distribution of valuations to be invariant to policy interventions or changes in the auction design. As we have shown, even something as simple as increasing the listing fee — for both auctions and posted-price listings — can change the distribution of buyers' willingness to pay.

At the same time, one could add realism by allowing heterogeneity among buyers or sellers. For instance, buyers could enter with differing deadlines or differing final values, or sellers could differ in their costs of production or discount rate. These extensions would complicate our steady state conditions, but the mechanisms underlying our solution would

still govern the final solution.

References

Akin, Ş. N. and B. C. Platt (2012): "Running Out Of Time: Limited Unemployment Benefits And Reservation Wages," *Review of Economic Dynamics*, 15, 149–170.

——— (2014): "A Theory of Search with Deadlines and Uncertain Recall," *Economic Theory*, 55, 101–133.

Ambrus, A., J. Burns, and Y. Ishii (2014): "Gradual Bidding in eBay-Like Auctions," *Mimeo, Harvard University.*

Athey, S. and P. A. Haile (2002): "Identification of Standard Auction Models," *Econometrica*, 70, 2107–2140.

Backus, M., T. Blake, D. Masterov, and S. Tadelis (2014a): "Is Sniping A Problem For Online Auction Markets?" *Mimeo.*

Backus, M., T. Blake, and S. Tadelis (2014b): "Cheap-Talk, Round-Numbers, and the Economics of Negotiation," *Mimeo.*

Backus, M. and G. Lewis (2012): "A Demand System for a Dynamic Auction Market with Directed Search," *Mimeo, Harvard University.*

Bauner, C. (2015): "Mechanism Choice and the Buy-It-Now Auction: A Structural Model of Competing Buyers and Sellers," *International Journal of Industrial Organization*, 38, 19–31.

Bodoh-Creed, A. (2012): "Optimal Platform Fees for Large Dynamic Auction Markets," *Mimeo, Cornell University.*

Budish, E. B. and L. N. Takeyama (2001): "Buy Prices In Online Auctions: Irrationality On The Internet?" *Economics Letters*, 72, 325–333.

Caldentey, R. and G. Vulcano (2007): "Online Auction And List Price Revenue Management," *Management Science*, 53, 795–813.

Einav, L., C. Farronato, J. D. Levin, and N. Sundaresan (2013): "Sales Mechanisms in Online Markets: What Happened to Internet Auctions?" .

Ellison, G. and A. Wolitzky (2012): "A Search Cost Model Of Obfuscation," *The RAND Journal of Economics*, 43, 417–441.

Ely, J. C. and T. Hossain (2009): "Sniping and squatting in auction markets," *American Economic Journal: Microeconomics*, 68–94.

Engelbrecht-Wiggans, R. (1994): "Sequential Auctions Of Stochastically Equivalent Objects," *Economics Letters*, 44, 87–90.

Etzion, H., E. Pinker, and A. Seidmann (2006): "Analyzing The Simultaneous Use Of Auctions And Posted Prices For Online Selling," *Manufacturing & Service Operations Management*, 8, 68–91.

Haile, P., H. Hong, and M. Shum (2003): "Nonparametric tests for common values at first-price sealed-bid auctions," Working Paper, Yale.

Hammond, R. G. (2013): "A Structural Model Of Competing Sellers: Auctions And Posted Prices," *European Economic Review*, 60, 52–68.

Hendricks, K., I. Onur, and T. Wiseman (2012): "Last-minute Bidding In Sequential Auctions With Unobserved, Stochastic Entry," *Review of Industrial Organization*, 40, 1–19.

Ingster, A. (2009): "A Structural Study Of A Dynamic Auction Market," *Mimeo, New York University*.

Jeitschko, T. D. (1999): "Equilibrium Price Paths In Sequential Auctions With Stochastic Supply," *Economics Letters*, 64, 67–72.

Julien, B., J. Kennes, and I. P. King (2001): "Auctions and Posted Prices in Directed Search Equilibrium," *Topics in Macroeconomics*, 1.

Kirkegaard, R. and P. B. Overgaard (2008): "Buy-out Prices In Auctions: Seller Competition And Multi-unit Demands," *The RAND Journal of Economics*, 39, 770–789.

Kultti, K. (1999): "Equivalence Of Auctions And Posted Prices," *Games and Economic behavior*, 27, 106–113.

Lucking-Reiley, D., D. Bryan, N. Prasad, and D. Reeves (2007): "Pennies from eBay: The determinants of price in online auctions*," *The Journal of Industrial Economics*, 55, 223–233.

Malmendier, U. and Y. H. Lee (2011a): "The Bidder's Curse," *The American Economic Review*, 749–787.

——— (2011b): "The Bidder's Curse," *The American Economic Review*, 749–787.

Milgrom, P. R. and R. J. Weber (1982): "A Theory Of Auctions And Competitive Bidding," *Econometrica: Journal of the Econometric Society*, 1089–1122.

——— (2000): "A Theory Of Auctions And Competitive Bidding II," in *The Economic Theory of Auctions*, ed. by P. Klemperer, Edward Elgar Publishing Ltd.

Myerson, R. B. (1998): "Population Uncertainty And Poisson Games," *International Journal of Game Theory*, 27, 375–392.

Nekipelov, D. (2007): "Entry Deterrence And Learning Prevention On Ebay," *Mimeo, Duke University.*

Peters, M. and S. Severinov (2006): "Internet Auctions With Many Traders," *Journal of Economic Theory*, 130, 220–245.

Said, M. (2011): "Sequential Auctions With Randomly Arriving Buyers," *Games and Economic Behavior*, 73, 236–243.

——— (2012): "Auctions With Dynamic Populations: Efficiency and Revenue Maximization," *Journal of Economic Theory*, 147, 2419–2438.

Salop, S. and J. Stiglitz (1977): "Bargains And Ripoffs: A Model Of Monopolistically Competitive Price Dispersion," *The Review of Economic Studies*, 493–510.

Schneider, H. (forthcoming): "The Bidder's Curse: Comment," *The American Economic Review.*

Sorensen, A. T. (2000): "Equilibrium Price Dispersion In Retail Markets For Prescription Drugs," *Journal of Political Economy*, 108, 833–850.

Stahl, D. O. (1989): "Oligopolistic Pricing With Sequential Consumer Search," *The American Economic Review*, 700–712.

Wang, R. (1993): "Auctions Versus Posted-price Selling," *The American Economic Review*, 838–851.

Zeithammer, R. (2006): "Forward-looking Bidding In Online Auctions," *Journal of Marketing Research*, 43, 462–476.

A Proofs

Derivation of Bellman Equations. Each of the continuous-time Bellman equations (Eq. 2, 15, and 16) in the model can be derived from a discrete-time formulation as follows. First, consider the expected profit of a seller in the auction market, Π_a. Let Δ be the length of a period of time, which we assume to be sufficiently short such that $\eta\Delta < 1$; this can then be interpreted as the probability of the auction closing during that period of time. The discrete time Bellman equation is thus:

$$\Pi_a = -\ell\Delta + \frac{1}{1+\rho\Delta}\left(\eta\Delta\left(1-e^{-\lambda}\right)(\theta-\gamma c) + \left(1-\eta\Delta\left(1-e^{-\lambda}\right)\right)\Pi_a\right). \tag{29}$$

The term $\ell\Delta$ is the listing fee incurred during the period of time. The term in parentheses computes the expected outcome in the next period of time: either the auction closes with at least one bidder, earning $\theta - \gamma c$, or it does not close or attracts no bidders, so the seller enters the next period with the same expected payoffs as the current period. These future payoffs are discounted by the factor $1/(1+\rho\Delta)$.

