Discounts and Deadlines in Consumer Search*

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

July 24, 2020

Abstract

We present a new equilibrium search model where consumers initially search among discount opportunities, but are willing to pay more as a deadline approaches, eventually turning to full-price sellers. The model predicts equilibrium price dispersion and rationalizes discount and full-price sellers coexisting without relying on ex-ante heterogeneity. We apply the model to online retail sales via auctions and posted prices, where failed attempts to purchase reveal consumers' reservation prices. We find robust evidence supporting the theory. We quantify dynamic search frictions arising from deadlines and show how, with deadline-constrained buyers, seemingly neutral platform fee increases can cause large market shifts.

JEL Classifications: C73, D44, D83

Keywords: Equilibrium search, deadlines, discount channels, mechanism choice, sequential auctions, endogenous valuations

*We thank Mark Armstrong, Lanier Benkard, Liran Einav, Chiara Farronato, Matt Gentzkow, Thomas Jeitschko, Jon Levin, Jason Lepore, Greg Lewis, Carol Lu, Daniel Quint, Yaron Raviv, Alan Sorensen, Steve Tadelis, Caio Waisman, Matthijs Wildenbeest, 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, the 2015 World Congress, the 2017 North American Winter Meetings of the Econometric Society, the Marketplace Innovation Conference, and the 2019 Consumer Search and Switching Costs Workshop for helpful comments. We thank Yuyang Jiang and Sharon Shiao for research assistance. Coey was an employee of eBay Research Labs while working on this project.

[†]Facebook, Core Data Science; coey@fb.com

[‡]Stanford University, Department of Economics and NBER, bjlarsen@stanford.edu

[§]Brigham Young University, Department of Economics, brennan_platt@byu.edu

Searching for the best price on a product requires time, and time can run out. Initially, a consumer may be willing to hunt for bargains on a given product, but if the search drags on with repeated failures, she may eventually turn to full-price retailers. Yet most of the search literature lacks this sense of urgency: the consumer will search indefinitely until finding a deal below her constant reservation price.

In this paper, we introduce *deadlines* into consumer search. A deadline can represent a specific date by which the consumer wants or needs the item she seeks, or simply a limit to the consumer's patience for continued search. Many consumer purchases involve a clear deadline, such as attire for an upcoming formal event, accessories for a planned vacation, clothing for an imminent change of season, supplies for an arriving newborn, or a gift for a birthday or anniversary. Even outside these event-driven purchases, search may be warranted to find a good deal on infrequently-purchased durable goods, and the consumer can easily grow frustrated as the search drags on unsuccessfully. Thus, the consumer behaves *as if* she has a deadline, even without a particular date in mind. In Section 2, we provide survey evidence that consumers explicitly recognize deadline pressures in their searching.

In our model, each buyer enters the market needing to acquire a good before her idiosyncratic deadline. The good is always available at a known posted-price outlet, but this is only used by buyers who have reached their deadline and literally need to buy it now. All other buyers are patient and search for opportunities to acquire the good through a discount channel. In equilibrium, a consumer will steadily increase her reservation price as her deadline approaches, eventually turning to the full-price outlet.

The model also allows sellers to offer their goods through either the discount channel (targeting patient buyers) or at full price (targeting desperate buyers). The latter yields a higher price but less frequent transactions; these effects exactly offset in equilibrium, making both channels equally profitable. Moreover, their goods are sold at a continuum of distinct prices. We document robust empirical evidence consistent with these predictions. We then discuss insights on consumer markets that would be missed by ignoring deadlines: in evaluating welfare consequences of discount markets and in analyzing market design decisions.

While the model can accommodate a variety of discount sales mechanisms, our main results depict the discount channel as an auction, both for expositional clarity and to match the data setting in which we test our theory: eBay auctions of popular, new-in-box items. These auctions potentially offer a low price but, from the buyer's perspective, have a low chance of success; meanwhile, the product is also available through posted-price listings. We focus on these eBay auctions to leverage a unique empirical advantage offered by auction data: buyers' bids indicate their willingness to pay *over a search spell*, even during *failed attempts* to acquire the good. We know of no other empirical work studying consumer search where changes in consumers' reservation prices are observable in the data.[1] We thus view eBay as an excellent

[1] In a similar vein, Genesove (1995) exploits failed offers in used-auto auctions to study a stationary partial equilibrium search model, but cannot observe repeat bids over time as we do. Panel data of unsuccessful

laboratory for studying time-sensitive search, which is likely to apply in other settings that are harder to measure.[2]

We demonstrate that the model's parameters are identified and can be estimated using functions of sample moments from eBay data—moments such as the number of bidders per auction or the number of auction attempts per bidder. 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.[3] Within this data, we focus our analysis on consumers who participate in multiple auctions, i.e. who *search* across auctions and reveal something about their reservation prices with each attempt. The data reveal a number of curious facts that find a unified explanation in our model, such as consumers increasing their reservation prices over time, equilibrium price dispersion, and coexistence of multiple sales channels. None of these patterns are exploited in fitting the model, and yet we find that the theoretical predictions for most of these facts are reasonably close to the magnitudes observed in the data.

We find that past losers tend to bid more in subsequent auctions — 1.2 percentage points more on average in the data, compared to 5.0 percentage points more in the fitted model. In the data, 70.6% of the auctions are won by the bidder who has been searching the longest (compared to 73.2% in the fitted model). In contrast, this would only occur 33.2% of the time if search length and bids were uncorrelated, as implicitly assumed in traditional search models. To our knowledge, this paper provides the first such evidence of time-sensitive search. Our model contributes to a small set of previous studies that also produce non-stationary search. We review this literature in Section 4.2.

The market response to buyer deadlines is also consistent with our model. First, deadlines can generate significant price dispersion even within homogeneous goods, as observed in the data. Despite sharing the same eventual utility from the good, buyers in the model differ in their time until deadline and hence submit distinct bids. We also see price dispersion between selling mechanisms, with auction prices averaging 15% lower than posted-price sales. Our model provides a plausible explanation for this price dispersion among identical goods, adding to a literature that, unlike our work, generally relies on ex-ante differences to generate dispersion (reviewed in Section 6.2).

Second, consumer search with deadlines rationalizes the coexistence of discount and non-discount mechanisms for identical products. Sellers use a mix of both auctions and posted prices, and the data and model have a similar distribution of this mix across products. The

transaction attempts may exist in markets for credit, housing, online bargaining, or online labor. For example, in studying search for auto loans, Argyle et al. (2017) observe a small fraction of consumers at a second financial institution after failing to secure a loan at the first.

[2]In Section D.1 of the Technical Appendix, we demonstrate that the theory extends to other discount sales mechanisms, including lotteries, bargaining, or discount posted prices.

[3]While eBay is popularly known as a avenue for buying and selling used goods, the platform sells over 80 million new-in-box items via auctions alone each year, totaling to 1.6 billion dollars annually in auction sales of new goods.

literature on competing mechanisms (reviewed in Section 6.3) only generates coexistence of multiple sales channels when there are exogenous differences among buyers or sellers, or under knife-edge conditions on parameters. Coexistence occurs for us with ex-ante homogeneity and under a robust set of parameters, as equilibrium selling speed compensates for difference in selling price.

Ignoring implications of consumers' idiosyncratic deadlines can skew the evaluation of market design and welfare. We demonstrate that, in the presence of deadlines, increases in platform fees can shift the market to more posted prices and fewer auctions, even if the fee increase is applied equally to both markets. This provides a possible micro-foundation for the decline in auction transactions relative to posted prices, particularly since 2009.[4] Indeed, the model predicts that sellers will completely abandon auctions if eBay commission fees rise an additional three percentage points.

We also show that the existence of the discount channel is always welfare improving when consumers have deadlines. The first-best solution would have sellers produce and sell the good to buyers at their deadline; but the presence of search frictions introduces an inefficiency, requiring sellers to produce well in advance. Auctions help conserve on sellers' waiting costs by closing more quickly, leading some sellers to choose auctions, and making the equilibrium auction usage constrained-efficient. Our model also allows us to quantify dynamic search frictions: the timing mismatch between the production and the consumption of the good produces roughly the same welfare cost as the static intermediation cost of bringing buyers and sellers together.

1 What Are Deadlines?

Before presenting the model, we first provide a discussion of what we mean by consumers searching with a *deadline*. A deadline can be a specific event for which a consumer needs a good, such as an air mattress or extra towels needed for arriving house guests, a birthday or anniversary gift, new running shoes for aspirational marathon training, a lantern for an upcoming campout, supplies for hosting a party, equipment for a soon-to-arrive baby, or supplies for a planned ski or beach trip. In some cases, a deadline could encompass a broader range than a specific date, as in the case of purchasing new clothes for an imminent change of season, larger clothes for a rapidly growing child, or a new baking dish for seasonal foods.[5]

The deadlines we have in mind are inherently *idiosyncratic*, not deadlines common to a large group of consumers, such as Christmas or Valentine's Day. In the presence of such common deadlines, both demand and supply change simultaneously as the deadline approaches,

[4] Einav et al. (2018) offered evidence of demand-side (and to a lesser extent, supply-side) shifts affecting the auction decline prior to 2009.

[5] All examples in this paragraph come from recent purchases in the household of one of the authors. The author reports that the marathon did not happen and the baby did arrive.

making it difficult to isolate the type of consumer behavior we model.[6] However, even these common deadlines may in practice generate the kind of idiosyncratic deadlines we model. For example, some consumers have idiosyncratic preferences for completing all Christmas shopping by early December, while others are willing to push the limits of Amazon Prime's on-time-delivery promise.

More generally, the deadlines we model represent a limit on how long a consumer is willing to spend procuring a good. For example, searching for a discount could become more difficult over time if consumers 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, consumers 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. In fact, while we model consumers as having a hard deadline, it can be shown that our model is isomorphic to a setting where consumers can search indefinitely but grow more impatient over time at an exogenously increasing discount rate. Despite the ubiquity of time-sensitive purchases in practice, deadlines have received sparse treatment in search literature, reviewed in Section 4.2.

To illustrate the prevalence of deadlines in search decisions, we conducted a survey of 1,210 random consumers from the Qualtrics consumer survey participation panel; survey details are provided in Technical Appendix A; see Coey et al. 2020b for the raw survey data. Each consumer identified a recent purchase for which she considered checking the price of multiple sellers. These responses remain stable across a wide price range and variety of product categories. First, we found that eBay plays a significant role in search: 28% of consumers reported checking the site as one of their options, compared to 25% searching Google Shopping and 68% searching Amazon.

Second, we asked consumers to indicate when they became aware that they wanted the item, and how long they would have been content without the item (had they not acquired it when they did). Only 2.5% of consumers reported unlimited patience; the remaining consumers had a potential search span averaging 70 days.

Third, many of the consumers reported motives that are consistent with the model. For instance, 32% of consumers needed the item for a specific event or gift, and 65% needed the item more over the course of their search. For 42% of consumers, the purchase was not urgent so long as it arrived in time for a particular deadline or use. Indeed, 65% of consumers indicated that they would have been willing to pay more later if they had been unable to purchase when they did, consistent with their reservation prices increasing over time.

Of course, many consumer purchases also fall outside the model's environment. For in-

[6] This simultaneous shift in supply and demand also occurs in perishable goods markets, as shown for NFL tickets in Waisman (2018) or in revenue management models (*e.g.* Deneckere and Peck, 2012; Board and Skrzypacz, 2016; Mierendorff, 2016; Dilme and Li, 2019), where a seller has multiple units of a good that expire at a known deadline.

stance, our model assumes that consumers know the product they want and are only searching on price, agreeing with half of survey respondents; the other half indicated that some portion of their search was to determine the right product as well. Also, 46% of consumers reported wanting the item as soon as possible, yielding a very brief duration of search.

Thus, roughly half to two-thirds of our survey respondents appear to be searching in a manner consistent with the deadline pressures we model. Of course, these answers rely on the respondents' imprecise memory and subjective evaluation of their own intentions, aggregated across widely varying items. In contrast, our eBay data in Section 4 records actual choices (bids) made with real consequences (potentially winning and having to pay) in seeking a homogeneous good. These observed actions in specific eBay markets are strongly consistent with the motives reported in the broad market surveyed here.

2 Buyers with Deadlines

We begin by modeling buyers' choices when faced with deadlines. Seller behavior is treated as exogenous until Section 3. In our continuous-time environment, buyers enter the market at a constant rate of δ, seeking one unit of the good that is needed within T units of time (*i.e.* the *deadline*).[7] The buyer receives βx dollars of utility at the time of purchase, while $(1-\beta)x$ dollars of utility are realized at the deadline, which is discounted at rate ρ. Thus, if the good is purchased with s units of time remaining until the buyer's deadline, her realized utility is:

$$u(s) \equiv \left(\beta + (1-\beta)e^{-\rho s}\right) x. \tag{1}$$

The extreme of $\beta = 0$ indicates that the good is literally of no use until the date of the deadline, so an early purchase provides no additional utility. When $\beta = 1$, the buyer immediately consumes the good when it is purchased, so further search comes at a cost of delayed consumption. 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 of mind from 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. For our empirical estimation, however, β is not separately identifiable from x because the two parameters always appear multiplied together in any of our equilibrium conditions. Indeed, on the margin, the equilibrium behavior reacts the same to greater overall utility x or more immediate consumption β.

The strategic questions for buyers are what price they are willing to bid in the auction channel and when to purchase from the posted-price channel. Let $V(s)$ denote the net present

[7]In our model, the event of a consumer entering the market is analogous to the consumer becoming *aware* that she needs/wants the good by some future date, and will now keep her eyes open for it as opportunities to search arrive. Such behavior is consistent with the findings of Blake et al. (2016), who document that consumers' web-browsing behavior consists of many searches over many non-consecutive days, well in advance of when the consumers actually purchase the item.

expected utility for a buyer with s units of time remaining until her deadline.

A buyer can obtain the good at any time through a non-discount option with posted price z, receiving utility $u(s) - z$. If the buyer instead waits until time $s = 0$ to purchase from the posted price option, she postpones (and discounts) the same expenditure z, obtaining a present discounted expected utility (*i.e.* in time s dollars) of $(x - z)e^{-\rho s}$. We assume throughout that $x \geq z$, so that buyers weakly benefit from purchasing via the posted-price option.[8] We also assume that $z > \beta x$, which ensures that $(x-z)e^{-\rho s} > u(s) - z$ for all $s > 0$. That is, buyers only purchase using the posted price when the deadline arrives, preferring to exhaust all discount opportunities first. Thus, the expected utility of a buyer who reaches her deadline is

$$V(0) = x - z. \tag{2}$$

Prior to her deadline, the buyer encounters a potential discount opportunity at rate α, and participates in it with exogenous probability τ, reflecting the possibility that a buyer can be distracted by other commitments or otherwise find it too much of a hassle to participate. Here, we treat each discount opportunity as a second-price sealed-bid auction that is executed immediately.[9] The highest bidder pays the second-highest bid and exits with the good, while losers continue their search. Let $W(b)$ denote the probability of winning with bid b, and $M(b)$ denote the expected payment under bid b.

A buyer's expected utility in state s can then be expressed in the following Hamilton-Jacobi-Bellman (HJB) equation:

$$\rho V(s) = \max_{b(s)} \tau\alpha \Big(W(b(s)) \left[u(s) - V(s)\right] - M(b(s)) \Big) - V'(s). \tag{3}$$

In this continuous-time formulation, the left-hand side of (3) represents the flow of expected utility that a buyer with s units of time remaining receives each instant while searching. The right hand side depicts potential changes in (net) utility times the rate at which those changes occur. When an auction occurs and the buyer participates in it (which occurs at rate of $\tau\alpha$), she will pay $M(b(s))$ on average. She also wins with probability $W(b(s))$, gaining consumption utility $u(s)$ but abandoning further search, which has expected utility $V(s)$. The derivative term accounts for the steady passage of time: remaining in the market for a unit of time reduces the buyer's state s by 1 unit, so her utility changes by $-V'(s)$.

The total stock of buyers in the market is denoted by H. The number of participants n in each auction is drawn from a Poisson distribution with mean λ. Indeed, $\lambda \equiv \tau H$ because over one unit of time, an average of $\tau\alpha H$ buyers participate in an auction while α auctions

[8] The consumption value x is treated as homogeneous across all buyers—in line with previous work that, like ours, focuses on commodity-like retail markets (*e.g.* Einav et al., 2018). Technical Appendix D.3 discusses ways homogeneity can be relaxed.

[9] Our empirical estimation approach accounts for the fact, pointed out by Song (2004), that in practice an eBay bidder may arrive at an auction after the current bid has passed her valuation, and hence she may not be observed bidding. Endogenous participation decisions are discussed in Technical Appendix D.4

occur, leading to τH participants per auction. If the cumulative distribution of bids is given by $G(b)$, then the probability of winning is:

$$W(b) \equiv \sum_{n=0}^{\infty} \frac{e^{-\lambda}\lambda^n}{n!} G(b)^n = e^{-\lambda(1-G(b))}. \tag{4}$$

To win the auction, the buyer must have the highest bid, which means all n other participants must have lower bids.[10] This occurs with probability $G(b)^n$. A participant's expected payment is the probability of winning times the expected second-highest bid when she wins:

$$M(b) \equiv e^{-\lambda}R + \sum_{n=1}^{\infty} \frac{e^{-\lambda}\lambda^n}{n!} \int_R^b nG(v)^{n-1} vG'(v)dv = e^{-\lambda}R + \int_R^b \lambda e^{-\lambda(1-G(v))} vG'(v)dv. \tag{5}$$

If there are no other participants (which occurs with probability $e^{-\lambda}$), the bidder pays the reserve price R. Inside the sum we find the probability of facing n opponents, while the integral computes the expected highest bid among those n opponents.

The first order condition of (3) yields $\lambda G'(b(s))e^{-\lambda(1-G(b(s)))}(u(s) - V(s) - b(s)) = 0$, so the optimal bid is thus:

$$b(s) = u(s) - V(s). \tag{6}$$

That is, a buyer in state s has a reservation price equal to the present value of the item minus the opportunity cost of skipping all future discount opportunities.[11]

Buyers randomly enter the market at differing times and thus will differ in their remaining time s, generating a distribution of valuations across buyers at any point in time. Let $F(s)$ denote the cumulative distribution of buyer types; therefore, $G(b(s)) = 1 - F(s)$. We assume that the reserve price is set to equal the lowest possible bid $b(T)$ (discussed in Technical Appendix D.2). Substituting for the optimal bid and the distribution of bidders into the HJB equation yields

$$\rho V(s) = -V'(s) + \tau\alpha \left(e^{-\lambda F(s)} [u(s) - V(s)] - b(T)e^{-\lambda} - \int_s^T b(t)\lambda e^{-\lambda F(t)} F'(t)dt \right). \tag{7}$$

[10] While the Poisson distribution literally governs the *total* number of participants per auction, it also describes the probability that n *other* bidders will participate. This convenient parallel between the aggregate distribution (in expected revenue and steady-state conditions) and the distribution faced by the individual (in her 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.

[11] One abstraction in our model is that bidders do not infer any information about their rivals from prior rounds. This approximates a large market, where the probability of repeat interactions are too low to justify tracking many opponents. In our data, a bidder has an 8.4% chance of encountering the same opponent in a subsequent auction.

2.1 Steady-State Conditions

In our model, the distribution $F(s)$ of buyer states is 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 she faces. We require that this distribution remain constant over time. As buyers exit the market, they are replaced by new consumers; as one group of buyers get closer to their deadlines, a proportional group follows behind.

Steady-state requirements are commonly used in equilibrium search models (*e.g.* Diamond, 1987; Albrecht et al., 2007) and more recently in dynamic auction models (*e.g.* Zeithammer, 2006; Said, 2011; Hendricks and Sorensen, 2018) for tractability, reducing the large state space that would be needed to track each entry or exit. This does not eliminate all uncertainty, such as the number or composition of bidders in each auction, but all shocks are transitory, as bidders in the next auction are independently drawn from a constant (though endogenous) distribution. Thus, steady-state conditions smooth out the short-run fluctuations around the average, and capture the long-run average behavior in a market.

Our environment ensures that the cumulative density function $F(s)$ is continuous on $[0, T]$. That is, there cannot be a positive mass (an *atom*) of buyers who share the same state s.[12] Conveniently, a continuous distribution also ensures that no two bids will tie with positive probability. Moreover, the probability density function, $F'(s)$, must also be continuous on $(0, T]$.[13] Indeed, the population of buyers changes according to

$$F''(s) = \tau\alpha F'(s)W(b(s)) = \tau\alpha F'(s)e^{-\tau HF(s)}. \tag{8}$$

That is, the relative density $F'(s)$ changes as buyers in state s participate in the discount sales channel (at a rate of $\tau\alpha$) and win (with probability $W(b(s))$), thereby exiting the market. Of course, a continuous distribution requires $F(0) = 0$ and $F(T) = 1$.

Finally, we ensure that the total population of buyers remains steady. Because 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.

[12] No stock of state 0 buyers can accumulate because all buyers who reach their deadline immediately purchase from a posted-price listing and exit the market. Similarly, no stock of state T buyers can accumulate because as soon as they enter the market, their clock begins steadily counting down. For interior states $s \in (0, T)$, exit can only occur by winning an auction; but the probability of participating in an auction at any given instant s is 0, thereby preventing a positive mass of buyers from exiting in the same state s.

