A Theory of Discounts and Deadlines in Retail Search*

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

December 23, 2016

Abstract

We present a new equilibrium search model where consumers face deadlines to obtain a retail product. Buyers initially search among discount opportunities, but are willing to pay more over time and eventually shift their search to full-price sellers. Even with homogeneous sellers and buyers, the model predicts equilibrium price dispersion and rationalizes the coexistence of both discount and non-discount sales mechanisms. We apply the model to online retail sales via auctions and posted prices; using data on one million listings for new-in-box items on eBay.com, we find robust evidence consistent with our model. As predicted, bidders increase their bids from one buying opportunity to the next, equilibrium price dispersion exists, and auctions and posted-price listings coexist. Fitting the model to the data, we find that the presence of the discount channel increases total welfare by 1.8% of the average retail price if intermediation is costly, but reduces welfare by 2.3% if intermediation is costless.

JEL Classifications: C73, D44, D83
Keywords: Mechanism choice, search, discount channels, sequential auctions, deadlines, endogenous valuations

*We thank Lanier Benkard, Liran Einav, Matt Gentzkow, Thomas Jeitschko, Jon Levin, Jason Lepore, Greg Lewis, Daniel Quint, Yaron Raviv, Steve Tadelis, Caio Waisman, and seminar and conference participants at Michigan, NC State, Penn State, Rice, Stanford, UC Davis, UCLA, UC Riverside, UT Austin, Wharton, eBay Research Labs, the Department of Justice, the 2013 Conference on Economic Design, the 2015 West Coast Search and Matching Workshop, the 2015 International Industrial Organization Conference, the 2015 Searle Center Conference on Internet Search and Innovation, the 2015 Western Economic Association International Conference, the 2015 NBER Summer Institute Joint IO/Digitization Meetings, and the 2015 World Congress of the Econometric Society for helpful comments.

†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

1 Introduction

Searching for the best price on a product requires patience, and patience can wear thin. A consumer may be willing to initially hunt for bargains on a given product, but there is no guarantee that one will be found. As the search drags on with repeated failures, the consumer can grow impatient and may eventually purchase the item outright at a standard posted price. This process of decaying patience is evident in many settings. For example, when a dress is needed for an upcoming formal event, the consumer can begin by combing through options at a discount store (e.g. Nordstrom Rack), but as the event draws nearer, she may need to purchase the dress for full price at a department store (e.g. Nordstrom). Similarly, the search for a new car may start with haggling at multiple dealerships, but the buyer becomes weary after failing several times to secure a discount, and may purchase the car at the sticker price.

We present a new non-stationary search model to investigate how consumers shift their strategy from discount to non-discount channels over the duration of their search. This theoretical framework allows us to investigate several important questions in retail search. First, how does the consumers' willingness to pay change over the duration of their search? Second, why do both discount and non-discount sales channels exist for identical goods? Finally, how is welfare affected by the coexistence of such discount and non-discount sales channels?

Specifically, our model depicts an online retail market for new-in-box items offered via auctions and posted prices. Auctions provide a discount sales mechanism, potentially offering a low price but with low chance of success; meanwhile, the same product is readily available at a higher posted price. This model showcases search dynamics that plausibly occur in many other retail settings; however, the auction marketplace offers a unique empirical advantage. In most retail markets, only successful transactions are observable; yet bids placed in an auction can indicate consumers' willingness to pay *over the course of the search*, even during *failed attempts* to acquire the good.[1] We focus on listings of new-in-box items, which eliminates unobserved heterogeneity in product characteristics and ensures a steady flow of listings.

We operationalize the notion of increasing time sensitivity through each buyer facing an individual-specific deadline by which he must obtain the good. At any given discount purchase opportunity, the buyer faces some probability of failing to acquire the good (e.g. not finding the item on discount clothing racks, failing to win an auction, or failing to reach a deal while haggling), and thus draws nearer his deadline.[2] With the non-discounted sales channel, on the other hand, the buyer obtains the good with probability one. Consistent with our application to auctions, we refer to a successful attempt to acquire the good through the discount sales channel as "winning."

The deadline faced by the buyer can be interpreted literally (*e.g.* the item is needed for

[1] This feature is also unique relative to clickstream data (De los Santos et al. 2012; Koulayev 2014). There, the researcher may observe what sites a consumer visits and what prices the consumer sees, but it is more difficult to detect changes in the consumer's willingness to pay during the search process.

[2] Throughout we will refer to a buyer as "he" and a seller as "she".

a particular event, like a vacation, birthday, birth of a newborn, or athletic/school/business event), but more generally, they represent a limit on how long a consumer is willing to spend procuring a good. For instance, searching for a discount could become more difficult over time if customers cannot sustain the same level of attention or become increasingly frustrated with repeatedly failing to acquire the item. 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 or bicycle tube) that hasn't yet failed but is increasingly likely to do so.

Deadlines introduce non-stationarity into the search problem and cause willingness to pay to increase with the duration of search. If the consumer reaches his deadline, he will buy at the posted price, but before his deadline, he might succeed at finding the item at a lower price through the discount sales channel. Thus, his willingness to pay early in his search is shaded down from the true valuation, with more time remaining providing more opportunity and hence greater shading. This time sensitivity also generates price dispersion. When a discount opportunity arrives, potential buyers differ in their willingness to pay because some buyers need the object today more than others (*i.e.* are closer to their deadlines). Rather than assuming that this heterogeneity is exogenously given, we derive this distribution of buyers' deadlines endogenously based on the rate at which buyers win and exit the market.

Our model also allows sellers to choose whether to list their product through a discount or non-discount option. The discount option earns less revenue but sells faster than the non-discount option. These competing effects are balanced in a mixed strategy equilibrium, allowing identical sellers to be indifferent between the two selling mechanisms.

We state the model in continuous time and analyze it in steady state, which keeps the model tractable and enables comparative static analysis. 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). The restriction to steady-state behavior implies that the distribution of buyers with respect to their deadline remains stable. Population dynamics are consistent 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 sales of standardized retail items.

We examine the model's predictions empirically using a new dataset of one million brand-new goods from 3,663 distinct products offered on eBay from October 2013 through September 2014. As predicted by the model, discount sales channels (auctions) and non-discount options (posted prices) coexist for these products, with auctions representing 47% of sales. We find significant price dispersion both within auction sales (an interquartile range equal to 32% of the average price) as well as between auctions and posted-price listings (auction prices are on average 15% lower than posted-price sales). We also find that, as the model predicts, past losers tend to bid more in subsequent auctions (2.0% more on average), with the rate of increase rising with the length of search. As further evidence of the time sensitivity we model, we find that repeat bidders are increasingly likely to bid on listings for which fast shipping options are available and are increasingly likely to win the auction relative to bidders who have not searched as long.

Having established these empirical facts, we use the data to estimate the model's parameters, each of which are conveniently pinned down by easily computable moments of the data. Using these parameters, we simulate bidders' willingness to pay. We find that the model performs strikingly well in fitting key data patterns that are not used as identifying moments. We also show that ignoring the effects of time-sensitivity on bidders' willingness to pay can strongly distort estimation of consumer surplus or demand from auction data.

A benefit of this tractable equilibrium search framework is that it allows evaluation of the welfare implications of the coexistence of options available to consumers during their search. Recent work on online retail markets by Einav et al. (2016) demonstrated that auction sales have declined relative to posted-price sales. We use our estimated model to evaluate whether auctions for standardized goods offer any welfare advantages when posted-price listings are already available. Winners in the discount market obtain the good before it is needed for consumption, causing some inefficiency; yet the discount market also shortens the time sellers spend in the market. The net welfare effect of these hinge on whether intermediation consumes real resources. If inventory holding fees merely cover the operating costs of market intermediaries, the discount mechanism reduces these operating costs and outweigh the consumption inefficiency, leading to higher welfare with both mechanisms in play. If holding fees are pure profit for intermediaries, most of the value of sellers' time on the market is transferred to the intermediaries rather than dissipated, and hence the consumption inefficiency typically dominates, yielding higher welfare if the discount channel is shut down.

We analyze discounts and deadlines in the stark case where all goods, buyers, and sellers are homogeneous, without the exogenous differences that would typically explain the variation across auction outcomes. This setting highlights the mechanisms at work and yields the cleanest predictions, which would still be at play even with exogenous differences among goods, buyers, and/or sellers.

While we see deadlines as a natural explanation for the key features of our data and of many other retail search settings, deadlines are not the only possible explanation for any one feature in isolation. Increasing willingness to pay over time, for example, can be also rationalized by a subtle story of buyers learning about the competition they face, which affects the option value of waiting for future buying opportunities and motivates buyers to revise their willingness to pay at each subsequent buying opportunity. We present empirical evidence showing that bidder learning is unlikely to fully explain the pattern of increasing bids over time. Moreover, while separate explanations might exist for some of the data patterns we document individually, buyer deadlines provide a single, unified explanation of all of the facts together.

1.1 Literature

Our paper contributes in several ways to the literature on sequential consumer search. First, we provide a framework for studying consumers with deadlines on their search process, and find strong evidence consistent with the predicted non-stationary search strategy. Canonical search models (*e.g.* Stigler, 1961; Rob, 1985; Diamond, 1987; Stahl, 1989) yield a constant reservation price for the duration of the search. Indeed, Kohn and Shavell (1974) show this always holds in static search: that is, when consumers sample from a fixed distribution, face the same search costs, and have at least one firm left to search. In our model, it is the last feature that varies over the search duration. Buyers always have a chance that the current discount opportunity will be their last, and this probability rises as they approach their deadline. Other features that can lead to non-stationary search include price matching guarantees (Janssen and Parakhonyak, 2013), costs incurred to recall past offers (Janssen and Parakhonyak, 2014), and the possibility that past quotes will not be honored (Akın and Platt, 2014).[3]

Our work also highlights deadlines as an interesting source of price dispersion. Typically, homogeneity of buyers and sellers leads to a single (monopoly) price being offered and thus eliminates the need for search, as shown in the seminal work of Diamond (1971). The equilibrium search literature has circumvented this Diamond Paradox by introducing exogenous differences among buyers' search costs (Salop and Stiglitz, 1977; Rob, 1985; Stahl, 1989; Backus et al., 2014; Schneider, 2014), valuations (Diamond, 1987), or distance to firms (Astorne-Figari and Yankelevich, 2014).[4] In contrast, our model delivers *pure price dispersion*, in the sense that sellers are identical and buyers are ex-ante identical in their valuation and time to search, yet their ex-post-differing deadlines create a continuum of dispersed prices. Indeed, deadlines generates dispersion across individual sales as well as dispersion in willingness to pay *within* individuals over time (as they did for unemployed workers in Akın and Platt, 2012).

Our work also rationalizes the coexistence of distinct sales mechanisms of a homogeneous good. Competing mechanisms have been well documented empirically in online vs. offline channels (Shriver and Bollinger, 2015), retail storefronts vs. factory outlets (Soysal and Krishnamurthi, 2015), mass discounters (i.e. "everyday low price" stores, such as Walmart) vs. typical supermarkets (Bell and Lattin, 1998; Ellickson et al., 2012), and auctions vs. negotiations (Bajari et al., 2008). Yet coexistence is difficult to sustain theoretically: in Wang (1993), Bulow and Klemperer (1996), Julien et al. (2001), and Einav et al. (2016), one mechanism is strictly preferred over the other except in "knife-edge" or limiting cases. Models in Etzion

[3] In the job search literature, Van Den Berg (1990) shows that reservation wages rise over the search duration if conditions are expected to deteriorate in the future (such as a worse distribution of wages, fewer offers, or lower unemployment benefits).

[4] Other work has augmented the exogenous differences of the Stahl (1989) model with endogenous choices of sellers to advertise (Robert and Stahl, 1993) or obfuscate (Ellison and Wolitzky, 2012). A more basic source of dispersion arises in auctions due to occasional lack of competition. Kultti (1999) and Julien et al. (2001) demonstrate this in the extreme case where a bidder wins an auction for free if he is alone, but if anyone else participates, they compete until the price equals their common valuation. In our model, this dispersion appears due to variation in the number and composition of competitors.

et al. (2006), Caldentey and Vulcano (2007), Hammond (2013), and Bauner (2015) rely on ex-ante buyer or seller heterogeneity to have both mechanisms operate simultaneously. In contrast, we show that both mechanisms may be used in equilibrium under generic parameters, even though buyers and sellers are homogeneous ex-ante.

In applying our search model to auctions vs. posted prices, our work also connects to the nascent literature on infinite sequential auctions (Zeithammer, 2006; Ingster, 2009; Said, 2011, 2012; Backus and Lewis, 2016; Bodoh-Creed et al., 2016; Hendricks et al., 2012; Hendricks and Sorensen, 2015), in which bidders shade their bids below their valuations, taking into account the continuation value of future participation. A distinguishing prediction of our model, and one for which we find strong empirical evidence, is that a bidder's bid will increase at subsequent attempts to acquire the item.

2 Buyer Behavior

We first model buyers' choices, treating seller behavior as exogenous until Section 3. Consider a market for a homogeneous good in a continuous-time environment. To fix ideas, we consider the search problem faced by consumers where the discount sales channel is an auction and the non-discount sales channel is a posted price. Throughout we comment on how alternative sales mechanisms would fit into this search framework.

Buyers enter the market at a constant rate of δ, seeking one unit of the good which is needed for consumption in T units of time and will provide utility x (measured in dollars). For expositional simplicity, we assume here that the good is only consumed at the intended deadline, not at the time of purchase. Section A of the Supplemental Appendix considers a generalization in which a portion of the utility is obtained at the time of purchase, with minimal impact on equilibrium behavior. We assume that buyers face a rate of time preferences ρ; thus, if the good is purchased with s units of time remaining until the buyer's deadline, his realized utility at the time of purchase is $xe^{-\rho s}$ minus the purchase price. It is important to note that this does *not* imply that in equilibrium the buyer will prefer getting the good later rather than earlier; equilibrium prices paid by buyers who obtain the good earlier will be lower than those who obtain the good later, leading to higher utility at the time of purchase for earlier purchasers.

Discount purchase opportunities occur at a Poisson rate α, from the individual buyer's perspective. When such an opportunity arises, the buyer participates with exogenous probability τ, reflecting the possibility that a buyer can be distracted from participation by other commitments. The outcome is resolved immediately, with the winner making payment while all losers continue their search.[5]

[5]Note that the model does not include an explicit cost of participating in the discount market; rather, there is an implicit time cost, since attempts in the discount market will fail with some probability whereas those in the posted-price market will not. The model can easily incorporate explicit participation costs, but with increased notational complexity and little gain in the model's insights.

Alternatively, at any time, a buyer can obtain the good directly from the non-discount option at a posted price of z.[6] We assume throughout that $x \geq z$, so that buyers weakly benefit from purchasing via the posted-price option. It is important to note that our model does not require that all buyers share the same final consumption utility x; this can vary across buyers so long as it is lies weakly above z.

Every buyer shares the same 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 total stock of buyers in the market, a finite mass denoted H. Thus, each auction fits into the symmetric independent private valuations format, although both $F(s)$ and H are endogenously determined.

The strategic questions for buyers are what reservation price to be willing to pay in the discount sales channel and when to purchase from the non-discount posted-price listings. This dynamic problem can be expressed recursively. Let $V(s)$ denote the net present expected utility of a buyer with s units of time remaining until his deadline; such a buyer is willing to pay up to his reservation price $b(s)$, which is the present value of the item minus the opportunity cost of skipping all future discount opportunities, yielding:

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

We assume that $b(s)$ is decreasing in s (i.e. willingness to pay increases as the deadline approaches), and later verify that this holds in equilibrium. A buyer's expected utility in state s can then be expressed in the following Hamilton-Jacobi-Bellman (HJB) equation:

$$\rho V(s) = -V'(s) + \tau\alpha\Big(\Pr(win|s)\left[xe^{-\rho s} - V(s)\right] - E[payment|s]\Big). \tag{2}$$

In this continuous-time formulation, the left-hand side of (2) represents the flow of expected utility that a buyer with s units of time remaining receives each instant. The right hand side depicts potential changes in (net) utility times the rate at which those changes occur.[7] The derivative term 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, so his utility changes by $-V'(s)$. When a discount opportunity arises and the individual participates in it — which occurs at a rate of $\tau\alpha$ per unit of time — the buyer wins with probability $\Pr(win|s)$. When this occurs, the term $xe^{-\rho s} - V(s)$ depicts the change in utility due to winning. $E[payment|s]$ represents the expected payment of the buyer (the expected price times the probability of winning and

[6] Section 3.2 discusses why it is an equilibrium for all posted-price sellers to choose the same z.

[7] Section C of the Supplemental Appendix indicates how this HJB formulation can be derived from a discrete-time Bellman equation. If buyers received any disutility from the search process itself, it would appear as a negative constant on the right-hand side. This would have negligible effect on the solution other than raising willingness to pay, since making a purchase would terminate this flow of disutility.

thus paying it).

The exact form of $\Pr(win|s)$ and $E[payment|s]$ depends on the specifics of the discount sales channel; in Section D of the Supplemental Appendix, we provide examples of search in the context of discount posted prices, haggling, and lotteries. Here, we focus on the case where the discount sales channel consists of second-price sealed-bid auctions: each participant submits a bid and immediately learns the auction outcome, with the highest bidder winning and paying the second highest bid. As in a static second-price auction, buyers find it weakly dominant to bid their reservation price $b(s)$.[8] The auctioneer is assumed to open the bidding at the lowest buyer reservation price, $b(T)$, which is only relevant in the unlikely event that just one bidder participates in the auction.[9]

We assume that the number of bidders in any given auction is drawn from a Poisson distribution with mean λ, so the probability that n bidders participate is $e^{-\lambda}\lambda^n/n!$. We require that $\lambda = \tau H$ in equilibrium, to ensure that participation in a given auction reflects the aggregate participation in the market. Thus, the probability that a buyer in state s wins a given auction is:

$$\Pr(win|s) = \sum_{n=0}^{\infty} \frac{e^{-\lambda}\lambda^n}{n!}(1 - F(s))^n. \tag{3}$$

To win the auction, the buyer must have the highest bid, which means all n other participants must have more than s units of time remaining.[10] This occurs with probability $(1 - F(s))^n$. The expected payment of the buyer is the average second-highest bid times the probability of winning and thus paying it, and is stated as:

$$E[payment|s] = e^{-\lambda}b(T) + \sum_{n=1}^{\infty} \frac{e^{-\lambda}\lambda^n}{n!} \int_s^T b(t)n(1 - F(t))^{n-1}F'(t)dt. \tag{4}$$

[8]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 $q \in (p, b(s))$. One abstraction in our model is 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 leads to further shading of bids (Backus and Lewis, 2016; Said, 2012). Leakage cannot occur in our model because encountering the same opponent in a future auction is a probability zero event (as in Zeithammer, 2006). This approximates a large market, where the probability of repeat interactions are too low to justify tracking hundreds of opponents. In our data, if a bidder encounters an opponent and later participates in another auction for the same item, he only has an 8.4% chance of encountering that opponent again.

[9]In Section B of the Supplemental Appendix we show that, with a reasonably short search length T, then $b(T)$ is in fact the optimal starting price due to competition among potential sellers.

[10]While the Poisson distribution literally governs the *total* number of participants per auction, it 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.

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 $n(1-F(t))^{n-1}F'(t)$.

Buyers also have the option to purchase from the non-discount posted-price option at any time, receiving utility $xe^{-\rho s} - z$. However, a buyer in state s can obtain a present expected utility of $(x-z)e^{-\rho s}$ by waiting until $s = 0$ to buy, 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 the consumer surplus from making the purchase:

$$V(0) = x - z. \tag{5}$$

2.1 Steady-State Conditions

In most models of demand, including 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, in contrast, the buyers' reservation prices $b(s)$ are endogenous, affected by the value of further search, $V(s)$, 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.