By moving $\Pi_a/(1+\rho\Delta)$ to the left-hand side, then dividing by Δ, this becomes:

$$\frac{\rho}{1+\rho\Delta}\Pi_a = -\ell + \frac{\eta\left(1-e^{-\lambda}\right)}{1+\rho\Delta}\left(\theta-\gamma c-\Pi_a\right),$$

and taking the limit as $\Delta \to 0$, we obtain Eq. 15.

The expected profit for posted-price sellers is derived similarly. Again, we assume that a period is short enough that $\zeta\Delta < 1$.

$$\Pi_p = -\ell\Delta + \frac{1}{1+\rho\Delta}\left(\zeta\Delta(z-\gamma c) + (1-\zeta\Delta)\,\Pi_p\right). \tag{30}$$

Like auction sellers, posted-price sellers incur the listing fee $\ell\Delta$. With probability $\zeta\Delta$, they encounter a buyer in the next period and earn $z-\gamma c$; otherwise they continue waiting. This rearranges as:

$$\frac{\rho}{1+\rho\Delta}\Pi_p = -\ell + \frac{\zeta}{1+\rho\Delta}\left(z-\gamma c-\Pi_p\right),$$

and taking the limit as $\Delta \to 0$, we obtain Eq. 16.

The derivation for the buyer's expected utility is similar, only with more sources of uncertainty if an auction occurs. Let the period length Δ be sufficiently short that $\tau\alpha\Delta < 1$. This can then be interpreted as the probability that an auction occurs and the buyer participates during the unit of time. A buyer's expected utility in state s can be expressed

as follows:

$$V(s) = \frac{1}{1+\rho\Delta}\Bigg(\left(1 - \tau\alpha\Delta\sum_{n=0}^{\infty}\frac{e^{-\lambda}\lambda^n}{n!}(1-F(s))^n\right)V(s-\Delta)$$
$$+\tau\alpha\Delta\Bigg(\sum_{n=0}^{\infty}\frac{e^{-\lambda}\lambda^n}{n!}(1-F(s))^n x e^{-\rho s} \tag{31}$$
$$-e^{-\lambda}b(T) - \sum_{n=1}^{\infty}\frac{e^{-\lambda}\lambda^n}{n!}\int_s^T b(t)n(1-F(t))^{n-1}F'(t)dt\Bigg)\Bigg).$$

On the left-hand side, all utility is discounted by factor $1/(1+\rho\Delta)$, meaning that the buyer does not receive any utility during the current period. By the next period, one of two outcomes could occur: either the buyer wins an auction and exits (second and third lines of Eq. 31), or he continues his search (first line, due to losing or not participating).

Specifically, the second line computes the probability of the individual participating in an auction ($\tau\alpha\Delta$) and winning (the first two terms of the summation), times the utility enjoyed from winning ($xe^{-\rho s}$). The third line compute the expected second-highest bid times the probability of winning and thus paying it. The first line considers when the buyer does not win or does not participate (the probability in parentheses), in which case the buyer will continue waiting for future auction opportunities. Yet, he will do so with less time remaining before his deadline, reflected in his state changing to $s-\Delta$.

To transform this to a continuous-time Bellman equation, we first multiply both sides by $(1+\rho\Delta)/\Delta$, then subtract $V(s)/\Delta$ from both sides, obtaining:

$$\rho V(s) = \frac{V(s-\Delta)-V(s)}{\Delta} + \tau\alpha\Bigg(\sum_{n=0}^{\infty}\frac{e^{-\lambda}\lambda^n}{n!}(1-F(s))^n\left(xe^{-\rho s} - V(s-\Delta)\right)$$
$$-e^{-\lambda}b(T) - \sum_{n=1}^{\infty}\frac{e^{-\lambda}\lambda^n}{n!}\int_s^T b(t)n(1-F(t))^{n-1}F'(t)dt\Bigg).$$

Then, by letting $\Delta \to 0$, we obtain the continuous-time Bellman Eq. 2.

□

Proof of Proposition 1. First, we note that the infinite sums in equations 2 and 4 can be readily simplified. In the case of the latter, it becomes:

$$F''(s) = \alpha\tau F'(s)e^{-\lambda F(s)}. \tag{32}$$

This differential equation has the following unique solution, with two constants of integration

k and m:

$$F(s) = \frac{1}{\lambda} \ln \left(\frac{\alpha\tau - e^{\lambda k(s+m)}}{\lambda k} \right). \tag{33}$$

The constants are determined by our two boundary conditions. Applying Eq. 6, we obtain $m = \frac{1}{\lambda k} \ln \left(\alpha\tau - \lambda k e^{\lambda} \right) - T$. By substituting this into Eq. 33, one obtains:

$$F(s) = \frac{1}{\lambda} \ln \left(\frac{\alpha\tau - e^{\lambda k(s-T)} \left(\alpha\tau - \lambda k e^{\lambda} \right)}{\lambda k} \right). \tag{34}$$

The other boundary condition, Eq. 5, requires that k satisfy:

$$\alpha\tau \left(1 - e^{-\lambda T k} \right) - \lambda k \left(1 - e^{\lambda - \lambda T k} \right) = 0. \tag{35}$$

From Eq. 7, we know that $H = \delta / F'(T)$, and using the solution for F in Eq. 34, this yields $H = \delta\lambda / \left(\lambda k - \alpha e^{-\lambda} \tau \right)$. We then substitute this into the fourth equilibrium requirement, $\lambda = \tau H$, and solve for k to obtain:

$$k = \frac{\tau}{\lambda} \left(\delta + \alpha e^{-\lambda} \right). \tag{36}$$

When we substitute this for k in Eq. 34, we obtain the equilibrium solution for F^* depicted in Eq. 9. Also, when Eq. 36 is used to replace k in the boundary condition in Eq. 35, we obtain the formula ϕ (Eq. 8) which implicitly solves for λ^*.

We now show that a solution always exists to $\phi(\lambda^*) = 0$ and is unique. Note that as $\lambda \to +\infty$, $\phi(\lambda) \to +\infty$. Also, $\phi(0) = -\delta \left(1 - e^{-\tau(\alpha+\delta)T} \right) < 0$. Since ϕ is a continuous function, there exists a $\lambda^* \in (0, +\infty)$ such that $\phi(\lambda^*) = 0$.

We next turn to uniqueness. The derivative of ϕ w.r.t. λ is always positive:

$$\phi'(\lambda) = \alpha e^{-\lambda} + \delta(e^{\lambda} + \alpha\tau T) e^{-\tau(\alpha e^{-\lambda} + \delta)T} > 0.$$

Thus, as an increasing function, $\phi(\lambda)$, crosses zero only one time, at λ^*.