[13] This is because buyers in state s become the buyers in state $s - \epsilon$ with the passage of time. Over ϵ units of time, they will participate in $\tau\alpha\epsilon$ auctions, but as $\epsilon \to 0$, the probability that a buyer of type s participates drops to zero, making it impossible to have a discontinuous drop in buyer density.

Because 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 equilibrium*[14] of this game as a bid function $b^* : [0,T] \to \mathbb{R}$, a distribution of buyers $F^* : [0,T] \to [0,1]$, and an average number of buyers $H^* \in \mathbb{R}^+$, such that

1. The distribution F^* satisfies the steady-state equation (8).
2. The average mass of buyers in the market H^* satisfies steady-state equation (9).
3. Bids b^* satisfy equations (2), (6), and (7), taking F^* and H^* as given.

The last condition requires buyers to bid optimally; the first two require that buyers correctly anticipate the distribution of competitors, consistent with steady state. Indeed, the equilibrium solution is separable: the first condition uniquely determines F^*, which then allows the second to determine H^*, which in turn combines with the third to determine b^*.

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 differential equations themselves have a closed-form analytic solution, but one of their boundary conditions does not; rather, the equilibrium H^* implicitly solves the boundary condition. If $\phi(H)$ is defined as

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

then the boundary condition is equivalent to $\phi(H^*) = 0$. This condition ensures that buyers newly entering the market exactly replace those who exit through winning an auction (the second term) or purchasing at the posted price (the third term). The solution H^* is unique because both terms are increasing in H — more buyers in the market will ensure that more auctions complete in a sale and that more buyers turn to posted prices. The rest of the equilibrium solution is expressed in terms of H^*.

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

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

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

[14] We will refer to the augmented equilibrium derived in Section 3, which takes into account sellers' decisions, as a *market equilibrium*.

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

$$b^*(s) = \beta x + (z - \beta x) \cdot \left(1 - \alpha \cdot \frac{\tau\kappa\,(1 - e^{\rho s})\,e^{-\rho T} + \rho\,(1 - e^{-\tau\kappa s})\,e^{\tau\kappa T}}{\tau\kappa\,(\delta e^{\tau H^*} + \alpha e^{-\rho T}) + \rho\,(\delta e^{\tau H^*} + \alpha e^{\tau\kappa T})}\right) \cdot e^{-\rho s}. \qquad (12)$$

The next result shows that this proposed solution is both necessary and sufficient to satisfy the equilibrium requirements.

Proposition 1. *Equations (11) and (12), together with $\phi(H^*) = 0$, 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 because $b''(s) > 0$. We state both results in the following proposition:

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

Bids increase as s falls for two reasons that can be seen in (12). The last exponential term $e^{-\rho s}$ simply reflects time discounting: buyers offer more because they are closer to enjoying the full utility of the good at the deadline.[15] Yet buyers shade their bid by the fractional term, which reflects the expected value of remaining search opportunities. As the deadline approaches, the buyer has fewer chances to win an auction, so the value of search falls and bidders shade less.

The convexity of bids is a consequence of backloaded search benefits. Buyers are unlikely to win early in their search due to low bids. Thus the benefit of search is nearly constant early on, in anticipation of future winning rather than from current winning. Later in their search, bidders have higher bids, and winning becomes more likely. But as the deadline approaches, fewer opportunities remain, so the value of search (and shading) quickly evaporates to zero.

Comparative statics for our model allow us to anticipate how the market would evolve if the underlying structure were to change. For example, if buyers were to become less patient or more auctions were to be offered, bidders' bidding profiles over search duration would become steeper. This would also occur if bidders were given more time to search (increasing T); this result is less obvious because if T were to increase there would be more chances to participate in auctions and also more participants per auction, but the former would always dominate. Although our equilibrium has no closed-form solution, these comparative statics are obtained by implicit differentiation, as derived and discussed in Technical Appendix B.

[15]If buyers are extremely patient ($\rho \to 0$), the bidding function approaches $b(s) = z$ regardless of time until deadline—even the fractional term of (12) approaches one. Impatience causes buyers to prefer postponing payment until closer to the time of consumption, and thereby creates some variation in willingness to pay. If impatience is eliminated, the variation disappears; everyone is willing to bid full price, so search does not offer a discount at all, in the spirit of the Diamond (1971) paradox.

3 Selling to Buyers with Deadlines

The most direct effect of buyer deadlines is seen as reservation prices increase with search duration. However, this behavior indirectly influences sellers as well; in this section, we develop the seller's side of the model and derive its implications for the market as a whole. This allows us to explain the coexistence of auction and posted-price sellers, predict the market reaction to increases in seller fees, and evaluate the welfare consequences of offering a discount selling mechanism.

We consider a continuum of sellers producing an identical good, and allow free entry to offer their product via either mechanism.[16] Each seller 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, effectively endogenizing α in the buyers' model. Since we consider goods that are readily available at numerous retailers, we take the posted price z to be exogenous, pinned down by the retail price at outlets outside of our model.

3.1 Seller Profit

Each seller can produce one unit of the good at a marginal cost of $c < z$, incurred when the seller enters the market. The seller also pays a listing fee, ℓ, which is a commission on the transaction price.[17] Each seller makes two choices: whether to enter, and upon entry, whether to join the discount or posted-price market. This choice resembles the tradeoff in directed search models. Buyers are more plentiful than sellers in the auction market, so sellers transact quickly, while buyers must wait to find a successful match. The reverse holds true in the posted-price market; however, the higher transaction price there compensates for the slower selling speed.

Examining first the posted-price market, let ζ denote the Poisson rate at which a posted-price seller encounters 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. Once a seller enters the market, the production cost c is sunk; thus the expected profit Π_p moving forward is

$$\rho\Pi_p = \zeta\left((1-\ell)z - \Pi_p\right). \tag{13}$$

When sellers encounter a buyer (at rate ζ), the purchase always occurs, with a net gain of

[16] 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 cost scales proportionately. Also, when sellers employ mixed strategies in our model, this can be interpreted literally as each seller randomizing which mechanism to use, or as sellers being divided into two groups playing distinct pure strategies in the proportion dictated by equilibrium.

[17] In addition to fees paid to the platform itself, this could also include costs of product storage, shipping, marketing, or customer service. This fee can also be modeled as a flat rate or as a flow of cost for the duration of the listing. See Coey et al. (2019) for a modification of this model with a flow cost. Also, while the marginal price z exceeds the marginal cost c, this markup does not contribute to any monopoly-like inefficiency because all buyers value the good at $x \geq z$.

$(1-\ell)z - \Pi_p$ after paying the commission.

Turning to the discount market, let η represent the Poisson rate of closing. This is taken as exogenous, representing a short time ($1/\eta$ on average) required for buyers to become aware of the listing. From the seller's perspective, a Poisson number of buyers (with mean λ) will participate in her auction at its conclusion, producing expected revenue θ:

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

Inside the parentheses, the first term applies when only one bidder participates (which occurs with probability $\lambda e^{-\lambda}$) and therefore wins at the opening price of $b(T)$.[18] 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. With probability $e^{-\lambda}$, no bidders participate and the item is relisted; dividing by $1 - e^{-\lambda}$ makes θ the expected revenue conditional on sale. From the perspective of a seller who has just listed an auction, the expected profit Π_a moving forward is

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

The listing closes at Poisson rate η, but if no bidders arrive then the seller re-lists the item (without incurring production cost c again) and waits for the new auction to close.[19] If at least one bidder participates, the seller's net gain is $(1-\ell)\theta - \Pi_a$.

From the perspective of a potential entrant, expected profits are net of the initial production cost: $\Pi_p - c$ or $\Pi_a - c$, respectively. We assume free entry into both markets, ensuring potential entrants will expect zero profits.

3.2 Steady-State Conditions

At any moment, both markets will have a stock of active listings waiting to close, denoted A for the measure of auction sellers and P for posted-price sellers. As with the population of buyers, the stock and flow of sellers are required to remain stable over time.

Each auction closes at rate η, and with A sellers in the market, this implies ηA auctions will close over a unit of time. From the buyer's perspective, α auctions close over a unit of time. These must equate in equilibrium:

$$\eta A = \alpha. \quad (16)$$

[18] If the reserve price were endogenous (as explored in Technical Appendix D.2), the seller would accept any bid greater than cost c, as in Levin and Smith (1994).

[19] It is possible to omit the waiting time for auction listings, in which case auctions occur instantaneously and $\Pi_a = (1-\ell)\theta$. Including η keeps the two seller problems parallel, and accommodates the empirical regularity that some listings receive no bidders and do not sell.

Each posted-price listing closes at rate ζ, so collectively the P sellers transact ζP units in one unit of time. Meanwhile, $HF'(0)$ buyers reach their deadline and make posted-price purchases over one unit of time. These must equate in equilibrium:

$$\zeta P \quad = \quad HF'(0). \tag{17}$$

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. Let σ be the fraction of entering sellers who join the auction market. Then $\sigma\delta$ new sellers list an auction over a unit of time. This must equal the flow of auction sellers who find at least one bidder over the same unit of time:

$$\sigma\delta = \alpha\left(1 - e^{-\lambda}\right). \tag{18}$$

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 who hit their deadlines:

$$(1-\sigma)\delta = HF'(0). \tag{19}$$

3.3 Market Equilibrium

With the addition of the seller's problem, we augment the equilibrium definition with three conditions. A *market 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. α^*, ζ^*, σ^*, A^*, and P^* satisfy the Steady-State equations (16) through (19).

3. Prospective entrants earn zero expected profits: $\Pi_p^* = c$, given ζ^*, and either $\Pi_a^* = c$ if $\alpha^* > 0$, or $\Pi_a^* \leq c$ if $\alpha^* = 0$.

The first requirement imposes that buyers behave optimally as developed in Section 2. The second imposes the steady-state conditions. The third ensures that sellers enter the market optimally, since free entry drives expected profits to zero in both markets. If the posted-price market were to offer positive profits, more sellers would enter, driving up the required time ζ for each seller to find a buyer, thus reducing expected profits. If auctions were to offer positive profits, more sellers would enter, increasing the frequency of auctions α. Buyers would then be more likely to win at a discount, so they would bid less and auction revenues would fall. Sellers are thus indifferent about which market they enter, allowing them to randomize according to the mixed strategy σ.

Note that the posted-price market will always operate in equilibrium. The uncertainty of winning in the discount market guarantees that a fraction of buyers will inevitably reach their deadlines. Thus, a sufficiently small stock of posted-price sellers can always break even. In contrast, a seller in the auction market might face low bids or a low number of bidders, insufficient to cover costs. If so, the auction market can shut down ($\alpha^* = \sigma^* = 0$), pushing all transactions to the posted-price market — referred to as a *degenerate equilibrium* in the language of equilibrium search theory.

In more moderate parameterizations, both markets will operate, which can an be called a *dispersed equilibrium*, because the homogeneous good is sold at a variety of prices and by multiple sales mechanisms. Note that buyers are always willing to purchase early if offered enough of a discount; the dispersed equilibrium occurs only when sellers can still cover their costs while providing that discount. This can be expressed by simplifying the third equilibrium requirement to:

$$(1-\ell)\theta^* = c\left(1 + \frac{\rho}{\eta\left(1 - e^{-\tau H^*}\right)}\right). \tag{20}$$

Intuitively, free entry requires that the expected post-commission revenue exactly equals the expected production cost, where the fractional term on the right is the expected interest incurred between production and transaction. We explore the efficiency of this dispersed equilibrium in Section 6.5.

While the market equilibrium conditions simplify considerably, they do not admit an analytic solution and we must numerically solve for both α^* and H^*. All other equilibrium objects can be expressed in terms of these. Proposition 3 in Technical Appendix C reports the simplification, while Proposition 4 proves that an equilibrium always exists and reports a precise condition for when that equilibrium will be degenerate (with $\alpha^* = 0$).

4 Empirical Evidence: eBay Auctions and Posted Prices

4.1 Data and Descriptive Statistics

The concept of consumer deadlines and increasing impatience during a consumer's search is likely to play out in a number of real-world settings. Among these, the eBay marketplace offers several advantages. Auctions (serving as the discount mechanism in our model) offer consumers repeated chances of obtaining the good, while posted-price sales (serving as the non-discount mechanism) offer consumers an identical good immediately at a higher price. For each auction we observe failed attempts at acquiring the good, including consumers' reservation price at each attempt.[20] By considering new-in-box products within a single platform, we

[20] This feature of our data provides a unique advantage even over the detailed clickstream data studied in De los Santos et al. (2012) or Blake et al. (2016), for example, where the authors observe a history of a buyer's

ensure product consistency across listings and across mechanisms.

Table 1: Descriptive Statistics

A. Transaction level	Posted Price		Auctions	
	Mean	Coefficient of Variation	Mean	Coefficient of Variation
Number of bidders	1	–	2.56	0.45
Revenue	106.82	0.32	97.27	0.30
Normalized revenue	1	–	0.85	0.30
Number of transactions	494,448		560,858	

B. Product level	Posted Price		Auctions	
	Mean	St. dev.	Mean	St. dev.
% of listings that transact	49.70	15.98	85.83	12.73
Transactions per product	134.98	220.82	153.11	343.63
Unique sellers per product	82.70	137.84	68.53	201.30
Unique buyers per product	129.03	208.02	334.79	921.59
Number of products	3,663			

Notes: Table displays descriptive statistics for our 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 the product-level coefficient of variation. Normalized revenue is computed by first dividing auction price by product-level average of posted-price sales. Panel B reports average and standard deviation of product-level averages and product-level counts.

Our data consist of auctions and posted-price sales on eBay.com for the year from October 1st, 2013, to September 30th, 2014.[21] As our model describes the sale of homogeneous goods, we restrict attention to brand new items that have been matched by the seller to a *product* in one of several commercially available catalogs. 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 refer to an individual attempt to sell the product as a *listing*. 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

web-browsing activity, but do not observe the buyer's reservation price at points during the search process.

[21]Instructions for accessing this proprietary dataset, Coey et al. (2020a), can be found in the supplemental material (Coey et al. 2020b).

popular items, principally electronics, media, or health/beauty products.

For each auction, our data report every bid and its timing. Each bidder has a unique identifier, allowing us to link a bidder across each auction attempt and observe any eBay posted-price purchase of the same product. We retain each bidder's highest bid in each auction. Each listing's shipping speed is also recorded, and the associated shipping fees are added to the posted price or bids throughout our analysis.

For each auction, we follow Bodoh-Creed et al. (2018) by keeping only *serious bids*, which include bids that are placed in the final hour of the auction, as well as the two highest bids prior to the last hour.[22] We impose this restriction because we find that, in our data, as in most eBay auctions, bids placed prior to the last hour of the auction tend to be low (averaging 49% of the good's winning price) with little or no chance of winning, and seem more like cheap talk than an informative signal of the bidder's true reservation price. While our model explains a number of empirical facts well, it is not well suited for explaining these low-ball bids. This sample restriction drops 52% of bids in our sample. Importantly, this restriction does not drop any *auctions* from our sample.[23]

Table 1 presents descriptive statistics for the listings that end in a sale. In all, there are 1,055,306 sales of 3,663 distinct products, split roughly evenly between auctions and posted prices. Panel A aggregates across transactions, then products; that is, we compute the mean and coefficient of variation across transactions of a given product, then average these results across all products. In our model, we treat the posted price for a given product as fixed (z in the model), whereas in reality posted prices can vary from listing to listing just as in auctions. However, the average selling price is higher with posted prices than auctions ($107 versus $97) with similar coefficients of variation. To adjust for differences across products, we follow Einav et al. (2018) and rescale all bids, dividing by the mean price of posted-price sales of that product. This rescaling is also consistent with our model, 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.

4.2 Bids Over Duration of Search

Our data allows us to follow each bidder across multiple auctions of the same product. We order these auctions 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

[22] The two highest bids from prior to the last hour are important only because, given the second-price nature of eBay auctions, these bids determine the standing bid at the beginning of the last hour.

[23] Technical Appendix G presents data and model estimates without imposing this sample restriction. We show that forcing the model to explain these low-ball bids results in an implausibly high discount rate.

sample. This yields 1,497,371 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. Our analysis in this section studies the behavior of bidders observed participating in at least two auctions for a given product (84% of bidders are only observed bidding in one auction).

We find that the average willingness to pay tends to increase from one auction to the next. Figure 1 displays the resulting trend across repeated auction attempts. In Panel (A), 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—whether the bidders participated in only two auctions, three auctions, or as many as six—the average bid steadily rises over time from the first to the last auction in the sequence.[24] Notably, the line for each sequence is successively higher as the sequence becomes shorter; they never cross. This is consistent with the model because, in the model, consumers who are observed in fewer total auctions are closer to their deadlines at the time when they are first observed participating; thus, their reservation price starts higher and rises more steeply. We find it striking that this feature is observed in the data as well. Panel (B) frames the trends from Panel (A) in terms of a regression result. Averaging across all sequence lengths and auction numbers, the bid increases by a statistically significant amount of 1.2 percentage points in each successive auction, or 2.9 percentage points from the first to last bid attempt.[25]

It is worth noting that canonical search models (*e.g.* Stigler, 1961; Diamond, 1987; Stahl, 1989) do not explain this empirical fact, and yield instead 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 constant search costs, and have at least one firm left to search.[26] 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.

While our primary interest is in the implications of deadlines for search more broadly, by applying our search model to auctions, our work also connects to the nascent literature on

[24] Note that in these figures the final bid in a sequence may be a winning bid, while by construction previous bids are not. This is not what drives the increase in bids across auctions, because bids tend to increase even before the final bid in a sequence. 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.

[25] As highlighted in Section 1, the deadlines we model are inherently idiosyncratic and unobservable, not common deadlines such as holidays. In our eBay data we found no clear evidence that consumers treat popular holidays as a common deadline; aggregate price trends exhibit no clear pattern leading up to such holidays. The average of *individual-specific* price trends, however, are consistent with the idiosyncratic deadlines we model.

[26] In the context of job search, Van Den Berg (1990) establishes that reservation wages will fall over the search spell in anticipation of worsening search costs, match rates, or wage distributions in the future. Aside from deadlines, 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), or the possibility that past quotes will not be honored (Akın and Platt, 2014).

Figure 1: Bids Over Search Duration — Data

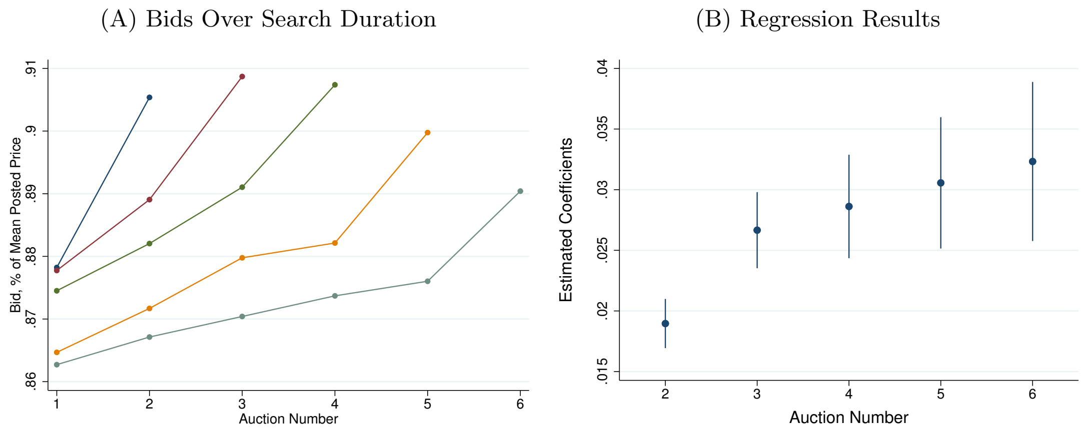

Notes: In Panel (A), a given line with m points corresponds to bidders who bid in m auctions total for a given product without winning in the first $m-1$ auctions. The horizontal axis represents an auction number within the sequence (from 1 to m) and the vertical axis represents the average normalized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

repeated sequential auctions (Zeithammer, 2006; Said, 2011; Hendricks et al., 2012; Backus and Lewis, 2016; Bodoh-Creed et al., 2018; Hendricks and Sorensen, 2018), in which bidders shade their bids below their valuations due to the continuation value of future search. Among this literature, our model is unique in its prediction that a bidder will increase her bid in subsequent attempts to acquire the item.[27]

We are also able to examine this bid increase for individual products (presented in Technical Appendix H), rather than averaging across all products. Within broad product categories, the average bid increase is quite similar to the results presented here, suggesting that deadline-like behavior may occur across a wide range of consumer goods.[28] Of course, many consumers may have idiosyncratic reasons for purchasing that are not driven by a specific timeline or growing impatience. Importantly for our empirical analysis, however, the presence of non-deadline-like behavior will likely work against us finding patterns consistent with the model.

4.3 Winners and Losers

Here we document additional patterns in the data concerning *who wins* and *what losers do*. First, the bidder in an auction with the longest observed time on the market (*i.e.* the time

[27]We offer a comparison to other auction models in Technical Appendix F.

[28]Our survey data discussed in Section 1 also includes a self-reported measure of consumers' individual deadlines. We analyze this measure in Technical Appendix A and find evidence of deadlines for a wide variety of products, with increasing reservation prices as the deadline approaches.

since the bidder's first observed bid) is frequently the winner, occurring in 70.6% of auctions.[29] This is consistent with the model.[30] In contrast, elapsed time and likelihood of winning would be *inversely* correlated if valuations were constant over time because high-valuation buyers would win shortly after entering the market while low-valuation bidders would require many repeated attempts to get lucky.