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 capture the long-run average behavior in a market.

To begin the steady-state analysis, consider the relative density of buyers 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:

$$F''(s) = F'(s)\tau\alpha \Pr(win|s) \tag{6}$$

where $\tau\alpha \Pr(win|s)$ is the rate at which buyers participate in and win through the discount sales channel (thus exiting the market). Recall that all buyers enter the market in state $s = T$; at all other states, buyers only exit the market (by winning, or running out of time at $s = 0$). Thus the buyer density $F'(s)$ must decrease as s falls, which holds because the right-hand side of (6) is always positive.

Equation (6) 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:

$$\lim_{s \to 0} F(s) = F(0) = 0 \tag{7}$$
$$\lim_{s \to T} F(s) = F(T) = 1. \tag{8}$$

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 also 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 stock of buyers in the market, $HF(s)$ depicts the mass 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:

$$\delta = H \cdot F'(T). \tag{9}$$

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

2.2 Buyer Equilibrium

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 HJB equations (1) through (5), taking F^* and λ^* as given.

2. The distribution F^* satisfies the Steady-State equations (6) through (8).

3. The average mass of buyers in the market H^* satisfies Steady-State equation (9).

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.

We now characterize the unique equilibrium of this game. 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. If $\phi(\lambda)$ is defined as:

$$\phi(\lambda) \equiv \delta - \alpha\left(1 - e^{-\lambda}\right) - \delta e^{\lambda - \tau T\left(\delta + \alpha e^{-\lambda}\right)}, \tag{10}$$

then the boundary condition is equivalent to $\phi(\lambda^*) = 0$. Note that δ is the mass of buyers who enter the market over a unit of time, while $\alpha\left(1 - e^{-\lambda}\right)$ is the mass of buyers who win an auction and thus exit over a unit of time. The last term in $\phi(\lambda)$ turns out to be $H \cdot F'(0)$ (derived below), which is the mass of buyers who exit because of hitting their deadline over a unit of time. Thus, $\phi(\lambda^*) = 0$ ensures that buyers newly entering the market exactly replace those who exit. The rest of the equilibrium solution is expressed in terms of λ^*.

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

$$F^*(s) = 1 - \frac{1}{\lambda^*}\ln\left(\frac{\delta + \alpha e^{-\lambda^*}}{\delta e^{\tau\left(\delta + \alpha e^{-\lambda^*}\right)(s-T)} + \alpha e^{-\lambda^*}}\right), \tag{11}$$

and its associated density function is:

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

If the probability of winning an auction were constant throughout a buyer's search, then buyers would be exponentially distributed. Indeed, we see an exponential density function in the numerator of $F'(s)$, but it decays as if buyers exit more often than auctions even occur (since $\tau\left(\delta + \alpha e^{-\lambda^*}\right) > \tau\alpha$ when $\phi(\lambda^*) = 0$). This exponential decay is slowed down by the fact that buyers only exit if they win an auction; this adjustment is reflected in the denominator of $F'(s)$.

Indeed, F' is always increasing in s but typically changes 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 decreases.

The average number of buyers in the market is simply:

$$H^* = \frac{\lambda^*}{\tau}. \tag{12}$$

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

$$b^*(s) = ze^{-\rho s} \cdot \frac{\tau\left(\delta + \alpha e^{-\lambda^*}\right)\left(\delta e^{\lambda^*} + \alpha e^{\rho(s-T)}\right) + \rho\left(\delta e^{\lambda^*} + \alpha e^{\tau(T-s)\left(\delta+\alpha e^{-\lambda^*}\right)}\right)}{\tau\left(\delta + \alpha e^{-\lambda^*}\right)\left(\delta e^{\lambda^*} + \alpha e^{-\rho T}\right) + \rho\left(\delta e^{\lambda^*} + \alpha e^{\tau T\left(\delta+\alpha e^{-\lambda^*}\right)}\right)}. \tag{13}$$

Alternatively, we can express the bidding function more succinctly, and with easier interpretation, as follows:

$$b^*(s) = z\left(1 - \frac{\rho\int_0^s g(t)dt}{g(T) + \rho\int_0^T g(t)dt}\right), \quad \text{where} \quad g(t) \equiv \frac{\tau\left(\delta + \alpha e^{-\lambda^*}\right)e^{-\tau\left(\delta+\alpha e^{-\lambda^*}\right)t}}{e^{-\lambda F(t)}}e^{-t\rho}.$$

To interpret the function $g(t)$, note that $e^{-\lambda F(t)}$ in the denominator is the probability of winning for a buyer who participates in the auction in state t. Thus, $1/\left(e^{-\lambda F(t)}\right)$ is the average number of auctions in which a buyer in state t would need to participate before winning. The numerator of $g(t)$ is the density function describing the likelihood of the next auction occurring in exactly t units of time. Finally, a win in state t is discounted by $e^{-t\rho}$ because the item is not needed until time $s = 0$.

Thus, the integral $\int_0^T g(t)dt$ is the average (discounted) number of auction attempts required to win before the deadline. The term $g(T)$ in the denominator of $b(s)$ accounts for the possibility that the buyer does not win any auction and is forced to buy at the posted price. The integral $\int_0^s g(t)dt$ in the numerator of $b(s)$ is the portion of those auction attempts that are still possible. The ratio of these integrals indicates the fraction of opportunities remaining. Buyers are effectively shading their bid below the retail price of z in accordance with the likelihood of winning between s and the deadline. As that window closes, they expect to have fewer opportunities and they draw closer to bidding the retail price.

The next result shows that this proposed solution is both necessary and sufficient to satisfy the equilibrium requirements. In the proof (given 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 (11) through (13) 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 $b''(s) > 0$.*

Time preference plays a key role in creating dispersion among the bidder valuations, even separately from the deadline T. For instance, if buyers become extremely patient ($\rho \to 0$), the bidding function approaches $b(s) = z$ regardless of time until deadline.[11] Impatience causes buyers to prefer postponing payment until closer to the time of consumption, and thereby creates some variation in willingness to pay. This makes it possible to win an auction at a discount relative to z, which further reduces the willingness to pay. But if impatience is eliminated, the variation disappears; everyone is willing to bid full price, so auctions do not offer a discount at all. Even so, the Proposition 2 prediction of bids increasing over time is not solely due to time preference; it also reflects reduced bid shading as buyers have fewer future opportunities to win an auction. We decompose these effects in Section 5.2.

The average time between auctions ($1/\alpha$) is of similar importance. 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, so a bidder's reservation price would simply equal $b(s) = xe^{-\rho s}$, his present value of the good.[12] For larger values of α, a bidder would optimally reduce his bid well below this, as he is likely to have many opportunities to win a deal before his deadline.

Comparative statics indicate that these insights consistently hold for all parameterizations. As buyers become less patient or more auctions are offered, their bidding profile over search duration becomes steeper. This also occurs when given more time to search (increasing T); this result is less obvious because there are more auctions but more participants per auction, but the former always dominates. Although our equilibrium has no closed form solution, these comparative statics can be obtained by implicit differentiation. Section E of the Supplemental Appendix reports all comparative statics and provides their proofs and additional discussion.

3 Seller Behavior

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 the discount or non-discount mechanism. We consider a continuum of sellers producing an identical good.[13] 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 discount opportunities are available. 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

[11] The fractional term of (13) approaches one. Note that the equilibrium λ^* from (10) is unaffected, as is the distribution of bidders in (11).

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

[13] 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 holding fees scale up proportionately. Also, when sellers employ mixed strategies, they can be interpreted literally as each seller randomizing which mechanism to use, or as dividing sellers into two groups playing distinct pure strategies in the proportion dictated by equilibrium.

when seller first enters the market (the *initial production cost*).[14] For either selling format, sellers also pay a holding fee of ℓ each unit of time from when the seller enters the market to when the good is sold. This flow cost of holding inventory includes any expenses required during an active listing, such as warehouse rental space to store inventory, fees paid to any intermediating platform, and costs of monitoring the listing, answering buyer questions, or paying employees to maintain a showroom. We assume that there are no barriers to entry for sellers. Upon entry, each seller must decide whether to join the discount or the posted-price market.

3.1 Discount Sellers

The advantage of the discount sector to a seller is that the sale occurs more quickly. Let θ denote the average revenue at the time of sale, and let η represent the Poisson rate of closing, meaning that the average time delay between listing and closing is $1/\eta$. Section D of the Supplemental Appendix indicates what these might look like in physical search, bargaining, or lottery environments. In the auction mechanism we focus on here, the listing length is treated as exogenous.[15] At its conclusion, the auction's expected revenue (conditional on at least one bidder participating) is denoted θ and computed as follows:

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

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 $n \geq 2$ 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, $1-e^{-\lambda}$. Note that here we assume the reserve (start) price is $b(T)$; in Section B of the Supplemental Appendix, we allow the seller's choice of reserve price to be endogenous and demonstrate that the optimal reserve price is indeed $b(T)$.

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 HJB equation:

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

[14] 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.

[15] Under the default setting at eBay, which is also the most common setting used in our data, an auction concludes one week after creating the listing, which provides time for bidders to examine the listing.

The right-hand side of (15) indicates that the seller incurs the holding 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 a discount seller, the expected profits from entry are net of the initial production cost: $\Pi_a - (1-\gamma)c$.

3.2 Non-discount Sellers

Sellers who charge the full retail price will earn more than discount sellers, since $b(s) < z$; the disadvantage of this format is that sellers may wait a considerable time before being chosen by a buyer.[16] 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 present-value expected profit of posted-price sellers already in the market, denoted Π_p, is computed in the following HJB equation:

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

Like auction sellers, posted-price sellers incur the holding fee of ℓ each unit of time that they await a buyer; for simplicity, we assume the same holding fees, but little changes if they differ across the two selling formats. 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 $\Pi_p - (1-\gamma)c$.

3.3 Steady-State Conditions

As with the population of buyers, the stock and flow of sellers are also assumed to remain stable over time. In 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.

In addition, the mass 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

[16] Although the model assumes that all posted-price sellers charge the same (exogenously determined) price z, this can also be supported as an equilibrium outcome even when each seller can endogenously choose their posted price. Specifically, if buyers anticipate that all sellers charge the same price z, they will expend no effort in searching among available sellers, but will choose one at random. Thus, a seller who deviates by posting a lower price will not sell any faster but would sacrifice some profit. Moreover, a seller who deviates by posting a higher price will always be rejected, since the buyer anticipates that another seller can immediately be found who charges price z. Although this is is an equilibrium, other equilibria are certainly possible, and exploring these would be an interesting avenue for future research.

a unit of time. This must equal the mass 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, her auction will close at rate η; but with A sellers in the market at any instant, there will be ηA auctions that close over 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, $HF'(0)$ posted price purchases occur over a unit of time (sold to buyers who reach their deadline). From the individual seller's perspective, she can sell 1 unit every $1/\zeta$ units of time; collectively, these sellers expect to sell ζP units in one unit of time. In equilibrium, the expected sales must equal the expected purchases:

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

3.4 Market Equilibrium

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 discount sector, $\sigma^* \in [0,1]$, such that:

1. Expected revenue θ^* is computed from Equation (14) 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 discount 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 2, given the endogenously determined auction arrival rate. The fourth imposes the steady-state conditions. The second and third requirements impose zero expected profits for both types of sellers, which is implied by the large, unrestricted pool of potential entrants. If either market offered positive profits, additional sellers would be attracted to that market and profits would fall: more posted-price sellers P would reduce the rate of selling ζ, and more auction sellers A would increase the auction arrival rate and decrease expected revenue θ. Together, these two requirements also ensure that sellers are indifferent about which market they enter, allowing them to randomize according to the 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 profit 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 if a sufficiently small stock of sellers serves these desperate buyers.

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

While the market equilibrium conditions simplify considerably, they do not admit an analytic solution and we must numerically solve for α^* and λ^*. All other equilibrium objects can be expressed in terms of these. The solution for these endogenous variables comes from equation (10) and the third market equilibrium requirement, which can be written:

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

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^*, (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 (10); here, we note that this equation involves both λ and α. Equilibrium is attained when both (10) 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:

$$\Pi_a^* = (1-\gamma)c \tag{22}$$
$$\Pi_p^* = (1-\gamma)c \tag{23}$$

$$
\begin{aligned}
A^* &= \frac{\alpha^*}{\eta} & (24)\\
P^* &= \frac{(z-c)\left(\delta - \alpha^*\left(1-e^{-\lambda^*}\right)\right)}{\ell + \rho(1-\gamma)c} & (25)\\
\zeta^* &= \frac{\ell + \rho(1-\gamma)c}{z-c} & (26)\\
\sigma^* &= \frac{\alpha^*\left(1-e^{-\lambda^*}\right)}{\delta}. & (27)
\end{aligned}
$$

It is 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 (11) through (14), and equations (21) through (27).*

The solution described in Proposition 3 can be called a *dispersed equilibrium*, to use the language of equilibrium search theory, as we observe the homogeneous good being sold at a variety of prices. By contrast, in a *degenerate equilibrium*, the good is always sold at the same price. This only occurs 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 (11) through (13) and equations (23) through (27) with $\alpha^* = 0$ and $\lambda^* = \tau\delta T$, exists if and only if*

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

Moreover, if this condition fails, a dispersed market equilibrium will exist. Thus, an equilibrium always exists.

The left-hand side of (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 (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. On the right-hand side, one can see that a high production cost or holding fee 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 exists 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. In that sense, they are always searching, giving sellers motivation to offer auctions, so long as (28) does not hold.

Proposition 4 proves that an equilibrium always exists; we further conjecture that the equilibrium is always unique. This claim would require that at most one dispersed equilibrium can occur, and that a dispersed equilibrium cannot occur when (28) holds — both of which are true if θ is a decreasing function of α (*i.e.*, more auctions always lead to lower revenue). The complicated expression for θ in the dispersed equilibrium precludes an analytic proof, but we have consistently observed this relationship between α and θ in numerous calculations across a wide variety of parameters.

Comparative statics for the market equilibrium also must be numerically evaluated; Section E of the Supplemental Appendix summarizes and discusses the typical effects. These provide insight on what drives more sellers to one market or the other. More sellers join the discount market (α increases) when buyers pay more attention, have less time, or are more patient; these are not surprising since they lead to higher bids (with a flatter bid profile). Higher seller costs (for holding or production) lead to negative profits, all else equal; the auction market compensates for this by offering fewer auctions and thereby forcing greater competition in the remaining auctions. Meanwhile, more buyers hit their deadline and must purchase in the posted-price market, restoring zero profits there. The same occurs if more production costs are paid upfront.

Regarding the causes of buyers' increasing willingness to pay, the market equilibrium comparative statics closely follow those of the buyer equilibrium. It is noteworthy, however, that production and holding costs of the seller can affect *buyer* search. With fewer discount opportunities, buyers see less option value from continued search and thus have a flatter profile of bids. These subtle responses to ℓ, c, and γ illustrate a potential hazard in market design if buyer valuations are not fundamental but rather the endogenous results of deeper factors. A seemingly neutral change applied to both markets (such as holding fees or timing of production costs) not only alters which market sellers use, but also warps the distribution of buyer valuations.

4 Empirical Evidence

4.1 Data and Descriptive Statistics

We examine the predictions of our model using online retail sales, where the discount sales channel is an auction and the non-discount channel is a posted price. Our data consist of auctions and posted-price sales on eBay.com for the year from October 1st 2013 to September 30th 2014.[17] As our model describes the sale of homogeneous goods, we restrict attention to brand new items which have been matched by the seller to an item in one of several commercially available catalogs; we will refer to these catalogued products as simply a *product*, and a individual attempt to sell the product will be referred to as a *listing*. These products 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 with outlier prices (defined as bids in the top or bottom 1% of bids for auctioned items of that product, and similarly for posted-price sales); products with under 25 auction or posted-price sales; and products that went more than 30 days without an auction. The products in our final sample are thus popular items, principally electronics, media, or health/beauty products.

Table 1 presents descriptive statistics for the listings that ended in a sale. In all, there are over one million sales of 3,663 distinct products, split roughly evenly between auctions and posted prices. Means and standard deviations in panel A are computed by taking the mean and standard deviation over all transactions of a given product and then taking the mean of these product-level means and standard deviations. The average number of bidders per sold auction is 5.3. The average selling price (calculated inclusive of shipping fees) is higher with posted prices than auctions ($107 versus $97). To adjust for differences across products, we follow Einav et al. (2016) and rescale all bids, dividing by the mean price of posted-price sales of that product. This rescaling is also consistent with our model above, in which bids scale multiplicatively with the posted price. The normalized revenue per auction sale is, on average, 85% of the posted price, reflecting the fact the auctions serve as a discount sales channel in this market. Panel B demonstrates that both auctions and posted-price sales contain a large number of transactions per product, with numerous distinct buyers and sellers involved in transactions of each product.

For the remainder of this section, we document several facts in the data that provide

[17] 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., 2015).

evidence in favor of the deadline model.

4.2 Bids Over Time

The principal advantage of using auctions to study the discount mechanism is that we observe each attempt to acquire the good, rather than only observing successful attempts as in physical search. Moreover, in each attempt, the optimal bid truthfully reveals the bidder's willingness to pay at that point in time. In our data, we follow each bidder across multiple auctions of the same product. While not all bidders make repeat attempts after failing to win, such bidders do 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 47% of all bids.

We find that bidders tend to increase their willingness to pay from one auction to the next. To compute this, we keep each bidder's highest bid in each auction in which he participates, yielding a sample of 4,077,410 bids. We order these bids in a chronological sequence for each bidder and product pair, ending when the bidder either wins an auction or does not participate in any more auctions in our sample. This yields 2,728,258 unique bidding-sequence and product pairs. We then compute the average of the normalized 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.[18] Averaging across all sequence lengths and auction numbers, the bid increases by 2.0 percentage points in each successive auction. These patterns of increasing willingness to pay are consistent with Proposition 2 and suggest a role for time sensitivity in search. These patterns stand in contrast to what would be expected if buyers held a stationary willingness to pay. An additional pattern observable in Figure 1 consistent with the model is the surprising feature that the bid sequence lines never cross, as discussed further in Section 5.

4.3 Price Dispersion

Our data reveal (and our model predicts) three forms of price dispersion over homogenous products. The first form is across mechanisms, in that auctions average 15% lower sales

[18] Figure A1 in the Supplemental Appendix contains similar results to those in Figure 1 in a regression framework, averaging over all auction sequence lengths. The 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. To see this, note that Figure 1 shows that even before the final bid in a sequence, bids tend to increase, and Figure A2 in the Supplemental Appendix presents the same figure excluding winning bids. Nor is it due to selection in the product mix across the auction number variable, as the sequences are constructed at the bidder-by-product level, so conditional on sequence length the product mix is constant across auction number.

prices than posted-price listings (see Table 1). The second form is that, even among repeated instances of discount sales of a given product, there is substantial dispersion in transaction prices. The interquartile range of the normalized second-highest bid across auctions is 32 percentage points. Some of this dispersion is due to low price items, which shows large variance in their normalized closing prices. Restricting attention to products with a mean posted price of over \$100, there remains a good deal of price dispersion, with an interquartile range of 20 percentage points. This dispersion remains even after controlling for seller and product fixed-effects: the residuals from the regression of the normalized second-highest bid on seller and product fixed effects have a standard deviation and interquartile range of 9 and 5 percentage points, respectively.[19] The third form of price dispersion is that a given individual participating in the discount sales channel systematically offers higher prices over time, as discussed in the preceding subsection. All three of these types of price dispersion are in line with the model.