We finally turn to the solution for the bidding function. Again, we start by simplifying the infinite sums in Eq. 2. The first sum is similar to that in Eq. 4. For the second, we first change the order of operation, to evaluate the sum inside the integral. This is permissible by the monotone convergence theorem, because $F(s)$ is monotone and $\sum \frac{e^{-\lambda}\lambda^n}{n!} b(t) n (1 - F(t))^{n-1}$ converges uniformly on $t \in [0, T]$. After evaluating both sums, we obtain:

$$\rho V(s) = -V'(s) + \alpha\tau \left(e^{-\lambda F(s)} \left(x e^{-\rho s} - V(s) \right) - e^{-\lambda} b(T) - \int_s^T \lambda e^{-\lambda F(t)} b(t) F'(t) dt \right).$$

Next, by taking the derivative of $b(s) = xe^{-\rho s} - V(s)$ (Eq. 1), we obtain $b'(s) = -\rho x e^{-\rho s} - V'(s)$. We use these two equations to substitute for $V(s)$ and $V'(s)$, obtaining:

$$(\rho + \alpha\tau e^{-\lambda F(s)})b(s) + b'(s) = \alpha\tau \left(e^{-\lambda} b(T) + \int_s^T \lambda e^{-\lambda F(t)} b(t) F'(t) dt \right). \tag{37}$$

This equation holds only if its derivative with respect to s also holds, which is:

$$(\rho + \alpha\tau e^{-\lambda F(s)})b'(s) + b''(s) = 0. \tag{38}$$

After substituting for $F(s)$ solved above, this differential equation has the following unique solution, with two constants of integration a_1 and a_2:

$$b(s) = a_1 \cdot \left(\frac{\delta e^{\lambda^* - \tau T\left(\delta + \alpha e^{-\lambda^*}\right)}}{\rho} + \frac{\alpha e^{-\tau s\left(\delta + \alpha e^{-\lambda^*}\right)}}{\rho + \tau\left(\delta + \alpha e^{-\lambda^*}\right)} \right) e^{-s\rho} + a_2. \tag{39}$$

This solves the differential equation, but to satisfy Eq. 37, a particular constant of integration must be used. We substitute for $b(s)$ in Eq. 37 using Eq. 39, and solve for a_2. This can be done at any $s \in [0, T]$ with equivalent results, but is least complicated at $s = T$ since the integral disappears: $(\rho + \alpha\tau e^{-\lambda F(T)})b(T) + b'(T) = \alpha\tau e^{-\lambda} b(T)$. After substituting $b(T)$, $b'(T)$, and $F(T)$, solving for a_2 yields:

$$a_2 = a_1 \frac{\alpha\tau\left(\delta + \alpha e^{-\lambda^*}\right)}{\rho\left(\rho + \delta\tau + \alpha\tau e^{-\lambda^*}\right)} e^{-\rho T - \tau T\left(\delta + \alpha e^{-\lambda^*}\right)}. \tag{40}$$

The other constant of integration is determined by boundary condition Eq. 3. If we translate this in terms of $b(s)$ as we did for the interior of the Bellman Equation, we get $b(0) = z$. We then substitute for $b(0)$ using Eq. 39 evaluated at 0, and substitute for a_2 using Eq. 40, then solve for a_1:

$$a_1 = \frac{\rho z\left(\rho + \delta\tau + \alpha\tau e^{-\lambda^*}\right) e^{\tau T\left(\delta + \alpha e^{-\lambda^*}\right)}}{\tau\left(\delta + \alpha e^{-\lambda^*}\right)\left(\delta e^{\lambda^*} + \alpha e^{-\rho T}\right) + \rho\left(\delta e^{\lambda^*} + \alpha e^{\tau T\left(\delta + \alpha e^{-\lambda^*}\right)}\right)}.$$

If the solutions for a_1 and a_2 are both substituted into Eq. 39, one obtains Eq. 11 with minor simplification. □

Proof of Proposition 2. The first derivative of $b^*(s)$ is:

$$b'(s) = -\frac{\rho z\left(\rho + \delta\tau + \alpha\tau e^{-\lambda^*}\right)\left(\delta e^{\lambda^*} + \alpha e^{\tau(T-s)\left(\delta + \alpha e^{-\lambda^*}\right)}\right)}{\tau\left(\delta + \alpha e^{-\lambda^*}\right)\left(\delta e^{\lambda^*} + \alpha e^{-\rho T}\right) + \rho\left(\delta e^{\lambda^*} + \alpha e^{\tau T\left(\delta + \alpha e^{-\lambda^*}\right)}\right)}.$$

Each of the parenthetical terms is strictly positive, thus the negative in front ensures that the derivative is negative.

The second derivative is:

$$b''(s) = \frac{\rho z \left(\rho + \delta\tau + \alpha\tau e^{-\lambda^*}\right) \left(\delta\rho e^{\lambda^*} + \alpha \left(\rho + \delta\tau + \alpha\tau e^{-\lambda^*}\right) e^{\tau(T-s)\left(\delta+\alpha e^{-\lambda^*}\right)}\right)}{\tau \left(\delta + \alpha e^{-\lambda^*}\right) \left(\delta e^{\lambda^*} + \alpha e^{-\rho T}\right) + \rho \left(\delta e^{\lambda^*} + \alpha e^{\tau T\left(\delta+\alpha e^{-\lambda^*}\right)}\right)} e^{-s\rho}.$$

Again, each parenthetical term is positive. Hence $b''(s) > 0$. □

Proof of Proposition 3. By Proposition 1, Eqs. 9 through 11 and $\phi(\lambda^*) = 0$ must be satisfied in order to be a buyer equilibrium, as required in the first condition.

The solutions to A^* and σ^* are simply restatements of Eq. 19 and 17, respectively.

The profits stated in Eqs. 22 and 23 are required by the third and second equilibrium conditions, respectively. From Eq. 16, profit solves as: $\Pi_p = \frac{\eta(z-\gamma c)-\ell}{\rho+\zeta}$, so for this to equal $(1-\gamma)c$, we require $\zeta^* = \frac{\ell+\rho c(1-\gamma)}{z-c}$ (Eq. 26). With this, Eq. 18 readily yields P^* as listed in Eq. 25.

The only remaining element regards expected auction profit. Eq. 15 solves as: $\Pi_a = \frac{\eta\left(1-e^{-\lambda}\right)(\theta-c\gamma)-\ell}{\eta\left(1-e^{-\lambda}\right)+\rho}$. By setting this equal to $(1-\gamma)c$ and solving for θ, we obtain Eq. 21.

To evaluate the integrals in Eq. 14, we first note that by interchanging the sum and integral and evaluating the sum, expected revenue simplifies to:

$$\theta = \frac{\lambda}{1-e^{-\lambda}} \left(e^{-\lambda} b(T) + \lambda \int_0^T b(s)F(s)F'(s)ds\right). \tag{41}$$

After substituting for $b(s)$ and $F(s)$ from the buyer equilibrium, this evaluates to:

$$\begin{aligned}\theta = \frac{z}{1-e^{-\lambda}} \quad \cdot \quad \Bigg(\quad & 1 + \frac{1}{(\rho+\kappa\tau)\left(\rho\delta + \tau(\kappa-\alpha)\left(\delta + \alpha e^{-\lambda-\rho T}\right)\right)} \cdot \\ & \Bigg((\alpha-\kappa)e^{-\lambda-\rho T}\left(\kappa\tau(\kappa\tau-\lambda\rho) - \lambda\rho^2\right) - \delta\rho(2\kappa\tau+\rho) \\ & +\kappa\rho\tau\left(\delta\Psi\left(1-\frac{\kappa}{\alpha}\right) + (\alpha-\kappa)e^{-\lambda-\rho T}\Psi\left(1-\frac{\kappa e^{\lambda}}{\alpha}\right)\right)\Bigg)\Bigg),\end{aligned}$$

where $\kappa \equiv \delta + \alpha e^{-\lambda}$ and $\Psi(q)$ is Gauss's hypergeometric function with parameters $a = 1$, $b = -1 - (\rho/\tau\kappa)$, $c = -\rho/\tau\kappa$, evaluated at q. Under these parameters, the hypergeometric function is equivalent to the integral:

$$\Psi(q) \equiv -\left(1 + \frac{\rho}{\tau\kappa}\right) \int_0^1 \frac{t^{-2-\frac{\rho}{\tau\kappa}}}{1-qt} dt.$$

While not analytically solvable for these parameters, Ψ is readily computed numerically. $\square$

Market Efficiency. As we consider the contribution to total welfare by this market, the degenerate equilibrium provides a good benchmark. For exposition purposes, first consider welfare taking the rate of auctions α as given, as in the buyer equilibrium (that is, ignore Eq. 21). If no auctions were available, all buyers would wait until the last minute and then purchase at the posted price. The net gain to society over a unit of time would be $\delta(x-c)$, as δ buyers are served, δ sellers incur completion costs γc, and δ other sellers incur initial production $(1-\gamma)c$ as they enter the market.