Second, when auction losers abandon auctions, we observe 11.3% of them later purchasing the same product via an eBay posted-price listing. Our model predicts that all losing auction participants will eventually turn to posted prices; and it is possible that many still do but through a different platform such as Amazon, or at a brick-and-mortar location. Alternatively, it is possible that some buyers may only intend to use auctions while others only use posted prices, which is a decision that our model does not seek to explain.

Regardless, among auction losers who do eventually purchase from an eBay posted-price listing, nearly all do so within a very short time of their last observed auction attempt: 57% do so within one day of their last losing attempt, 73% within 5 days, 80% within 10 days, etc. Indeed, the cumulative probability of switching to a posted-price listing is concave in the time elapsed since the last losing attempt, as shown in Figure 2.A. This is as the model predicts: the buyer's last observed auction attempt should be close to the buyer's deadline, and thus a posted-price purchase is most likely to occur close to the last auction attempt.

Third, the data report the time elapsed between a buyer's bids. We find that the average duration between bids is shorter for bidders with more bidding attempts. This decrease is consistent with the model: all buyers have the same available time T, so those fortunate enough to participate in more auctions must have placed bids closer together. We return to this point in Section 5.4

4.4 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.[31] Consider the case where bidders are uncertain about the degree of competition they face, and form different estimates of its intensity. A bidder who underestimates the number of competitors or the bids of

[29] This percentage is computed conditional on auctions in which the winner has bid at least once before. Without this conditioning, the number is similar (70.5%).

[30] Note that the model does not predict that will occur 100% of the time. The buyer only enters our data when her first bid is placed, even though she may have started her search earlier. Thus, in both the data and the model, the bidder with the longest *observed* time since first bid should frequently (but not always) win.

[31] 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 her information about its value. Learning stories can also generate *decreasing* reservation prices, as in De los Santos et al. (2017).

Figure 2: Switching Rates, Expensive Products, and Experienced Bidders

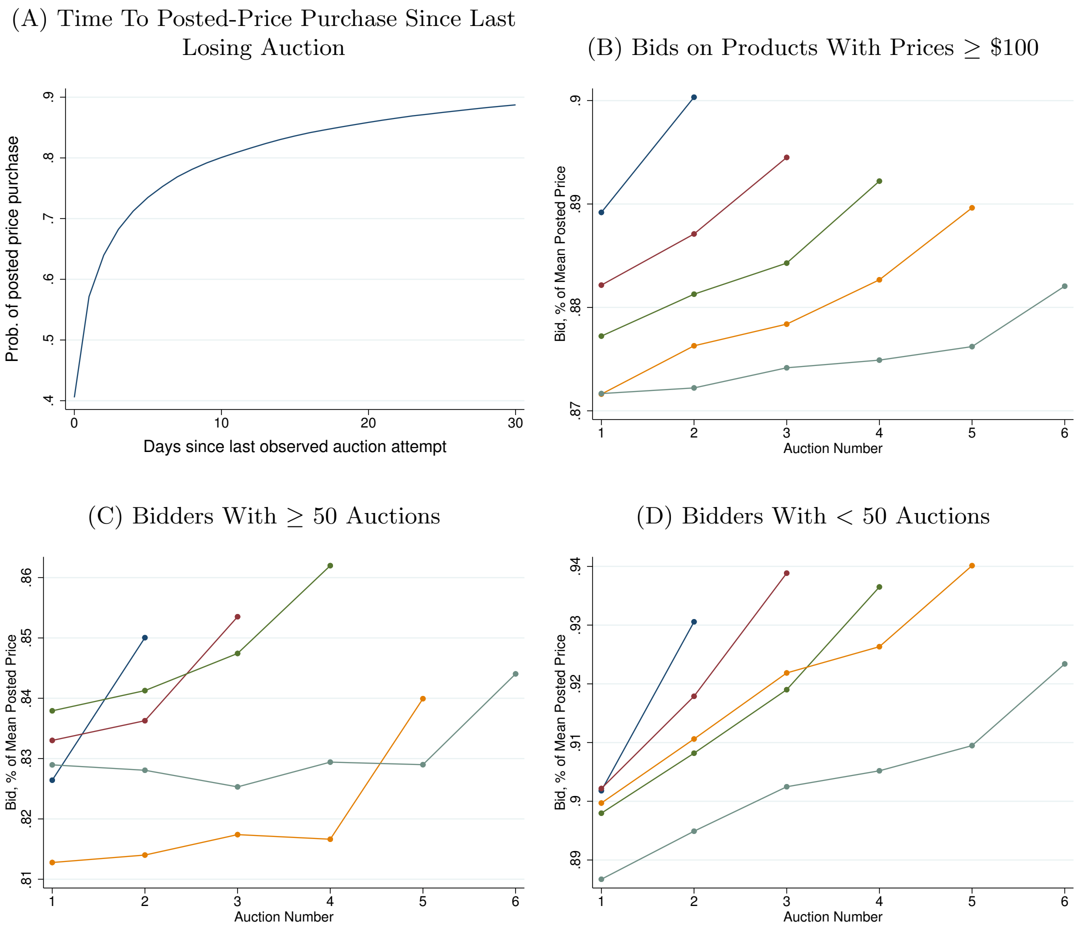

Notes: Panel (A) displays cumulative density of the time difference between the last observed auction attempt and the posted-price purchase conditioning on bidders who attempted an auction and did not win and were later observed purchasing the good on an eBay posted-price listing. The remaining panels report average bids over duration of search as in Panel (A) of Figure 1. Panel (B) limits to products with average transaction price $\geq$ \$100; Panel (C) limits to bidders who have bid in at least 50 auctions; Panel (D) limits to bidders who bid in less than 50 auctions.

competitors will overestimate her likelihood of winning in future auctions; this raises her continuation value and causes her to shade her bid lower. Such a bidder will gradually revise her estimates upwards as she 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. Lauermann et al. (2017) provide a theory of this form, though in the context of first-price auctions and without empirical testing; they refer to this pattern of losers raising their bids due to learning as the *loser's curse*. While such learning likely occurs in practice (and our model abstracts away from it), there are several reasons why bidder learning is unlikely to be the sole driver of increasing bids in our setting.

First, users can easily learn prices and bid histories for current and past listings by selecting the "Sold Listings" checkbox on the eBay search results page; this is far quicker and more informative than auction participation.

Second, learning by participation is more costly with expensive products, due to the danger of bidding too high and winning when initially uninformed. We therefore expect buyers of expensive products to be more inclined to gather information before bidding (such as through searching eBay sold listings) rather than using repeated bid attempts to discover the market price. Figure 2.B shows the same increasing bid pattern for products with a median price over \$100, although the bid increase across auctions is smaller than in the Figure 1.A (0.5 percentage points on average from auction to auction, and 1.3 percentage points from the first to last bid).

Third, 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 greatly. We define experience as having participated in at least 50 auctions prior to the current auction (even outside our sample).[32] Experienced bidders place 28.6% of bids in our sample, and we observe a similar (if noisier) pattern of increasing bids among experienced bidders, shown in panel (C) of Figure 2, and inexperienced bidders, shown in panel (D). Experienced bidders bid on average 8 percentage points lower than inexperienced bidders, but the bids increase over nearly the same range of 5 percentage points. Also, 12.0% of experienced bidders turn to posted prices after abandoning auctions, compared to 11.1% of the inexperienced.

Fourth, other facts have no clear connection to a bidder learning story. For instance, in Technical Appendix E, we show that bidders with more attempts gravitate toward auctions that offer fast shipping and that close soon. While the model does not directly speak to choosing among available auctions, this ancillary evidence is consistent with increasing impatience.

[32]For each bidder in each auction, our data reports the number of previous auctions the bidder has participated in, including listings outside of our sample period and for products outside of our sample.

We wish to emphasize here that we do not attempt to (nor do we believe it would be feasible to) rule out the possibility of bidder learning; indeed, some degree of bidder learning seems plausible and intuitive.[33] The appeal of our model of time-sensitive buyers is that it provides a single, unified explanation of a number of facts together, even while alternative explanations may generate some of the patterns we observe. For example, one alternative explanation for the increase in bids at the end of bidding sequences in Figure 1.A is that, from one auction to the next, a bidder's valuation is determined as an *iid* random draw, 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 valuations would fail to explain the pattern of increasing bids *prior* to the final auction in the sequence. While learning and other alternative explanations likely play some role in this market, the bulk of the evidence also seems to indicate some role for time sensitivity.

5 Taking the Model to the Data

Here we describe the process of fitting our model to the data and compare the resulting estimates to facts from the preceding section. The approach is generally straightforward, as each parameter corresponds directly to a transformation of sample moments. In computing these moments, we normalize the unit of time to be one month and normalize $z = 1$. This latter normalization is done by dividing bids for a given product by the average posted price for that product. This rescaling is equivalent to "homogenizing" bids (Balat et al., 2016). This has no effect on the distribution $F(s)$ and or the bids $b(s)$, as bids scale proportionally with z. We set $\beta = 0$ throughout this estimation because it is not separately identifiable from x — the two parameters always appear multiplied together in any of our equilibrium conditions. Setting $\beta > 0$ would decrease shading in our model, and would only affect our estimation by increasing ρ and decreasing c.

The first column of Table 2 displays each of the sample moments we exploit. We deliberately avoid selecting moments described in Section 4, allowing us to later compare these observed behaviors to the predictions of the fitted model. The target data moments are computed as averages across product-level averages; thus, this exercise should be interpreted as fitting the model for the average product.[34] The second column of Table 2 displays the corresponding theoretical equivalent for each moment, and the third column the resulting parameter estimates. We discuss each sample moment and estimation step in turn in the following subsection.

[33] 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 she increases her bid in subsequent auctions. Yet our deadline model also generates this positive correlation, because low bidders (early in their search) will typically lose by the largest margins, and these buyers are rarely observed until much later in the search process when their bid has increased substantially.

[34] Technical Appendix H applies this same procedure separately product-by-product.

5.1 Data Moments and Model Equivalents

We now discuss the moments used to identify the parameters. The first parameter in Table 2 is the number of participants per auction (λ), which plays an outsized role in the model because it determines the degree of competition among bidders and thus influences the value of searching in auctions. Our identification of λ explicitly addresses the issue raised by Song (2004): a participant may arrive at the auction after the standing bid has passed her reservation price and thus will not be observed placing a bid. Platt (2017) shows that the Poisson mean λ of the number of bidders arriving at the auction is identified by $\lambda P(\lambda)$, which is the average number of *observed* bidders, where $P(\lambda)$ is the probability that an arriving participant can successfully place a bid.[35] Thus, each theoretic moment in Table 2 involving the number of participants is multiplied by $P(\lambda)$ to state them in terms of observed bidders, as in the data.

Table 2 shows that the second parameter, α, is identified by the frequency with which auctions occur. Auction frequency is computed conditional on sale in the data; thus, in the theoretical equivalent, α is multiplied by $1 - e^{-\lambda}$, the probability that the listing receives at least one bid.[36] Note that we also condition on completed auctions in the first moment and the last three.

The third parameter in Table 2 is τ, which is identified by the frequency with which a bidder participates in a specific item's auctions. In our theory, the average bidder will participate in $\tau\alpha$ auctions per month, but her bids will only be observed in fraction $P(\lambda)$ of them. Furthermore, there will be some data months in which she cannot participate because she has not entered or has concluded her search. Thus, we measure participation per month, conditional on being observed in at least one auction that month (in the theoretical equivalent, the denominator accounts for this same conditioning). Importantly, this moment is the only one estimated at the *bidder* level, and thus the fit of the bidder-level patterns we document below (such as the average bid increase over time) is by no means baked into the model estimation exercise.

The fourth and fifth parameters in Table 2 are δ and T, which are identified jointly from a data moment and a steady state condition. The data moment, shown in the fourth row of Table 2, is the flow of new buyers entering the market for a given item. We compute this

[35] This probability is given by $P(\lambda) \equiv \frac{1}{\lambda}\left(2\left(\ln(\lambda) - \Gamma'(1) - \Gamma(0,\lambda)\right) - 1 + e^{-\lambda}\right)$, 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 upper incomplete gamma function. Platt (2017) assumes that participants in an auction arrive in random order, which is compatible with our model because the payoff of losing participants is the same whether or not they were able to actually place a losing bid. The assumption implies that a buyer with low willingness to pay will only be observed if the buyer happens to have arrived early compared to other bidders in that auction, as in Hendricks and Sorensen (2018).

[36] In the model, auctions only fail to transact when no one arrives, which happens with 5% probability at our estimated parameters. In the data, a much larger fraction of auctions (14%) fail to transact—despite having an average highest bid of 93.7% of the posted price. Some of these fail because eBay allows sellers to impose a secret reserve price, and it can happen that no bidder bids above this value. In that sense, an auction with a sufficiently high secret reserve price functions effectively like a posted-price listing, won only on rare occasion by high-value bidders. Others fail because the transaction is canceled by the buyer or seller. We effectively exclude all failed auctions from our analysis.

Table 2: Key Data Moments and Matching Parameter Values

	Observed in Data	Theoretical Equivalent	Fitted Parameter
Bidders per completed auction	2.57	$\frac{\lambda \cdot P(\lambda)}{1-e^{-\lambda}}$	$\lambda = 3.01$ (0.020)
Completed auctions per month	12.76	$\alpha\left(1-e^{-\lambda}\right)$	$\alpha = 13.42$ (0.548)
Auctions a bidder is observed in per month	1.11	$\frac{\tau \alpha P(\lambda)}{1-e^{-\tau \alpha P(\lambda)}}$	$\tau = 0.019$ (0.00066)
New bidders per month who never win	16.33	$(\delta - \alpha)(1 - e^{-\tau \alpha T P(\lambda)})$	$\delta = 41.46$ (2.56)
—	—	Eq. (10)	$T = 4.25$ (0.142)
Average revenue per completed auction	0.853	θ	$\rho = 0.056$ (0.0024)
Average listing fee paid	0.116	ℓ	$\ell = 0.116$ (0.0029)
Average duration of an auction listing (months)	0.156	$1/\eta$	$\eta = 6.39$ (0.028)
—	—	Eq. (20)	$c = 0.748$ (0.0036)

Notes: Table displays the model parameter estimates in the last column, obtained by setting the theoretical equivalent (second column) equal to the observed value in the data (the first column) and solving for a given parameter. 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.

monthly flow conditional on bidders who never win to ensure a complete search spell for these bidders. The theoretical object we match uses this same conditioning, given by $\delta - \alpha$ (the flow per month of new arrivals minus new winners) multiplied by the probability of being observed in at least one auction over the full search span $(1 - e^{\tau \alpha T P(\lambda)})$.

Note that the search span T cannot be directly observed in the data. Instead, we identify T using the buyer steady state condition (10). In the model, this condition determines the endogenous number of buyers in the market, H^*, which in turn determines participants per auction, λ, so as to ensure that buyers win and exit the market at precisely the same rate that new entrants follow behind them. Our procedure has already identified λ and τ, forcing $H^* = \frac{\lambda}{\tau}$, so we solve for the search span T that is consistent with that population size. Too small of a value for T, for instance, would leave too few buyers in the market to sustain the estimated level λ of per-auction competition.

The sixth parameter in Table 2 is ρ, which is identified from the average second-highest bid per completed auction (*i.e.* average auction revenue θ). Note that Eq. (14), which defines θ, depends only on parameters of the buyer's problem. A larger discount rate ρ creates a steeper bid function for each bidder, and thus leads to a lower second-highest bid. This allows us to identify ρ from the average auction revenue, and thus preserve the average *increase between bids* as a check for the goodness of fit.

The seller's problem requires three additional parameters: listing fees, auction duration, and production costs, which are the last three parameters shown in Table 2. The first two are directly observed in the data, providing immediate estimates for ℓ and η. The seller's production cost parameter c is not observable, however. Instead, we identify c as the cost that makes sellers indifferent between selling via auctions or posted prices, which is the market equilibrium requirement (20), as shown in the last row of Table 2.

5.2 Discussion of Parameter Estimates

Many of the parameter estimates in the final column of Table 2 seem reasonable (such as α, λ, δ, η, ℓ), in part because they simply reflect features that are directly observable in the eBay marketplace. We obtain an estimate of $P(\lambda) = 0.81$, so the number of underlying bidders arriving at the average auction is 23% larger than the number of observed bidders. Our estimate of $\tau = 0.019$ may seem somewhat low, but this is simply capturing the data fact that bidders are participating in relatively few of the available auctions. The stock of buyers in the market is $H = 160$ under these parameters, while only $\lambda = 3.01$ serious bidders participate per auction. Other parameters, including c, T, and ρ, merit further discussion as they speak to supply and demand features that relate to other markets.

The estimated production cost c ensures that both mechanisms earn the same expected profit. Posted-price sellers make a markup of $(1 - \ell)z - c = 13.7\%$, which seems reasonable for retail sales. Auction sellers make an average $(1 - \ell)\theta - c = 0.7\%$ markup, but with the resulting $\zeta = 0.31$ (the rate at which a posted-price listing finds a buyer) and $\eta = 6.39$ (the rate at which auctions close), auction transactions occur 21 times faster. Thus, auction sellers are compensated for the lower markup by selling more quickly.

The estimate of T implies that, for the average product in our sample, consumers become aware of their need for the product 4.25 months in advance of their private deadline. However, the model predicts a much shorter observed search. Early in their search, buyer bids are more often lower than the standing price and thus cannot be placed. As a result, the model predicts an average search span of 1.6 months from the first to last observed auction. Furthermore, the model predicts that 77% of auction winners win in the last 2 months of their search span. These predicted durations are in the ballpark of the average estimated length of time on the market from our survey (2.3 months), but longer than the observed average time in the market in our eBay data (0.4 months).

Our estimate of ρ, a monthly discount rate of 0.056, is clearly higher than implied by a strict interest-rate interpretation of discount rates. One possible cause of this high estimated rate is that it may absorb other relevant behavior omitted from the model (as discussed in Frederick et al., 2002). In the context of our model, for instance, discount rates could be higher that the pure rate of time preference due to bidders' risk aversion about future bidding opportunities, opportunity costs of watching for auction listings, or increasing frustration with losing auctions, none of which are explicitly addressed in the model.

To see what features of the model and data lead to this high estimate of ρ, note that ρ is estimated to match the expected revenue per auction between the model and data. An increase in ρ leads to a bidding profile (the bid at a given point s relative to the deadline) that is lower and steeper, reducing the average winning price. The estimate of ρ is highly sensitive to the number of bidders per auction (λ); a decrease in the number of bidders will automatically lower the expected revenue and (by increasing the value of continued search, $V(s)$) will also lower bids. The flow of new bidders (δ) is also of particular importance in determining ρ. If δ is lower (holding α constant), then a higher fraction of buyers will win an auction before their deadlines, thereby raising the value of search and lowering bids. Thus, a lower λ or δ would reduce expected revenue, and in turn reduce the implied discount rate ρ.

In what follows, we simulate from the model using these fitted parameters and compare simulated outcomes to those in the data. The sample moments we have exploited in estimating the model parameters do not correspond to the outcomes we explore in the remainder of this section or in Section 6 below, and thus comparing the patterns observed to those predicted by the model provides a number of different dimensions on which to judge the model's goodness of fit. We will also highlight features of the data that the model does not explain.

5.3 Bids Over Duration of Search: Model Predictions

Figure 3.A provides the same analysis as in Figures 1.A but using data simulated from the model under the fitted parameters. Again, we see the average bid steadily rises within each sequence. As highlighted for the data in Section 4.2, the lines in Figure 3.A do not cross.

The data and model patterns certainly do not agree perfectly quantitatively. For example, the model range in Figure 3.A extends about 8 percentage points lower and 4 percentage points higher than the data range in Figure 1.A. We also note that the average final bid in Figure 3.A is incrementally about 0.7 percentage points higher for each longer sequence (of length 4, 5, or 6), while the last bid of each sequence is 0.5 percentage points lower for each of the longest sequences in Figure 1.A. At the same time, these longer sequences only constitute 2.6% of the data (per Figure 4.A), so these estimates are less precise.

The underlying bidding strategy (as a function of time remaining) is depicted in the solid line of Figure 3.B. Initially (for s near T), the price path is more or less linear, but as the

Figure 3: Bids Over Search Duration — Model

(A) Bids over Search Duration (B) Shading by Bidders

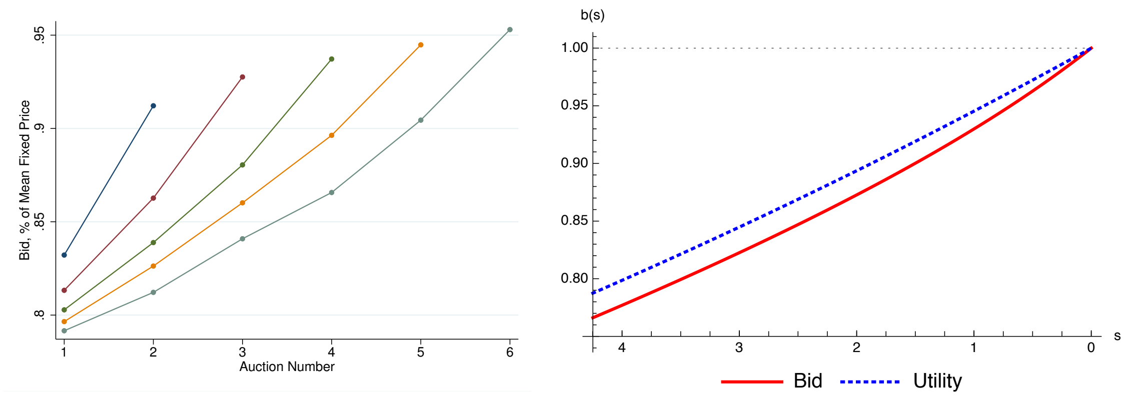

Notes: Panel (A) reproduces Figure 1.A from simulated data under the fitted parameters. Likewise, Panel (B) reports bids (solid line) and utility (dotted line) as a function of time remaining s. Because $z = 1$, these may be read as percentages relative to the retail price.

deadline approaches, greater curvature is introduced. On average, a buyer increases her bid at a rate of 5.5 percentage points per month. Because the average bidder participates in 1.11 auctions per month, this translates to an increase of 5.0 percentage points between each auction of a given product—which is 4 times larger than the 1.2 percentage point gain seen in the data (Figure 1.B).