4.4 Coexistence of Auctions and Posted Price Sales

Our data indicates that a mixture of both mechanisms are simultaneously used for frequently listed homogenous goods, rather than relying almost exclusively on auctions or on posted prices. This finding necessarily holds in our preferred sample (which only includes products with at least 25 transactions of each type), and therefore to document evidence of the mix of sales mechanisms, we turn to a larger sample nesting our preferred sample. This broader sample includes all products sold at least 50 times in our sample period without regard to listing method, yielding 12,368 products. For each product, we compute the fraction of listings sold via auction. The histogram in Figure 2 indicates the distribution of this fraction across products. Under 1% of products 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 products have between 10% and 90% of sales by auction. Averaging across products in this broader sample, the average rate of auction use is 47%. This coexistence of the two sales channels for homogeneous goods is again consistent with the model, and contrasts with many other models of competing mechanisms which predict only one mechanism being adopted in equilibrium.

Transactions are typically completed faster through the discount (auction) mechanism relative to posted-price listings. 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 3 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 (21% vs 24%), and are over twice as likely to sell within 10 days (83% vs 41%). This is

[19]The corresponding numbers for all products, rather than only those with a mean posted price over \$100, are 17 and 13 percentage points.

consistent with logic of the model: a seller is willing to sell through auctions or posted prices, with the former mechanism selling considerably faster but at a lower price than the latter.

4.5 Winners, Losers, and Shipping Speed

Here we document several other patterns relating time and bidder behavior that are consistent with our model. First, the bidder in an auction with the longest observed time on the market (i.e the time since the bidder's first observed bid) is frequently the winner, occurring in 40% of auctions. In Section 5.2, we find that this is consistent with our model's predictions. In contrast, if elapsed time and likelihood of winning were orthogonal, the longest duration bidder would only have a 7.7% chance of winning, since each of the bidders are equally likely to win.[20] Moreover, elapsed time and likelihood of winning would be inversely correlated if valuations were constant over time, since high valuation buyers would win shortly after entering the market while low valuations bidders would require many repeated attempts to get lucky.

Second, buyers who lose in an auction and then purchase the same good from one of eBay's posted-price listings tend to do so shortly after their last observed auction attempt: 52% do so within one day of their last losing attempt. Furthermore, the cumulative probability of switching to a posted-price listing is concave in the time elapsed since the last losing attempt (see Figure A3 in the Supplemental Appendix). This is as the model predicts: the buyer's last observed auction attempt must be close to the buyer's deadline, or else the buyer would have attempted more auctions, and thus a posted-price purchase is most likely to occur close to the last auction attempt.

Third, after repeated losses, buyers are increasingly more likely to participate in auctions where expedited shipping is available, also consistent with the time sensitivity we model. The overall fraction of buyers bidding in auctions with fast shipping available is 0.45, and this fraction increases on average by 0.6 percentage points per auction attempt (see Figure A4 of the Supplemental Appendix). Each of these features is consistent with the nonstationary search model we propose.

4.6 Bidder Learning and Alternative Explanations

Deadlines provide a single explanation for multiple data patterns, one of which is the robust pattern observed in our data of bidders increasing their bids over time. Another possible explanation for that particular fact might involve bidder learning.[21] Consider the case where bidders are uncertain about the degree of competition they face, and form different esti-

[20] As explained in Section 5, the average number of participants per auction is 13, and $\frac{1}{13} \approx 7.7\%$.

[21] We note that learning does not necessarily imply increases in bids across auctions. In Jeitschko (1998), bidders can learn their opponents types from their bids in the first auction, but in equilibrium, they reach the same expected price in the second auction. The model in Iyer et al. (2014) generates bids that rise on average, but there learning occurs only for the auction winner, who needs to experience the good to refine his information about its value.

mates 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. Bidders who overestimate the amount of competition will bid more aggressively than those who underestimate. However, their initial aggressive bidding makes them likely to win 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. In practice, however, there are several reasons why bidder learning is unlikely to be the sole driver of increasing bids.

First, users can easily learn prices and bid histories for current and past listings by selecting the "Sold Listings" checkbox on the search results page; this is far more informative and quicker than auction participation. Second, experienced bidders should have more familiarity with the auction environment and with alternative means for gathering information, so learning by participation should not affect them. Yet the same pattern of increasing bids appears among experienced bidders, as shown in the left panel of Figure A5 in the Supplemental Appendix for bidders who participated in at least 50 auctions prior to the current auction and in the right panel for bidders who, over the past year, participated at least 10 auctions in the same product grouping as the reference auction.[22] Third, learning by participation is more costly with expensive products, due to the danger of bidding too high and winning when initially uninformed. We would expect buyers to be more cautious with such products and use alternative methods of learning. Yet Figure A6 in the Supplemental Appendix shows the same increasing bid pattern for products with a median price over \$100.[23]

We emphasize here that we do not attempt to entirely rule out the possibility of bidder learning or other alternative explanations for individual pieces of the patterns we observe. However, the appeal of our model of time-sensitive buyers is that it provides a single, unified explanation of all of these facts together. For example, one alternative explanation for the increase in bids at the end of bidding sequences in Figure 1 is that from one auction to the next bidders receive random shocks to their valuations and that the increase at the end of the sequence is caused by bidders winning and exiting after a positive shock. However, a story of random valuation shocks would fail to explain the pattern of increasing bids *prior* to the final auction in the sequence, nor would it speak to the coexistence of auctions and posted-price sales or the relationships discussed above between time and bidder behavior. While these alternative explanations could play a role, the bulk of the evidence also seems to indicate a

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

[23]One might be tempted to test the learning story by looking for positive correlation between the amount by which a bidder loses an auction and the amount he increases his bid in subsequent auctions. Yet such a positive correlation is also consistent with the deadlines model, because low bidders (early in their search) will typically lose by the largest margins. As discussed in the next section, these buyers are rarely observed until much later in the search process when their bid has increased substantially.

role for time-sensitivity.

5 Parameter Estimation and Counterfactuals

The preceding section documented that the data facts are qualitatively similar to the model predictions. Here, we illustrate the equilibrium behavior under estimated parameter values. In most dimensions, we find that the model closely matches the data facts, and the parameter values provide insight into the forces shaping this market behavior. The estimated parameters also allow us to compute the welfare impact of eliminating the discount (auction) sales channel in this market.

5.1 Estimation Procedures

Fitting our model to the data is relatively straightforward, as each parameter either corresponds directly to a sample moment or to a known transformation of sample moments. The moments we match and the corresponding model parameter estimates for the buyer equilibrium are shown in the upper half of Table 2, while additional moments and parameters for the market equilibrium are shown in the lower half.

Several of the moments used in this procedure relate to bidders' participation in auctions, and these participation estimates require an adjustment to match the data and theory. In our model, auctions proceed by sealed bid, yet eBay conducts ascending-like auctions. As pointed out by Song (2004), this can prevent a willing participant from submitting a bid if the auction's standing price passes his valuation before his arrival. While this does not affect the eventual winner or final price, it will cause the observed number of bidders to underestimate the true number of participants. We adjust for this bias using the approach proposed in Platt (2015). Under the assumption that participants in an auction arrive in random order,[24] Platt (2015) derives the probability that a random participant places a bid as:

$$P(\lambda) \equiv \frac{1}{\lambda}\left(2\left(\ln(\lambda) + \Gamma(1) - \Gamma(0,\lambda)\right) - 1 + e^{-\lambda}\right), \tag{29}$$

where $\Gamma(1) \approx 0.57721$ is Euler's constant and $\Gamma(0,\lambda) \equiv \int_{\lambda}^{\infty} \frac{e^{-t}}{t} dt$ is the incomplete gamma function. Thus, the expected number of bidders per auction is a strictly increasing function $\lambda \cdot P(\lambda)$, yielding a one-to-one relation between the average number of observed bidders and the average number of would-be auction participants. Also, $\tau\alpha P(\lambda)$ is the average rate at which buyers in the market place bids. These observations allow for identification and estimation of λ, α, and τ, as shown in Table 2.

[24] This assumption does not affect our model, since the payoff of losing participants is the same whether or not they were able to actually place a losing bid. The assumption implies that buyers with low willingess to pay will only be observed placing bids when the buyer happens to have arrived early compared to other bidders in that auction, as in Hendricks and Sorensen (2015).

We can also compute the number of bids a given buyer might place over the full duration of his search. This is given by:

$$D \equiv \tau\alpha \int_0^T \left(\frac{2\left(1 - e^{-\lambda F(s)}\right)}{\lambda F(s)} - e^{-\lambda F(s)} \right) ds. \tag{30}$$

The integrand is the probability that a participant in state s can place a bid (Platt, 2015). This object can then be used to estimate δ as shown in Table 2.

Several parameters can be directly observed in the data, such as ℓ (where we assume that, in this setting, inventory holding fees are primarily the listing fees paid to the platform), θ (which allows for identification of ρ), and η (where $1/\eta$ is the average duration of an auction). The remaining parameters—T, γ, and c—can be pinned down by equilibrium conditions. Section F of the Supplemental Appendix describes the identification and estimation of each parameter of the model in more detail.

In computing these moments, we normalize the posted price as $z = 1$. As discussed in Section 4.1, this has no effect on the distribution; $F(s)$, and $b(s)$ will scale proportionally. Therefore we similarly transform bidding data: for each item, we compute the average price (including shipping costs) among all sold posted-price listings (the analog of z) of that product, then divide all bids (including shipping costs) for that item by this product-level average. This rescaling is equivalent to "homogenizing" bids (Haile et al., 2003). All monetary values (θ, ℓ, c) are therefore in units of fractions of the average posted price. We consider one unit of time to be a month, which is also a normalization and it merely adjusts the interpretation of T and ρ. The observed values in the data reported in the second column of Table 2 are computed by taking averages across product-level averages; thus, this exercise should be interpreted as fitting the model for the average product.

The resulting parameter estimates are displayed in the final column of Table 2, which we discuss in the following subsections.

5.2 Willingness to Pay Over Buyers' Search

In our model, three factors drive buyers to shade down their bids early in the search and raise them later. First is their discounting of future consumption, as they do not want to pay for something far in advance of its use. Our monthly rate of time preference (ρ) of 5.5% falls roughly in the middle of empirical estimates surveyed in Frederick et al. (2002). They explain that these estimated rates are likely absorbing factors omitted from the estimated model, pushing them well above market interest rates. In the context of our model, for instance, discount rates could be higher that the pure rate of time preference due to risk aversion about future bidding opportunities, omitted opportunity costs of watching for auction listings, or increasing frustration with losing auctions.

The second factor driving bid shading is the intensity of competition is in each auction, since the buyer is unlikely to win early in his search with many competitors. We estimate an average of $\lambda = 13.10$ bidders per auction. On average, 21.13 new participants (δ) arrive per month. This ensures robust competition in each auction, and gives low probability of winning for anyone in the first two thirds of their search.

The third factor driving bid shading is the expected number of auctions that the buyer will encounter during his search. We estimate an arrival rate (α) of 12.76 auctions per month, but buyers only pay attention to a small fraction of these arriving auctions ($\tau = 0.064$). Even so, buyers begin their search with sufficient advance notice ($T = 10.30$ months) that they have a non-trivial chance of winning an auction before reaching their deadline.

We now use the estimated parameters to illustrate how buyers in the model systematically are willing to offer higher bids as their search continues. The solid line in the right panel of Figure 4 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 bids are dispersed across a range equal to 47% of the posted price; this dispersion becomes larger with greater impatience or longer deadlines. Initially (for s near T), the price path is more or less linear, but as the deadline approaches, greater curvature is introduced. The dashed line in the right panel of Figure 4 indicates the utility that the buyer gets if he purchases at time s, which increases as the deadline approaches purely due to time preference. The gap between the lines indicates shading relative to the bidder's current utility. This also highlights the fact that the increasing bids pattern predicted by the model is not solely due to impatience, but also reflects the reduced option value of future auction opportunities.

On average, a buyer increases his bid at a rate of 4.3 percentage points per month. As the average bidder participates in 1.18 auctions per month, that translates to an increase of 3.6 percentage points between each auction of a given item. As reported in Section 4.2, in the data we see an increase of 2.0 percentage points between each auction attempt. Bear in mind that in estimating the model's parameters we did not exploit any details about average bids over time.

Our model also explains why bidders with longer sequences start with a lower initial bid and why bidding sequences of different lengths should not cross, as Figure 1 indicates for the data. 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 estimated parameters and then summarize the data similarly, the resulting Figure 5 is qualitatively similar to the corresponding patterns in the data in Figure 1.

Next, consider the distribution of bids. The left panel of Figure 4 illustrates the equilibrium density of bidders. Note that $F'(s)$ is nearly constant from $s = 4$ to 10. With an average of $\lambda = 13$ participants per auction, those with lower valuations (hence longer time remaining) are highly unlikely to win. Yet the relative density cuts in half between $s = 4$ and $s = 1$, and then nearly does so again before $s = 0$. Those closest to their deadline are far more likely to win and exit.

Figure 6 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 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, with an average of $\lambda = 13$ bidders per auction, the highest bid 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 5.9 percentage points. As reported in Section 4.3, in the data the interquartile range of closing prices, after controlling for seller and product fixed effects, is 5 percentage points, serving as further evidence of the model's fit given that this price dispersion was not exploited in the estimation.

Our model predicts that the auction participant with the most time in the market will always win; however, in auction data, we do not observe when the participant entered the market, but only the first time he placed a bid. Indeed, the winner of a given auction may have been in the market the longest but only recently have been observed participating in a prior auction for the first time, giving him a short *observed* duration. We therefore measure the observed time in the market as the elapsed time since the buyer's first observed bid. Under that metric, the winner has the longest search duration in 45.7% of the simulated auctions, which is remarkably close to the 40% reported in Section 4.5 for our data.

The notable departure of the fitted 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 12.5% place a single bid. However, after conditioning on bidders who participated in three or more auctions, the model prediction is much closer to the observed distribution of bidding sequence lengths (see Figure A8 in the Supplemental Appendix).

In Section G of the Supplemental Appendix, we provide a comparison of this model to two other models: a static model and a stationary dynamic model (as opposed to our nonstationary dynamic model). In both cases bidders' willingness to pay is exogenously determined and stationary, unlike in our deadlines model. We demonstrate that ignoring the effects of time-sensitivity on bidders' willingness to pay can lead to miscalculations of objects of interest such as demand and consumer surplus: ignoring time-sensitivity in a dynamic model and treating bidders as having a constant willingness to pay in subsequent auctions substantially overestimates bidder surplus, and ignoring bidding dynamics altogether underestimates surplus. Incorrect estimates of the demand curve or consumer surplus could easily distort calculations needed for profit maximization, price discrimination, regulation, and many other applications.

5.3 Competing Mechanisms

Next, we consider what our model reveals concerning the simultaneous use of discount and non-discount mechanisms. Section 4.4 reports that 47% of sales occur via auction, while our fitted model predicts a somewhat higher $\sigma^* = 60\%$. Average auction revenue (θ) is 0.853. A sale in the posted-price market therefore generates $z - \theta = 14.6\%$ more revenue than a sale in the auction market, but this is offset in that the sale occurs after 1.8 months on average, which is 11 times longer than the average auction (0.156 months, implying an η of 6.39 auctions ending per month).

While informative, these parameters are essentially direct reflections of the data. However, the model uncovers the cost structure that sellers must face to optimally split between auctions in this way. With a production cost of $c = 0.840$, discount sellers are only paid a 1.3% markup on average. Moreover, nearly all the costs of production are incurred at the time of sale, with $\gamma = 0.984$, suggesting that a large portion of costs are shipping fees or that sellers may source their inventory at the time of the sale rather than at the time of listing. In Supplemental Appendix H, we find evidence that γ has increased in recent years, and we discuss possible implications or causes of this change.

5.4 Welfare Consequences of the Discount Sales Channel

Here, we consider whether the availability of the discount sales channel enhances or detracts from total expected welfare, comparing the dispersed (auctions and posted prices) versus degenerate (posted prices only) equilibria. For this exercise, we ignore whether sellers prefer posted prices over auctions (Eq. 28), since the latter could be inefficient even if sellers find them to be individually rational; thus, the counterfactual exercise asks whether welfare would increase if auctions were forbidden.[25]

Total expected welfare is derived from three sources in our environment: buyers, sellers, and *intermediaries* (any recipient of holding fees). A newly entering buyer expects utility of $V(T)$ by participating in the market (measured in terms of dollars, and net of any payments to sellers). We multiply this expected utility by $e^{\rho T}$ to measure it as of the time of consumption of the good, rather than the time of entry in the market. A newly entered seller earns zero expected profit due to free entry into the market.

For intermediaries, we consider two stark cases. In the first case, we suppose that holding fees merely cover the real resources that are consumed in the process of connecting buyers and sellers, such as maintaining warehouse space for inventory, building the intermediating platform, staffing showrooms, etc. Thus, intermediaries obtain zero surplus, total welfare is

[25] Auctions also have the potential to be welfare-improving by bringing in additional customers who value the good at less than the retail price z, though this is beyond the scope of our model.

simply equal to the buyers' expected utility, and auctions are socially beneficial, as shown below.

Proposition 5. *If holding costs are socially wasteful, then a dispersed equilibrium produces total welfare $V(T)e^{\rho T}$, which strictly exceeds the total welfare $x-z$ in a degenerate equilibrium.*

To appreciate this result, one should note that there are two avenues through which sellers are being driven to zero profit: one is to receive less revenue through auctions, and the other is to incur more costs by waiting longer for buyers in the posted-price market. But the latter is inefficient in this setting because it burns up real resources. Auctions also create some inefficiency as buyers pay for the product earlier than they need it, but most of the dissipated seller profit is transferred to buyers rather than destroyed. In this scenario, under our estimated parameters, the existence of auctions increase welfare compared to a market with only posted prices by 1.8% of the product's retail price.

In the second setting, suppose that the holding fees represent a pure transfer to the intermediaries, which would hold if their activities are costless or costs were previously sunk. In this case, the holding fees are all profit and contribute to social welfare. If a product is listed for t periods, then the present value of those holding fees is $\left(1-e^{-\rho t}\right)\ell/\rho$. With a posted-price listing, the duration is exponentially distributed with parameter ζ, whereas with an auction, the parameter is $\eta\left(1-e^{-\lambda}\right)$. Recall that fraction σ of products are listed as auctions, so the present value of expected profit to intermediaries is:

$$\sigma\int_0^\infty \eta\left(1-e^{-\lambda}\right)e^{-\eta\left(1-e^{-\lambda}\right)t}\frac{\left(1-e^{-\rho t}\right)\ell}{\rho}dt+(1-\sigma)\int_0^\infty \zeta e^{-\zeta t}\frac{\left(1-e^{-\rho t}\right)\ell}{\rho}dt. \tag{31}$$

After evaluating the integrals and substituting for the equilibrium values of σ and ζ, this becomes:

$$\frac{\ell}{\delta}\left(\frac{(z-c)(\kappa-\alpha)}{\rho(z-c\gamma)+\ell}+\frac{\alpha\left(1-e^{-\lambda}\right)}{\eta\left(1-e^{-\lambda}\right)+\rho}\right) \tag{32}$$

In the degenerate equilibrium, $\alpha=0$, so the expected intermediary profit becomes: $\frac{(z-c)\ell}{\rho(z-c\gamma)+\ell}$.

Comparing these, the change in expected profit from shutting down auctions is:

$$\frac{\ell\alpha\left(1-e^{-\lambda}\right)}{\delta}\left(\frac{z-c}{\rho(z-c\gamma)+\ell}-\frac{\theta-c}{\rho(\theta-c\gamma)+\ell}\right), \tag{33}$$

after substituting for η using (21). Since $z>\theta$, this expression is positive, so intermediaries always collect more revenue without auctions.