In addition, we must consider the listing fees these sellers incur in our welfare calculations. If we include the market host (eBay) in total welfare, then listing fees are merely a transfer between sellers and the host. Yet the listing fee reflects a sort of inefficiency in the market caused by the matching friction. For instance, with perfect coordination, a buyer and seller could instantly be matched on entry so that no stock P of posted-price listings needed to be maintained. This does not occur because of the random arrival and matching of buyers to sellers in P. Through their excessive entry, sellers increase the expected wait for a buyer, so that any eventual profit is consumed in listing fees. In the degenerate case, total listing fees would be ℓP.

Now consider the dispersed equilibrium. For sellers, total costs incurred over a unit of time are $\Gamma \equiv \ell(A+P)+\delta c$, which includes the listing cost for the stock of uncompleted auctions. One can substitute for the equilibrium stock of both sellers, but the intuition can be seen without doing so. Because posted-price listings require more time to sell than auctions do, any change in the market that decreases P by 1 will increase A by less than 1, though still serving δ buyers per unit of time. In other words, the market is more efficient as it serves the same population while reduces the volume of listings.

Buyers in the dispersed equilibrium will experience different utility depending on when they buy: obtaining the good with s units of time remaining provides $xe^{-\rho s}$ utils. Because of this, early acquisition is inefficient; the production occurs immediately but the consumption is discounted. Of course, in a given auction, the highest valuation bidder will always win, who will also be closest to consumption. However, changes in the auction environment will alter the distribution and number of bidders in each auction, and thus change how close the average winner is from his deadline.

The distribution of the state s of the highest bidder in a given auction (conditional on

having any bidders show up) is represented by density function $G'(s)$, as follows:

$$\begin{aligned} G'(s) &\equiv \frac{1}{1-e^{-\lambda}} \sum_{n=1}^{\infty} \frac{e^{-\lambda}\lambda^n}{n!} n(1-F(s))^n F'(s) \\ &= \frac{\lambda e^{-\lambda F(s)} F'(s)}{1-e^{-\lambda}}. \end{aligned}$$

Auctions occur at rate α per unit of time, and attract at least one buyer with probability $1-e^{-\lambda}$. In addition, $H \cdot F'(0)$ buyers will purchase at the posted price per unit of time, enjoying the full utility x. Summing across all buyers, the total utility U generated by the entire market over each unit of time is:

$$U \equiv x\left(\delta - \alpha\left(1-e^{-\lambda}\right)\right) + \alpha \int_0^T x e^{-\rho s} \lambda e^{-\lambda F(s)} F'(s) ds. \tag{42}$$

The integral in Eq. 42 evaluates with similar complexity to expected revenue, precluding analytic analysis. However, with significant algebraic manipulation, one can show that for the cumulative distribution $G(s)$, $\frac{\partial G(s)}{\partial \alpha} \leq 0$ for all $s \in [0, T]$, meaning that any increase in the number of auctions will decrease the fraction of winners who are close to their deadline. That is, the distribution with fewer auctions first-order stochastically dominates the distribution with more auctions. Since the utility enjoyed at a given s is unaffected by α, more auctions unambiguously decreases average consumer utility.

Intuitively, this happens because frequent auctions provide more opportunities to win, so bidders exit earlier in their search spell; in addition, it leads to fewer participants per auction. While both of these are beneficial to the consumer, they ensure that the average winner is further from his deadline. Indeed, by this logic, the degenerate equilibrium produces the most utility from the goods by ensuring that they are fully enjoyed at the time of purchase.

The total welfare added through this market is $W \equiv U - \Gamma$. The two terms move in opposite directions, with more auctions reducing costs but increasing early acquisition; the net effect can only be compared in numeric examples. However, across a wide variety of parameterizations, the cost effect consistently dominates the utility effect. That is, more auctions are strictly beneficial because they decrease listing costs by more than they reduce average utility. This remains true even when the listing fee is quite low or the discount rate is quite high. Indeed, in every computation we have performed where a dispersed equilibrium existed, total welfare was greater than if it were forced into a degenerate market (by outlawing auctions, for instance).

The preceding analysis has focused on how the auction rate affects welfare. In the market equilibrium, this rate is endogenously determined. To evaluate evaluating the welfare

consequences of changing the fundamental parameters, one can simply compute how these parameters affects the equilibrium auction rate. For instance, a larger listing fee (ℓ) will reduce the auction rate and total welfare, while deferring more costs the time of sale (γ) will increase the auction rate and total welfare. In short, total welfare seems to track perfectly with the endogenous auction rate. □

B Reserve Prices

We now relax the assumption that auction sellers always set their reserve price equal to $b(T)$, the lowest bid any buyer might make in equilibrium. There is clearly no incentive to reduce the reserve price below that point: doing so would not bring in any additional bidders, but would decrease revenue in those instances where only one bidder participates.

Now consider a seller who contemplates raising the reserve price to $R > b(T)$, taking the behavior of all others in the market as given. This will only affect the seller when a single bidder arrives or the second highest bid is less than R; with this higher reserve price, the seller closes the auction without sale in these situations and re-lists the good, a strategy which has a present discounted value of Π_a. Of course, the seller gives up the immediate revenue and completion cost, which is no more than $R - \gamma c$.

Since $\Pi_a = (1-\gamma)c$ in equilibrium, deviating to the reserve price R is unprofitable if $R - \gamma c \leq (1-\gamma)c$, or rearranged, $R \leq c$. In words, the optimal seller reserve price should equal the total cost of production. Thus, in our context, $b(T)$ is the optimal seller reserve price so long as $b^*(T) \geq c$.

If $b^*(T) < c$, then the seller would prefer to set a reserve price of c. One can still analyze this optimal reserve price in our model by endogenizing the buyer deadline, T. For instance, suppose that buyers who enter six months before their deadline are only willing to bid below the cost of production. By raising the reserve price, these bidders are effectively excluded from all auctions; it is as if they do not exist. They only begin to participate once they reach time S such that $b^*(S) = c$. In other words, it is as if all buyers enter the market with S units of time until their deadline. To express this in terms of our model, we make T endogenous, requiring $b^*(T^*) = c$ in equilibrium. All else will proceed as before.

Even with optimal reserve prices, the entry and exit of sellers will ensure that expected profits from entering the market are zero. Any gains from raising the reserve price are dissipated as more auctions are listed. To consider the absence of this competitive response, consider if one seller had monopoly control of both markets. The optimal choice would be to shut down the auction market, forcing all buyers to purchase at the highest price z. When there are numerous independent sellers, however, they cannot sustain this degenerate equilibrium (at least when Equation 28 does not hold). There is always an advantage to

offering an auction if all other sellers offer fixed price listings: your product sells faster, even if at a slightly lower price.