The dotted line in Figure 3.B indicates the utility that the buyer gets by purchasing at time s (under the fitted parameters); this increases as the deadline approaches purely due to time preference. The gap between the dashed and solid lines indicates shading relative to the bidder's current utility. Note that the gap is essentially constant up until 2 months before the deadline; this is because, in this early phase of the search, bids are so low that only 11% of bidders win in this phase. In contrast, 36% of bidders win in the last 2 months (or 77% of all auction winners). Over this latter phase, bidding opportunities are running out, causing the gap to shrink. This is precisely the cause of curvature in the bidding function, and 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.

5.4 Winners and Losers: Model Predictions

As highlighted in Section 4.3, we observe in the data that 70.6% of the time the winner is the bidder with the longest observed time in the market. Under the fitted parameters, the model predicts a similar frequency of 73.2%, and this moment is not exploited in fitting the model's parameters. In contrast, if elapsed time and likelihood of winning were completely orthogonal, as assumed in standard models of consumer search, the likelihood of this event

Figure 4: Repeat Bidding Behavior

(A) Distribution of Sequence Lengths (B) Duration Between Bids, by Sequence Length

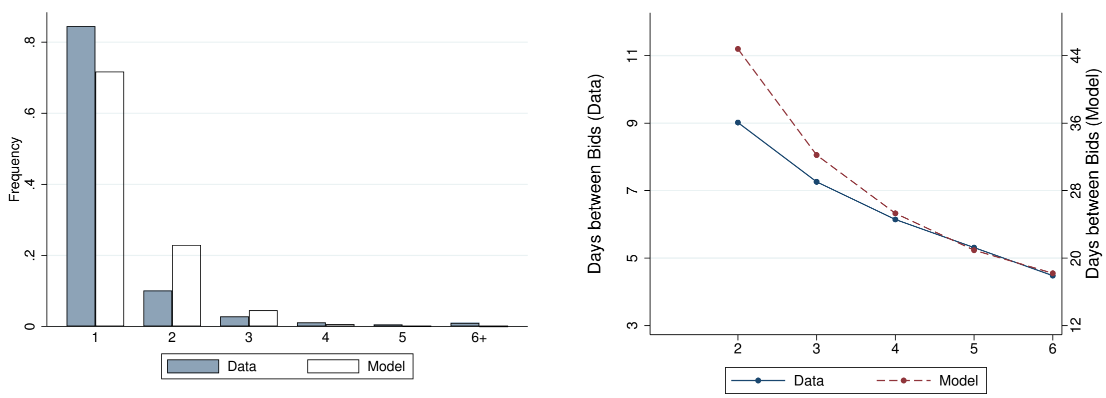

Notes: Panel (A) shows a histogram of the number of bids a buyer places on listings of the same product, for data (shaded) and model (white). Panel (B) shows the average time between consecutive bids depending on the total sequence length, for data (solid) and model (dashed).

would be drastically lower, given by $\frac{1}{\lambda} = 33.2\%$, because such orthogonality would make each bidder equally likely to win regardless of her time spent searching so far.

Figure 4.A displays the distribution of the number of auction attempts by bidders in the model and data. The model somewhat under-predicts the fraction of bidders who are observed in only one auction (72%, as opposed to 84% in the data) and over-predicts the fraction who are observed in two auctions (23%, as opposed to 10% in the data).

The data and model match qualitatively in their prediction of the average time between bids being lower for bidders observed bidding in more auctions, although the scale of the average time between bids is roughly four times larger in the model than in the data, as shown in Figure 4.B. This discrepancy is driven by our estimated T and the observed time in the market it implies (1.6 months) being larger than the observed time in the market in the data (0.4 months), as discussed in Section 5.2 above. However, the duration between bids falls at the same rate in the model and the data.

6 Market Implications of Deadlines

The market equilibrium of our model yields three clear predictions, each of which are strongly evident in the data: discounts should yield sales faster than posted prices, discount and full-price mechanisms should coexist in the market, and price dispersion should be evident both within the discount mechanism and between the mechanisms. These results are even more stark due to our focus on homogeneous products (both in the model and in the data): despite being a market of identical, new-in-box products, transaction prices vary widely and sellers use

Figure 5: Sales Rates and Price Dispersion

(A) Cumulative Fraction Sold Since Listing (B) Distribution of Closing Prices

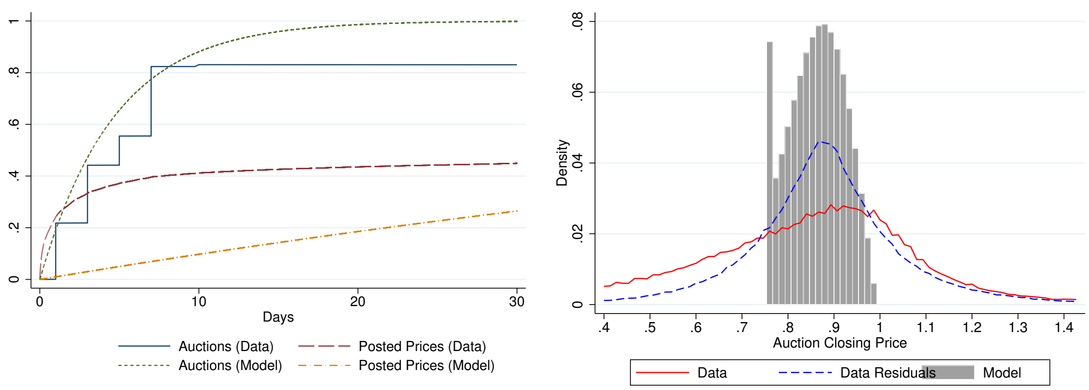

Notes: Panel (A) displays the cumulative fraction of listings sold (vertical axis) against the number of days since the listing was posted (horizontal axis) for auctions and posted-price listings, as observed in the data and as predicted from the fitted model. Panel (B) plots the density of bids predicted by the fitted model (bars), observed in the normalized eBay bids (solid line) winsorized at the 5th and 95th percentile, and observed in residuals after controlling for seller and product fixed effects (dashed line) winsorized at the 1st and 99th percentile.

multiple sales mechanisms. We then examine how underlying market design changes may have contributed in part to a recent trend toward more posted-price sales.[37] We also discuss market welfare and quantify dynamic search frictions that are unique to consumers with deadlines.

6.1 Selling Time

As predicted in the model, transactions in our data are typically completed faster through the discount (auction) mechanism than through posted-price listings. On eBay, the seller explicitly chooses the auction length for either 1, 3, 5, 7 or 10 days, whereas posted-price listings are available until a buyer purchases it and can be renewed if not purchased after 30 days. Figure 5.A plots the cumulative fraction of listings sold against the number of days after listing the item for sale. In the data, auctions (solid line) sell at the same rate as posted prices (long dash) for the first day, but the posted-price rate slows considerably thereafter.

The model predicts a similar rate for auction sales (dotted) as in the data, which is no surprise because this moment was used in estimating η. For posted prices (dash dotted), the model under-predicts the sales rate and shows less curvature than the data.

[37]Deadlines also have implications for correctly estimating demand. We show in Technical Appendix F that ignoring deadlines in a static or dynamic model would understate or overstate demand.

6.2 Price Dispersion

Our data reveal (and our model predicts) three forms of price dispersion over homogeneous products. The first form is across mechanisms, in that auctions average 15% lower sales prices than posted-price listings (see Table 1). The second form is dispersion across transaction prices *within* the discount mechanism. The distribution of the normalized second-highest bid across auctions is depicted in Figure 5.B (solid line), which has an interquartile range of 32 percentage points.[38] Some of this dispersion is due to low-price items, which show 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 in a regression of the normalized second-highest bids; the dashed line in Figure 5.B reports the distribution of the resulting residuals, which have an interquartile range of 13 percentage points, or 6 percentage points when restricted to products with a mean posted price over $100. The third form of price dispersion is that a *given individual* participating in the discount sales channel systematically offers higher prices over time, as seen in Section 4.2.

Figure 5.B reports the fitted model's predicted distribution of auction closing prices (depicted with bars). Closing prices are dispersed from 77 to 100% of the posted price. The small spike at 0.77 arises from auctions in which only one serious bidder participates (predicted to happen in 15% of auctions), and thus the auction closes at its opening price. While the data has a wider range of prices than the model, they are in closer agreement when considering the interquartile range, which is 9.0 percentage points in the model despite the assumption of homogeneous goods.

Thus, deadlines can be seen 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 overcome this result by introducing exogenous differences among buyers' search costs (e.g. Stahl, 1989) or valuations (e.g. Diamond, 1987). 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.[39] It is only after randomly arriving to the market that buyers differ ex-post, leading to a continuum of dispersed prices.[40]

[38] Lach (2002) finds levels of dispersion that are nearly this high in grocery commodity prices such as flour and frozen chicken, with an interquartile range of 15 to 19% of the average price.

[39] Directed search models (surveyed in Wright et al., 2019) can also generate pure price dispersion if buyers are indifferent about seeking lower prices but with less chance of success, and similarly for sellers offering higher prices. Sellers in our setting are likewise indifferent between fast discount sales and slow full-price sales; buyers, on the other hand, strictly prefer the discount mechanism until their deadline arrives.

[40] Deadlines had a similar effect for declining reservation wages among unemployed workers in Akın and Platt (2012); although there, workers passively responded to posted job offers, rather than buyers actively selecting bid strategies here. Also, labor markets lack the auction mechanism to extract and record reservation wages throughout a search spell.

Figure 6: Mechanism Dispersion and Listing Fees

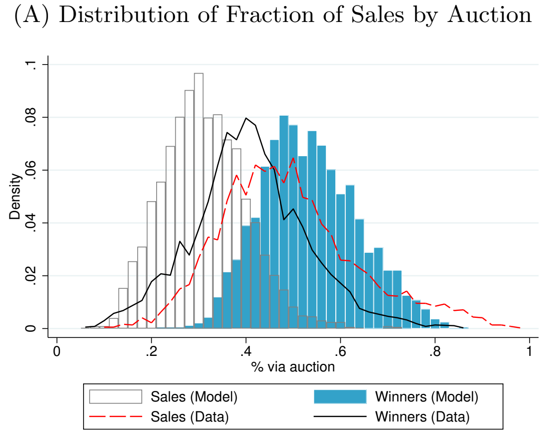

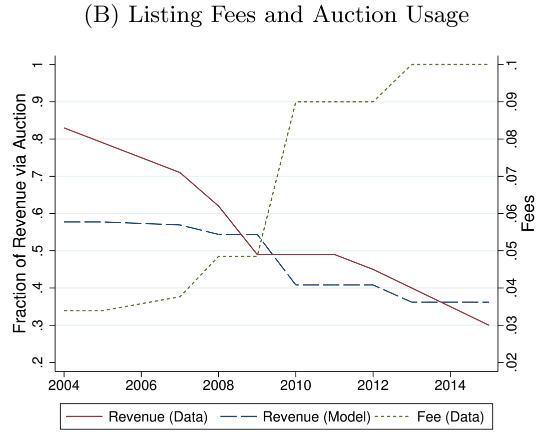

Notes: The dashed line in Panel (A) is a histogram for the fraction of eBay sales occurring via auction (versus posted prices) for each of our 3,663 products, while the solid line does the same for the fraction of eBay bidders who eventually win an auction for that product. White and shaded bars provide the corresponding model predictions, derived from product-level estimation. The solid line in Panel (B) indicates the fraction of revenue from auction (as opposed to posted-price) sales over time as reported by Einav et al. (2018). The dashed line indicates the model prediction of the fraction of revenue from auctions, holding all parameters as in Table 2 except platform fees (ℓ), which eBay adjusted over time as reported in the dotted line (on the right axis); see also Table A8.

6.3 Coexistence of Auctions and Posted-Price Sales

In our model, sellers use a mixed strategy, and thus sellers are observed transacting through both auctions and posted prices. At the estimated parameters, the fraction of sellers who sell through an auction is $\sigma = 30.8\%$. In the data, we see that 49.9% of transactions occur through auctions.[41] We also examine the fraction of sales through auctions across our various products. For this exercise, we estimate the model separately for each product, as described in Technical Appendix H. Figure 6.A provides a histogram across the products of the fraction sold by auction in the data (dashed line) and the resulting model estimate (white bars). The latter distribution is shifted about 20 percentage points lower than the former but has the same shape. The fraction of auction sales in the model and data varies somewhat across products; even so, 90% of products in the data have an auction-fraction lying between 29% and 79%, suggesting that coexistence is widespread.

One possible reason the model under-predicts auction usage is that sellers (and buyers) may be turning to other posted-price platforms such as Amazon, which we are unable to

[41]Of course, products are only included in our sample if at least 25 transactions occurred under both mechanisms. To document coexistence more broadly, we replicated these results in a larger sample that includes all products sold at least 50 times in our sample period regardless of listing method. We found very similar coexistence patterns in that broader data to those described in this section.

measure in our data. However, our model offers a dual approach to test mechanism dispersion using only auction data. In equilibrium, the fraction of sellers using auctions, σ, equals the fraction of buyers who eventually win an auction, $(\delta - H \cdot F'(0))/\delta$, and we can compare this latter object in the model and data. The caveat is that we only see buyers who register a bid in the data; for the theoretical equivalent, we divide σ by $1 - e^{-\tau\alpha P(\lambda)T}$, which is one minus the probability that a buyer never registers a bid over T periods (due to not encountering an auction or entering it after the standing price passes her valuation). At our main parameter estimates, the model predicts that 52.5% of observed bidders will win an auction during their search span; in our eBay data, we see that 41.6% do. Thus, the model over-predicts auction winners by roughly 11 percentage points. We compare the distribution of these measures across products in Figure 6.A, where the shaded bars show a histogram for the model's prediction and the solid line shows the same for the data.

While discount and non-discount sales channels frequently offer the same good in practice, such coexistence is difficult to sustain theoretically: in Wang (1993), Bulow and Klemperer (1996), Julien et al. (2001), and Einav et al. (2018), one mechanism is strictly preferred over the other except in "knife-edge" or limiting cases. Models in 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, both mechanisms are active in our model over a wide range of parameters. This is because our ex-ante identical buyers become different ex-post as they reach their deadlines. Thus, sellers can obtain a higher price at the cost of a longer wait, and free entry ensures that these forces offset each other.

6.4 Market Design: Equilibrium Effects of an eBay Listing Fee Change

We now demonstrate that even simple adjustments to the eBay marketplace may have unexpected consequences when consumers search with a deadline. From 2004 to 2015, eBay increased its fees almost every year, either by raising its percentage commission or raising its cap on the commission, as reported in Table A8 of the Technical Appendix. The majority of fee increases have occurred in final value fees—a percent commission of the final price, in line with our parameter ℓ.[42] In most of these years, the same fee structure applied to auctions and posted prices, which would seem to make any fee increase neutral across sales channels. However, in the context of our model, a higher listing fee will incentivize more sellers to choose to list through the posted-price channel.

The mechanics behind this distortion are as follows. All else equal, a fee increase makes both channels unprofitable in expectation, but this decrease in expected profits is largest for auctions, where seller profit margins are always smaller. The equilibrium is restored in the auction market as sellers offer fewer auctions, which reduces the value from continued search, $V(s)$, raising equilibrium bids and expected auction revenue. The reduction in the flow of

[42] Other fees, such as insertion fees shown in Table A8, would be included as up-front costs, c, in our model.

auctions forces more buyers to use posted-price listings. This shortens the time before a posted-price listing receives a buyer, restoring profitability for posted-price listings. In the process, the balance of transactions shifts toward more posted prices.

To quantify the impact of listing fees on auction usage, we evaluate the model predictions over this time period as ℓ changes, using the actual eBay fees in different years from Table A8 and holding other model parameters fixed at our main estimates. We show the level of ℓ in each year using the dotted line in Figure 6.B (with levels corresponding to the right axis of the figure) and the model's predicted share of auction revenue $\frac{\sigma\theta}{1-\sigma+\sigma\theta}$ for each year in the dashed line in Figure 6.B (levels shown on the left axis). For comparison, the solid line in Figure 6.B shows the auction revenue share in each year over this time frame, which declined by about 53 percentage points from 2004 and 2015.[43] Through the lens of our model, fee increases could explain nearly half of this decline (22 percentage points).

Recent work by Einav et al. (2018) focused on the first half of this period (2003–2009); their model ascribes the majority of the auction-revenue decline to reduced demand for auctions and the remainder to changing supply conditions.[44] Figure 6.B offers complementary evidence to their finding: our model suggests that fee increases explain only a small fraction of the auction decline over this early period. In contrast, fee increases appear as a potentially strong explanation for nearly all of the auction decline in the *post*-2009 period. We emphasize here that our finding is by no means causal evidence, as many other market factors may have changed over this time period.

In our model, the endogenous parameter σ (the fraction of auctions) falls more quickly as ℓ rises, and this leads to a nonlinear impact of fee increases on the auction share. For example, a 1 percentage point commission increase in 2008 decreased the auction share of revenue by 2.5 percentage points; but a similar increase in 2013 had nearly twice the impact because fees were already high. The model predicts that auctions would be completely eliminated (the degenerate equilibrium, derived in Proposition 4 in the Technical Appendix) if listing fees were to reach a level of $\ell = 14.7\%$. On the other hand, if listing fees were completely eliminated ($\ell = 0$), auctions would reach maximum usage at 70.0% revenue share.

While the market response to higher listing fees helps explain the declining use of auctions, it also illustrates a potential hazard of ignoring deadlines in market design. If buyer valuations in a given transaction are not fundamental but rather are the endogenous results of deeper factors, even a seemingly neutral change in listing fees (applied to *both* the auction and posted-price markets) not only alters which market sellers use, but also warps the distribution of buyer

[43]The auction revenue share comes from Figure 1 of Einav et al. (2018); see also Figure 1 of Backus et al. (2018), as well as our Table A8. Note that eBay's internal categorization of new vs. used items became available in 2010, so the products underlying the solid line are potentially more heterogeneous (with a mixture of new and used items) than in our sample. Also, we do not observe the exact timing of fee changes, as our data on fees comes from historical snapshots of the eBay site.

[44]The sample that Einav et al. (2018) used to estimate their model lacks some of the data required to estimate our model (such as repeat bidding), preventing us from estimating our model using their sample, and our sample lacks some variables required for estimating theirs (such as auction starting prices).

valuations and changes buyer behavior.

6.5 Welfare Consequences of Deadlines

Our model rationalizes the seemingly redundant coexistence of discount and posted-price channels when goods, sellers, and buyers are ex-ante homogeneous—an uncommon result, as noted in Section 6.3. A question that naturally follows is what social value does the discount mechanism offer? The structure of our model allows us to investigate this question. Note first that, in our model, all buyers eventually obtain and consume the good, which yields utility $x - z$ at the deadline. The remaining welfare calculations quantify the *additional* utility or costs in this marketplace relative to this utility of $x - z$.

The first-best outcome in this setting would be for the buyer to produce her own good at cost c at the time of consumption. This would increase the total welfare $x - z$ by $z - c$, or by 25.2% of the retail price under our estimated parameters. However, our environment necessitates that buyers *search* for sellers. Total welfare in the market equilibrium is simply the expected consumer surplus, because sellers earn zero expected profit. A newly-entering buyer in this market expects utility of $V(T)$ from the full span of participation in both mechanisms (measured in terms of dollars, and net of any payments to sellers). If we measure utility relative to the time of deadline (so as to be comparable to the first-best computation), this increases total welfare $x - z$ by $V(T)e^{\rho T} - (x - z)$, or by 2.6% of the retail price in the fitted model. The gap between first-best (25.2%) and equilibrium welfare (2.6%) reflects additional costs that sellers incur to connect with buyers, which come in three forms.

First, sellers must pay listing fees to the intermediary, which averages $\ell(\sigma\theta + (1-\sigma)z)$ per buyer, or 11.1% at our estimated parameters. Sellers are willing to pay these fees precisely because the intermediary connects them with a broad market of buyers, and this intermediation can be viewed as a variant of traditional search costs.[45]

Second, sellers must produce the good before it is sold. This reduces total welfare by the time cost of those resources, computed as the interest accrued during the expected wait for a buyer: $\rho c \left(\frac{\sigma}{\eta(1-e^{-\lambda})} + \frac{1-\sigma}{\zeta} \right)$. At our estimated parameters, this equals 9.7% of the retail price, most of which comes from the lengthy wait endured by posted-price sellers.