To determine the welfare effect of shutting down auctions, one must weigh the loss of expected utility to consumers against the increase in profits to intermediaries. This requires numeric evaluation due to the endogenous α, λ, and θ involved; however, it is typically the case that auctions are inefficient in the pure transfer setting. Under our estimated parameters (and assuming $x=z$), the intermediaries' profit falls from 14.4% to 6.5% of the retail price if auctions are allowed; even with the 5.6% benefit to buyers, auctions creates a net welfare loss equal to 2.3% of the retail price.

This is essentially a competition between two forms of inefficiency: in the discount sales mechanism, buyers are paying for the item earlier than they need it, but in posted-price listings, sellers are entering the market earlier than needed. The latter inefficiencies tend to be smaller; exceptions only occur when the holding costs ℓ are quite small or the flow of buyers δ is large.

Even so, the posted-price listings do create inefficiency, falling short of a first-best solution. For instance, if buyers could produce their own unit, total welfare would be $x - c$, which is 16.0% of the retail price under our estimated parameters. To see this formally, we compare this first-best welfare to expected welfare in the pure transfer degenerate equilibrium:

$$x - c > x - z + \frac{(z-c)\ell}{\rho(z - c\gamma) + \ell} \quad \iff \quad \frac{\rho(z-c\gamma)(z-c)}{\rho(z-c\gamma)+\ell} > 0.$$

Intermediation requires sellers to wait for buyers to arrive. This delay is socially costly, but unavoidable in this search environment.

In sum, the auction discount mechanism is welfare improving in a setting with costly intermediation, even with homogeneous goods that are readily available elsewhere. With costless intermediation, auctions are often wasteful. In reality, one would expect that a platform such as eBay can sustain some profits (for instance, network externalities make it beneficial for sellers and buyers to rely on a single platform), but it is also unlikely to be costless to facilitate the listings. An intermediate level of platform profits could result in auctions being welfare neutral.

In Tables A3 through A5 of the Supplemental Appendix, we examine whether the welfare implications of offering auctions for these new-in-box goods have changed over the period from 2010–2014. Einav et al. (2016), studying instead the pre-2010 period, document that posted-price listings grew in popularity, overtaking auctions by 2009.[26] The authors' model suggests that the primary cause for this decline pre-2010 was a shift in buyer preferences, with buyers becoming more time-sensitive and less interested in auctions. Using the same sample restrictions as in our main data sample, we recompute the model's parameters and the corresponding welfare measures in other years. We find that the rising use of posted prices relative to auctions has continued post-2010, and the flow of buyers δ has substantially declined, despite minimal change in auction prices. Moreover, unlike the pre-2010 period, auction participants in the post-2010 period have not become more time sensitive, as the search time has grown longer and buyers are more patient (larger T and smaller ρ).[27] The

[26] Einav et al. (2016) do not limit to new-in-box items as we do, as they study the period from 2003–2009 and eBay's categorization of new vs. used items is only available from 2010 onward.

[27] Even so, this could be driven by a changing composition of buyers: the most time-sensitive buyers could be entirely avoiding online auctions more than in the past and hence the remaining participants would look less time sensitive. Similarly, the remaining buyers in the market appear to be paying greater attention (τ) to

shift toward posted prices in the post-2010 period appears to have been driven instead by a change in how sellers produce/source/ship their inventory. In particular, the parameter γ has risen dramatically in the post-2010 period, consistent with a shift toward "just-in-time" production, which erodes some of the advantages of the auction mechanism. Assuming a pure transfer setting, we find that allowing auctions in 2010 increased welfare by 0.6% of the retail price; however, by 2014, auctions reduced welfare by 2.1%. Section H of the Supplemental Appendix provides further details.

6 Conclusion

This work examines consumer retail search in a new light, modeling decisions in a non-stationary environment where consumers grow less willing to search for a deal the longer they have been searching. Consumers are time sensitive and have deadlines by which they must obtain the good, leading to an increasing willingness to pay as consumers approach their deadlines. The model also rationalizes the coexistence of discount and full-price sales channels selling the same item, since transactions occur quickly in the former but at a lower price.

We apply the model to online retail sales, where consumers may acquire the good through auctions (the discount channel) or a posted price (the non-discount channel). We document a variety of reduced-form findings consistent with the time sensitivity we model. In particular, buyers offer 2% more in each successive attempt to win a discount, are more likely to win, and are more likely to pursue items with fast shipping. Under estimated parameters, our model can reasonably approximate several key moments of the data not used for estimation, and indicates a key role for buyer impatience and for the timing of production costs. Despite the redundancy of discount markets when the same good is readily available at a posted price, auctions are unambiguously welfare enhancing if the inventory holding fees merely cover the cost of intermediation. However, if intermediaries collect these fees as pure profit, welfare would typically increase by eliminating the discount sales channel.

While we have focused on new, homogenous goods sold online, the same lessons are equally applicable for impatient repeat buyers on imperfectly interchangeable items. Indeed, we anticipate similar results from other sales mechanisms where buyers must make repeated attempts, such as bargaining or shopping at physical discount outlets: time-sensitive buyers will adjust their strategy as they approach their deadlines and eventually resign themselves to the posted-price market. As in the online markets we study, this could generate diverse behavior in otherwise homogenous markets.

each auction.

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.

Astorne-Figari, C. and A. Yankelevich (2014): "Asymmetric Sequential Search," *Economics Letters*, 122.

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

Backus, M., T. Blake, B. Larsen, and S. Tadelis (2015): "Price Formation in Bilateral Trade: Evidence from Online Bargaining," Mimeo, Stanford University.

Backus, M. and G. Lewis (2016): "A Demand System for a Dynamic Auction Market with Directed Search," Working paper, Harvard University.

Backus, M. R., J. U. Podwol, and H. S. Schneider (2014): "Search Costs and Equilibrium Price Dispersion in Auction Markets," *European Economic Review*, 71, 173–192.

Bajari, P., R. McMillam, and S. Tadelis (2008): "Auctions Versus Negotiations in Procurement: An Empirical Analysis," *Journal of Law, Economics, and Organization*, 25, 372–399.

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.

Bell, D. R. and J. M. Lattin (1998): "Shopping Behavior and Consumer Preference for Store Price Format: Why "Large Basket" Shoppers Prefer EDLP," *Marketing Science*, 17, 66–88.

Bodoh-Creed, A., J. Boehnke, and B. Hickman (2016): "How Efficient are Decentralized Auction Platforms?" Working paper, UC Berkeley.

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

Bulow, J. and P. Klemperer (1996): "Auctions Versus Negotiations," *American Economic Review*, 86, pp. 180–194.

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

De los Santos, B., A. Hortaçsu, and M. R. Wildenbeest (2012): "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior," *American Economic Review*, 102, 2955–2980.

Diamond, P. (1987): "Consumer Differences And Prices In A Search Model," *Quarterly Journal of Economics*, 102, 429–436.

Diamond, P. A. (1971): "A Model Of Price Adjustment," *Journal Of Economic Theory*, 3, 156–168.

Einav, L., C. Farronato, J. D. Levin, and N. Sundaresan (2016): "Auctions versus Posted Prices in Online Markets," *Journal of Political Economy*, forthcoming.

Ellickson, P. B., S. Misra, and H. S. Nair (2012): "Repositioning Dynamics and Pricing Strategy," *Journal of Marketing Research*, 49, 750–772.

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

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.

Frederick, S., G. Loewenstein, and T. O'Donoghue (2002): "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, 40, 351–401.

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.

Hendricks, K. and A. Sorensen (2015): "The Role of Intermediaries in Dynamic Auction Markets," Working paper, University of Wisconsin.

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

Iyer, K., R. Johari, and M. Sundararajan (2014): "Mean Field Equilibria of Dynamic Auctions with Learning," *Management Science*, 60, 2949–2970.

Janssen, M. C. W. and A. Parakhonyak (2013): "Price Matching Guarantees and Consumer Search," *International Journal of Industrial Organization*, 31, 1–11.

——— (2014): "Consumer Search Markets with Costly Revisists," *Economic Theory*, 55, 481–514.

Jeitschko, T. D. (1998): "Learning in Sequential Auctions," *Southern Economic Journal*, 65, 98–112.

Julien, B., J. Kennes, and I. P. King (2001): "Auctions and Posted Prices in Directed Search Equilibrium," *Topics in Macroeconomics*, 1, 1–16.

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

KOHN, M. G. AND S. SHAVELL (1974): "The Theory of Search," *Journal of Economic Theory*, 9, 93–123.

KOULAYEV, S. (2014): "Search for Differentiated Products: Identification and Estimation," *RAND Journal of Economics*, 45, 553–575.

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," *Journal of Industrial Economics*, 55, 223–233.

MILGROM, P. R. AND R. J. WEBER (1982): "A Theory Of Auctions And Competitive Bidding," *Econometrica*, 50, 1089–1122.

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

PLATT, B. (2015): "Inferring Ascending Auction Participation from Observed Bidders," Working paper, Brigham Young University.

ROB, R. (1985): "Equilibrium Price Distributions," *Review of Economic Studies*, 52, 487–504.

ROBERT, J. AND D. O. STAHL (1993): "Informative Price Advertising in a Sequential Search Model," *Econometrica*, 657–686.

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," *Review of Economic Studies*, 44, 493–510.

SCHNEIDER, H. (2014): "The Bidder's Curse: Comment," *American Economic Review*, 106, 1182–1194.

SHRIVER, S. AND B. K. BOLLINGER (2015): "Structural Analysis of Multi-Channel Demand," *Columbia Business School Research Paper*.

SONG, U. (2004): "Nonparametric Estimation of an eBay Auction Model with an Unknown Number of Bidders," Working paper, University of British Columbia.

SOYSAL, G. AND L. KRISHNAMURTHI (2015): "How Does Adoption of the Outlet Channel Impact Customers' Spending in the Retail Stores: Conflict or Synergy?" *Management Science*.

STAHL, D. O. (1989): "Oligopolistic Pricing With Sequential Consumer Search," *American Economic Review*, 79, 700–712.

Stigler, G. J. (1961): "The Economics of Information," *Journal of Political Economy*, 213–225.

Van Den Berg, G. J. (1990): "Nonstationarity in Job Search Theory," *Review of Economic Studies*, 57, 255–277.

Wang, R. (1993): "Auctions Versus Posted-price Selling," *American Economic Review*, 83, 838–851.

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

Figure 1: Bids Over Time

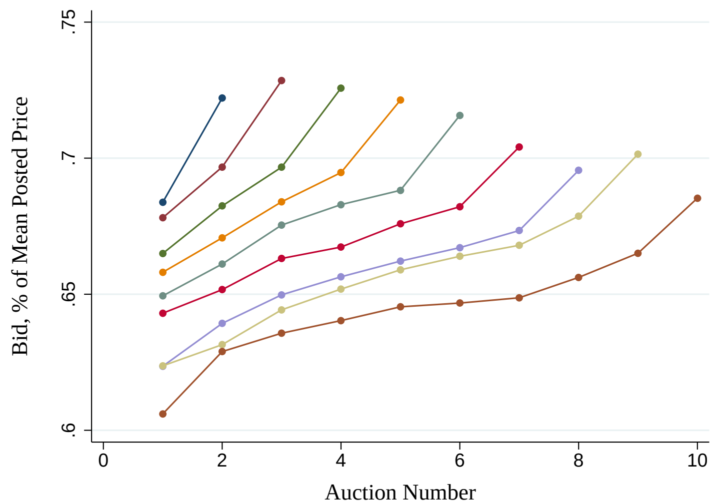

Notes: In the figure, a given line with n points corresponds to bidders who bid in n auctions total for a given product without winning in the first $n-1$ auctions. Horizontal axis represents auction number within the sequence (from 1 to n) and vertical axis represents the average normalized bid.

Figure 2: Distribution of Fraction of Sales By Auction Across Items

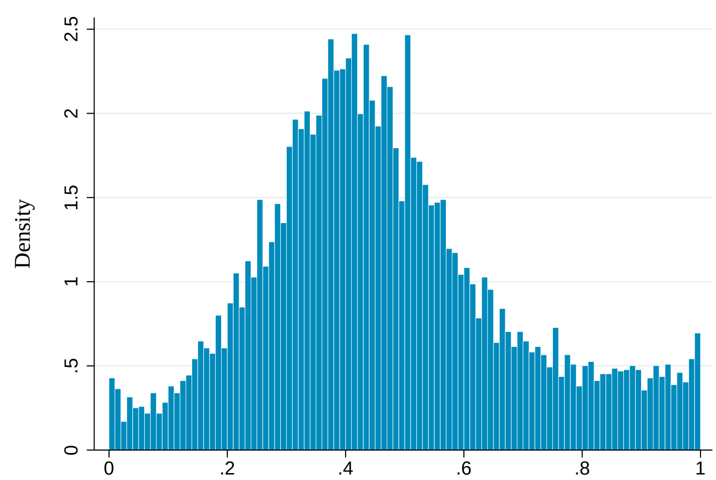

Notes: Figure shows a histogram at the product level of the fraction of listings sold by auction (rather than posted price) for a given product. The sample used to generate this figure is a superset of our main sample, containing the 12,368 distinct products with at least 50 transactions observed in the sample period without regards to listing method.

Figure 3: Cumulative Fraction Sold by Days Since Listing

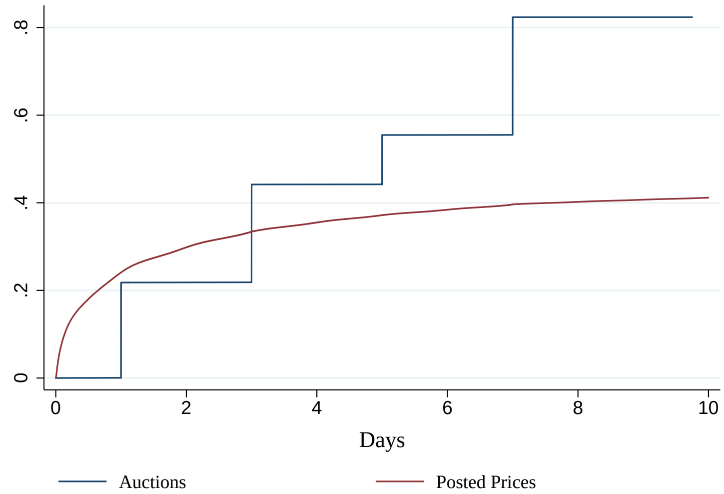

Notes: Figure displays the cumulative fraction of listings sold (vertical axis) against the number of days since the listing was posted (horizontal axis) for auctions and posted-price listings.

Figure 4: Bidding under Fitted Parameters

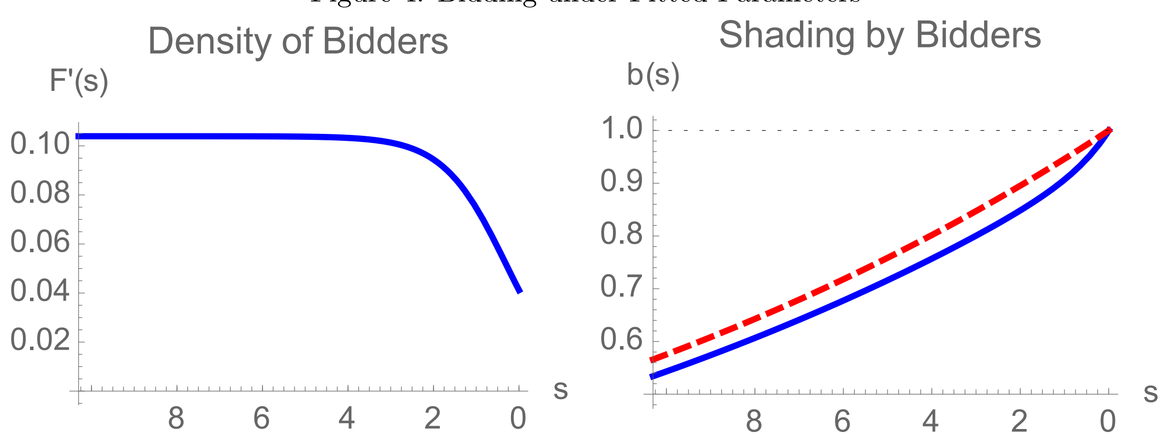

Notes: The left panel displays the density of bidders with s time remaining until deadline. The right panel indicates the bids (solid line) and utility (dashed line) as a function of time remaining s. Since $z = 1$, these may be read as percentages relative to the retail price.

Figure 5: Bids Over Time, Simulated Data

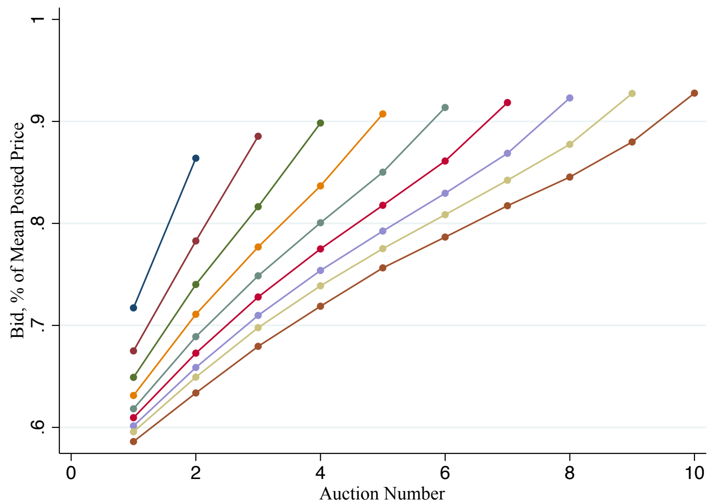

Notes: Figure displays average bids at each auction number for different auction sequence lengths, simulated from the model at the estimated parameter values. In the figure, a given line with n points corresponds to bidders who bid in n auctions total for a given product without winning in the first $n-1$ auctions. Horizontal axis represents auction number within the sequence (from 1 to n) and vertical axis represents the average normalized bid.

Figure 6: Price Density under Fitted Parameters

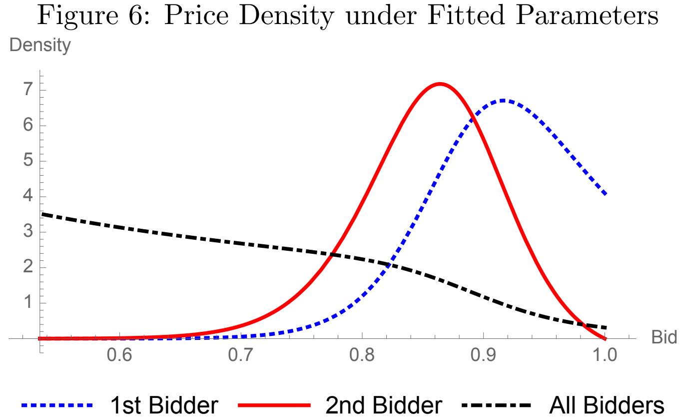

Notes: Figure displays the equilibrium density of the highest bid in an auction (dotted), the second highest bid (solid), and all bids (dot-dashed), simulated from the model at the estimated parameter values.