Figures

Figure 1: Bids Over Time

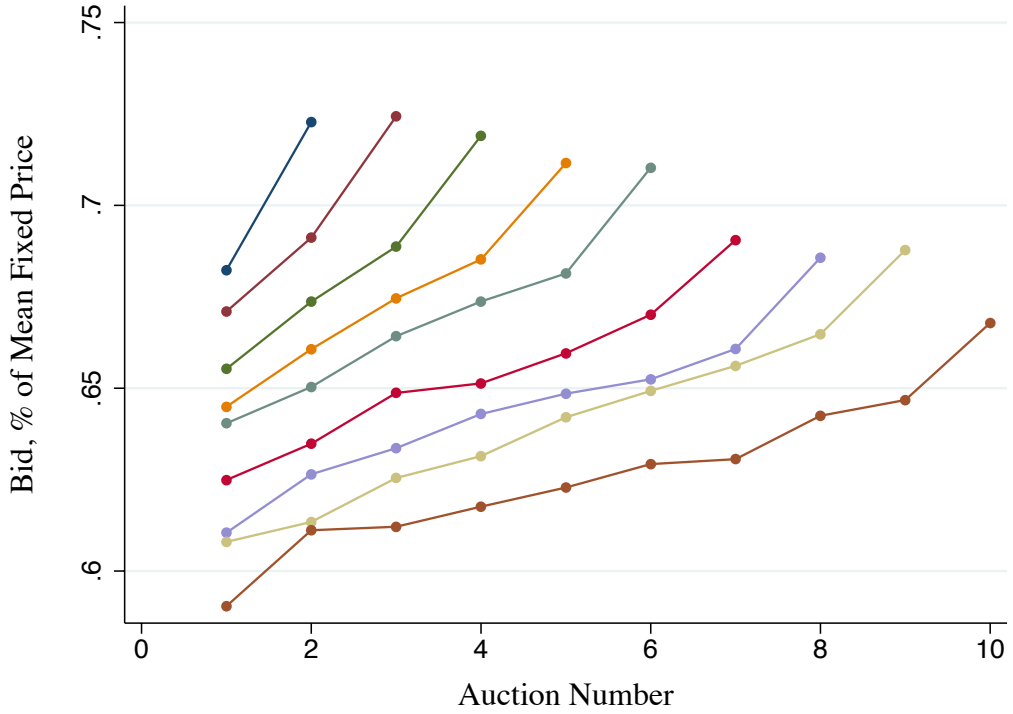

Figure 2: Bids Over Time, Regression Results

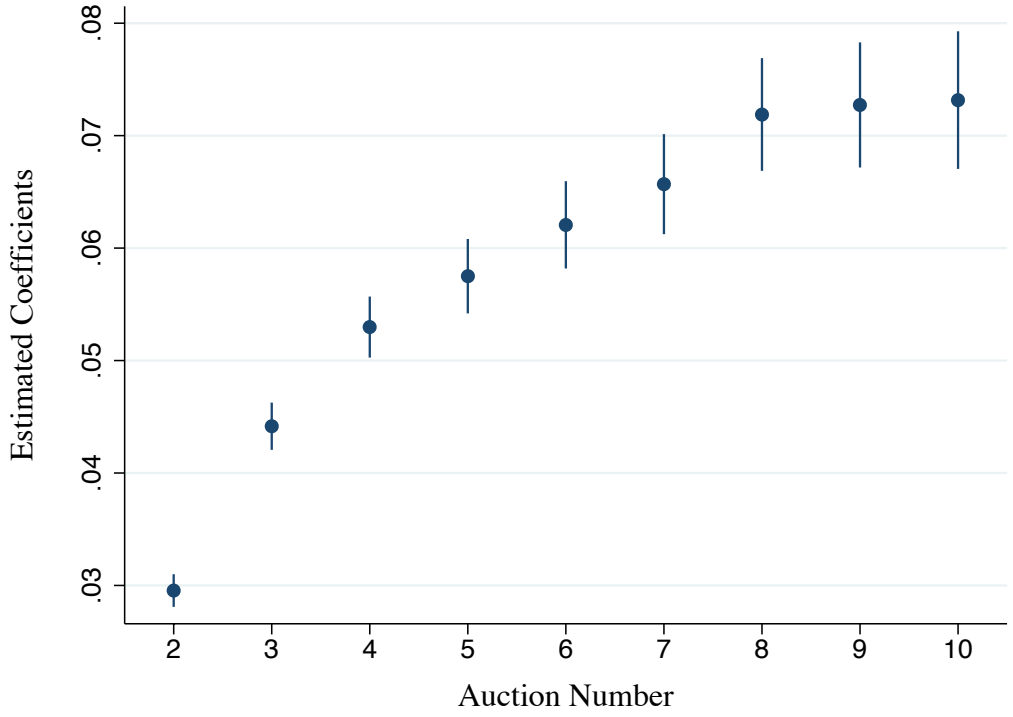

Figure 3: Distribution of Normalized Second-Highest Bids

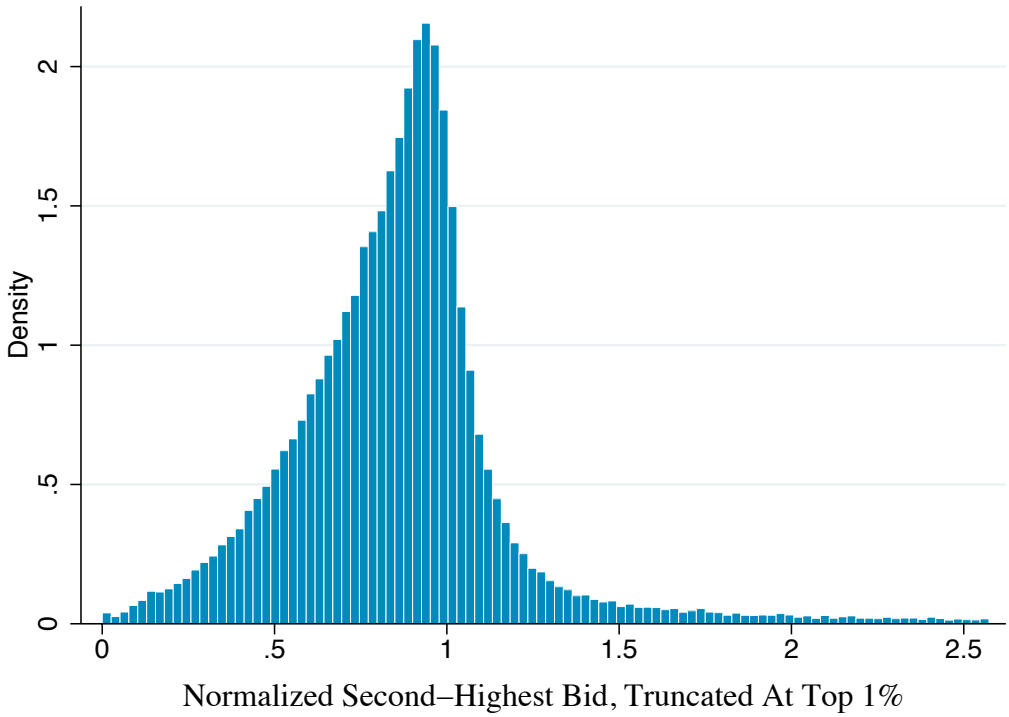

Figure 4: Distribution of Normalized Second-Highest Bids, Products with Mean Price over $100