Third, even after it is sold in an auction, the buyer waits until the deadline to consume the good.[46] Purchasing the good before it is needed sacrifices the interest she could have had on her money. This is most easily computed as the residual gap after the other two costs are deducted, which is 2.0% under our fitted parameters.

These latter two costs are dynamic search frictions, reflecting the opportunity cost of tying up resources before they are needed for consumption. Note that together, these dynamic costs

[45] If these fees are merely a transfer between seller and intermediary, they would be included in equilibrium welfare; otherwise, they reflect real resources that are consumed in creating the platform for buyers and sellers.

[46] Here, we evaluate this with $\beta = 0$. This welfare cost is lessened but not eliminated when some consumption takes place at the time of purchase, $\beta > 0$. Other welfare results are unaffected.

are slightly larger than the static cost of intermediation. The third friction, where buyers may purchase earlier than needed, is a cost unique to this deadline setting, and constitutes one tenth of the total welfare costs.

We note that buyers efficiently sort across mechanisms, because the highest-valuation buyers go where they will be served immediately, while others wait for auctions. Similar inter-temporal sorting happens within a posted-price market in Deneckere and Peck (2012); when supply is insufficient in a period, low-valuation buyers delay their purchase, allowing high-valuations buyers to consume. The allocation is also efficient within a given auction in our model, because the highest bidder also has the highest valuation. However, sorting across auctions is imperfect in our model because of search frictions: buyers can miss an auction despite having a higher valuation than the winner. Related models of directed search are frequently *constrained efficient*, as in Albrecht et al. (2014). They model endogenous entry by buyers and sellers into a static, one-shot auction, and find that informational rents by buyers are exactly offset by the negative externality of an additional auction on other sellers, yielding constrained efficiency. Similar forces are present in our model, but by incorporating buyer deadlines, our model adds a dynamic component to welfare computations, and it is precisely the *timing mismatch* between production and consumption that creates the second and third welfare costs above.

This mismatch suggests the potential to raise welfare by shutting down the discount market while still allowing the posted-price market to operate freely. This would seem to have the virtue of ensuring that all purchases take place at the time of consumption, eliminating the third inefficiency. However, this shift to all posted-price listings would increase the amount of platform fees paid to $z\ell = 11.6\%$ instead of 11.1%. The second inefficiency would also increase (to 13.7%) as sellers who previously used auctions shift to the lengthy posted-price mechanism. As a result, consumer surplus (and total welfare) is simply $x - z$; thus, shutting down the discount market only takes away from total welfare. More generally, it is straightforward to show that the equilibrium consumer surplus with both markets operating ($V(T)e^{\rho T}$) is *always* higher than it would be in a market with only posted prices ($x - z$).

7 Conclusion

This work examines consumer 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 reservation price as consumers approach their deadlines. The model also rationalizes the coexistence of discount and full-price sales channels selling the same item, because transactions occur more quickly in the former but at a lower price.

While the idea that buyers would be willing to pay more as a deadline draws near is intuitive, it has far-ranging logical consequences: *e.g.* who wins auctions, how buyers are distributed in the market, and which market sellers will enter. In answering these questions, the model is consistently disciplined with deadlines as the single source of ex-post heterogeneity. Even with this rigid structure, the model replicates many key features of the observed data, including moments that were not used in fitting the parameters. By omitting exogenous differences that would typically explain the variation across auction outcomes, this setting yields the cleanest predictions and highlights the mechanisms at work, which would still be at play even if we were to introduce exogenous differences among discount rates, valuations of goods, seller costs, etc. See Technical Appendix D for further discussion.

In our empirical application we document a variety of reduced-form findings consistent with the time sensitivity we model. In particular, buyers offer more and are more likely to win in each successive attempt to win a discount. These conclusions from observational eBay data are also consistent with evidence we present from directly surveyed consumers. We also estimate the model's parameters and demonstrate that buyer deadlines have implications for rates of sales, price dispersion, mechanism coexistence, welfare calculations, and market design.

While our empirical application focused on new, homogeneous goods sold online, the lessons we learn are equally applicable for impatient repeat buyers on imperfectly interchangeable items. Indeed, we anticipate similar results for consumer search in the presence of 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.

References

Akın, Ş Nuray and Brennan C Platt, "Running Out of Time: Limited Unemployment Benefits and Reservation Wages," *Review of Economic Dynamics*, 2012, *15* (2), 149–170.

_ and _, "A Theory of Search with Deadlines and Uncertain Recall," *Economic Theory*, 2014, *55* (1), 101–133.

Albrecht, James, Axel Anderson, Eric Smith, and Susan Vroman, "Opportunistic Matching in the Housing Market," *International Economic Review*, 2007, *48* (2), 641–664.

_, Pieter Gautier, and Susan Vroman, "Efficient Entry in Competing Auctions," *American Economic Review*, 2014, *104* (10), 3288–3296.

Argyle, Bronson, Taylor Nadauld, and Christopher Palmer, "Real Effects of Search Frictions in Consumer Credit Markets," 2017. MIT Sloan Working Paper 5242-17.

Backus, Matt and Gregory Lewis, "Dynamic Demand Estimation in Auction Markets," 2016. NBER Working Paper 22375.

_, Tom Blake, Brad Larsen, and Steve Tadelis, "Sequential Bargaining in the Field: Evidence from Millions of Online Bargaining Interactions," 2018. NBER Working Paper 24306.

Balat, J., P. Haile, H. Hong, and M. Shum, "Nonparametric Tests for Common Values at First-price Sealed-bid Auctions," 2016. Working paper, Yale.

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

Berg, Gerard J. Van Den, "Nonstationarity in Job Search Theory," *Review of Economic Studies*, 1990, *57* (2), 255–277.

Blake, Thomas, Chris Nosko, and Steven Tadelis, "Returns to Consumer Search: Evidence from eBay," in "Proceedings of the 2016 ACM Conference on Economics and Computation" ACM 2016, pp. 531–545.

Board, Simon and Andrzej Skrzypacz, "Revenue Management with Forward-Looking Buyers," *Journal of Political Economy*, 2016, *124* (4), 1046–1087.

Bodoh-Creed, Aaron, Jorn Boehnke, and Brent Hickman, "How Efficient are Decentralized Auction Platforms?," *Review of Economic Studies*, 2018, *forthcoming.*

Bulow, Jeremy and Paul Klemperer, "Auctions Versus Negotiations," *American Economic Review*, 1996, *86* (1), pp. 180–194.

Caldentey, René and Gustavo Vulcano, "Online Auction and List Price Revenue Management," *Management Science*, 2007, *53* (5), 795–813.

Coey, Dominic, Bradley Larsen, and Brennan Platt, "Discounts and Deadlines in Consumer Search," 2019. NBER Working Paper 22038.

_ , _ , and _ , "Confidential Data for: Discounts and Deadlines in Consumer Search," 2020. Data available for other researchers via contract with eBay, Inc. (file "CoeyLarsen-Platt_2020_eBay_Data.zip." Created 2020-05-11.

_ , _ , and _ , "Data and Code for: Discounts and Deadlines in Consumer Search," 2020. Nashville, TN: American Economic Association [publisher]. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 2020-07-24. https://doi.org/10.3886/E119387V1.

De los Santos, Babur, Ali Hortaçsu, and Matthijs R Wildenbeest, "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior," *American Economic Review*, 2012, *102* (6), 2955–2980.

_ , _ , and _ , "Search with Learning for Differentiated Products: Evidence from E-commerce," *Journal of Business & Economic Statistics*, 2017, *35* (4), 626–641.

Deneckere, Raymond and James Peck, "Dynamic Competition with Random Demand and Costless Search: A Theory of Price Posting," *Econometrica*, 2012, *80* (3), 1185–1247.

Diamond, Peter, "A Model of Price Adjustment," *Journal Of Economic Theory*, 1971, *3* (2), 156–168.

_ , "Consumer Differences and Prices in a Search Model," *Quarterly Journal of Economics*, 1987, *102*, 429–436.

Dilme, Francesc and Fei Li, "Revenue Management Without Commitment: Dynamic Pricing and Periodic Fire Sales," *Review of Economic Studies*, 2019, *86* (5), 1999–2034.

Einav, Liran, Chiara Farronato, Jonathan D Levin, and Neel Sundaresan, "Auctions versus Posted Prices in Online Markets," *Journal of Political Economy*, 2018, *126* (1), 178–215.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue, "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, 2002, *40* (2), 351–401.

Genesove, David, "Search at Wholesale Auto Auctions," *Quarterly Journal of Economics*, 1995, *110* (1), 23–49.

Hammond, Robert G, "A Structural Model of Competing Sellers: Auctions and Posted Prices," *European Economic Review*, 2013, *60*, 52–68.

Hendricks, Kenneth and Alan Sorensen, "Dynamics and Efficiency in Decentralized Online Auction Markets," 2018. NBER Working Paper 25002.

_ **, Ilke Onur, and Thomas Wiseman**, "Last-minute Bidding in Sequential Auctions with Unobserved, Stochastic Entry," *Review of Industrial Organization*, 2012, *40* (1), 1–19.

Iyer, Krishnamurthy, Ramesh Johari, and Mukund Sundararajan, "Mean Field Equilibria of Dynamic Auctions with Learning," *Management Science*, 2014, *60*, 2949–2970.

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

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

Jeitschko, Thomas D, "Learning in Sequential Auctions," *Southern Economic Journal*, 1998, *65*, 98–112.

Julien, Benoit, John Kennes, and Ian Paul King, "Auctions and Posted Prices in Directed Search Equilibrium," *Topics in Macroeconomics*, 2001, *1*, 1–16.

Kehoe, Patrick J, Bradley J Larsen, and Elena Pastorino, "Dynamic Competition in the Era of Big Data," 2018. Working paper, Stanford University.

Kohn, Meir G. and Steven Shavell, "The Theory of Search," *Journal of Economic Theory*, 1974, *9*, 93–123.

Lach, Saul, "Existence and Persistence of Price Dispersion: An Empirical Analysis," *Review of Economics and Statistics*, 2002, *84* (3), 433–444.

Lauermann, Stephan, Wolfram Merzyn, and Gábor Virág, "Learning and Price Discovery in a Search Market," *Review of Economic Studies*, 2017, *85* (2), 1159–1192.

Levin, Dan and James Smith, "Equilibrium in Auctions with Entry," *American Economic Review*, 1994, *84* (3), 585–599.

Mierendorff, Konrad, "Optimal Dynamic Mechanism Design with Deadlines," *Journal of Economic Theory*, 2016, *161*, 190–222.

Myerson, Roger B, "Population Uncertainty and Poisson Games," *International Journal of Game Theory*, 1998, *27* (3), 375–392.

Platt, Brennan, "Inferring Ascending Auction Participation from Observed Bidders," *International Journal of Industrial Organization*, 2017, *54*, 65–88.

Said, Maher, "Sequential Auctions with Randomly Arriving Buyers," *Games and Economic Behavior*, 2011, *73* (1), 236–243.

Song, Unjy, "Nonparametric Estimation of an eBay Auction Model with an Unknown Number of Bidders," 2004. Working paper, University of British Columbia.

Stahl, Dale O, "Oligopolistic Pricing with Sequential Consumer Search," *American Economic Review*, 1989, *79*, 700–712.

Stigler, George J, "The Economics of Information," *Journal of Political Economy*, 1961, *69* (3), 213–225.

Waisman, Caio, "Selling Mechanisms for Perishable Goods: An Empirical Analysis of an Online Resale Market for Event Tickets," 2018. Working paper, Stanford University.

Wang, Ruqu, "Auctions Versus Posted-price Selling," *American Economic Review*, 1993, *83*, 838–851.

Wright, Randall, Philipp Kircher, Benoit Julien, and Veronica Guerrieri, "Directed Search and Competitive Search Equilibrium: A Guided Tour," *Journal of Economic Literature*, 2019, *forthcoming.*

Zeithammer, Robert, "Forward-looking Bidding in Online Auctions," *Journal of Marketing Research*, 2006, *43* (3), 462–476.

Appendix: Proofs

Proof of Proposition 1. This differential equation (8) is an application of the Kolmogorov forward equation (with no stochastic component or time trend), and has the following unique solution with two constants of integration k and m:

$$F(s) = \frac{1}{\tau H} \ln \left(\frac{\alpha\tau - e^{\tau H k(s+m)}}{\tau H k} \right). \tag{21}$$

The constants are determined by our two boundary conditions. Applying $F(T) = 1$, we obtain $m = \frac{1}{\tau H k} \ln\left(\alpha\tau - \tau H k e^{\tau H}\right) - T$. The other boundary condition, $F(0) = 0$, requires that k satisfy:

$$\alpha\tau \left(1 - e^{-\tau H T k}\right) - \tau H k \left(1 - e^{\tau H(1-Tk)}\right) = 0. \tag{22}$$

From (9), we know that $HF'(T) = \delta$, and using the solution for F in (21), this yields:

$$k = \frac{\delta + \alpha e^{-\tau H}}{H}. \tag{23}$$

When we substitute for m and k in (21), we obtain the equilibrium solution for F^* depicted in (11). Also, (23) is used to replace k in the boundary condition in (22), we obtain the formula for ϕ in (10) which implicitly solves for H^*.

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

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

$$\phi'(H) = -\tau \left(\alpha e^{-\tau H} + \delta(e^{\tau H} + \alpha\tau T) e^{-\tau(\alpha e^{-\tau H} + \delta)T}\right) < 0.$$

Thus, as a decreasing function, $\phi(H)$, crosses zero only one time, at H^*.

We finally turn to the solution for the bidding function. Again, we start by simplifying the infinite sums in (4) and (5). The first sum is similar to that in (8). 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:

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

Next, by taking the derivative of $b(s) = \left(\beta + (1-\beta)e^{-\rho s}\right) x - V(s)$ in (6), we obtain $b'(s) = -\rho(1-\beta) x e^{-\rho s} - V'(s)$. We use these two equations to substitute for $V(s)$ and $V'(s)$, obtaining:

$$(\rho + \alpha\tau e^{-\lambda F(s)})b(s) + b'(s) = \rho\beta x + \alpha\tau\left(e^{-\lambda}b(T) + \int_s^T \lambda e^{-\lambda F(t)} b(t)F'(t)dt\right). \tag{24}$$

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

$$(\rho + \alpha\tau e^{-\lambda F(s)})b'(s) + b''(s) = 0. \tag{25}$$

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

$$b(s) = a_1 \cdot \left(\frac{\delta e^{\tau H^* - \tau\kappa T}}{\rho} + \frac{\alpha e^{-\tau\kappa s}}{\rho + \tau\kappa}\right) e^{-s\rho} + a_2. \tag{26}$$

One constant of integration is pinned down by (24). We substitute for $b(s)$ in (24) using (26), and solve for a_2. This can be done at any $s \in [0, T]$ with equivalent results, but is least complicated at $s = T$ because the integral disappears: $(\rho + \alpha\tau e^{-\lambda F(T)})b(T) + b'(T) = \rho\beta x + \alpha\tau e^{-\lambda}b(T)$. After substituting $b(T)$, $b'(T)$, and $F(T)$, solving for a_2 yields:

$$a_2 = \beta x + a_1 \frac{\alpha\tau\kappa}{\rho\,(\rho + \tau\kappa)} e^{-(\rho+\tau\kappa)T}. \tag{27}$$

The other constant of integration is determined by boundary condition (2). 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 (26) evaluated at 0, and substitute for a_2 using (27), then solve for a_1:

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

If the solutions for a_1 and a_2 are both substituted into (26), one obtains (12) with minor simplification. □

Proof of Proposition 2. Note that $e^{\rho s}$ and $e^{-\tau\kappa s}$ both equal 1 at $s = 0$, causing the fractional term in (12) to become 0 and yielding $b(0) = z$. The first derivative of $b^*(s)$ is:

$$b'(s) = -\frac{(z - \beta x)\rho(\rho + \tau\kappa)\left(\delta e^{\tau H^*} + \alpha e^{\tau\kappa(T-s)}\right)}{\tau\kappa\left(\delta e^{\tau H^*} + \alpha e^{-\rho T}\right) + \rho\left(\delta e^{\tau H^*} + \alpha e^{\tau\kappa T}\right)} \cdot e^{-\rho s} < 0,$$

where the sign holds because each of the parenthetical terms is strictly positive. The second derivative is:

$$b''(s) = \frac{(z - \beta x)\rho(\rho + \tau\kappa)\left(\rho\delta e^{\tau H^*} + (\rho + \tau\kappa)\alpha e^{\tau\kappa(T-s)}\right)}{\tau\kappa\left(\delta e^{\tau H^*} + \alpha e^{-\rho T}\right) + \rho\left(\delta e^{\tau H^*} + \alpha e^{\tau\kappa T}\right)} \cdot e^{-\rho s} > 0.$$

Again, each parenthetical term is strictly positive. □

Online Appendix to
"Discounts and Deadlines in Consumer Search"

Dominic Coey Bradley Larsen Brennan C. Platt[47]

A Survey of Deadlines in Consumer Search

From September 27th to November 1st, 2018, Qualtrics administered a survey on our behalf to a panel of consumers (see Coey et al. 2020b for the raw survey data). Qualtrics is a survey administration company that recruits survey participants through a variety of means, including websites, member referrals, targeted email lists, gaming sites, social media, and other sources. Panelists are incentivized to complete the survey through some small monetary compensation or through points toward a particular product loyalty program. These panelists are thus likely to be comfortable with online activity.

Members of the Qualtrics panel were selected at random to receive an email offering them the opportunity to participate in our survey. Consumers who opted to start the survey were given the following screening question to identify participants who could recall an item for which they had searched:

> Can you think of a recent purchase for which you considered searching at multiple locations (either online or offline) in order to find a good price? Note: Think back only on non-food items. Examples might include a phone/tablet/laptop (or other consumer electronic item), a toy, an item of clothing or accessory, a sporting good, a book, an appliance or other household item, or even a car.
>
> - Yes
> - No

Consumers who responded "No" were given no further questions. Consumers who responded "Yes" entered into our sample and were given the following survey. Respondents were required to make a response to all questions. Questions 1, 2, 3, 5, and 6 were free-response questions. Questions 4 and 7 were check-box questions, and the respondents were allowed to select as many of the options as desired, but were required to select at least one. Questions 8–12 were radio-button questions, and the respondents were required to select one and only one option.

1. What was the item you purchased? Describe it in just a few a words.
2. About how much money (in dollars) did you pay for it?
3. About how much money (in dollars) do you think you saved by searching around?

[47]Coey: Facebook, Core Data Science; coey@fb.com. Larsen: Stanford University, Department of Economics and NBER, bjlarsen@stanford.edu. Platt: Brigham Young University, Department of Economics, brennan_platt@byu.edu

4. Where did you search? Select ALL that apply:
   (a) Amazon
   (b) eBay
   (c) Google
   (d) Large retailer's physical store
   (e) Small retailer's physical store
   (f) Other
5. How many times did you visit a physical store in attempting to find the item?
6. How many times did you visit an online retail site in attempting to find the item?
7. Select ALL that apply to the item you purchased: [Respondents were allowed to select as many of the following as desired, but were required to select at least one.]
   (a) The item was a gift for someone
   (b) I wanted/needed this item for an upcoming event
   (c) I wanted/needed this item more as time went by
   (d) I knew where I could find this item for sure at a high price, but I searched around to find a low price
   (e) None of the above
8. Which of the following best describes the urgency with which you wanted/needed the item? [Respondents were required to select one and only one of the following]
   (a) I wanted/needed this item as soon as possible
   (b) It wasn't urgent that I get the item as soon as possible, just as long as it came in time for a particular deadline or a particular use of the item I had in mind
   (c) None of the above
9. If you hadn't found/purchased the item when you did, which of the following best describes what you would have done next in your attempt to get it? [Respondents were required to select one and only one of the following]
   (a) Given up searching.
   (b) Kept trying to find a good price, and eventually purchased it even if it had cost a little more than (respondent's answer to Q.2)
   (c) Kept trying to find a good price, and eventually purchased it only if it had cost (respondent's answer to Q.2) or less
10. Which response best completes the following sentence? "If I hadn't purchased this item when I did, I would have been fine getting this item anytime within the next ______."
    (a) one day
    (b) one week
    (c) two weeks
    (d) month
    (e) two months
    (f) four months
    (g) six months
    (h) one year

(i) century (in other words, anytime would have been fine — I had no timeline for getting this item)

11. Which response best completes the following sentence? “I was aware that I wanted/needed to eventually buy this item about ______ before I purchased it.”
    - (Same options as prior question except the last)

12. Select the answer that best describes what you were trying to learn from your search:
    (a) I was only trying to find the best price; I knew exactly what item I wanted
    (b) I was mainly trying to find the best price, but I was also trying to find which product was the best fit for me
    (c) Price and product fit were equally important to me in my search
    (d) I was mainly trying to find which product was the best fit for me, but I was also trying to find the best price
    (e) I was only trying to find which product was the best fit for me, independent of price
    (f) None of the above

Qualtrics screens for non-serious responders in several ways. First, the company collects responses until 50 consumers have completed the survey. The company then computes a speed threshold (by computing the median time taken on the survey among those first 50 completers, and setting the threshold to half of that time); any respondent (or subsequent respondent) who completes the survey faster than that threshold (which in our case is 1.15 minutes) is not considered a serious respondent. Second, Qualtrics allowed us to examine responses to identify those in which the free response questions were non-serious (*e.g.* an answer of 0 for Q.2; answers such as “I don’t know” or “none” for Q.1; or answers for Q.1 that describe food, which violates the screening question.).