Table 1: Descriptive Statistics

A. Transaction level	Posted Price		Auctions	
	Mean	Std. dev.	Mean	Std. dev.
Bidders per transaction	1	–	5.30	2.20
Revenue per transaction	106.82	21.74	97.27	16.60
Revenue per transaction, normalized by avg posted price	1	–	0.85	0.17
Number of transactions	494,448		560,861	

B. Product level	Posted Price		Auctions	
	Mean	Std. dev.	Mean	Std. dev.
Number of transactions per product	134.98	220.82	153.12	343.63
Unique sellers per product	82.70	137.84	68.53	201.30
Unique buyers per product	129.03	208.02	606.08	1,365.60
Number of products	3,663			

Notes: Table displays descriptive statistics for primary data sample: transactions from October 1, 2013 through September 30, 2014 meeting the sample restrictions described in the text. All values are computed for completed (sold) listings. In panel A, values reported are means of product-level means and means of product-level standard deviations. Normalized revenue is computed by first dividing auction price by product-level average of posted-price sales. Panel B reports average and standard deviation, taken across all products, of the number transactions of a given product, the number of unique sellers selling a given product, and the number of unique buyers bidding in an auction for a given product or purchasing the product through a posted-price listing.

Table 2: Key Data Moments and Matching Parameter Values

	Observed Value in Data	Theoretical Equivalent	Fitted Parameter
Bidders per completed auction	5.30	$\frac{\lambda \cdot P(\lambda)}{1-e^{-\lambda}}$	$\lambda = 13.10$ (0.243)
Completed auctions per month	12.76	$\alpha\left(1 - e^{-\lambda}\right)$	$\alpha = 12.76$ (0.525)
Auctions a bidder tries per month	1.18	$\frac{\tau\alpha P(\lambda)}{1-e^{-\tau\alpha P(\lambda)}}$	$\tau = 0.064$ (0.002)
New repeat bidders per month who never win	7.41	$(\delta - \alpha)\left(1 - (1+D)e^{-D}\right)$	$\delta = 21.13$ (1.163)
—	—	Eq. (10)	$T = 10.30$ (0.255)
Average revenue per completed auction	0.853	θ	$\rho = 0.055$ (0.003)
Average listing fee paid	0.087	ℓ	$\ell = 0.087$ (0.0003)
Average duration of an auction listing (months)	0.156	$1/\eta$	$\eta = 6.39$ (0.028)
Average % of posted-price listing sold in 30 days	48.1%	$1 - e^{-\frac{\ell+\rho(1-\gamma)c}{z-c}}$	$\gamma = 0.985$ (0.039)
—	—	Eq. (21)	$c = 0.840$ (0.003)

Notes: Table displays observed moments in data and corresponding theoretical equivalent for additional market equilibrium parameters. Standard errors, from 200 bootstrap replications at the product level, are contained in parentheses. Data moments are averaged for each product (and month, where noted), then averaged across these.

Appendix: Proofs

Proof of Proposition 1. First, we note that the infinite sums in equations (2) and (6) can be readily simplified. In the case of the latter, it becomes:

$$F''(s) = \alpha\tau F'(s)e^{-\lambda F(s)}. \tag{34}$$

This differential equation has the following unique solution, with two constants of integration k and m:

$$F(s) = \frac{1}{\lambda}\ln\left(\frac{\alpha\tau - e^{\lambda k(s+m)}}{\lambda k}\right). \tag{35}$$

The constants are determined by our two boundary conditions. Applying (8), we obtain $m = \frac{1}{\lambda k}\ln\left(\alpha\tau - \lambda k e^{\lambda}\right) - T$. By substituting this into (35), one obtains:

$$F(s) = \frac{1}{\lambda}\ln\left(\frac{\alpha\tau - e^{\lambda k(s-T)}\left(\alpha\tau - \lambda k e^{\lambda}\right)}{\lambda k}\right). \tag{36}$$

The other boundary condition, (7), requires that k satisfy:

$$\alpha\tau\left(1 - e^{-\lambda T k}\right) - \lambda k\left(1 - e^{\lambda - \lambda T k}\right) = 0. \tag{37}$$

From (9), we know that $H = \delta/F'(T)$, and using the solution for F in (36), 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:

$$k = \frac{\tau}{\lambda}\left(\delta + \alpha e^{-\lambda}\right). \tag{38}$$

When we substitute this for k in (36), we obtain the equilibrium solution for F^* depicted in (11). Also, (38) is used to replace k in the boundary condition in (37), we obtain the formula for ϕ in (10) 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 (2). The first sum is similar to that in (6). 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:

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

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

$$(\rho + \alpha\tau e^{-\lambda F(s)})b(s) + b'(s) = \alpha\tau\left(e^{-\lambda}b(T) + \int_s^T \lambda e^{-\lambda F(t)} b(t)F'(t)dt\right). \tag{39}$$

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

$$(\rho + \alpha\tau e^{-\lambda F(s)})b'(s) + b''(s) = 0. \tag{40}$$

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:

$$b(s) = a_1 \cdot \left(\frac{\delta e^{\lambda^* - \tau T\left(\delta + \alpha e^{-\lambda^*}\right)}}{\rho} + \frac{\alpha e^{-\tau s\left(\delta + \alpha e^{-\lambda^*}\right)}}{\rho + \tau\left(\delta + \alpha e^{-\lambda^*}\right)}\right) e^{-s\rho} + a_2. \tag{41}$$

This solves the differential equation, but to satisfy (39), a particular constant of integration must be used. We substitute for $b(s)$ in (39) using (41), 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:

$$a_2 = a_1 \frac{\alpha\tau\left(\delta + \alpha e^{-\lambda^*}\right)}{\rho\left(\rho + \delta\tau + \alpha\tau e^{-\lambda^*}\right)} e^{-\rho T - \tau T\left(\delta + \alpha e^{-\lambda^*}\right)}. \tag{42}$$

The other constant of integration is determined by boundary condition (5). If we translate this in terms of $b(s)$ as we did for the interior of the HJB equation, we get $b(0) = z$. We then substitute for $b(0)$ using (41) evaluated at 0, and substitute for a_2 using (42), 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 (41), one obtains (13) with minor simplification. $\square$

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, equations (11) through (13) 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 (19) and (17), respectively.

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

The only remaining element regards expected auction profit. Equation (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 (21).

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

$$\theta = \frac{\lambda}{1-e^{-\lambda}}\left(e^{-\lambda} b(T) + \lambda \int_0^T b(s)F(s)F'(s)e^{-\lambda F(s)} ds\right). \tag{43}$$

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

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

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. $\square$

Proof of Proposition 4. The proposed Buyer and Market Equilibria still apply when $\alpha^* = 0$, bearing in mind that as $\alpha \longrightarrow 0$, the solution to $\phi(\lambda^*) = 0$ approaches $\lambda^* = \tau\delta T$. In the absence of auctions, the distribution of bidders is uniformly distributed across $[0, T]$, since none of them exit early; so $F^*(s) = s/T$ and $H^* = \delta T$. Moreover, the buyer's willingness to bid (if an auction unexpectedly occurred) reduces to: $b(s) = ze^{-\rho s}$.

For $\alpha^* = 0$ to be a market equilibrium, we need $\Pi_a^* \leq \Pi_p^*$. To prevent further enter, $\Pi_p^* = (1-\gamma)c$ is still required. From (15), if an auction were unexpectedly offered, the seller would generate $\Pi_a^* = \frac{\eta(\theta - \gamma c) - \ell}{\rho + \eta\left(1 - e^{\tau\delta T}\right)}$. Thus, the expected profit comparison simplifies to: $\theta \leq c + \frac{\ell + \rho c(1-\gamma)}{\eta\left(1 - e^{-\tau\delta T}\right)}$. This is equivalent to (28), where the left-hand side is evaluated from (43):

$$\begin{aligned}
\theta &= \frac{\tau\delta T}{1 - e^{-\tau\delta T}} \left(e^{-\tau\delta T} b(T) + \int_0^T b(s) F(s) F'(s) e^{-\tau\delta T F(s)} ds \right) \\
&= \frac{\tau\delta T}{1 - e^{-\tau\delta T}} \left(e^{-\tau\delta T} z e^{-\rho T} + \int_0^T z e^{-\rho s} \frac{s}{T^2} e^{-\tau\delta s} ds \right) \\
&= \frac{\tau\delta}{1 - e^{-\tau\delta T}} \cdot \frac{\tau\delta + \left(\rho T(\rho + \tau\delta) - \tau\delta\right) e^{-(\rho+\tau\delta)T}}{(\rho + \tau\delta)^2} z.
\end{aligned}$$

Thus, if (28) holds, then the profit from offering an auction is never greater than continuing to offer a posted-price listing, making $\alpha^* = 0$ an equilibrium. If (28) fails to hold, then $\alpha^* = 0$ cannot be an equilibrium, since some firms will earn greater profit by deviating and offering an auction.

To prove the last claim, first note that in a buyer equilibrium, $\lambda \to 0$ as $\alpha \to \infty$. In addition, $b(s) \to 0$ for all $s > 0$, because auctions occur every instant, in which the buyer faces no competition. Thus, expected revenue is 0 in the limit, yielding profit $\Pi_a < 0$ for $\alpha \to \infty$. At the same time, the violation of (28) is equivalent to $\Pi_a > 0$ for $\alpha = 0$. Since expected revenue is continuous in α, by the intermediate value theorem there must exist an $\alpha^* > 0$ such that $\Pi_a(\alpha^*) = 0$, which will constitute a dispersed equilibrium. $\square$

Proof of Proposition 5. Total expected welfare is simply the expected utility of the new entrant, measured at the time of consumption, $V(T)e^{\rho T}$. This is compared to the welfare that would occur if all buyers were forced to use the posted price, $x - z$. The former will be greater if:

$$
\begin{aligned}
x - \frac{z\kappa(\tau\kappa + \rho)}{\delta(\tau\kappa+\rho) + (\rho e^{\tau\kappa T} + \tau\kappa e^{-\rho T})\,\alpha e^{-\lambda}} > x - z \quad &\iff \\
\kappa(\tau\kappa+\rho) < \delta(\tau\kappa+\rho) + \left(\rho e^{\tau\kappa T} + \tau\kappa e^{-\rho T}\right)\alpha e^{-\lambda} \quad &\iff \\
\alpha e^{-\lambda}(\tau\kappa+\rho) < \left(\rho e^{\tau\kappa T} + \tau\kappa e^{-\rho T}\right)\alpha e^{-\lambda} \quad &\iff \\
\rho\left(e^{\tau\kappa T} - 1\right) - \tau\kappa\left(1 - e^{-\rho T}\right) > 0 &
\end{aligned}
$$

The l.h.s. of the last line is strictly increasing in T, with derivative: $\tau\kappa\rho\left(e^{(\tau\kappa+\rho)T} - 1\right)e^{-\rho T}$. Moreover, at $T = 0$, it evaluates to 0. Therefore, the expression is always greater than 0 for $T > 0$, and expected welfare is strictly greater with auctions than without.

□

Supplemental Appendix to "A Theory of Discounts and Deadlines in Retail Search"

Dominic Coey Bradley Larsen Brennan C. Platt[28]

A Model with Immediate Consumption

In the paper, we assumed that the good is only needed at the time of the deadline, as in the purchase of a birthday gift, travel arrangements, or seasonal sports equipment. Suppose instead that utility x is split between the immediate value from purchase, βx, and the value realized only at the deadline, $(1-\beta)x$, where the latter utility is discounted at rate ρ. The extreme of $\beta = 0$ indicates that the good is literally of no use until the date of the deadline, while $\beta = 1$ indicates that it starts providing the same flow of value regardless of when it is purchased. The intermediate case seems reasonable for many deadlines: for instance, a gift is not needed until the birthday, but the giver may enjoy some peace from mind of having it secured early. A spare automobile part provides similar insurance even if it is not literally needed until the failure of the part it replaces. Thus, if the good is purchased with s units of time remaining until the buyer's deadline, his realized utility is $(\beta + (1-\beta)e^{-\rho s})\, x$ minus the purchase price.

We assume throughout that $\beta x < z$; otherwise, buyers would strictly prefer using the posted-price option in equilibrium, as the possible savings from winning at auction would not compensate for delaying the immediate benefits of an outright purchase.[29]

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

$$b(s) = \left(\beta + (1-\beta)e^{-\rho s}\right) x - V(s). \tag{44}$$

In light of this, a buyer's HJB equation is revised to be:

$$\rho V(s) = -V'(s) + \tau\alpha\Big(\Pr(win|s)\left[\left(\beta + (1-\beta)e^{-\rho s}\right) x - V(s)\right] - E[payment|s]\Big). \tag{45}$$

Buyers also have the option to purchase from the posted-price listings at any time, receiving utility $(\beta + (1-\beta)e^{-\rho s})\, x - z$. However, a buyer in state s can obtain a present expected utility of $e^{-\rho s}(x - z)$ by waiting until $s = 0$ to make the purchase, and this is strictly preferred assuming that $\beta x < z$. 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

[28] Coey: eBay Research Labs, dcoey@ebay.com. Larsen: Stanford University, Department of Economics and NBER, bjlarsen@stanford.edu. Platt: Brigham Young University, Department of Economics, brennan_platt@byu.edu

[29] To see this in our equilibrium solution, note that buyers are willing to pay more than z (that is, $b^*(s) > z$ for $s > 0$) if and only if $\beta x \geq z$.

making the purchase:

$$V(0) = x - z. \tag{46}$$

Note that the steady-state conditions and equilibrium definitions are unaffected by the immediate consumption of the good by auction winners. Indeed, in the equilibrium solution, we obtain the same distribution of buyers, $F(s)$ and H. The only alteration is to the equilibrium bidding function:

$$b^*(s) = z - (z - \beta x)\frac{\delta(\kappa\tau + \rho)e^{\lambda^*}\left(1 - e^{-\rho s}\right) + \alpha\rho e^{\tau\kappa T}\left(1 - e^{-s(\rho+\tau\kappa)}\right)}{\delta\rho e^{\lambda^*} + \tau\kappa\left(\delta e^{\lambda^*} + \alpha e^{-\rho T}\right) + \alpha\rho e^{\tau\kappa T}}, \tag{47}$$

where $\kappa \equiv \delta + \alpha e^{-\lambda^*}$.

Proof. This closely follows the proof of Proposition 1. The HJB equation simplifies to:

$$\begin{aligned}\rho V(s) = -V'(s) + \alpha\tau\Bigg(& e^{-\lambda F(s)}\left(\left(\beta + (1-\beta)e^{-\rho s}\right)x - V(s)\right) \\ & - e^{-\lambda}b(T) - \int_s^T \lambda e^{-\lambda F(t)}b(t)F'(t)dt\Bigg).\end{aligned} \tag{48}$$

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

$$(\rho + \alpha\tau e^{-\lambda F(s)})b(s) + b'(s) = \rho\beta x + \alpha\tau\left(e^{-\lambda}b(T) + \int_s^T \lambda e^{-\lambda F(t)}b(t)F'(t)dt\right). \tag{49}$$

This equation holds only if its derivative with respect to s also holds. Yet this delivers the same differential equation as when $\beta = 0$:

$$(\rho + \alpha\tau e^{-\lambda F(s)})b'(s) + b''(s) = 0. \tag{50}$$

As in the paper, this differential equation has the following unique solution, with two constants of integration a_1 and a_2:

$$b(s) = a_1 \cdot \left(\frac{\delta e^{\lambda^* - \tau T\left(\delta + \alpha e^{-\lambda^*}\right)}}{\rho} + \frac{\alpha e^{-\tau s\left(\delta + \alpha e^{-\lambda^*}\right)}}{\rho + \tau\left(\delta + \alpha e^{-\lambda^*}\right)}\right)e^{-s\rho} + a_2. \tag{51}$$

What changes are the constants of integration. We substitute for $b(s)$ in (49) using (51) and evaluate it at $s = T$: $(\rho + \alpha\tau e^{-\lambda F(T)})b(T) + b'(T) = \rho\beta x + \alpha\tau e^{-\lambda}b(T)$. After substituting

$b(T)$, $b'(T)$, and $F(T)$, solving for a_2 yields:

$$a_2 = \beta x + a_1 \frac{\alpha\tau\left(\delta + \alpha e^{-\lambda^*}\right)}{\rho\left(\rho + \delta\tau + \alpha\tau e^{-\lambda^*}\right)} e^{-\rho T - \tau T\left(\delta + \alpha e^{-\lambda^*}\right)}. \tag{52}$$

The other constant of integration is determined by boundary condition (5), which is still $b(0) = z$. We then substitute for $b(0)$ using (51) evaluated at 0, and substitute for a_2 using (52), then solve for a_1:

$$a_1 = \frac{\rho(z - \beta x)\left(\rho + \delta\tau + \alpha\tau e^{-\lambda^*}\right) e^{\tau T\left(\delta + \alpha e^{-\lambda^*}\right)}}{\delta\rho e^{\lambda^*} + \tau\left(\delta + \alpha e^{-\lambda^*}\right)\left(\delta e^{\lambda^*} + \alpha e^{-\rho T}\right) + \alpha\rho e^{\tau T\left(\delta + \alpha e^{-\lambda^*}\right)}}. \tag{53}$$

If the solutions for a_1 and a_2 are both substituted into (51), one obtains (47) with minor simplification. □

Qualitatively, the model with immediate consumption behaves quite similarly to the model in the paper. As β increases, the bid $b(s)$ rises, all else equal. This will flatten the bidding function (since $b(0) = z$ as before), so that the bid increases by a smaller amount each unit of time. It also compresses the range of bids and reduces the variance of expected revenue. Moreover, immediate consumption will reduce auction inefficiencies as only a portion of the utility is delayed.

One difficulty in including immediate consumption in the estimation exercise is that β is not separately identifiable from x, as these two parameters always appear multiplied together in any of our equilibrium conditions. However, one could exogenously set $x = z$, and then estimate β by matching another moment of the distribution of auction revenue. We attempted this with median revenue as the additional moment. When $\beta = 0$, the model predicts a median revenue that is 85.4% of the posted price, while the data reports a median revenue of 83.2% . However, any increase in β would increase the predicted median. Thus, the corner solution of $\beta x = 0$ is the closest fit possible. The same occurs if instead the variance of revenue is used as the moment to match.

B Optimal 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 that has a present expected profit of Π_a. Of course, the seller gives up the immediate revenue and completion cost, which is no greater than $R - \gamma c$.

Since $\Pi_a = (1-\gamma)c$ in equilibrium, deviating to the reserve price R is certain to be unprofitable if $R - \gamma c < (1-\gamma)c$, or rearranged, $R < c$. 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 Proposition 4 in the paper does not apply). There is always an advantage to offering an auction if all other sellers offer posted-price listings: the product sells faster, even if at a slightly lower price.

C Discrete Time Derivation of Bellman Equations

Each of the continuous-time Bellman equations—(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:

$$\Pi_a = -\ell\Delta + \frac{1}{1+\rho\Delta}\left(\eta\Delta\left(1-e^{-\lambda}\right)(\theta-\gamma c) + \left(1-\eta\Delta\left(1-e^{-\lambda}\right)\right)\Pi_a\right). \tag{54}$$

The term $\ell\Delta$ is the holding 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:

$$\frac{\rho}{1+\rho\Delta}\Pi_a = -\ell + \frac{\eta\left(1-e^{-\lambda}\right)}{1+\rho\Delta}\left(\theta - \gamma c - \Pi_a\right), \tag{55}$$

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

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

$$\Pi_p = -\ell\Delta + \frac{1}{1+\rho\Delta}\left(\zeta\Delta(z-\gamma c) + (1-\zeta\Delta)\,\Pi_p\right). \tag{56}$$

Like auction sellers, posted-price sellers incur the holding 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:

$$\frac{\rho}{1+\rho\Delta}\Pi_p = -\ell + \frac{\zeta}{1+\rho\Delta}\left(z - \gamma c - \Pi_p\right), \tag{57}$$

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

The derivation for the buyer's expected utility is similar, only with more sources of uncertainty when 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:

$$\begin{aligned} V(s) = \frac{1}{1+\rho\Delta}\Bigg[&(1-\tau\alpha\Delta\Pr(win|s))\,V(s-\Delta) \\ &+ \tau\alpha\Delta\Bigg(\Pr(win|s)xe^{-\rho s} - E[payment|s]\Bigg)\Bigg]. \end{aligned} \tag{58}$$

On the right-hand side, all utility is discounted by factor $1/(1+\rho\Delta)$, since the buyer does not receive any current-period utility during the search. By the next period, one of two outcomes could occur: either the buyer wins an auction and exits (second line of (58)), 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 ($\Pr(win|s)$) times the utility enjoyed from winning ($xe^{-\rho s}$), minus the expected payment required. 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 Hamilton-Jacobi-Bellman equation, we first multiply both sides by $(1 + \rho\Delta)/\Delta$, then subtract $V(s)/\Delta$ from both sides, obtaining:

$$\rho V(s) = \frac{V(s-\Delta) - V(s)}{\Delta} + \tau\alpha\Big(\Pr(win|s) \left[xe^{-\rho s} - V(s-\Delta)\right] - E[payment|s]\Big). \tag{59}$$

Then, by letting $\Delta \to 0$, we obtain (2).