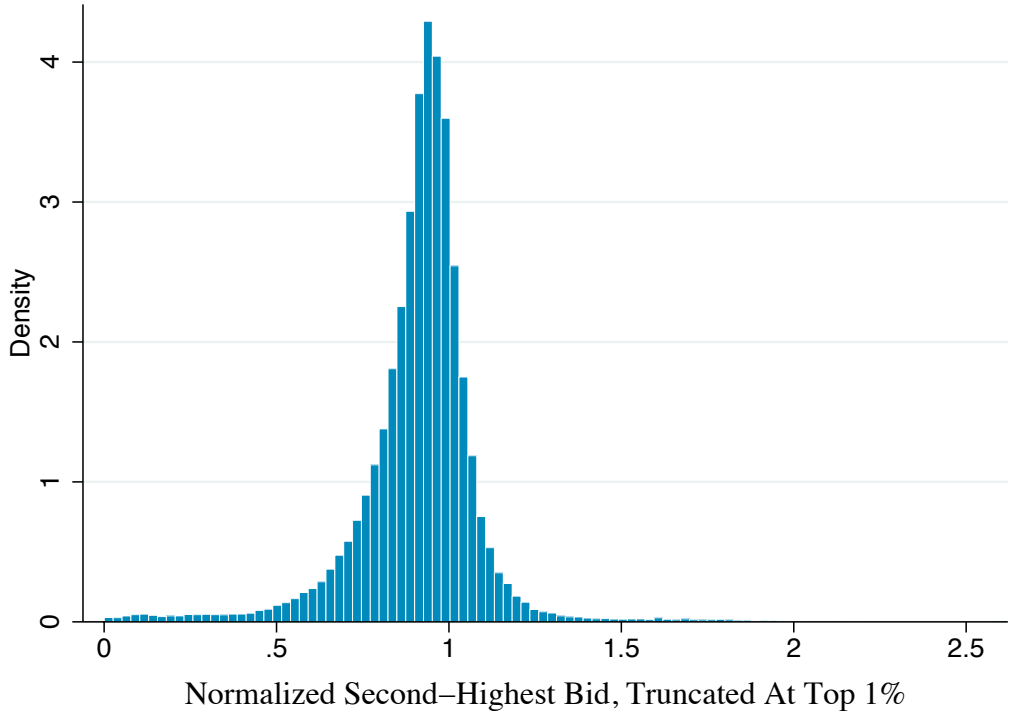

Figure 5: Distribution of Fraction of Sales By Auction Across Items

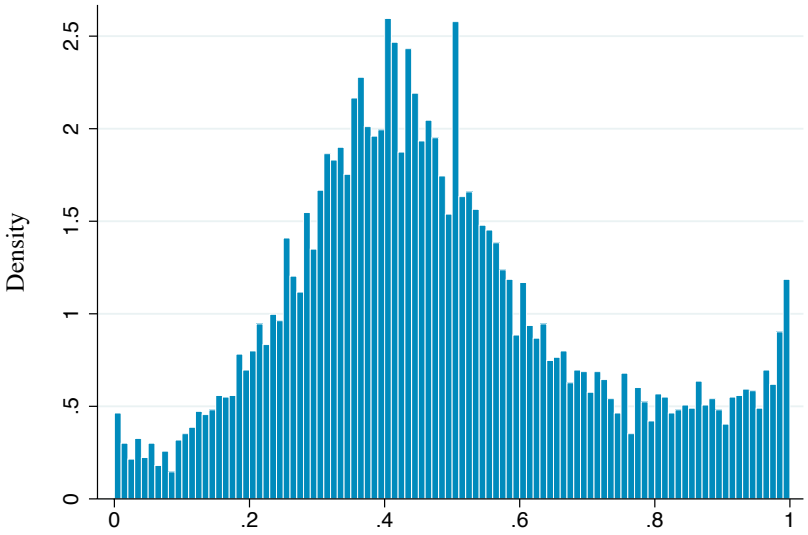

Figure 6: Cumulative Fraction Sold by Days Since Listing

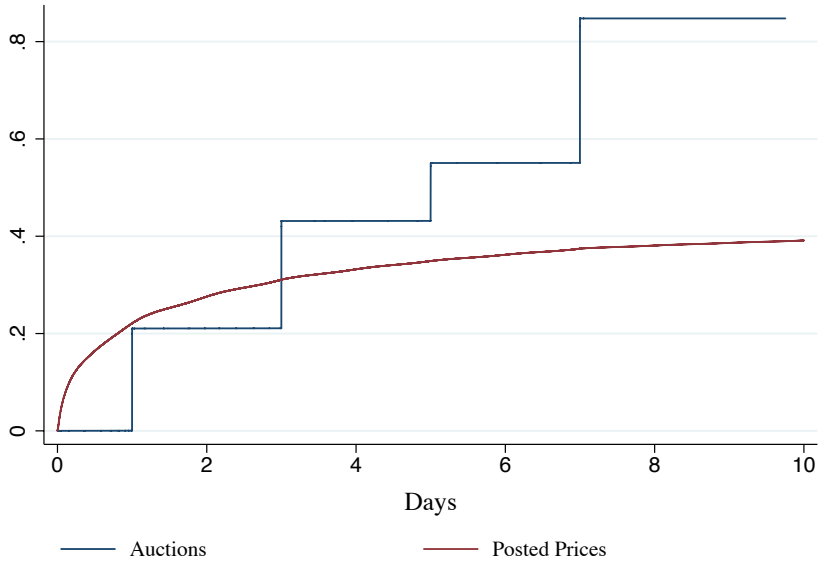

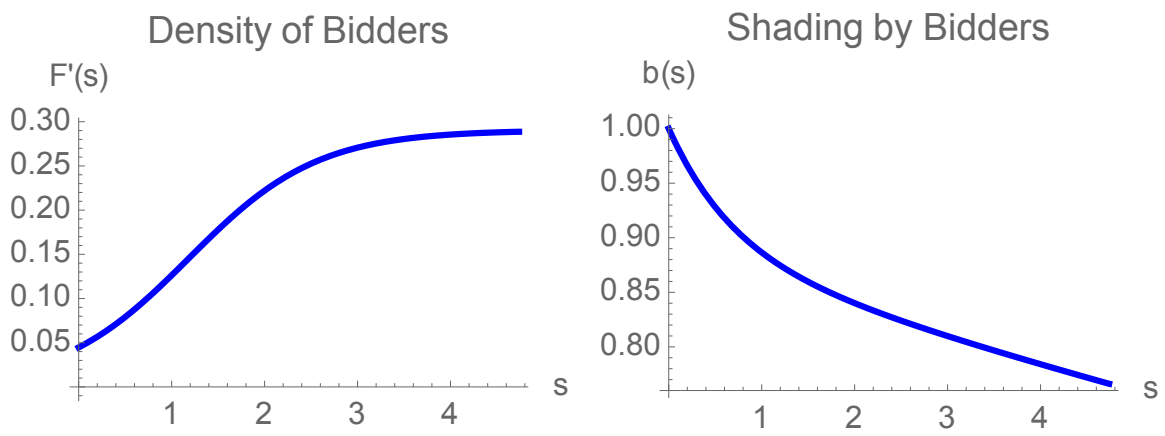

Figure 7: Bidding under Calibrated Parameters: the distribution of bidder states (left panel) and the % by which bidders shade relative to z (right panel)