The survey responses are summarized by price range in Table A1 and by product categories in Table A2. Categories were determined from respondents’ free-response item descriptions (Q.1) as follows: Automotive (vehicles and parts), Technology (computers, TVs, phones, game consoles), Entertainment (video games, books, sports equipment, toys), Household (appliances, furniture), Clothing (clothes, jewelry), and Other. The responses show remarkable consistency across the various products and prices. A notable exception is with automotive purchases, which are much more expensive, are rarely motivated by a special event, are less likely to be needed more over the search spell, and have more searches occur but at specialized websites rather than popular consumer websites.

Using the respondents’ estimated savings, we consider whether those who completed their purchase relatively early in their search span saved more, consistent with our model’s prediction. To account for the wide price range and differing potential search spans, we measure both variables in percentage rather than absolute terms. Table A3 reports the regression results. Despite heterogeneous goods and potentially imprecise guesses from respondents on savings and potential search span, we find a positive correlation between early purchases and greater savings. The estimate is quite noisy in the first column. The point estimate and its precision increase as we narrow the sample to those whose reported motives most closely fit

Table A1: Survey Summary Statistics by Price Range

			≤ \$33	> \$33 & ≤ \$150	> \$150	Total
		N	416	397	397	1210
Q2	Purchase price	(mean)	16	77	2600	884
		(sd)	9	35	7213	4299
Q3	% saved	(mean)	39	29	22	30
Q10&11	Potential search span	(mean)	46	67	99	70
		(sd)	66	85	135	101
	% of search remaining	(mean)	50	49	45	48
	Unlimited potential span	(%)	3.1	1.8	2.5	2.5
Q5	# of physical searches	(mean)	2.4	1.8	2.	2.1
Q6	# of online searches	(mean)	3.1	3.8	5.5	4.1
Q4	a. Searched Amazon	(%)	74	73	59	69
	b. Searched eBay	(%)	31	28	25	28
	c. Searched Google	(%)	24	25	27	25
Q7	a–b: For a special event	(%)	36	38	23	32
	a–c: Needed more over time	(%)	65	66	64	65
	d. Knew high-price option	(%)	43	47	50	47
Q8	a. Needed ASAP	(%)	40	44	53	46
	b. Needed by deadline	(%)	45	42	38	42
Q9	b. Willing to pay more in future	(%)	66	63	64	64
Q12a	a. Only searching on price	(%)	52	48	46	49

Notes: Table provides means and standard deviations for a participants' survey responses. The first column denotes the question number and, in some cases, the response letter corresponding to the survey questions described in the text of Technical Appendix A. The second column provides an abbreviated explanation of the survey question. The final column contains statistics for the full sample. The columns labeled with monetary amounts (*e.g.* "≤ \$33,") report statistics for a particular subsample based on the participant's reported purchase price.

Table A2: Survey Summary Statistics by Category

			Automotive	Technology	Entertainment	Household	Clothing	Other	Total
		N	52	329	110	210	183	326	1210
Q2	Purchase price	(mean)	15,613	398	56	375	63	92	884
		(sd)	14,213	476	98	6194	109	455	4,299
Q3	% saved	(mean)	14	27	35	29	33	32	30
Q10&11	Potential search span	(mean)	122	78	55	91	56	54	70
		(sd)	149	117	79	118	72	74	101
	% of search remaining	(mean)	41	47	50	50	49	48	48
	Unlimited potential span	(%)	3.8	2.1	3.6	1.9	2.2	2.8	2.5
	Span > 20 days	(%)	85	74	72	77	78	65	73
Q5	# of physical searches	(mean)	2.7	1.8	1.3	1.7	4.1	1.6	2.1
Q6	# of online searches	(mean)	6.5	4.5	3.4	4.2	3.4	3.8	4.1
Q4	a. Searched eBay	(%)	13	29	37	21	25	33	28
	b. Searched Google	(%)	7.7	29	27	23	26	25	25
	c. Searched Amazon	(%)	17	73	84	66	65	71	69
Q7	a–b. For a special event	(%)	1.9	30	47	25	43	34	32
	a–c. Needed more over time	(%)	42	66	65	63	65	67	65
	d. Knew high-price option	(%)	56	46	45	45	51	45	47
Q8	a. Needed ASAP	(%)	54	52	34	44	38	48	46
	b. Needed by deadline	(%)	35	36	50	40	50	42	42
Q9	b. Willing to pay more in future	(%)	62	64	57	66	64	66	64
Q12	a. Only searching on price	(%)	46	46	65	40	50	50	49

Notes: Table provides descriptive statistics for the same survey responses as in Table A1, but broken down by product category (based on the participants' responses to survey Q1).

the model assumptions, such as having several weeks or more to search, or being willing to pay more over time, or searching purely for the best price rather than across competing products.

Table A3: Dependent Variable: Percentage Savings, Self-Reported

	(1)	(2)	(3)	(4)
% Remaining Search Time	1.82	4.66	7.71	13.92
	(3.15)	(3.65)	(4.78)	(6.81)
Constant	29.2	28.5	27.1	23.8
	(1.61)	(1.82)	(2.31)	(3.22)
N	764	534	347	162
Willing to pay more in future	X	X	X	X
Span > 20 days		X	X	X
Exclude clothing and household			X	X
Only searching on price				X

Notes: Table displays results of a regression of the percent saved by the consumer (computed as the response to Q3 divided by the response to Q2) regressed on the percent of search time remaining (computed as number of days corresponding to the response to Q10 divided by the sum of the days corresponding to Q10 and Q11), with progressively more restrictive samples used in Columns (1) through (4). Column (1) limits the sample to those respondents who indicated a willingness to pay more in the future (Q9b); column (2) adds a restriction that search span be greater than 20 days; column (3) excludes clothing and household items; and column (4) only includes those participants who were searching only for a good price (Q12a). Robust standard errors are displayed in parentheses.

B Comparative Statics

B.1 Comparative Statics in the Buyer Equilibrium

In this section we discuss comparative statics results for the model parameters. Although our equilibrium has no closed-form solution, these comparative statics can be obtained by implicit differentiation of $\phi(k)$, which allows for analytic derivations reported below.

Table A4 reports the sign of the derivatives of four key statistics in the buyer equilibrium. 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 measure of buyers who never win an auction and must use the posted-price listings; this crucially affects the profitability of the posted-price market in the market equilibrium. 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

Table A4: Comparative Statics on Key Statistics: Buyer Equilibrium

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

Notes: $**$ 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.

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 she 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 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 equation (10) or (11). Intuitively, this is because the rate at which bidders exit is determined by how often auctions occur, which is exogenous here. Also, who exits depends on the ordinal ranking of their valuations, which does not change even if the cardinal values are altered. Their bids react as one would expect: buyers offer less when their utility from future consumption is valued less. By the same token, a decrease in β has no effect on the distribution of bidders, but will reduce their bids because more utility from consumption is delayed until the deadline.

We can also consider the effect (not shown in Table A4) of the parameter change on the expected revenue generated in an auction. For the first four 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 farther away; there, the additional participants override the lower initial bid, driving up expected revenue.

Table A5: Comparative Statics on Key Statistics: Market Equilibrium

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

Notes: Reported signs are numeric computations under estimated parameters.

B.2 Comparative Statics in the Market Equilibrium

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. Table A5 summarizes these typical effects. We are particularly interested in how parameter changes affect the distribution of sellers across mechanisms. We find that more sellers join the discount market when buyers are more attentive (τ), less patient (ρ), or have less time (T). Higher seller costs (whether in listing fee, ℓ, or production, c) also shift sellers from auctions to posted prices.

To examine the effects in greater depth, first consider an increase in τ. In the buyer equilibrium, this leads to more participants per auction, who then are willing to bid more. In the market equilibrium, however, more attentive buyers also induce sellers to offer more auctions. This more than offsets the effect of more participants per auction, producing a net decline in bids. On net, however, expected revenue slightly increases.

Next, an increase in ρ reduces bids but had no effect on the distribution of buyers in the buyer equilibrium. 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. The comparative statics for listing fees ℓ behave similarly, as discussed in the Section 6.4 of the text.

B.3 Derivation of Buyer Equilibrium Comparative Statics

Because we do not have a closed-form solution for the endogenous number of participants per auction, we use implicit differentiation of $\phi(H^*) = 0$ from (10) to determine the effect of the exogenous parameters on H^*. In fact, we find it convenient to express this implicit differentiation in terms of the participants per auction, $\lambda^* \equiv \tau H^*$; so with slight abuse of notation, we refer to $\phi(\lambda)$ when literally it would be $\phi(\lambda/\tau)$. In preparation for implicit differentiation, we note that $\phi'(\lambda) < 0$ for all λ:

$$\frac{\partial \phi}{\partial \lambda} = -\alpha e^{-\lambda} - (\tau T\alpha + e^{\lambda})\delta e^{-\tau T\kappa} < 0, \tag{28}$$

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:

$$b(T) = ze^{-\rho T} \cdot \frac{\kappa\left(\tau\kappa + \rho\right)e^{\lambda^*}}{\tau\kappa\left(\delta e^{\lambda^*} + \alpha e^{-\rho T}\right) + \rho\left(\delta e^{\lambda^*} + \alpha e^{\tau T\kappa}\right)}. \tag{29}$$

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:

$$b(T) = \frac{ze^{-\rho T}}{\delta} \cdot \frac{\left(\tau\kappa + \rho\right)\left(\kappa - \alpha\right)}{\tau\left(\kappa - \alpha + \alpha e^{-(\rho - \tau\kappa)T}\right) + \rho}. \tag{30}$$

B.3.1 Auction Rate, α

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

$$\begin{aligned}
\frac{\partial \phi}{\partial \alpha} &= -1 + e^{-\lambda} + \tau T\delta e^{-\tau T\kappa} &(31)\\
&= -1 + e^{-\lambda}\left(1 + \left(\frac{\delta + \alpha e^{-\lambda} - \alpha}{\delta + \alpha e^{-\lambda}}\right)\ln\left(\frac{\delta e^{\lambda}}{\delta + \alpha e^{-\lambda} - \alpha}\right)\right). &(32)
\end{aligned}$$

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:

$$\ln\left(\frac{\delta e^{\lambda}}{\delta + \alpha e^{-\lambda} - \alpha}\right) - \left(e^{\lambda} - 1\right)\frac{\delta + \alpha e^{-\lambda}}{\delta + \alpha e^{-\lambda} - \alpha} \leq 0 \tag{33}$$

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:

$$-\frac{\left(e^{\lambda}-1\right)\left(\alpha+\delta e^{\lambda}\right)\left(2\alpha+e^{\lambda}(\delta-\alpha)\right)}{\left(\alpha+(\delta-\alpha)e^{\lambda}\right)^{2}} \tag{34}$$

Each parenthetical term is strictly positive for all $\lambda > 0$, so the left-hand side of (33) is strictly decreasing in λ. Thus, (33) 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,

$$\frac{\partial \lambda}{\partial \alpha} = -\frac{1-(1+\tau T(\kappa-\alpha))e^{-\lambda}}{\kappa-\alpha+(1+\tau T(\kappa-\alpha))\alpha e^{-\lambda}}. \tag{35}$$

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

$$\frac{\partial F'(0)}{\partial \alpha} = e^{-\lambda} - 1 + \alpha \cdot \frac{\partial \lambda}{\partial \alpha}. \tag{36}$$

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.

B.3.2 Attention, τ

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

$$\frac{\partial \phi}{\partial \tau} = \delta \kappa T e^{-\tau T \kappa} > 0. \tag{37}$$

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

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

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:

$$\frac{\partial F'(0)}{\partial \tau} = \frac{\partial F'(0)}{\partial \lambda} \cdot \frac{\partial \lambda}{\partial \tau} = -\alpha e^{-\lambda} \cdot \frac{\partial \lambda}{\partial \tau} \tag{39}$$

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 (30) rather than examining it in terms of $g(t)$. Because $\kappa'(\tau) = \alpha e^{-\lambda} \lambda'(\tau)$, the comparative static on $b(T)$ works out to:

$$\frac{\partial b(T)}{\partial \tau} = \frac{z \alpha e^{\lambda} \psi}{(\kappa - \alpha)\left(\tau\alpha + (\tau(\kappa - \alpha) + \rho) e^{T(\rho + \tau\kappa)}\right)^2 \left(\alpha + (\kappa - \alpha)\left(\tau \alpha T + e^{\lambda}\right)\right)}. \tag{40}$$

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 $\psi > 0$ because the remaining terms in $\frac{\partial b(T)}{\partial \tau}$ are always positive.

To verify the sufficient conditions listed under Table A4 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, because $\kappa > \alpha$ 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:

$$\rho\left(e^{\lambda}(\kappa - \alpha) + \alpha\right) - T(\kappa\tau + \rho)\left(\tau(\kappa - \alpha)^2 + \kappa\rho\right) \quad > \quad \rho e^{\lambda}(\kappa - \alpha) - T\left(\tau\kappa + \rho\right)^2 \kappa$$

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

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

B.3.3 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:

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

Because $t \le T$, 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.

B.3.4 Immediate Consumption, β

The fraction of immediate consumption has no impact on (10), so λ^* will not change even if consumers obtain more utility at the time of purchase. Thus the number and distribution of buyers in the market are unaffected. The bid function is thus directly impacted as

$$\frac{\partial b(T)}{\partial \beta} = x \cdot \frac{(1-e^{-\rho T})\delta e^{\lambda^*}\left(\tau\kappa+\rho\right)+\rho\alpha\left(e^{\tau\kappa T}-e^{-\rho T}\right)}{\tau\kappa\left(\delta e^{\lambda^*}+\alpha e^{-\rho T}\right)+\rho\left(\delta e^{\lambda^*}+\alpha e^{\tau\kappa T}\right)} > 0. \tag{41}$$

The inequality holds because $e^{\tau\kappa T} > 1 > e^{-\rho T}$.

B.3.5 Deadline, T

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

$$\frac{\partial \phi}{\partial T} = \delta\kappa\tau e^{\lambda^*} e^{-\tau T\kappa}, \tag{42}$$

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,

$$\frac{\partial \lambda}{\partial T} = \frac{\delta\tau\kappa}{\delta\left(1+\tau T\alpha e^{-\lambda^*}\right)+\alpha e^{\tau T\kappa - 2\lambda^*}} \tag{43}$$

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:

$$\frac{\partial F'(0)}{\partial T} = \frac{\partial F'(0)}{\partial \lambda} \cdot \frac{\partial \lambda}{\partial T} = -\alpha e^{-\lambda} \cdot \frac{\partial \lambda}{\partial T} \tag{44}$$

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 $t \in [0,T]$. Because $\mu(0) > \mu(T) = 0$ and $\mu''(t) < 0$, $\mu(t) > 0$ for all $t \in [0,T)$.

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

C Market Equilibrium Model

The equilibrium solution to the market model is presented here, along with the propositions for the existence of degenerate and dispersed equilibria and their proofs.

C.1 Equilibrium Solution

While the market equilibrium conditions simplify considerably, they do not admit an analytic solution and we must numerically solve for both α^* and H^*. Equilibrium is attained when both (10) and (20) simultaneously hold. 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. Once α^* and H^* are found, the remaining equilibrium objects are easily solved as follows:

$$\Pi_a^* = c \tag{45}$$

$$\Pi_p^* = c \tag{46}$$

$$A^* = \frac{\alpha^*}{\eta} \tag{47}$$

$$P^* = \frac{(z(1-\ell)-c)\left(\delta - \alpha^*\left(1-e^{-\tau H^*}\right)\right)}{\rho c} \tag{48}$$

$$\zeta^* = \frac{\rho c}{z(1-\ell)-c} \tag{49}$$

$$\sigma^* = \frac{\alpha^*\left(1-e^{-\tau H^*}\right)}{\delta}. \tag{50}$$

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

Proposition 3. *A market equilibrium with an active discount channel ($\alpha^* > 0$) must satisfy $\phi(H^*) = 0$, equation (20), equations (11) through (14), and equations (45) through (50).*

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 and by multiple sales mechanisms. 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.[48]

Proposition 4. *The degenerate market equilibrium, described by equations (11) and (12) and equations (45) through (50) with $\alpha^* = 0$ and $H^* = \delta T$, exists if and only if*

$$\beta x + \frac{(z-\beta x)\tau\delta}{1-e^{-\tau\delta T}} \cdot \frac{\tau\delta + (\rho T(\rho+\tau\delta) - \tau\delta)\, e^{-(\rho+\tau\delta)T}}{(\rho+\tau\delta)^2} \leq \frac{c}{1-\ell} \cdot \left(1 + \frac{\rho}{\eta\,(1-e^{\tau\delta T})}\right). \tag{51}$$

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

The left side of (51) calculates the expected revenue θ that a seller would earn by offering an auction when no one else does ($\alpha = 0$). For this equilibrium to exist, the expected revenue must be lower than the expected cost of entering the market (the right side of (51)). We can consider such a deviation because buyers still wait until their deadline before purchasing via the posted-price listing, and are willing to bid their reservation price $b(s) = \beta x + (z - \beta x)e^{-\rho s}$ if given the chance.

Equation (51) indicates that auctions are not viable when expected costs are high, such as high production costs or listing fees, or long delays before closing (small η). In contrast, the posted-price market can compensate for these costs by keeping a low stock of sellers so that the item is sold very quickly. Auctions can also be undermined by weak competition among bidders producing low expected revenue, which occurs with a small flow of buyers (δ) or few of them paying attention (τ).

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

[48] 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 when (51) does not hold.

can occur, and that a dispersed equilibrium cannot occur when (51) 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.

C.2 Proofs

Proof of Proposition 3. By Proposition 1, equations (11) and (12) and $\phi(H^*) = 0$ 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 (16) and (18), respectively. It is apparent that $\sigma^* \geq 0$. To see that $\sigma^* < 1$, note that the equilibrium condition $\phi(H^*) = 0$ requires that $\alpha\left(1 - e^{-\tau H}\right) < \delta$. This also ensures that $P^* > 0$.

The profits stated in (45) and (46) are required by the third and second equilibrium conditions, respectively. From (13), profit solves as: $\Pi_p = \frac{\zeta z(1-\ell)}{\rho+\zeta}$, so for this to equal c, we require $\zeta^* = \frac{\rho c}{z(1-\ell)-c}$ as in equation (49). With this, (19) readily yields P^* as listed in (48).

The only remaining element regards expected auction profit. Equation (15) solves as: $\Pi_a = \frac{\eta\left(1-e^{-\tau H}\right)(1-\ell)\theta}{\eta\left(1-e^{-\tau H}\right)+\rho}$. By setting this equal to c and solving for θ, we obtain (20).

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

$$\theta = \frac{\lambda}{1-e^{-\lambda}}\left(e^{-\lambda}b(T) + \lambda\int_0^T b(s)F(s)F'(s)e^{-\lambda F(s)}ds\right). \tag{52}$$

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

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

where $\kappa \equiv \delta + \alpha e^{-\tau H}$ 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 $\alpha \to 0$, the solution to $\phi(H^*) = 0$ approaches $H^* = \delta T$. In the absence of auctions, the distribution of bidders is uniformly distributed across $[0, T]$ because 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: $b(s) = \beta x + (z - \beta x)e^{-\rho s}$.

For $\alpha^* = 0$ to be a market equilibrium, we need $\Pi_a^* \leq \Pi_p^*$. To prevent further entry, $\Pi_p^* = c$ is still required. From (15), a seller would earn $\Pi_a^* = \frac{\eta\left(1-e^{\tau\delta T}\right)(1-\ell)\theta}{\rho+\eta\left(1-e^{\tau\delta T}\right)}$ by offering an auction unexpectedly. Thus, the expected profit comparison simplifies to: $\theta \leq \frac{c}{1-\ell} \cdot \left(1 + \frac{\rho}{\eta\left(1-e^{\tau\delta T}\right)}\right)$. This is equivalent to (51), where the left-hand side is evaluated from (52):

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

Thus, if (51) 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 (51) fails to hold, then $\alpha^* = 0$ cannot be an equilibrium because some firms will earn greater profit by deviating and offering an auction.

To prove the last claim, first note that in a buyer equilibrium, $H \to 0$ 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 (51) is equivalent to $\Pi_a > 0$ for $\alpha = 0$. Because 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. □

D Extensions

D.1 Alternative Mechanisms: Physical Search, Bargaining, or Lotteries

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 search problems could be formulated, changing the discount mechanism in (3) while maintaining the deadlines embedded in the $-V'(s)$ term and the full price option z.

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

D.1.1 Physical Search

First, consider physical search for a homogeneous good where sellers post a price, but discovering these sellers is time consuming. At each encounter, the buyer learns a specific seller's price but has to purchase immediately or lose the opportunity. The buyer in state s formulates 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:

$$W(s) = 1 - G(s), \tag{53}$$

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

$$M(s) = \int_s^T b(t) dG(t). \tag{54}$$

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 $G(s)$).

We now consider physical search from the seller's perspective. A deeper discount results in lower revenue but a higher likelihood of sale because it will be acceptable to more buyers. 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:

$$\rho\Pi_a = \eta F(s) \left((1-\ell)b(s) - \Pi_a\right). \tag{55}$$

To obtain price dispersion, each targeted price $b(s)$ must yield the same expected profit Π_a. The equilibrium in this environment is closely related to the labor market model of Akın and Platt (2012).