D Alternative Discount Mechanisms

Our model of non-stationary search for discounts can be readily adapted for settings beyond auctions. Here, we briefly outline several examples of how the buyer's and seller's search problems could be formulated.

First, consider physical search for a homogenous good where sellers post a price, but discovering these sellers is time consuming. At each encounter, the buyer will learn a specific seller's price but has to purchase immediately or lose the opportunity. The buyer in state s will formulate a reservation price $b(s)$, purchasing if and only if the quoted price is at or below $b(s)$. Let $G(s)$ depict the cumulative distribution of sellers offering a price at or above $b(s)$. One could say that a firm charging $b(s)$ is targeting buyers of type s, and will only sell to those who have s or less time remaining. In this case, the probability that a buyer "wins" the discount is:

$$\Pr(win|s) = 1 - G(s), \tag{60}$$

since the buyer will reject any discount targeted at buyers more desperate than himself. The expected payment would be:

$$E[payment|s] = \int_s^T b(t)dG(t). \tag{61}$$

When offered, the buyer accepts any price between $b(T)$ and $b(s)$, but pays nothing if a higher price is offered (which occurs with probability $G(s)$).

From the seller's perspective, deeper discounts result in lower revenue but a higher likelihood of sale. A seller who targets buyers with s time remaining will only complete the sale to fraction $F(s)$ of buyers but will be paid $b(s)$ when the sale is completed. Thus, the discount mechanism generates an expected profit of:

$$\rho\Pi_a = -\ell + \eta F(s)\left(b(s) - \gamma c - \Pi_a\right). \tag{62}$$

The equilibrium in this environment is characterized in Akin and Platt (2011).

Alternatively, consider an environment in which buyers are randomly paired with sellers and enter Nash bargaining. Again, let $G(s')$ denote the distribution of seller states, where a seller in state s' is willing to accept any price at or above $b(s')$. Upon meeting, their private states are revealed. Matches with negative surplus are ignored, while matches with positive surplus lead to a sale with a price $\omega b(s) + (1-\omega)b(s')$, where ω is the Nash bargaining power of the seller. Here, a buyer in state s will only make a purchase if the seller is willing to accept a lower price than $b(s)$, which occurs if $s' > s$; so the buyer "wins" the discount with probability:

$$\Pr(win|s) = 1 - G(s). \tag{63}$$

The expected payment would be:

$$E[payment|s] = \int_s^T \big(\omega b(s) + (1-\omega)b(s')\big)\, dG(s'). \tag{64}$$

A bargaining seller of type s' would only find a mutually agreeable price with buyers of type $s < s'$, which occurs in a random match with probability $F(s')$. The exact price depends on the type of the buyer, so we integrate over all possibilities.

$$\rho\Pi_a(s') = -\ell + \eta\left(\int_0^{s'} \big(\omega b(s) + (1-\omega)b(s') - \gamma c\big)\, dF(s) - F(s')\Pi_a(s')\right). \tag{65}$$

Finally, consider a lottery setting. Here, buyers are occasionally presented with a lottery as the discount option, with the freedom to buy as many tickets $k(s)$ as desired, with one being selected at random to win. If the number of lottery tickets purchased by others are distributed according to $G(k')$, then the probability of winning would be:

$$\Pr(win|s) = \int_0^T \frac{k(s)}{k(s)+k'} dG(k'). \tag{66}$$

If p denotes the price of one lottery ticket, then the expected payment would be:

$$E[payment|s] = pk(s). \tag{67}$$

For a lottery seller, the revenue is simply the number of tickets sold, while the lottery will result in a winner for sure at its close. The expected profit would then be:

$$\rho\Pi_a = -\ell + \eta\left(\int_0^T pk' dG(k') - \Pi_a\right). \tag{68}$$

To our knowledge, these non-stationary bargaining and lottery problems have not been studied before, but present interesting settings for future work.

E Comparative Statics

In this section we discuss the comparative statics results from Tables A1 and A2. Although our buyer equilibrium has no closed form solution, these comparative statics can be obtained

Table A1: Comparative Statics on Key Statistics: Buyer Equilibrium

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

Note: Derivations are provided in Section E of the Supplemental Appendix. $**$ indicates that the sign depends on parameter values. Sufficient conditions for a positive sign are $\delta\tau T > 1$ and $\tau(\kappa - \alpha) > \rho > \tau(2\kappa - \alpha)\sqrt{\tau\kappa T e^{-\lambda}}$. An exact condition is provided in the proof.

by implicit differentiation of $\phi(k)$, which allows for analytic derivations reported below. For the market equilibrium, 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 discuss several of these results not discussed in the body of the paper, and then provide analytic derivations where available.

Table A1 reports the sign of the derivatives of four key statistics. The first and second are the average number of participants per auction, λ^*, which reflects how competitive the auction is among buyers, and the average mass of buyers in the market, H^*, which is always proportional to λ^*. Third is the flow of buyers who never win an auction and must use the posted-price listings; in the next section we will see that this crucially affects the profitability of the posted-price market. Fourth 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 simplest computation occurs at $s = T$. This comparative static also captures price dispersion, both within auctions and between auctions and posted prices. The posted price z is fixed, so a lower $b^*(T)$ indicates greater dispersion.

Changes in α have an intuitive impact. With more frequent auctions (reduced search frictions) the value of continued search is greater as there are more opportunities to bid. The increase in auctions creates more winners, reducing the stock of bidders and the number of competitors per auction. Both of these effects lead bidders to lower reservation prices.

Changes in τ have nearly the reverse effect from that of α, though there are opposing forces at work. A higher likelihood of participating also reduces the search friction of a given bidder, as he will participate in more of the existing auctions. However, all other bidders are more likely to participate as well. The net result is typically higher bids, because the greater

Table A2: Comparative Statics on Key Statistics: Market Equilibrium

		$\partial/\partial\tau$	$\partial/\partial\rho$	$\partial/\partial T$	$\partial/\partial c$	$\partial/\partial\ell$	$\partial/\partial\gamma$
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.

number of competitors dominates the increased auction participation to reduce the value of search. However, this does depend on parameter values; in particular, when τ or ρ are very close to zero, extra participation dominates extra competitors, leading to lower bids.

The rate of time preference has no impact on the number or distribution of bidders, as ρ does not enter into equations (10) through (12). 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 less when their utility from future consumption is valued less. Changes in ρ will play a greater role once we endogenize sellers in the market equilibrium.

We can also consider the effect (not shown in Table A1) of the parameter change on the expected revenue generated in an auction. For the first three parameters, 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 is further away; there, the additional participants override the lower initial bid, driving up expected revenue.

Table A2 displays comparative statics in the market equilibrium. First, we note that for increases in τ under the market equilibrium, the bid response is typically opposite of that in the buyer equilibrium. The difference is that when buyers are more attentive, sellers are willing to offer more auctions in the market equilibrium. Additional auctions will improve the continuation value of buyers, and thus reduce their bids. There will still be 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 market equilibrium.

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 discount buying opportunities. This drives up bidders' reservation prices, but not enough to prevent a smaller flow of auction sellers.

A higher holding fee, ℓ, has a similar effect, but this is surprising because the holding 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 market design if buyer valuations are not fundamental but rather the endogenous results of deeper factors. A seemingly neutral change in the holding fee not only alters which market sellers use, but also warps the distribution of buyer valuations.

E.1 Preliminary Steps for Comparative Statics Derivations

Here we derive the signs of comparative statics that are reported in Table 1 of the paper. Because we do not have a closed form solution for the endogenous number of participants per auction, we use implicit differentiation of $\phi(\lambda^*) = 0$ from (10) to determine the effect of the exogenous parameters on λ^*. In preparation for this, we initially note that $\phi'(\lambda) < 0$ for all λ:

$$\frac{\partial \phi}{\partial \lambda} = -\alpha e^{-\lambda} - (\tau T \alpha + e^{\lambda})\delta e^{-\tau T \kappa} < 0, \tag{69}$$

where $\kappa \equiv \delta + \alpha e^{-\lambda}$ is used for notational convenience, though we treat κ as a function of α and λ when taking derivatives.

Also note that $H = \frac{\lambda^*}{\tau}$ and $F'(0) = \kappa - \alpha$, while the lowest bid is:

$$b(T) = ze^{-\rho T} \cdot \frac{\kappa\left(\tau\kappa + \rho\right) e^{\lambda^*}}{\tau\kappa\left(\delta e^{\lambda^*} + \alpha e^{-\rho T}\right) + \rho\left(\delta e^{\lambda^*} + \alpha e^{\tau T \kappa}\right)}. \tag{70}$$

Because this is always evaluated at the equilibrium λ^*, we can substitute for e^{λ^*} using

$\phi(\lambda^*) = 0$, which is $\delta e^{\lambda} = (\kappa - \alpha)e^{\tau T \kappa}$, thus obtaining:

$$b(T) = \frac{ze^{-\rho T}}{\delta} \cdot \frac{(\tau\kappa + \rho)(\kappa - \alpha)}{\tau\left(\kappa - \alpha + \alpha e^{-(\rho - \tau\kappa)T}\right) + \rho}. \tag{71}$$

E.2 Auction Rate, α

Using implicit differentiation, we compute the effect of α on λ^*.

$$\begin{aligned} \frac{\partial \phi}{\partial \alpha} &= -1 + e^{-\lambda} + \tau T \delta e^{-\tau T \kappa} && (72)\\ &= -1 + e^{-\lambda}\left(1 + \left(\frac{\delta + \alpha e^{-\lambda} - \alpha}{\delta + \alpha e^{-\lambda}}\right)\ln\left(\frac{\delta e^{\lambda}}{\delta + \alpha e^{-\lambda} - \alpha}\right)\right). && (73) \end{aligned}$$

The second equality comes from substituting for T using a rearrangement of $\phi(\lambda^*) = 0$, which is $T = \frac{1}{\tau\kappa}\ln\left(\frac{\delta e^{\lambda}}{\kappa - \alpha}\right)$.

By rearrangement, $\frac{\partial \phi}{\partial \alpha} \leq 0$ if and only if:

$$\ln\left(\frac{\delta e^{\lambda}}{\delta + \alpha e^{-\lambda} - \alpha}\right) - \left(e^{\lambda} - 1\right)\frac{\delta + \alpha e^{-\lambda}}{\delta + \alpha e^{-\lambda} - \alpha} \leq 0 \tag{74}$$

As $\lambda \longrightarrow 0$, the left-hand side approaches 0. If we take the derivative of the left-hand side w.r.t. λ, we obtain:

$$-\frac{\left(e^{\lambda} - 1\right)\left(\alpha + \delta e^{\lambda}\right)\left(2\alpha + e^{\lambda}(\delta - \alpha)\right)}{\left(\alpha + (\delta - \alpha)e^{\lambda}\right)^2} \tag{75}$$

Each parenthetical term is strictly positive for all $\lambda > 0$, so the left-hand side of (74) is strictly decreasing in λ. Thus, (74) strictly holds for any $\lambda > 0$, including the equilibrium λ^*. Therefore, $\frac{\partial \phi}{\partial \alpha} < 0$, and $\frac{\partial \lambda}{\partial \alpha} = -\left(\frac{\partial \phi}{\partial \alpha}\right) / \left(\frac{\partial \phi}{\partial \lambda}\right) < 0$. Specifically,

$$\frac{\partial \lambda}{\partial \alpha} = -\frac{1 - (1 + \tau T(\kappa - \alpha))e^{-\lambda}}{\kappa - \alpha + (1 + \tau T(\kappa - \alpha))\alpha e^{-\lambda}}. \tag{76}$$

Next, consider the impact on the fraction purchasing from posted-price listings, which is affected both directly by α and indirectly through λ:

$$\frac{\partial F'(0)}{\partial \alpha} = e^{-\lambda} - 1 + \alpha \cdot \frac{\partial \lambda}{\partial \alpha}. \tag{77}$$

This is strictly negative because $e^{-\lambda} < 1$ and $\frac{\partial \lambda}{\partial \alpha} < 0$.

To demonstrate the effect to α on the bidding function, we use the alternate depiction in terms of the function $g(t)$:

$$b(T) = \frac{g(T)}{g(T) + \rho\int_0^T g(t)dt},$$

recalling that

$$g(t) \equiv \tau e^{-\rho t}\left(\kappa - \alpha\left(1 - e^{-t\tau\kappa}\right)\right).$$

Of course, $g(t)$ is a function of α (including its effect on κ), so let $g_\alpha(t)$ denote its derivative with respect to α. Thus,

$$g_\alpha(t) = \tau e^{-\rho t}\left(e^{-\tau\kappa t} + \frac{\kappa\left(1 - \alpha\tau t e^{-\tau\kappa t}\right)}{\alpha + (\kappa - \alpha)\left(e^\lambda + \alpha\tau T\right)} - 1\right).$$

When we take the derivative of $b(T)$ w.r.t. α, we obtain:

$$\frac{\partial b(T)}{\partial \alpha} = z\rho \frac{\int_0^T \left(g(t)g_\alpha(T) - g(T)g_\alpha(t)\right)dt}{\left(g(T) + \rho\int_0^T g(t)dt\right)^2}.$$

The denominator is clearly positive. The numerator is always negative; in particular, at each $t \in [0, T]$, the integrand is negative. This integrand simplifies to:

$$-\frac{\kappa\tau^2 e^{-(t+T)(\kappa\tau+\rho)}\left(\alpha^2\tau(T-t) + (\kappa-\alpha)\left(\alpha\tau(T-t)e^{\kappa\tau T} + e^\lambda\left(e^{\kappa\tau T} - e^{\kappa t\tau}\right)\right)\right)}{\alpha + (\kappa-\alpha)\left(e^\lambda + \alpha\tau T\right)} < 0.$$

The inequality holds that because $T \geq t$ and $\kappa > \alpha$, making each parenthetical term in the expression positive.

E.3 Attention, τ

Using implicit differentiation, we compute the effect of τ on λ^*.

$$\frac{\partial \phi}{\partial \tau} = \delta\kappa T e^{-\tau T\kappa} > 0. \tag{78}$$

All of these terms are strictly positive. Since $\frac{\partial \phi}{\partial \lambda} < 0$, then by implicit differentiation, $\frac{\partial \lambda}{\partial \tau} = -\left(\frac{\partial \phi}{\partial \tau}\right) / \left(\frac{\partial \phi}{\partial \lambda}\right) > 0$. Specifically,

$$\frac{\partial \lambda}{\partial \tau} = \frac{\delta\kappa T e^\lambda}{\alpha e^{\tau T\kappa} + \delta e^\lambda\left(e^\lambda + \alpha\tau T\right)}. \tag{79}$$

Next, consider the impact on the fraction purchasing from posted-price listings. The probability of participation τ has no direct effect on $F'(0)$, but affects it only through λ:

$$\frac{\partial F'(0)}{\partial \tau} = \frac{\partial F'(0)}{\partial \lambda} \cdot \frac{\partial \lambda}{\partial \tau} = -\alpha e^{-\lambda} \cdot \frac{\partial \lambda}{\partial \tau} \tag{80}$$

which is always negative.

Finally, consider the effect on the lowest bid. Here, the sign of the derivative will depend on parameter values, so it is more convenient to take comparatives on (71) rather than examining it in terms of $g(t)$. Since $\kappa'(\tau) = \alpha e^{-\lambda}\lambda'(\tau)$, the comparative static on $b(T)$ works out to:

$$\frac{\partial b(T)}{\partial \tau} = \frac{z\alpha e^{\lambda}\psi}{(\kappa-\alpha)\left(\tau\alpha + (\tau(\kappa-\alpha)+\rho)e^{T(\rho+\tau\kappa)}\right)^2\left(\alpha + (\kappa-\alpha)\left(\tau\alpha T + e^{\lambda}\right)\right)}. \tag{81}$$

where

$$\begin{aligned}\psi \quad \equiv \quad & e^{\lambda}(\kappa-\alpha)^2\left(\rho(\tau\delta T - 1) + \delta\kappa\tau^2 T - \frac{\alpha e^{-\lambda}\rho}{\kappa-\alpha}\right) \\ & + \delta e^{\lambda+\rho T}\left(\rho\left(e^{\lambda}(\kappa-\alpha)+\alpha\right) - T(\kappa\tau+\rho)\left(\tau(\kappa-\alpha)^2+\kappa\rho\right)\right).\end{aligned}$$

The lowest bid is increasing in τ if and only if $\psi > 0$, since the remaining terms in $\frac{\partial b(T)}{\partial \tau}$ are always positive.

To verify the sufficient conditions listed under Table A1 in the paper, note that $\tau\delta T > 1$ ensures that the first term in the first line is positive. For the remaining terms of the first line, note that $\delta\kappa\tau^2 T > \kappa\tau$ by the same assumption. Moreover, since $\kappa > \alpha$ and $1 > e^{-\lambda}$, then $\delta\kappa\tau^2 T > \alpha\tau e^{-\lambda}$. Thus, the sufficient condition $\tau(\kappa-\alpha) > \rho$ ensures that $\delta\kappa\tau^2 T > \frac{\alpha e^{-\lambda}\rho}{\kappa-\alpha}$.

For the second line, we note that by omitting the first and last α in the first step, then applying the second sufficient condition twice in the second, we get:

$$\begin{aligned}\rho\left(e^{\lambda}(\kappa-\alpha)+\alpha\right) - T(\kappa\tau+\rho)\left(\tau(\kappa-\alpha)^2+\kappa\rho\right) \quad & > \quad \rho e^{\lambda}(\kappa-\alpha) - T\left(\tau\kappa+\rho\right)^2\kappa \\ & > \quad \frac{\rho^2 e^{\lambda}}{\tau} - T\left(\tau(2\kappa-\alpha)\right)^2\kappa.\end{aligned}$$

The third sufficient condition, $\rho > \tau(2\kappa-\alpha)\sqrt{\tau\kappa T e^{-\lambda}}$, ensures that this last term is positive.

E.4 Impatience, ρ

The rate of time preference ρ does not enter into ϕ, so therefore $\frac{\partial\phi}{\partial\rho} = 0$ and $\frac{\partial\lambda}{\partial\rho} = 0$. Similarly, ρ has no direct effect on $F'(0)$ or indirect effect through λ.