Figure 8: Bids Over Time, Simulated Data

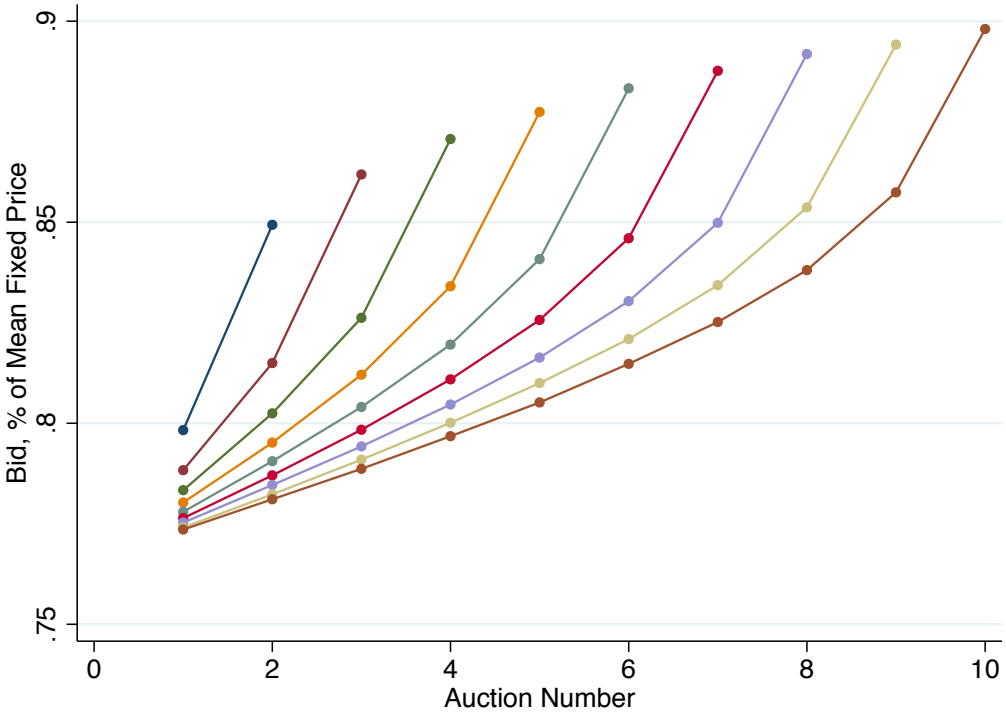

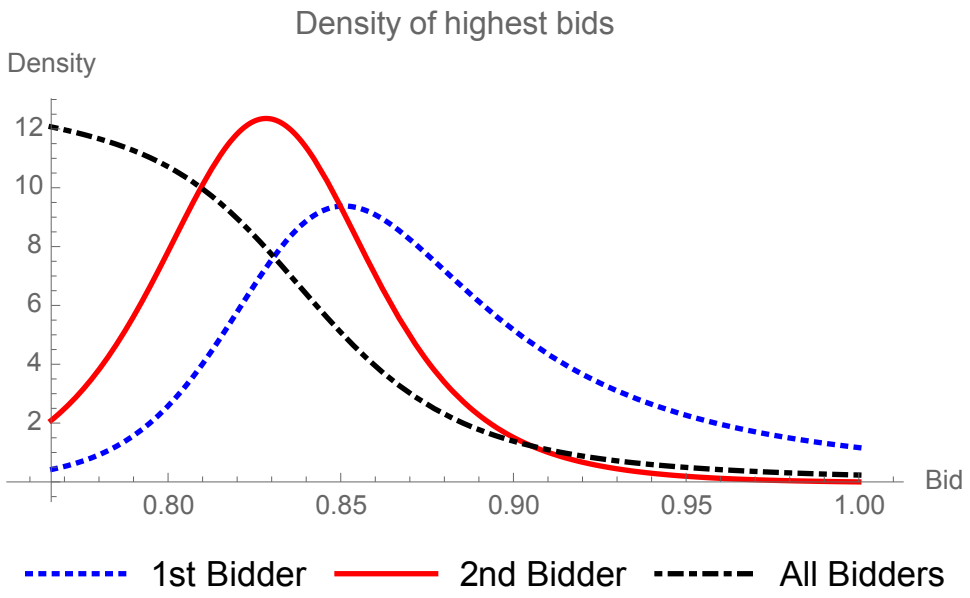

Figure 9: Price Density under Calibrated Parameters: the equilibrium density of the highest bid in an auction (dotted), the second highest bid (solid), and all bids (dot-dashed).

Figure 10: Bids Over Time, Products With Average Transaction Price of At Least $100

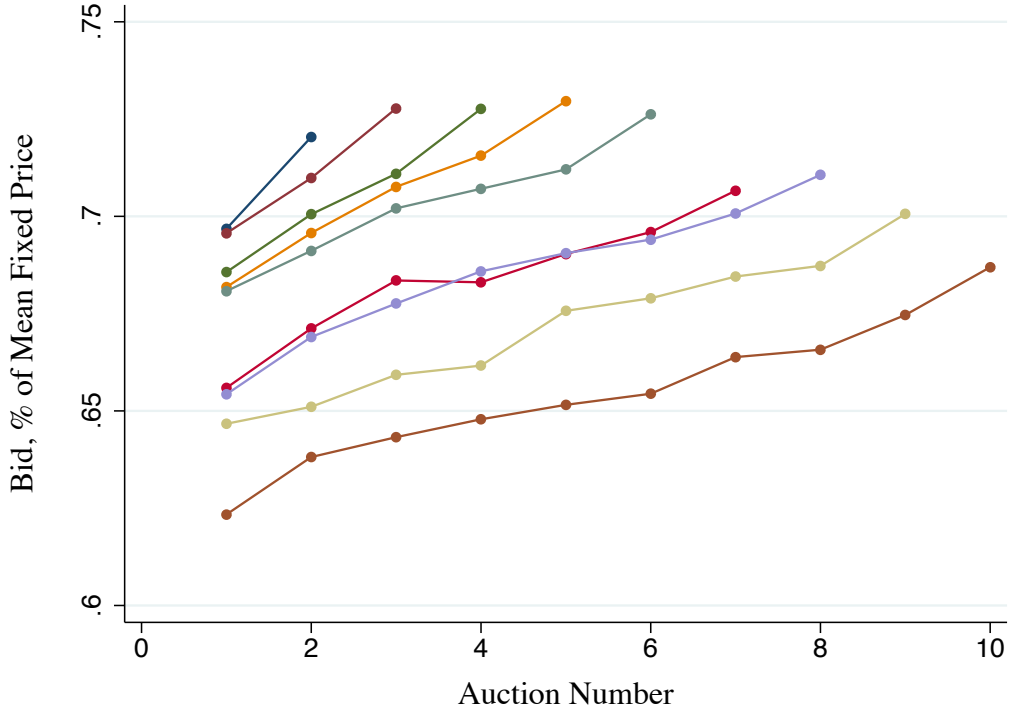

Figure 11: Bids over Time, Bidders Participating in at least 50 Auctions in the Past Year

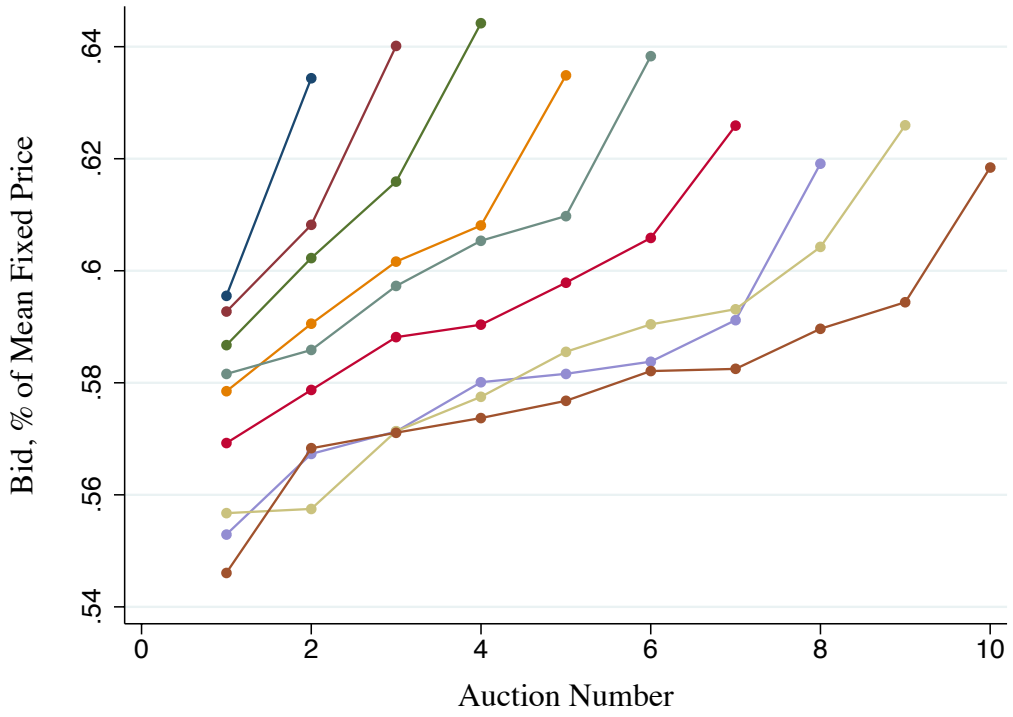

Figure 12: Bids over Time, Bidders Participating in at least 10 Auctions in the Past Year for Products in the Same Product Grouping