D.1.2 Bargaining

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 dissolved, 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:

$$W(s) = 1 - G(s). \tag{56}$$

The expected payment would be:

$$M(s) = \int_s^T \big(\omega b(s) + (1-\omega)b(s')\big)\, dG(s'). \tag{57}$$

Now consider Nash bargaining from the seller's perspective. A 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 $\omega b(s) + (1-\omega)b(s')$ depends on the type of the buyer, so we integrate over all possibilities.

$$\rho \Pi_a(s') = \eta \left((1-\ell) \int_0^{s'} \big(\omega b(s) + (1-\omega)b(s')\big)\, dF(s) - F(s')\Pi_a(s') \right). \tag{58}$$

D.1.3 Lottery

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 other buyers collectively are distributed according to $G(k')$, then the probability of winning would be:

$$W(s) = \int_0^T \frac{k(s)}{k(s) + k'} dG(k'). \tag{59}$$

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

$$M(s) = pk(s). \tag{60}$$

A seller's revenue in a lottery setting 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:

$$\rho \Pi_a = \eta \left((1-\ell) \int_0^T pk' dG(k') - \Pi_a \right). \tag{61}$$

D.2 Endogenous Posted Price and Reserve Price

The model assumes that all posted-price sellers charge the same exogenous price z. If the model were to be expanded to allow each seller to endogenously choose her own posted price, there would still exist an equilibrium in which all sellers would choose the same z. Specifically, if buyers anticipate that all sellers charge the same posted 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 does not sell any faster but sacrifices some profit. Moreover, a seller who deviates by posting a higher price will always be rejected because the buyer anticipates that another seller can immediately be found who charges price z. Of course, other equilibria are certainly possible, posing an interesting avenue for future research.

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 if a single bidder arrives or if the second-highest bid is less than R. With this higher reserve price, the seller would close the auction without sale in these situations and would re-list the good, a strategy that has a present expected profit of Π_a. Because $\Pi_a = c$ in equilibrium, deviating to the reserve price R is certain to be unprofitable if $R < c$. 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 would make T endogenous, requiring $b^*(T^*) = c$ in equilibrium. All else would proceed as before.

Even with sellers using 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, imagine one seller has monopoly control of both markets. The optimal market design for this monopolist 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 the conditions for degeneracy from Proposition 4 are not satisfied). An individual seller always has an incentive to offer an auction if all other sellers offer posted-price listings: the product sells faster through auctions, even if at a slightly lower price.

D.3 Buyer and Seller Heterogeneity

The baseline model assumes ex-ante homogeneity of buyers and sellers. This focus is intentional in order to discipline the model and allow us to isolate the effect of consumer deadlines on repeated bidding, price dispersion, and sales channel decisions rather than confounding these effects with differences among the market participants. However, the model can accommodate certain types of heterogeneity among buyers or sellers with minimal impact on the overall behavior. For example, some sellers might have stronger preferences than others for posted-prices over auctions; this would determine which sellers would participate in each mechanism, though the overall mix would be determined by the marginal seller, as in the baseline model. The same would occur if some buyers were to have a stronger distaste for auction participation.

Another potential extension would be to allow buyers to differ in their raw consumption utility, which is particularly straightforward when $\beta = 0$ (all consumption utility is realized at the deadline).[49] Suppose x is a random variable drawn for each buyer, similar to the exogenously-given valuations in traditional auction models. If x is bounded below by z, all of the model's results carry through without modification, as bids are chosen relative to the posted price (which all bidders have as their common outside option), rather than relative to their idiosyncratic consumption utility.[50]

D.4 Endogenous participation

A final group of extensions endogenizes when buyers start or conclude their participation in the discount mechanism. First, suppose that a buyer incurs some cost while searching for auctions. This would lead her to postpone her search until closer to her deadline in an effort to avoid the search cost while the chances of winning are exceptionally low. Relative to our baseline model, this would be a simple extension that would effectively endogenize T; buyers would be aware of their need earlier, but search would really begin only once the expected utility from search is equal to the cost of search.

Second, consider a case where buyers must also search to find a posted-price listing. This is in contrast to the baseline model, in which a posted-price option is always readily available. If such search were required, some buyers would abandon the discount market prior to their deadline to increase their chances of securing the good in the posted-price market (depending on the penalty for missing the deadline). This would effectively endogenize participation at the end of the search spell. This extension, and the costly search extension discussed in the previous paragraph, would affect when discount search would begin or end (and must be solved for numerically), but bids would still rise during the search spell and sellers would still find it profitable to utilize both mechanisms.

An alternative adjustment to participation would be to introduce exogenous heterogeneity

[49] If $\beta > 0$, some of the utility x is immediately obtained on purchase, and becomes relevant in the bidding function. This disrupts analytic tractability of the equilibrium bidding function, but we have found that numeric solutions under this extension preserve the same qualitative features as the baseline model.

[50] The behavior is more nuanced if x can be less than z; in such a setting, some bidders would be worse off purchasing at the posted price, and extending the model in this case would require specifying the consequences of missing the deadline.

in the initial time-until-deadline T or attention given to discount opportunities τ. For the latter, a buyer might increase her attention $\tau(s)$ as her deadline approaches. Unlike the heterogeneity extensions in the preceding subsection, this type of heterogeneity would disrupt the analytic tractability of the solution; however, we have found that numerical solutions under this extension produce similar qualitative results to our baseline model.

At the same time, we note that *observed participation* already increases over the search duration in our baseline model, even though attention is assumed to be constant throughout the search. Song (2004) first noted that a buyer who arrives after the auction's current bid exceeds her reservation price will be precluded from submitting a bid and will remain unobserved. In our empirical application, we account for the feature of our model that buyers closer to their deadline have higher reservation prices, and increasing reservation prices also lead to a higher frequency of being observed. We use methods from Platt (2017) to explicitly account for unobserved participation in the structural estimation of the model, as described in Section 5.1 of the paper.

E Shipping Speeds and Closing Times

We now present two empirical patterns that provide strong ancillary evidence that buyers grow more time-sensitive over the duration of their search. First, after repeated losses, buyers are increasingly likely to participate in auctions where expedited shipping is available, consistent with the time sensitivity we model. Here we define fast shipping as any shipping option with guaranteed delivery within 96 hours. The overall fraction of buyers bidding in auctions with fast shipping available is 44%, and this fraction rises with the number of auctions a buyer has attempted. This can be seen in Figure A1.A. The horizontal axis indicates the total number of auctions a bidder participates in, and the vertical axis indicates the fraction of cases where the *last auction* the bidder participates in offers fast shipping. We find that those who have participated in more prior auctions gravitate toward fast shipping (roughly 2% more for each additional auction).

The choice of which auction to use is beyond the scope of the model, but we would expect that fast shipping would be most relevant to buyers within a week of their deadline. Of course, deadlines are not observed in the data, and so Figure A1.A proxies for closeness to the deadline by how many attempts a bidder has made. To give a sense of the magnitude of the effect, we use simulated data from the model to determine the relationship between bidder attempts and closeness to the deadline. This is reported in Figure A1.B, which shows the fraction of bidders who are in their last week at the time of their last bid.

Note that, in the data, participation in fast-shipping auctions is much more prevalent than would be suggested by the model if bidders only join such auctions in their last week. Yet fast-shipping participation still rises 10 percentage points from those who bid once to those who bid six times. In the model, the fraction who are in their last week grows 19 percentage

Figure A1: Shipping Speed

(A) Fraction in Fast Shipping Auction (B) Fraction in Last Week (Model)

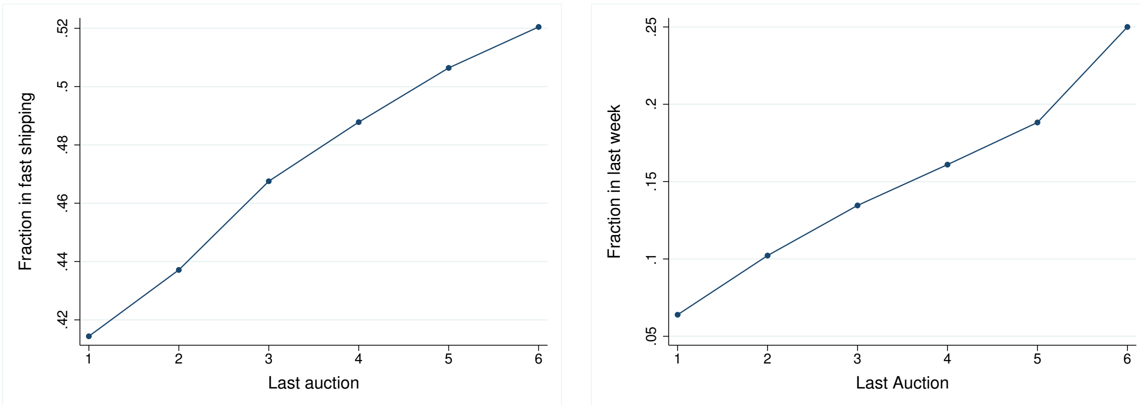

Notes: Panel (A) reports the fraction of bidders in the data participating in an auction with fast shipping during the last bid attempt on a product. Panel (B) reports the fraction of bidders in the simulated data who, at the time of their last bid attempt, are within a week of their deadline.

points. Thus, the participation in fast shipping rises about as much as the rise in last-week bidders. We see this as favorable evidence that buyers with longer auction sequences are feeling greater time pressures, though clearly this is not the only reason they participate in fast-shipping auctions.

Second, we find that as bidders move farther along in their search process they are increasingly likely to participate in auctions that are ending soon. Our main sample examines primarily bidders who participate just before the auction closes, so we broaden our analysis here to include non-serious bids. In this broader sample, a buyer's highest bid in a given auction is, on average, placed when there are 1.34 days remaining. Figure A2.A demonstrates that this number decreases steadily and significantly across auction attempts (with the average time until the auction closes falling by 2.43% per auction attempt), again consistent with growing time sensitivity during the search process. Hendricks and Sorensen (2018) report a similar fact in their data: high-value bidders tend to prefer auctions that end soon. While this preference toward soon-to-close auctions is not explicitly micro-founded in either model, deadlines provide one motivation: in the deadlines model, high-value bidders are precisely those who need the item sooner.

F Consumer Surplus and Demand

Online retail markets are a rich source of data about consumer demand. However, demand data has wildly different interpretations depending on the model in which it is analyzed. For example, if consumers grow increasingly time sensitive over the duration of their search, ignor-

Figure A2: Days Left in Auction Regression (A) and Derived Demand Curve (B)

(A) Days Left in Auction (B) Derived Demand Curve

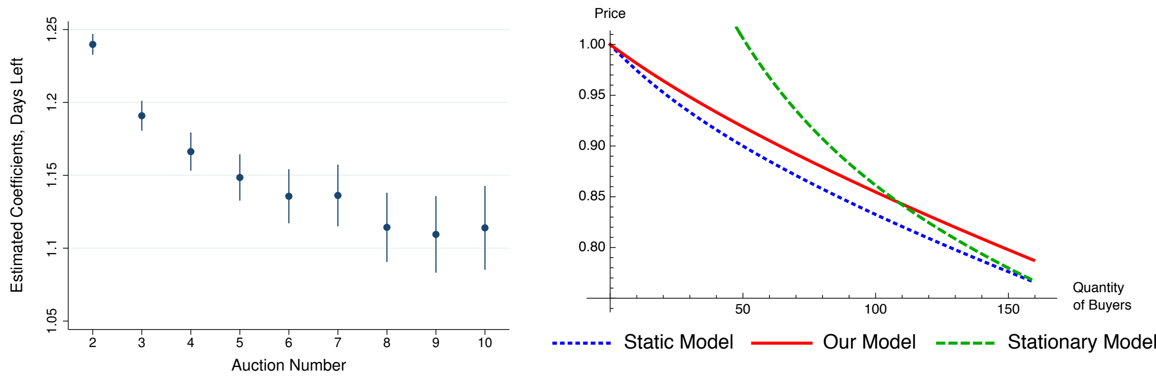

Notes: Panel (A) displays estimated coefficients for dummy variables for each auction number (*i.e.* where the auction appears in the sequence) from a regression of a dependent variable on these auction number dummies and on dummies for the length of auction sequence. The dependent variable is the number of days left in the auction when the bidder bid. Regressions are performed after removing outliers in the auction number variable (defined as the largest 1% of observations). 95% confidence intervals are displayed about each coefficient. Panel (B) reports the demand curve inferred from bids reported in Panel (C) of Figure 3 using our deadlines model (solid) vs. treating the data as though it came from a static model (dotted) or a stationary dynamic model (dashed). The dashed line is truncated, but would intersect the vertical axis at a price of 1.41.

ing this non-stationarity would lead to mis-measurement of demand and consumer surplus.[51] To demonstrate this, we consider two alternatives to our non-stationary dynamic search model: a static model and a stationary dynamic model. Buyers in the static model only make one purchase attempt, while the stationary dynamic model allows multiple attempts; but in both, buyer valuations are exogenously given and constant.

For the static model, assume that the valuation of bidder type s is denoted $x(s)$, 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, $x(s)$ still denotes the valuation of bidder type 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 con-

[51]Incorrect estimates of the demand curve could potentially distort calculations needed for profit maximization, price discrimination, regulation, and other applications. Moreover, individual-level estimates of willingness to pay are essential in providing individualized product recommendations, targeted advertising, and personalized pricing. One implication of consumer-specific deadlines is that firms engaged in personalized pricing based on consumer data (*e.g.* Kehoe et al. 2018) might benefit by including in their models a measure of a given consumer's observed search duration.

tinuation value of a bidder is:

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

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

$$x(s) \equiv b(s) + \frac{\tau\alpha}{\rho}\left(e^{-\lambda F(s)} b(s) - e^{-\lambda}\left(b(T) + e^{\lambda}\int_s^T b(t)\lambda e^{-\lambda F(t)} F'(t)\, dt\right)\right). \tag{62}$$

In the static model, buyers reveal their valuations in their single truthful bid, so the econometrician can estimate demand by inverting the empirical CDF of bids. By way of comparison, if bidding data were generated by our model, but the data is then used to estimate demand under a static model, we obtain the dotted line in Figure A2.B, in a parametric plot of $(H \cdot F(s), b(s))$.

However, in our paper's environment, 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 discount opportunities provide positive expected surplus. When observed bids are adjusted to determine the valuations,[52] it generates the true demand curve, depicted as the solid line in Figure A2.B. The static interpretation of data generated from a dynamic process will *underestimate* demand — on average by 1.4% of the retail price.

Of course, other dynamic models (Zeithammer, 2006; Said, 2011; Backus and Lewis, 2016; Bodoh-Creed et al., 2018; Hendricks and Sorensen, 2018) can make a similar critique because the option to participate in future discount opportunities reduces buyers' willingness to bid. However, these stationary dynamic models predict that the highest-valuation bidders have the greatest option value from search and thus shade their bids the most aggressively. This is not true in our model, where the highest-valuation bidders are about to abandon the discount mechanism and thus do not shade their bids. If bids were generated by deadline-motivated buyers but interpreted using a *stationary* dynamic model, it would *overstate* demand by 27.9% of the retail price (the dashed line in Figure A2.B). In the stationary model, low-valuation buyers are unlikely to win in current *or* future auctions and thus they are willing to pay nearly their full valuation. Meanwhile, high-valuations buyers are most likely to win in current and future auctions, so they shade their bids aggressively (by as much as 41%). In our non-stationary model, however, high-valuation buyers are closer to their deadline and hence shade *less* than low-valuation bidders.

[52] Here, we set $x = z$, which creates the smallest difference between the static model and ours.

G Including Non-Serious Bids: Data and Model Results

Our main sample includes bids submitted in the last hour of the auction and the two highest bids prior to that time ("serious" bids). This screens out extremely low bids that have no chance of winning and yet are never raised later in that auction.[53] Here we repeat the key analysis from the paper when these non-serious bids are included in the sample.[54]

In Figure A3, we replicate the data facts reported in Figures 1 and 2 of the paper. In Table A6, we report the parameter estimates obtained in this expanded sample compared to the main estimates from the paper. Figure A4 then replicates the comparison of model fit from the paper on key graphs where they are affected.

We note that including non-serious bids leads to more long sequences; thus we report sequences of length up to 10. We observe the same pattern of increasing average bids in the data among all bidders (Figure A3.A and Figure A3.B), those bidding on expensive items (Figure A3.D), experienced bidders (Figure A3.E), and inexperienced bidders (Figure A3.F). We also observe line sequences that rarely cross. However, including non-serious bids pulls down the average bid amount by almost 20 percentage points, leading to a gap between average bids in the data (in Figure A3.A) and the model equivalent (Figure A4.A), illustrating the better fit of the serious bids sample. In Figure A4.A, the model prediction shows the longest sequence line crossing the shorter lines, but this is due the small number of observations in the simulated data reaching eight auctions (literally only a single bidder participates in eight auctions in the simulated data). The rate of switching to posted prices (Figure A3.C) is essentially identical.

Some other comparisons to the data and model are similar even when non-serious bids are included, such as the distribution of sequence lengths (Figure A4.B) and the duration between bids (Figure A4.C). We find that, with non-serious bids included, the fitted model predicts fewer auction sales than are observed in the data, but the fit between the fraction of bidders who eventually win in the data and model is quite close (Figure A4.D).

Table A6 demonstrates that including non-serious bids has the largest impact in increasing the number of participants per auction (λ) and the flow of participants entering the market (δ). Changes in these fitted parameters then lead to slightly shorter implied time frame for search (T) than in the main model (decreasing from 4.3 to 2.5 months) and a (unrealistically large) estimate for the discount factor (ρ), which increases from 0.056 in the paper to 0.380 when non-serious bids are included. The reason the model yields this large estimate for ρ is as follows: the model rationalizes these extremely low-ball bids by treating these bidders as

[53] Note that dropping these *bids*, for the most part, drops particular *bidders* who do not appear to be ever bidding seriously. Only 16% of *bidders* are observed in the data having a serious bid in one auction and a non-serious bid in another auction. All other bidders place only serious bids or only non-serious bids. This fact, along with the fact that non-serious bids do not affect final prices, suggests that non-serious bidding is unlikely to be an important strategic or outcome-driving feature of the marketplace.

[54] An alternative way to expand the sample would be to lengthen the window for a bid to qualify as serious; not surprisingly, such an approach yields results a mixture of the paper results and those presented here.

Figure A3: Data Facts, Including Non-Serious Bids

(A) Bids Over Search Duration (cf. Fig 1.A)

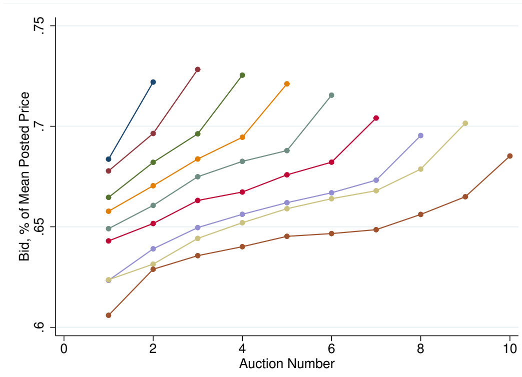

(B) Regression Results (cf. Fig 1.B)

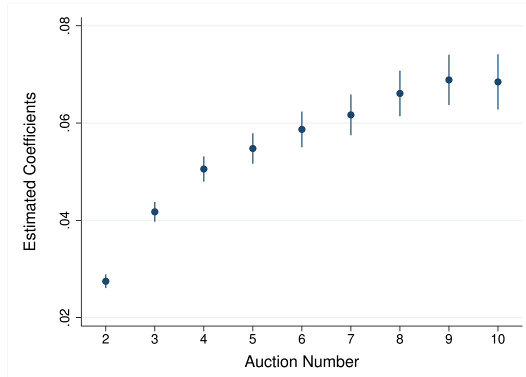

(C) Time To Posted-Price Purchase Since Last Losing Auction (cf. Fig 2.A)

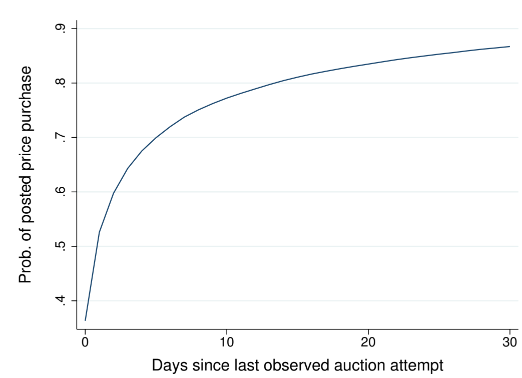

(D) Bids on Products With Prices $\geq$ \$100 (cf. Fig 2.B)

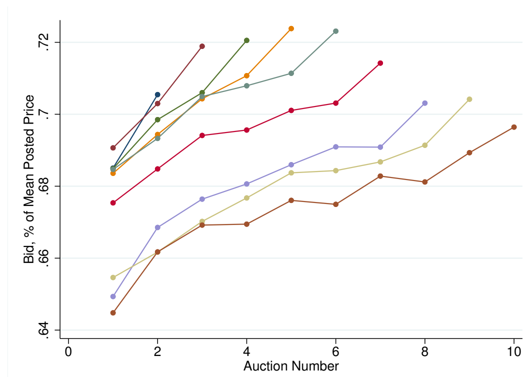

(E) Bidders With ≥ 50 Auctions (cf. Fig 2.C)

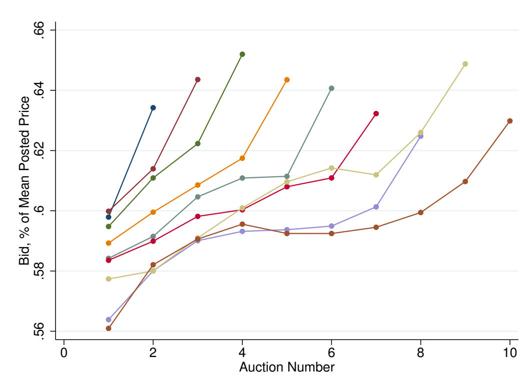

(F) Bidders With < 50 Auctions (cf. Fig 2.D)

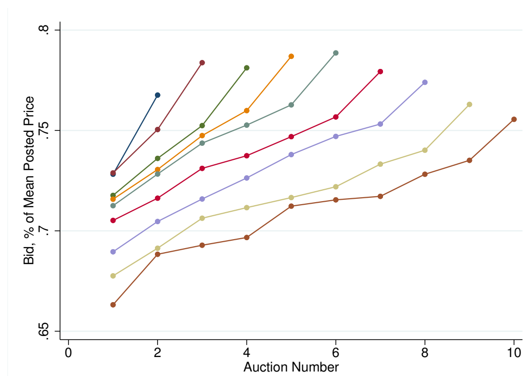

Notes: Figure displays the equivalents of Figures 1 and 2 with non-serious bids included.