To demonstrate the effect to ρ on the bidding function, we use the alternate depiction in terms of the function $g(t)$:

$$b(T) = \frac{g(T)}{g(T) + \rho\int_0^T g(t)dt},$$

recalling that

$$g(t) \equiv \tau e^{-\rho t}\left(\delta + \alpha\left(e^{-\lambda} + e^{-t\tau\left(\delta+\alpha e^{-\lambda}\right)} - 1\right)\right).$$

Of course, $g(t)$ is a function of ρ, so let $g_\rho(t)$ denote its derivative with respect to ρ. Thus,

$$g_\rho(t) = -t\tau e^{-\rho t}\left(\delta + \alpha\left(e^{-\lambda} + e^{-t\tau\left(\delta+\alpha e^{-\lambda}\right)} - 1\right)\right).$$

Therefore, when we take the derivative of $b(T)$ w.r.t. ρ, we obtain:

$$\frac{\partial b(T)}{\partial \rho} = z\frac{\int_0^T \left(\rho g(t) g_\rho(T) - \rho g(T) g_\rho(t) - g(t) g(T)\right) dt}{\left(g(T) + \rho \int_0^T g(t) dt\right)^2}.$$

The denominator is necessarily positive. We will show that the integrand is negative for all t, implying that $\frac{\partial b(T)}{\partial \rho} < 0$. The integrand simplifies to:

$$\frac{\tau^2(\rho(t-T)-1)}{e^{(t+T)\left(\tau\left(\alpha e^{-\lambda}+\delta\right)+\rho\right)}} \cdot \left(\left(\alpha(1-e^{-\lambda})-\delta\right) e^{\tau t\left(\alpha e^{-\lambda}+\delta\right)} - \alpha\right) \cdot \left(\left(\alpha(1-e^{-\lambda})-\delta\right) e^{\tau T\left(\alpha e^{-\lambda}+\delta\right)} - \alpha\right).$$

Since $t \leq T$, the numerator is always negative, and the exponential term in the denominator is always positive. Finally, we note that $\alpha\left(1-e^{-\lambda}\right) - \delta < 0$ because $\delta - \alpha\left(1-e^{-\lambda}\right) - \delta e^{\lambda - \tau T\left(\delta + \alpha e^{-\lambda}\right)} = 0$ in equilibrium. This ensures that second and third parenthetical terms are negative.

E.5 Deadline, T

Using implicit differentiation, we compute the effect of T on λ^*.

$$\frac{\partial \phi}{\partial T} = \delta \kappa \tau e^{\lambda^*} e^{-\tau T \kappa}, \tag{82}$$

which is clearly positive. Then by implicit differentiation, $\frac{\partial \lambda}{\partial T} = -\left(\frac{\partial \phi}{\partial T}\right) / \left(\frac{\partial \phi}{\partial \lambda}\right) > 0$. Specifically,

$$\frac{\partial \lambda}{\partial T} = \frac{\delta \tau \kappa}{\delta\left(1 + \tau T \alpha e^{-\lambda^*}\right) + \alpha e^{\tau T \kappa - 2\lambda^*}} \tag{83}$$

Moreover, the number of buyers H^* is not directly affected by T, so it increases only because λ^* increases.

Next, consider the impact on the fraction purchasing from posted-price listings. The deadline T has no direct effect on $F'(0)$, but affects it only through λ:

$$\frac{\partial F'(0)}{\partial T} = \frac{\partial F'(0)}{\partial \lambda} \cdot \frac{\partial \lambda}{\partial T} = -\alpha e^{-\lambda} \cdot \frac{\partial \lambda}{\partial T} \tag{84}$$

which is always negative.

To demonstrate the effect of T on the bidding function, we again use the definition of $b(T)$ in terms of $g(t)$, but to distinguish between an intermediate time t and the initial time T, we write it as:

$$b(T) = \frac{g(T,T)}{g(T,T) + \rho \int_0^T g(t,T) dt},$$

where

$$g(t,T) \equiv \tau e^{-\rho t}\left(\kappa - \alpha\left(1 - e^{-t\tau\kappa}\right)\right),$$

where T only affects the expression by changing λ and hence changing κ.

The derivative of $b(T)$ w.r.t. T is thus:

$$\frac{\partial b(T)}{\partial T} = -\frac{z\rho\left(\int_0^T\left(\frac{g(T,T)^2}{T} - g(t,T)g_t(T,T)\right)dt + \int_0^T\Big(g(T,T)g_T(t,T) - g(t,T)g_T(T,T)\Big)dt\right)}{\left(g(T)+\rho\int_0^T g(t)dt\right)^2},$$

where g_t and g_T are derivatives with respect to the first and second terms, respectively. Specifically, these evaluate to:

$$g_t(T,T) = \left(\rho\tau(\alpha-\kappa) - \alpha\tau(\kappa\tau+\rho)e^{-\tau\kappa T}\right)e^{-T\rho}$$

and

$$g_T(t,T) = \alpha\tau\left(\alpha t\tau - e^{\kappa t\tau}\right)e^{-\lambda - t(\kappa\tau+\rho)}\lambda'(T).$$

Because $\kappa > \alpha$, we know that $g_t(T,T) < 0$ and $g(t,T) > 0$ for all t. Thus, the first integral in the numerator is always positive.

The integrand of the second integral simplifies to $\mu(t)\alpha^2\tau^2\lambda'(T)e^{-\lambda-(t+T)(\tau\kappa+\rho)}$, where:

$$\mu(t) \equiv e^{\tau\kappa t}(\tau T(\alpha-\kappa)-1) + e^{\tau\kappa T}(t\tau(\kappa-\alpha)+1) + \alpha\tau(t-T).$$

We have already shown that $\lambda'(T) > 0$; thus, to show that the integral is positive, we only need to show that $\mu(t) \geq 0$ for all t. First note that $\mu(T) = 0$ and $\mu(0) = e^{\tau\kappa T} - \tau\kappa T - 1 > 0$. To see the latter inequality, note that this has the form $e^x - x - 1$, which is equal to 0 at $x = 0$ and has a positive derivative $e^x - 1 \geq 0$ for all x.

Next, note that $\mu''(t) = -(1+\tau T(\kappa-\alpha))\tau^2\kappa^2 e^{\tau\kappa t} < 0$ for all $t \in [0,T]$. Since $\mu(0) > \mu(T) = 0$ and $\mu''(t) < 0$, then $\mu(t) > 0$ for all $t \in [0,T)$.

Thus, the integrand of the second integral is always positive. Thus $\frac{\partial b(T)}{\partial T} < 0$.

F Model Estimation Details

We first discuss the moments used to identify the parameters of the buyer equilibrium and then discuss those used for estimating the full market equilibrium. The buyer equilibrium depends on one endogenous variable (λ) and five parameters (α, δ, τ, T, and ρ); the market equilibrium depends on the additional parameters ℓ, η, γ, and c. These are estimated one at a time, in the following sequence.

The first moment in Table 2 is the number of observed bidders per completed auction. The theoretical equivalent of this is $\lambda P(\lambda)$ bidders per auction, since $P(\lambda)$ is the probability that a participant is able to place a bid and be observed, as described in Section 5. In the data, we restrict ourselves to auctions that resulted in a sale, meaning that at least one bidder arrived. In the model, this occurs with probability $1-e^{-\lambda}$, so we divide $\lambda P(\lambda)$ by this to get the average number of observed bidders per *completed* auctions.

The second moment is the number of completed auctions per month. The model predicts α auctions per month, but this includes auctions that close without a bidder. Thus, $(1-e^{-\lambda})\alpha$ provides the average number of completed auctions per month.

In the third moment, we compute the number of auctions per month that a buyer places a bid. Again, this is measured conditional on having at least one bid by that buyer during the month. The number of observed bids per month is Poisson distributed with parameter $\tau\alpha P(\lambda)$, since opportunities arise for the bidder at rate α and are used with probability τ, but will only register as a bid with probability $P(\lambda)$. To obtain the predicted conditional average, we divide this by $1 - e^{-\tau\alpha P(\lambda)}$.

To get a measure of the market size, we measure the number of new bidders per month. We focus this by conditioning on bidding at least twice during the span of the data; this is done through the last term, since D is the Poisson rate of bids per buyer over the duration of his search (as described in Section 5). However, D is computed for a complete search spell (*i.e.* all T periods); thus, we multiply by $\delta-\alpha$ rather than δ to get the average flow per month of newly observed repeat bidders who never win, and similarly condition the data moment on never winning.

The deadline length T is not directly observable, since we cannot observe when a buyer entered or exited the market, but only the first and last times they bid on an item. Instead, we select T so that the equilibrium condition $\phi(\lambda) = 0$ holds. Effectively, we are choosing the deadline T such that the observed number of bidders per auction (which is endogenously determined) will occur in equilibrium.

The rate of time preference is determined by matching the mean of auction revenue θ, whose formula is reported in the Appendix with the proof of Proposition 3. Note that θ is already conditional on having at least one bid, consistent with the computed data moment.

The parameters in the full market equilibrium can also be computed using simple empirical averages or functions of these averages. First, for simplicity, we treat inventory holding costs as equaling the listing fees paid to the platform. These fees are directly observable, yielding an estimate of ℓ. The average time for which an auction is listed is also observable in the data, yielding an estimate of $1/\eta$.

We also observe the fraction of posted-price listing that sell within the 30 day window of their listing. Since the model predicts that posted-price sellers exit at Poisson rate ζ, $1-e^{-\zeta}$ will have exited within one month; the table simply inserts the equilibrium value for ζ. We also note that α, which was computed in the second row of Table 2 for the buyer equilibrium, is now endogenous. In the market equilibrium, we can use the equilibrium condition for α in (21) to determine the underlying cost of production. The last two conditions are solved jointly to recover γ and c.

We now summarize what sources of variation in the data leads to the most change in the estimate of each parameter. Variation in many of the parameters (λ, α, τ, δ, ℓ, and η) maps directly back to variation in the data moments found in the corresponding rows of Table 2. The effect is less obvious for those involving the equilibrium conditions (T, ρ, γ, and c), which we now consider in turn.

The parameter T relies on the requirement for a buyer equilibrium found in (10), which effectively balances buyer entry and exit in steady state. It is most sensitive to increases in the number of bidders per auction (the first data moment in Table 2). If the number of bidders per auction increases, buyers are more likely to hit their deadline without winning, causing too much exit relative to entry; but the estimation restores balance by keeping buyers in the market longer, giving them more chances to win. While α, τ, and δ also appear in (10), the effect of their corresponding moments on the estimated T are somewhat muted.

The parameter ρ is directly determined by revenue per auction, holding the five preceding parameters fixed. Higher revenue occurs if buyers are more patient, since they bid closer to the posted price. The number of bidders per auction also indirectly affects ρ because, as explained above, an increase in the number of bidders per auction leads to an increase in the inferred T, and as T increases, the average bidder will be further from his deadline and bid less.

The parameters γ and c are jointly determined in the last two rows of Table 2. Recall that the market equilibrium requirement in (21) determines the endogenous number of auctions so that both mechanisms will be equally profitable. As a consequence, c is closely tied to average auction revenue, keeping ex-post profits low. Thus, γ is left to reconcile the time it takes to sell a posted-price listing and the number of auctions listings. The latter has a particularly large impact, as seen in Table A5. While the auction revenue and speed of selling have hardly changed, the number of auctions declined precipitously; this can be justified if γ increased, because posted-price sellers will have fewer up-front costs before the long wait for a buyer.

G Alternative Models for Interpreting Auction Data

G.1 Overview of Results

Our model considers buyers with repeat opportunities to bid while facing a deadline; but to what extent do these features matter in empirical uses of auction data? We explore this question by comparing how data generated in our setting would be interpreted under two standard auction models. If estimates of the demand curve or consumer surplus are biased under these standard approaches, they would distort calculations needed for profit maximization, price discrimination, regulation, and many other applications.

Second-price auctions are empirically convenient because bids truthfully reveal the bidders' valuation, so the observed distribution of bids traces a demand curve for the product. Truthful revelation also allows a precise measurement of consumer surplus, since the first price truthfully reveals the winner's valuation, while the second price indicates what winner actually paid. Yet these interpretations of bidding data are only correct in a *static* second-price auction, with private valuations independently drawn from an exogenous distribution; estimates of demand or consumer surplus can be wildly skewed if the same interpretation is applied to bidding data generated in a setting of dynamic bidding by time-sensitive buyers.

In our setting, the buyer's value, $xe^{-\rho s}$, is no longer the same as willingness to pay, $b(s) = xe^{-\rho s} - V(s)$. Buyers are truthful about their willingness to pay, but they they do not bid their full value because tomorrow's bidding opportunities provide positive expected surplus. They would rather lose and continue their search than pay a price above $b(s)$, even if they would have gotten positive surplus. Thus, the bid only reveals how much the bidder values *the current auction*, not the product itself. Thus, the static interpretation of our data will underestimate demand — on average by 5.3% of the retail price using our estimated parameter values, as illustrated in the dotted line in Figure A7. Similarly, auction winners in our model enjoy a consumer surplus equal to 8.7% of the retail price; yet the static interpretation would estimate a surplus of 5.4%, underestimating by 39%.

Of course, other models with an infinite sequence of auctions (Zeithammer, 2006; Ingster, 2009; Said, 2011; Backus and Lewis, 2016; Bodoh-Creed et al., 2016) can make a similar critique, since the option to participate in future auctions reduces willingness to bid in those models as well. Yet the non-stationary search process in our model still leads to significantly different interpretations of bid data. For illustration, we compare our model to a stationary dynamic auction environment in which a bidder of type s (where s is no longer the time until deadline but instead a fixed type) enjoys utility $x(s)$ from purchase, and this is constant throughout his search. Types are still distributed according to $F(s)$, and bidders still participate at rate $\tau\alpha$ with an average of λ bidders per auction. For low bids, this model predicts practically no shading, while high bidders shade aggressively (by 60%). Thus, if the data were generated by in our non-stationary environment, but then interpreted using the stationary dynamic model, demand would be overstated by an average of 2.5% of the retail price (the dashed line in Figure A7), while consumer surplus would be highly overestimated at a value of 52.4% of the retail price.

G.2 Details on Alternative Models

Here we describe the static and the stationary dynamic models that could be used for the interpretation of bidding data. We highlight how product demand and consumer surplus would be determined in each.

For the static model, assume that bidder types s determine their valuation $x(s)$, which is a decreasing function of s. Types are independently drawn from an exogenous distribution $F(s)$. Each bidder has only one opportunity to bid. In such a model, the optimal bid will be $b(s) = x(s)$, so that bids precisely reveal the underlying utility of bidders.

For the stationary dynamic model, bidder types s still determine their valuation $x(s)$, and these valuations are persistent throughout their search. Types in a given auction are distributed by $F(s)$, which could be endogenously determined. Bidders participate in auctions at rate $\tau\alpha$ with an average of λ bidders per auction. In this dynamic environment, the continuation value of a bidder is:

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

The optimal bid is $b(s) = x(s) - V(s)$; so after substituting this into the HJB equation, it simplifies to:

$$x(s) \equiv b(s) + \frac{\tau\alpha}{\rho}\left(e^{-\lambda F(s)}b(s) - e^{-\lambda}\left(b(T) + e^{\lambda}\int_s^T b(t)\lambda e^{-\lambda F(t)}F'(t)\,dt\right)\right). \tag{85}$$

That is, by observing all bids, $b(s)$, and their distribution, $F(s)$, one can infer the underlying utility of the bidders.

To estimate demand in the static model, the econometrician inverts the empirical CDF of bids. We employed a parametric plot of $(H \cdot F(s), b(s))$ to create the dashed line in Figure A7. In the stationary dynamic model, one would use the empirical CDF of bids in (85) to determine the valuation associated with each $b(s)$. The demand curve plots $(H \cdot F(s), x(s))$. Finally, in our deadline model, one would estimate the parameter values as indicated in Section 5, with the demand curve plotting $(H \cdot F(s), xe^{-\rho s})$. Note that x is not identifiable from the model, as it drops out of the equilibrium solution; here, we set $x = z$, which creates the smallest difference between the static model and ours.

Next, for each model we provide the theoretical expressions for consumer surplus in the average auction, conditional on having any participants. These can then be computed using the estimates from the preceding paragraph. In our deadline model, this would be:

$$CS_{\text{model}} = \frac{\int_0^T xe^{-\rho s}\lambda e^{-\lambda F(s)}F'(s)ds - \int_0^T b(s)\lambda^2 e^{-\lambda F(s)}F(s)F'(s)ds - e^{-\lambda}b(T)}{1 - e^{-\lambda}}. \tag{86}$$

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 the static model, consumer surplus is the difference between the first price (which truthfully reveals the winner's valuation) and the second price (which the winner actually pays). Therefore, we replace $xe^{-\rho s}$ in (86) with $b(s)$, which simplifies to:

$$CS_{\text{static}} = \frac{\lambda}{1-e^{-\lambda}} \left(\int_0^T b(s) e^{-\lambda F(s)} (1 - \lambda F(s)) F'(s) ds - e^{-\lambda} b(T) \right). \tag{87}$$

In the stationary dynamic model, the expected consumer surplus replaces $xe^{-\rho s}$ in (86) with $x(s)$:

$$CS_{\text{stationary}} = \frac{\int_0^T x(s) \lambda e^{-\lambda F(s)} F'(s) ds - \int_0^T b(s) \lambda^2 e^{-\lambda F(s)} F(s) F'(s) ds - e^{-\lambda} b(T)}{1-e^{-\lambda}}. \tag{88}$$

When the static model is applied to deadline data, it consistently underestimates demand and consumer surplus. This is because it ignores the continuation value of further search, which causes the observed bid to be less than the true value. When the stationary dynamic model is applied to deadline data, it rotates the true demand curve, estimating it to be steeper than it really is. This is because buyers with a low s are almost certain to win in any auction, and thus can drastically shade their bid in hopes of a good deal. As s increases, the bid still falls, but the bid shading becomes less extreme; those with a very high s hardly shade at all because they are so unlikely to win given that $V(s)$ is nearly zero. In other words, the stationary model interprets the true valuations as being nearly the same as the lowest bids, but rising much steeper for those who bid more. The overestimates of consumer surplus on high bids far outweigh the underestimates on low bids.

H Estimation Across Years

This section considers how online retail markets have evolved over a five-year period, and examines how this affects our estimated parameters. Using one-year intervals, we apply the same sample restriction as in the paper to eBay data from October 1st, 2010 through September 30th, 2015. Table A3 reports the descriptive statistics from each year, while Table A4 summarizes the data moments used to estimate our parameters. Note that the number of auction transactions has fallen dramatically. The number of unique auction buyers has also fallen significantly.

We then recompute the model's parameters and the corresponding welfare measures for each year, as reported in Table A5. In reviewing the estimated parameters and data moments, it is noteworthy that some have changed very little. In particular, the average auction revenue has held fairly steady at $\theta \approx 85\%$ of the retail price (see the fifth row of Table A4). In Table A5, the production cost remains about 1.4% below expected revenue. Thus, over this time frame, it is not that auctions have become less profitable. Even so, the flow of buyers into the market (δ) has decreased by nearly half, and similarly for the rate at which auctions are offered (α).

On its face, time sensitivity does not seem to be the driving force, as the search time T

has grown longer and buyers are more patient (smaller ρ). Indeed, the remaining buyers in the market appear to be paying greater attention (τ) to each auction. Even so, this could be driven by a changing composition of buyers: δ has fallen dramatically, and if those who no longer participate were the most time-sensitive buyers, those who remain would look less time sensitive.

One of the most dramatic changes over time is found in γ, the fraction of production costs that can be delayed until the time of sale. We estimate $\gamma = 35\%$ in 2010, but found it to be above 85% in the last three years of data. The primary feature of the data driving this estimated decline is that auction revenue and the selling speed of posted-price listings both stayed fairly constant over this period, implying that costs were fairly constant; yet the endogenous arrival rate of auctions decreased, which can be justified if production costs are incurred later.[30] This could be due to changes in the relative importance of shipping costs, changes in how sellers source their inventories, or other movements toward a "just-in-time" production setting. As this occurs, the long wait (for sellers) associated with posted-price listings becomes less costly and more beneficial for overall welfare.