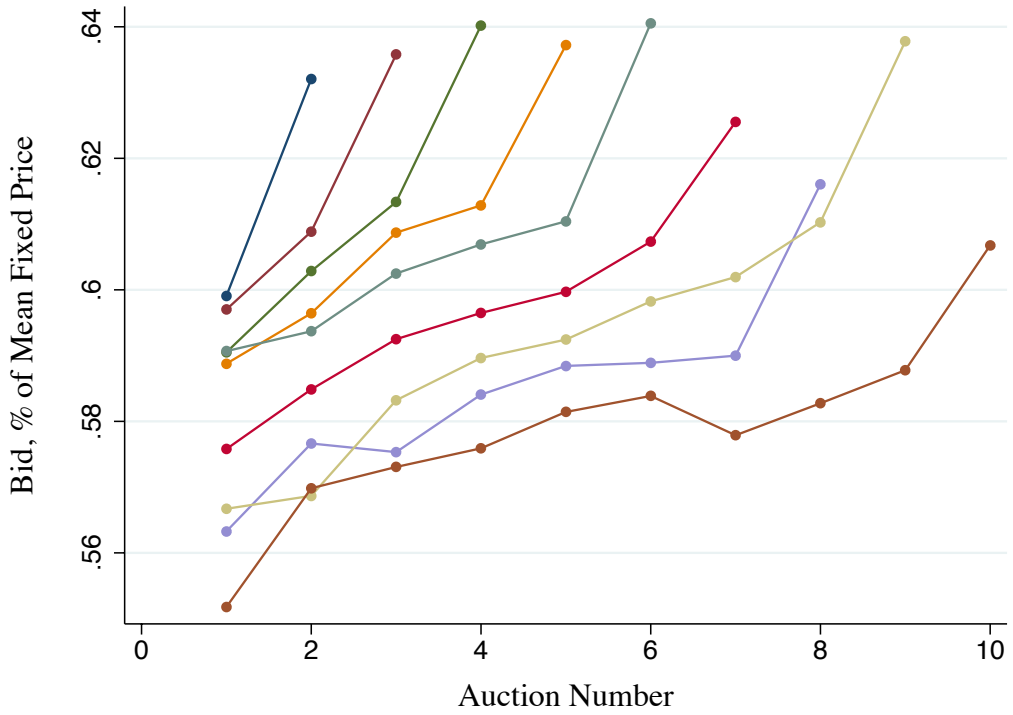

Tables

Table 1: Descriptive Statistics

	Auctions	Posted Price
Number of Transactions	584,755	540,194
Mean Revenue	82.57	100.06
Std. Dev. Revenue	154.93	172.76
Mean Bidders per Transaction	5.54	1
Std. Dev. Bidders per Transaction	3.75	—
Number of Product Categories	6,708	

Table 2: Calibration process and parameter values

	Observed Value in Data	Theoretical Equivalent	Calibrated Parameter
Bidders per completed auction	5.0	$\lambda/\left(1-e^{-\lambda}\right)$	$\lambda = 4.97$
Completed auctions per month	17.2	$\alpha\left(1-e^{-\lambda}\right)$	$\alpha = 17.3$
Auctions a bidder tries per month	1.2	$\tau\alpha$	$\tau = 0.071$
Auctions attempted per bidder	5.8	$\tau\alpha T$	$T = 4.74$
—	—	Eq. 8	$\delta = 20.33$
Average revenue per completed auction	0.838	θ	$\rho = 0.031$

Table 3: Comparative statics on key statistics: Buyer Equilibrium

		$\partial/\partial\alpha$	$\partial/\partial\tau$	$\partial/\partial\rho$	$\partial/\partial x$	$\partial/\partial T$
Participants per Auction	λ^*	–	+	0	0	+
Buyers using Posted Price	$F'(0)H^*$	–	–	0	0	–
Lowest Bid	$b^*(T)$	–	+	–	+	–

Note: Derivations are provided in the Online Appendix.

Table 4: Additional Calibration for Market Equilibrium

	Observed Value in Data	Theoretical Equivalent	Calibrated Parameter
Average listing fee paid	0.084	ℓ	$\ell = 0.454$
Average duration of an auction listing (months)	0.163	$1/\eta$	$\eta = 6.15$
Average % of posted-price listing sold in 30 days	47.8%	$\frac{\ell+\rho(1-\gamma)c}{z-c}$	$\gamma = 0.188$
—	—	Eq. 21	$c = 0.838$

Table 5: Comparative statics on key statistics: Market Equilibrium

		$\partial/\partial c$	$\partial/\partial\gamma$	$\partial/\partial\ell$	$\partial/\partial\tau$	$\partial/\partial T$	$\partial/\partial\rho$
Auction Rate	α^*	$-$	$+$	$-$	$+$	$-$	$-$
Participants per Auction	λ^*	$+$	$-$	$+$	$+$	$+$	$+$
% Buying via Posted Price	$F'(0)H^*$	$+$	$-$	$+$	$-$	$+$	$+$
Stock of Posted Price Sellers	P^*	$+$	$+$	$+$	$-$	$-$	$+$
Lowest Bid	$b^*(T)$	$+$	$-$	$+$	$-$	$-$	$-$
Expected Revenue	θ^*	$+$	$-$	$+$	$-$	$-$	$+$

Note: Reported signs are for numeric computations on the example.

Table 6: Bid Changes by Group

	N	Bid Increases	Bid Constant	Bid Decreases
All Bidders	1,224,062	48.3%	12.4%	39.4%
Products With Average Transaction Price $\geq$ \$100	431,619	49.9%	8.3%	41.9%
Bidders Entering $\geq$ 50 Auctions in Past Year	606,199	45.5%	15.5%	39.0%
Bidders Entering $\geq$ 10 Auctions in Same Product Grouping in Past Year	476,257	44.3%	16.6%	39.1%

Notes: Table shows the fraction of consecutive auctions in which bidders' highest bids increase, remain constant, or decrease from one auction to the next. N is the number of pairs of consecutive auctions entered by bidders for the same product. Different rows correspond to different samples of bidders.

Table 7: Bid change from previous to current auction regressed on amount previously outbid

Dep. var: Bid change, previous to current auction	(1)	(2)	(3)	(4)
Auction price minus bid, last auction	0.136 0.002	0.184 0.004	0.205 0.005	0.212 0.006
Auction price minus bid, 2nd to last auction	-	−0.094 0.029	−0.075 0.004	−0.060 0.005
Auction price minus bid, 3rd to last auction	-	-	−0.057 0.004	−0.055 0.005
Auction price minus bid, 4th to last auction	-	-	-	−0.039 0.005
R^2	0.047	0.055	0.063	0.065
N	518,291	209,842	113,244	70,931

Notes: Table shows regressions, for $n \in \{1, 2, 3, 4\}$, of bidders' bid increases from the n to $n + 1^{th}$ auctions in their bidding sequences on the amounts by which the auction revenue exceeded their bids in the previous n auctions. The dependent variable is winsorized at the 1st and 99th percentiles. Standard errors are clustered at the bidder level.