Table A6: Data Moments and Parameter Values, Including Non-Serious Bids

	Observed in Data		Theoretical Equivalent	Fitted Parameter	
	Paper	With Non-Serious		Paper	With Non-serious
Bidders per completed auction	2.57	5.30	$\frac{\lambda \cdot P(\lambda)}{1-e^{-\lambda}}$	$\lambda = 3.01$ (0.020)	13.10 (0.243)
Completed auctions per month	12.76		$\alpha\left(1-e^{-\lambda}\right)$	$\alpha = 13.42$ (0.548)	12.76 (0.525)
Auctions a bidder is observed in per month	1.11	1.17	$\frac{\tau\alpha P(\lambda)}{1-e^{-\tau\alpha P(\lambda)}}$	$\tau = 0.019$ (0.00066)	0.064 (0.0024)
New bidders per month who never win	16.33	39.10	$(\delta-\alpha)\cdot\left(1-e^{-\tau\alpha T P(\lambda)}\right)$	$\delta = 41.46$ (2.56)	81.55 (3.45)
—	—		Eq. (10)	$T = 4.25$ (0.142)	2.54 (0.050)
Average revenue per completed auction	0.853		θ	$\rho = 0.056$ (0.0024)	0.380 (0.012)
Average listing fee paid	0.116		ℓ	$\ell = 0.116$ (0.0029)	
Average duration of an auction listing (months)	0.156		$1/\eta$	$\eta = 6.39$ (0.028)	
—	—		Eq. (20)	$c = 0.748$ (0.0036)	0.712 (0.0038)

Notes: Table displays the equivalent of Table 2 with non-serious bids included in the data sample moments.

Figure A4: Model Fit, Including Non-Serious Bids

(A) Bids over Search Duration (cf. Fig 3.A)

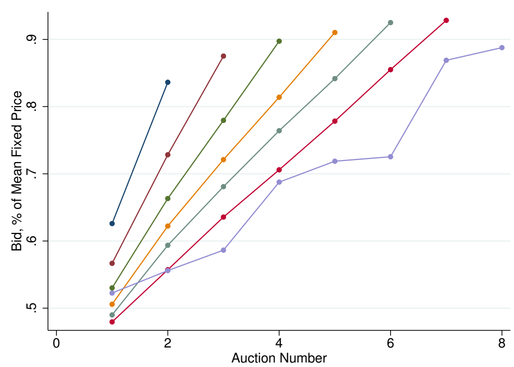

(B) Distribution of sequence lengths (cf. Fig 4.A)

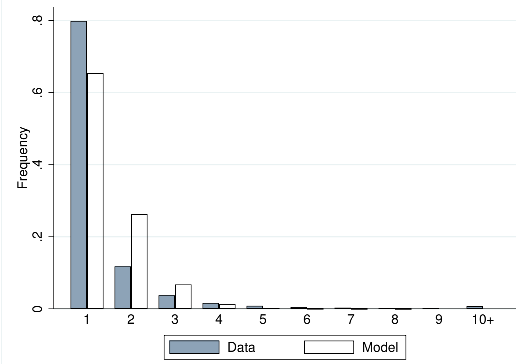

(C) Duration between bids, by sequence length (cf. Fig 4.B)

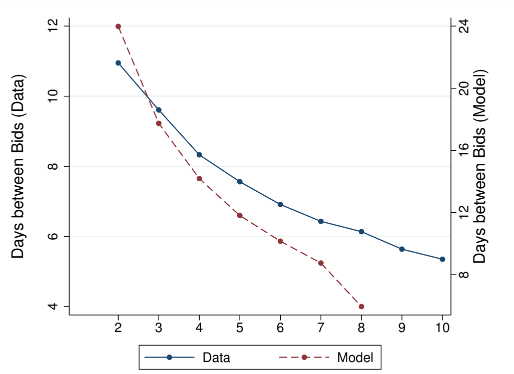

(D) Distribution of Fraction of Sales by Auction (cf. Fig 6.A)

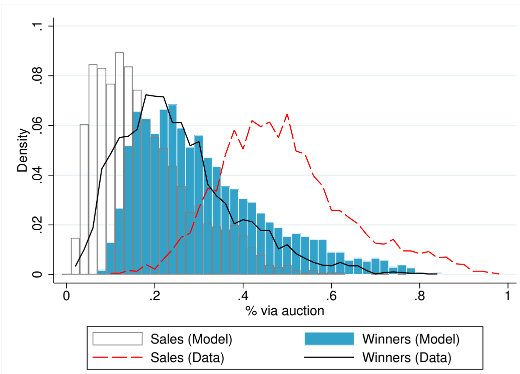

Notes: Figure displays the equivalents of Figures 3.A, Figure 4.A–B, and Figure 6.A with non-serious bids included in the data sample moments and with the model-fitting exercise performed using this expanded sample.

agents who will eventually be willing to pay the full price upon reaching their deadline, thus interrupting the large fraction of non-serious bids as though they (and all bids) must have been generated by agents who steeply discount the future.

H Differences Across Products in Behavior and Parameter Estimates

Not surprisingly, behavior in our data varies across products. We illustrate this here by examining the amount of repeat bidding and the rate of bid increase. For the average product, repeat bidders place 23% of bids, with an interquartile range of 15 percentage points across all products. The highest rates of repeat bidding occur with Computers/Tablets (26.3%) and DVDs/Movies (25.0%). We note that items with higher rates of repeat bidding typically have lower estimates for T. The rate of bid increase also varies across products. For the average product, bidders raise their bid on average by 1.9% per observed attempt, with an interquartile range of 2.5 percentage points. This rate is nearly the same across the major categories with the exception of Video Games (4.9%).

We now explore estimation of the model's parameters separately for each product. The estimates presented in the paper (with the exception of Figure 6.B) uses aggregate data moments, aggregated across all products, yielding a fit that is representative of the average product in the market. Here, instead, we estimate the model's parameters product-by-product, matching the data moments for a given product to the theoretical moment to obtain product-specific parameter estimates.

Table A7 summarizes the mean, standard deviation, and median of these product-level parameter estimates. The first column in Table A7 displays our main estimates for comparison. For the parameters λ, α, δ, η, ℓ and c, the mean product-level estimates are nearly the same as the main estimates. In contrast, the average τ, ρ, and T are somewhat larger. For these three parameters, the median products estimates are smaller (and, for τ and T, are in closer agreement to our main estimates). A few factors contribute to the distinctions between our main estimates and these product-by-product estimates.

First, some targeted moments are not normally distributed across products. For instance, the monthly flow of new buyers is highly skewed, with 75% of the products below the mean of 16, while the top 1% of products reach into triple digits. Attempted auctions per month are similarly skewed, though to a lesser degree. Second, the estimation procedure tends to add or exacerbate skewness in λ, due to the non-linearity of $P(\lambda)$.

Together, these factors lead to disproportionately skewed product-level estimated parameters. The aggregated targets are necessarily kept away from extremes, but any given product target could be an extreme, and such outliers have a large influence on the average of the product-level parameters. This skewness also explains why the median estimates are in closer

Table A7: Comparison of Parameter Estimates

		Product-Level Parameter Estimates			
	Main Parameter Estimates	Mean	Median	Std. Dev.	Coeff. Var.
λ	3.01	3.13	3.04	1.29	0.41
α	13.42	13.33	7.57	29.24	2.19
τ	0.019	0.033	0.021	0.038	1.15
δ	41.46	41.23	22.26	127.94	3.10
T	4.25	7.68	5.42	7.29	0.95
ρ	0.056	0.064	0.040	0.088	1.36
η	6.39	6.91	6.24	2.46	0.36
ℓ	0.12	0.11	0.10	0.17	1.48
c	0.75	0.73	0.76	0.14	0.19

Notes: Main parameter estimates come from Table 2 (which were estimated by fitting model moments to the average of product-level averages of data moments). The mean, median, and standard deviation columns display the mean, median, and standard deviation across products of parameters estimated separately for each product. Coefficient of variation displays the standard deviation over the mean.

agreement for some parameters.

We also note that 11% of individual products cannot fit the model under any parameters. Most of these misfits are due to cases where the average auction revenue is greater than the average posted price, which our model cannot rationalize. However, for about 2% of products, the lack of a solution is because bidders on those products are never observed bidding in more than one auction (which is the data moment used to identify τ).

While our data permits us to classify equivalent products together through an anonymized product id, it does not allow us to see *what* the product actually is. For example, we cannot tell whether a given product id corresponds to an X-box or a PlayStation. This limits our ability to consider whether particular parameters seem appropriate for a specific product. However, we can analyze heterogeneity using a broad category identifier and the average price level of the product.

In Figure A5, we show how estimates for T, ρ, and c vary across these classifications and within them. We focus on these three parameters because they are easily interpretable even beyond the eBay context (unlike, for example, λ). The box indicates the 25th, 50th, and 75th percentiles for products within a given classification, while the whiskers extend to the 5th and 95th percentile. The categories we display include at least 100 products from our sample, while the price ranges split our products into roughly four quartiles.

While some groups of products have systematically higher T (in the first row), such as toys or items under \$12, the variation within each group is very large, with overlapping confidence intervals for all groups. Similarly, estimates for ρ (in the second row) can be high for categories such as health, but all confidence intervals overlap. For the estimates of c (in the third row), note that a lower estimated cost c is equivalent to a higher percentage markup. We find that this markup appears to be higher among toys, movies, and health products, as well as lower-priced items.

Figure A5: Distribution of Product-Level Estimates

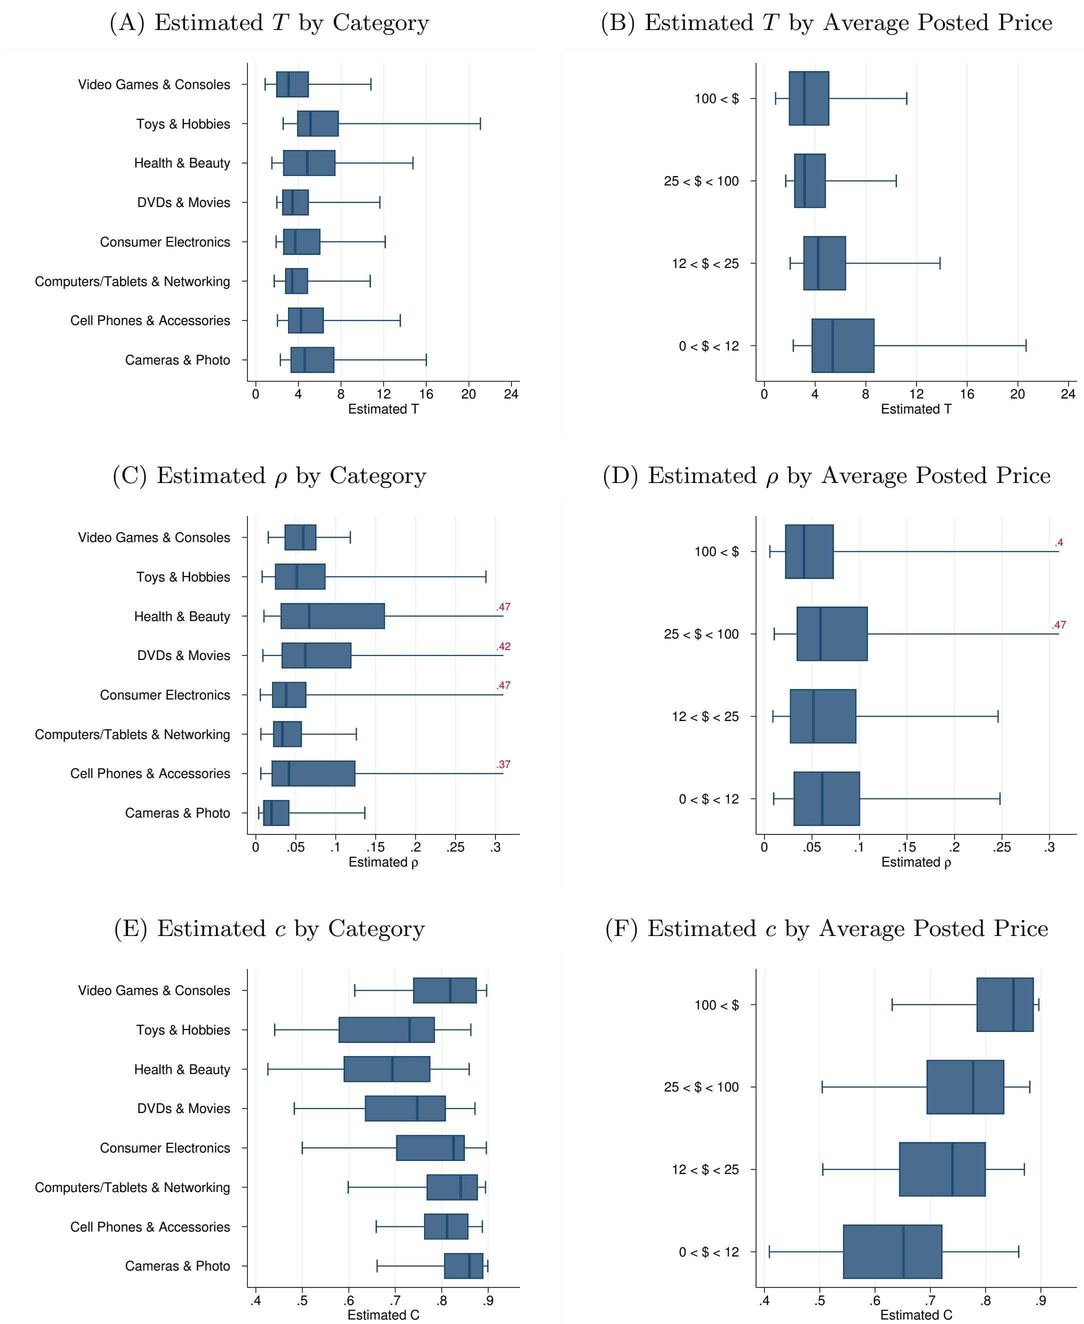

Notes: Panels (A), (C), and (E) display model parameters estimated separately product-by-product and then aggregated by product category for each category containing at least 100 products. In each panel, boxes indicate the 25th, 50th, and 75th percentiles of the parameter estimate for products within a given category, while the whiskers extend to the 5th and 95th percentile. For categories with large values for the 95th percentile, the value of the 95th percentile is shown in red type. Panels (B), (D), and (F) display similar results but aggregated by average posted price level of each product rather than by product category. Panels (A) and (B) display estimates of T, panels (C) and (D) display estimates of ρ, and panels (E) and (F) display estimates of c.

Table A8: eBay Fees Over Time

				2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015
Auction	Insertion fee (by starting or reserve price)		0.01-0.99	$0.30	$0.25	$0.20	$0.20	$0.15	$0.15	$0.10	$0.10	$0.10	$0.30	$0.30	$0.30
			1.00-9.99	$0.35	$0.35	$0.35	$0.35	$0.35	$0.35	$0.25	$0.25	$0.25	$0.30	$0.30	$0.30
			10.00-24.99	$0.60	$0.60	$0.60	$0.60	$0.55	$0.55	$0.50	$0.50	$0.50	$0.30	$0.30	$0.30
			25.00-49.99	$1.20	$1.20	$1.20	$1.20	$1.00	$1.00	$0.75	$0.75	$0.75	$0.30	$0.30	$0.30
			50.00-199.99	$2.40	$2.40	$2.40	$2.40	$2.00	$2.00	$1.00	$1.00	$1.00	$0.30	$0.30	$0.30
			200.00-499.99	$3.60	$3.60	$3.60	$3.60	$3.00	$3.00	$2.00	$2.00	$2.00	$0.30	$0.30	$0.30
			500.00+	$4.80	$4.80	$4.80	$4.80	$4.00	$4.00	$2.00	$2.00	$2.00	$0.30	$0.30	$0.30
	Final value fee (% of closing price, cumulative)		0.01-25.00	5.25	5.25	5.25	5.25	8.75	8.75	9.00	9.00	9.00	10.00	10.00	10.00
			25.01-1000.00	2.75	2.75	3.00	3.25	3.50	3.50	9.00	9.00	9.00	10.00	10.00	10.00
			1000+	1.50	1.50	1.50	1.50	1.50	1.50	9.00	9.00	9.00	10.00	10.00	10.00
			Maximum charge							$50.00	$100.00	$250.00	$250.00	$250.00	$750.00
Posted Price	Final value fee (% of posted price, cumulative)	Consumer Electronics	$0.01 - $50.00	same as Auction style	same as Auction style	same as Auction style	same as Auction style	8.00	8.00	8.00	7.00	7.00	same as Auction style	same as Auction style	same as Auction style
			$50.01 - $1,000.00					4.50	4.50	5.00	5.00	5.00			
			$1,000.00+					1.00	1.00	2.00	2.00	2.00			
		Computers & Networking	$0.01 - $50.00					6.00	6.00	8.00	7.00	7.00			
			$50.01 - $1,000.00					3.75	3.75	5.00	5.00	5.00			
			$1,000.00+					1.00	1.00	2.00	2.00	2.00			
		Clothing, Shoes & Accessories	$0.01 - $50.00					12.00	12.00	12.00	10.00	10.00			
			$50.01 - $1,000.00					9.00	9.00	9.00	8.00	8.00			
			$1,000.00+					2.00	2.00	2.00	2.00	2.00			
		Media	$0.01 - $50.00					15.00	15.00	15.00	13.00	13.00			
			$50.01 - $1,000.00					5.00	5.00	5.00	5.00	5.00			
			$1,000.00+					2.00	2.00	2.00	2.00	2.00			
		All Other Categories	$0.01 - $50.00					12.00	12.00	12.00	11.00	11.00			
			$50.01 - $1,000.00					6.00	6.00	6.00	6.00	6.00			
			$1,000.00+					2.00	2.00	2.00	2.00	2.00			
	Insertion fee (by posted price)	Media	$1.00+					$0.15	$0.15	$0.50	$0.50	$0.50	$0.05	$0.05	
		Other Categories	$1.00+					$0.35	$0.35	$0.50	$0.50	$0.50	$0.30	$0.30	
(1) Final Value Fee (%) at Average Price in Auction Sample ($97)				3.39	3.39	3.58	3.77	4.85	4.85	9.00	9.00	9.00	10.00	10.00	10.00
(2) Final Value Fee (%) at Median Price in Auction Sample ($31)				4.77	4.77	4.81	4.86	7.73	7.73	9.00	9.00	9.00	10.00	10.00	10.00
(3) Einav et al. (2018) Fraction Revenue from Auctions				0.83	0.79	0.75	0.71	0.62	0.49	0.49	0.49	0.45	0.40	0.35	0.30
(4) Backus et al. (2018) Fraction Revenue from Auctions				--	0.82	0.76	0.73	0.66	0.56	0.52	0.48	0.42	0.36	0.32	0.27

Notes: Fees come from archived eBay.com pages on Wayback Machine (one snapshot per year), accessed on October 18, 2019; the dates and URLs for each snapshot are found in Coey et al. (2020b). No final value fee charged if item not sold. Starting in 2011, final value includes shipping fee. First 50 listings per month have no insertion fees starting in 2011 for auctions and 2013 for posted prices. "Media" category nests Books, Music, DVDs & Movies, Video Games. "Consumer Electronics" category nests Consumer Electronics, Video Game Systems, Cameras & Photo. Some additional category-specific exceptions are omitted from table, as are other optional promotion or listing add-on fees. Insertion and final value fees after 2015 are relatively constant and are omitted from table. Rows (1) and (2) at bottom of table show commission based only on auction final value fees evaluated at the same price in different years: $97 for row (1) (the mean of auction price plus shipping in the paper sample) and $31 for (2) (the median). Rows (3) comes from Figure 1 of Einav et al. (2018) and row (4) from Figure 1 of Backus et al. (2018).