Indeed, these delayed costs have greatly improved market efficiency. Between 2010 and 2014, total welfare in the dispersed equilibrium has grown by 7% from 0.112 to 0.120 of the retail price. However, the welfare under a degenerate, posted-price-only market has grown even more (33%), from 0.106 to 0.141 of the retail price. This created a large enough shift to reverse the welfare comparison: In 2010, allowing auctions increased welfare 6% above what posted-prices alone could offer (0.112 vs. 0.106), but by 2014, a market with auctions provides 15% less welfare than the posted-prices alone (0.120 vs. 0.141). This assumes that holding fees are a pure transfer (profit) to the platform, which is the least favorable scenario for judging auction efficiency.

[30] The end of Section F examines the sensitivity of each parameter to specific data moments.

Figure A1: Bids Over Time, Regression Results

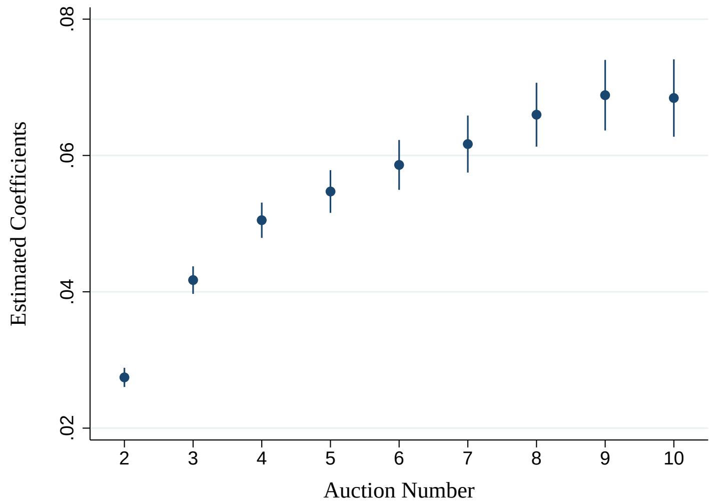

Notes: Figure displays estimated coefficients for dummy variables for each auction number (i.e. where the auction appears in the sequence) from a regression of normalized bid on these dummies and on dummies for the length of auction sequence. This regression is performed after removing outliers in the auction number variable (defined as the largest 1% of observations). 95% confidence intervals are displayed about each coefficient.

Figure A2: Bids Over Time, Excluding Winning Bids

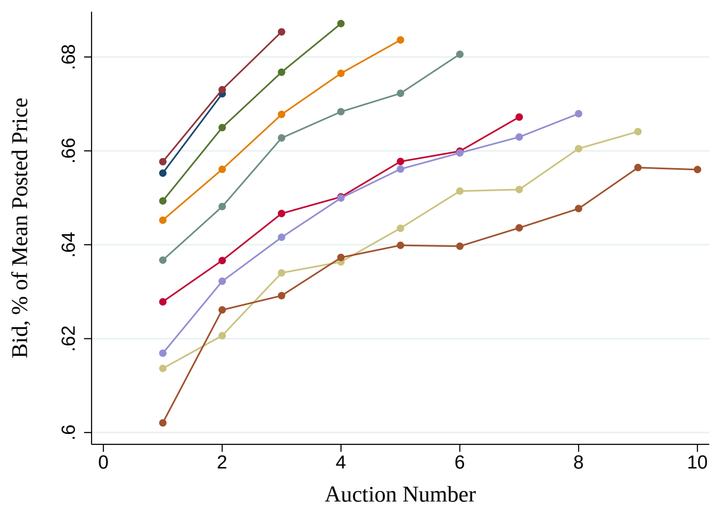

Notes: Figure displays the same analysis as in Figure 1 from the paper, but with winning bids excluded. In the figure, a given line with n points corresponds to bidders who bid in n auctions total for a given product without winning in the first $n-1$ auctions. Horizontal axis represents auction number within the sequence (from 1 to n) and vertical axis represents the average normalized bid.

Figure A3: Time To Posted-Price Purchase Since Last Losing Auction

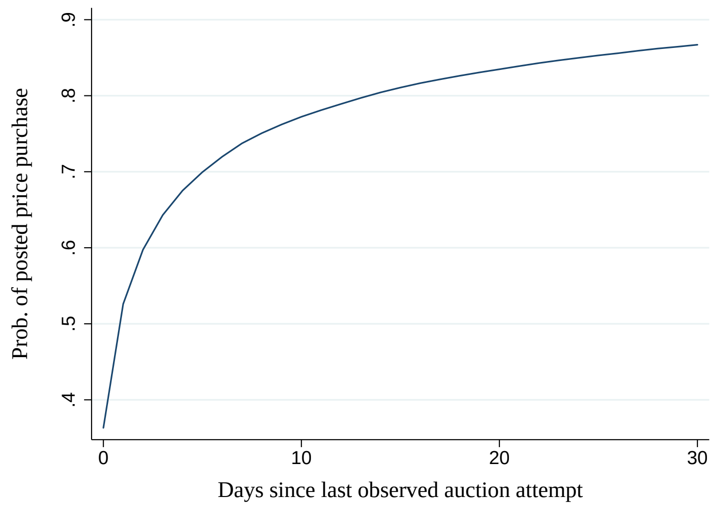

Notes: Figure displays cumulative density of the time difference between the last observed auction attempt and the posted-price purchase conditioning on bidders who attempted an auction and did not win and were later observed purchasing the good on an eBay posted-price listing.

Figure A4: Fraction Bidding on Listings with Fast Shipping

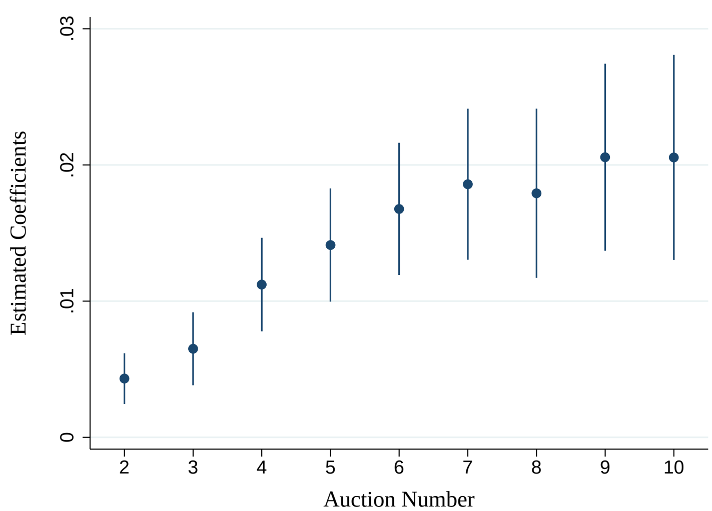

Notes: Figure displays estimated coefficients for dummy variables for each auction number (i.e. where the auction appears in the sequence) from a regression of a fast-shipping dummy (an indicator for whether the listing offered a shipping option guaranteed to arrive within 96 hours) on these auction number dummies and on dummies for the length of auction sequence. This regression is performed after removing outliers in the auction number variable (defined as the largest 1% of observations). 95% confidence intervals are displayed about each coefficient. The regression constant (overall mean of the fast shipping dummy) is 0.43. Averaging over all regression coefficients yields an estimate of 0.0028, or 0.6% of the overall mean.

Figure A5: Bids over Time, Experienced Bidders

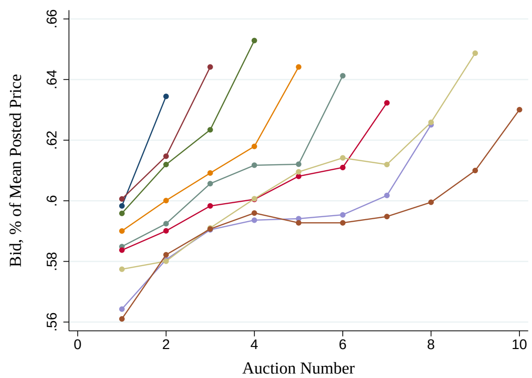

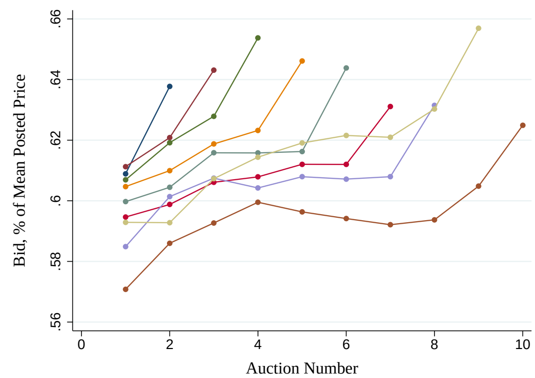

Notes: Figures limit to bidders who have bid in at least 50 auctions (left panel) or bidders who have bid in at least 10 auctions for products in the same product grouping (right panel) in the past year prior to the current auction. In the figures, a given line with n points corresponds to bidders who bid in n auctions total for a given product without winning in the first $n-1$ auctions. Horizontal axis represents auction number within the sequence (from 1 to n) and vertical axis represents the average normalized bid.

Figure A6: Bids Over Time, Products With Average Transaction Price of At Least \$100

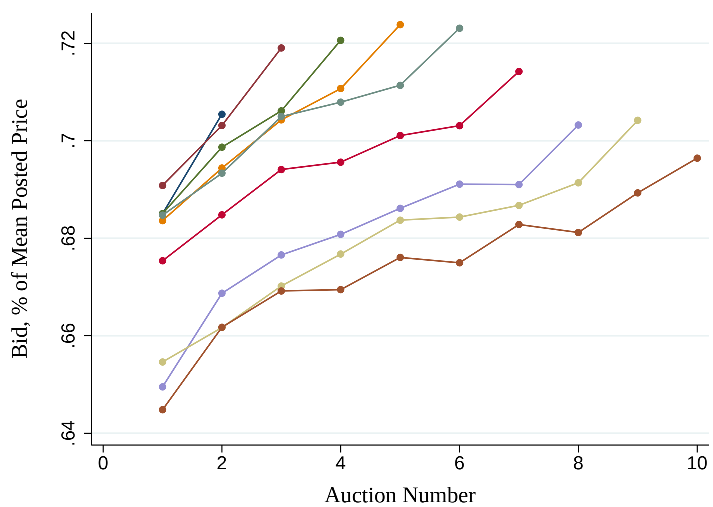

Notes: Figure limits to products with average transaction price $\geq$ \$100. In the figure, a given line with n points corresponds to bidders who bid in n auctions total for a given product without winning in the first $n-1$ auctions. Horizontal axis represents auction number within the sequence (from 1 to n) and vertical axis represents the average normalized bid.

Figure A7: Auction Demand Curve

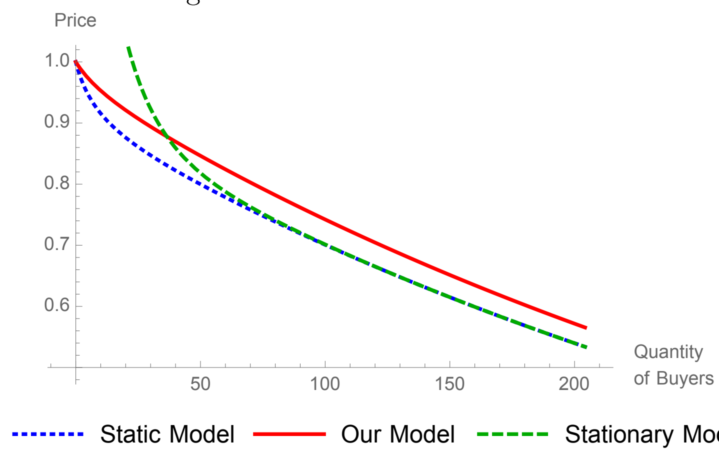

Notes: Figure shows inferred auction demand curve using the deadlines model (solid red line) vs. treating the data as though it came from a static model (dotted blue line) or a stationary dynamic model (dashed green line). The dashed line is truncated, but would intersect the vertical axis at a price of 2.5.

Figure A8: Participation by Number of Auctions

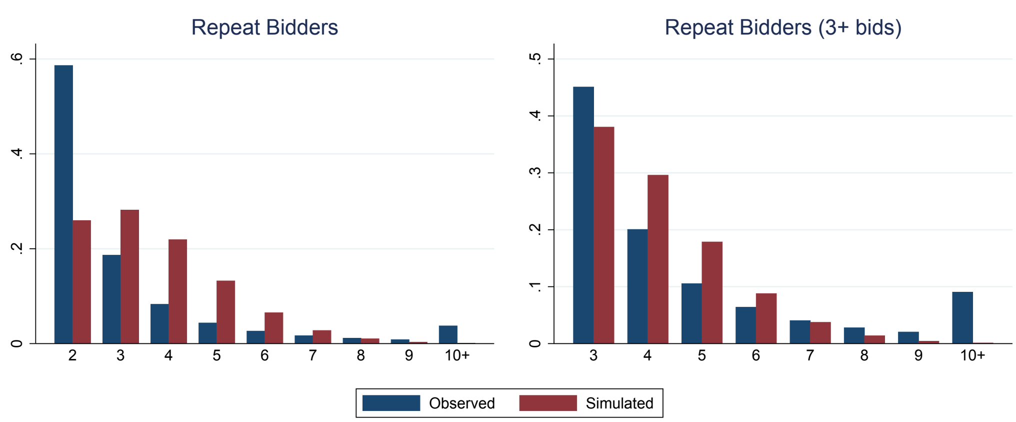

Notes: Each histogram reports the fraction of bidders who participate in a given number of auctions, conditional on participating in at least two (left panel) or three (right panel) auctions, for observed data (blue) and simulated data (red).

Table A3: Descriptive Statistics From 2010–2014 Data Samples

	2010	2011	2012	2013	2014
Number of products	4,060	3,704	4,231	3,663	2,102
Posted prices					
Transactions	335,691	416,300	537,441	494,448	269,263
Revenue	104.55 (23.89)	115.54 (23.45)	108.24 (23.22)	106.82 (21.74)	127.08 (23.22)
Transactions per product	82.68 (116.63)	112.39 (149.47)	127.02 (166.57)	134.98 (220.82)	128.10 (142.06)
Unique sellers per product	45.47 (57.09)	66.62 (85.24)	75.86 (92.86)	82.70 (137.84)	78.49 (79.18)
Unique buyers per product	79.77 (108.52)	106.32 (138.21)	117.65 (149.02)	129.03 (208.02)	122.42 (130.46)
Auctions					
Transactions	914,219	795,699	818,223	560,861	268,001
Revenue	93.14 (16.40)	105.15 (16.86)	99.20 (16.71)	97.27 (16.60)	116.52 (19.09)
Normalized revenue	0.83 (0.17)	0.85 (0.14)	0.87 (0.17)	0.85 (0.17)	0.86 (0.18)
Bidders per transaction	4.99 (2.21)	5.27 (2.29)	5.41 (2.28)	5.30 (2.20)	5.41 (2.05)
Transactions per product	225.18 (383.68)	214.82 (331.60)	193.39 (320.73)	153.12 (343.63)	127.50 (177.81)
Unique sellers per product	102.36 (144.91)	95.40 (126.33)	82.97 (109.57)	68.53 (201.30)	60.49 (69.04)
Unique buyers per product	792.02 (1468.70)	783.63 (1299.21)	734.97 (1212.34)	622.18 (1394.80)	508.69 (632.59)

Notes: Table displays descriptive statistics computed as in Table 1 using data samples from 2010–2014. Data was constructed using same sample restrictions as in main sample. Each year label corresponds to one year of data beginning in Oct. of that year and continuing through Sep. of the following year (i.e. the 2013 column corresponds to the main data sample used in the body of the paper). Values in table are computed as described in notes of Table 1. In Revenue, Normalized revenue, and Bidders per transaction rows, values reported are means of product-level means, with means of product-level standard deviations in parentheses. In all rows specifying per-product measures, values reported are the average values across all products, with standard deviations across products in parentheses.

Table A4: Data Moments Used in Estimation From 2010–2014 Data Samples

	2010	2011	2012	2013	2014
Bidders per completed auction	4.988	5.267	5.406	5.301	5.408
Completed auctions per month	18.765	17.902	16.116	12.760	10.625
Auctions a bidder tries per month	1.215	1.215	1.201	1.176	1.150
New repeat bidders per month who never win	10.635	10.522	9.552	7.406	5.261
Average revenue per completed auction	0.835	0.852	0.868	0.853	0.855
Average listing fee paid	0.070	0.076	0.081	0.087	0.089
Average duration of an auction listing (months)	0.154	0.154	0.157	0.156	0.159
Average % of posted-price listing sold in 30 days	0.512	0.556	0.521	0.481	0.515

Notes: Table displays observed data used for estimation (as in Table 2) for data sampless from 2010–2014. Data was constructed using same sample restrictions as in main sample. Each year label corresponds to one year of data beginning in Oct. of that year and continuing through Sept. of the following year (i.e. the 2013 column corresponds to the main data sample used in the body of the paper). Displayed moments were used to obtain parameter estimates displayed in Table A5.

Table A5: Parameter Values and Welfare Estimates in 2010–2014

	2010	2011	2012	2013	2014
λ	11.196 (0.201)	12.877 (0.250)	13.807 (0.225)	13.101 (0.243)	13.817 (0.311)
α	18.765 (0.533)	17.902 (0.462)	16.116 (0.395)	12.760 (0.525)	10.625 (0.340)
τ	0.048 (0.001)	0.055 (0.001)	0.060 (0.001)	0.064 (0.002)	0.069 (0.002)
δ	31.033 (1.065)	29.862 (0.909)	26.879 (0.795)	21.131 (1.163)	16.357 (0.557)
T	8.125 (0.148)	8.394 (0.144)	9.154 (0.146)	10.301 (0.255)	13.233 (0.267)
ρ	0.069 (0.002)	0.068 (0.002)	0.058 (0.002)	0.055 (0.003)	0.041 (0.001)
ℓ	0.070 (0.000)	0.076 (0.000)	0.081 (0.000)	0.087 (0.000)	0.089 (0.000)
η	6.475 (0.023)	6.476 (0.026)	6.372 (0.023)	6.395 (0.028)	6.300 (0.033)
γ	0.351 (0.036)	0.463 (0.029)	0.889 (0.038)	0.985 (0.039)	0.855 (0.071)
c	0.818 (0.003)	0.836 (0.002)	0.855 (0.003)	0.840 (0.003)	0.841 (0.004)
	Welfare comparison (pure transfer case)				
Degenerate equil.	0.106 (0.001)	0.106 (0.001)	0.124 (0.001)	0.144 (0.001)	0.141 (0.001)
Dispered equil.	0.112 (0.001)	0.105 (0.001)	0.107 (0.001)	0.121 (0.001)	0.120 (0.001)
First best	0.182 (0.003)	0.164 (0.002)	0.145 (0.003)	0.160 (0.003)	0.159 (0.004)

Notes: Table displays estimates, using data from 2010–2014, of parameter values and welfare estimates for the degenerate (posted prices only) equilibrium, dispersed (auctions and posted prices) equilibrium, and first best settings when the holding fees are a pure transfer (profit) to the platform. Data was constructed using same sample restrictions as in main sample. Each year label corresponds to one year of data beginning in Oct. of that year and continuing through Sep. of the following year (i.e. the 2013 column corresponds to the main data sample used in the body of the paper). Units for parameters are as in Table 2 and units for welfare measures are fraction of retail (posted) price. Standard errors, from 200 bootstrap replications at the product level, are contained in parentheses.