NBER WORKING PAPER SERIES

DISCOUNTS AND DEADLINES IN CONSUMER SEARCH

Dominic Coey
Bradley Larsen
Brennan Platt

Working Paper 22038
http://www.nber.org/papers/w22038

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2016, Revised September 2019

Previously circulated as "A Theory of Bidding Dynamics and Deadlines in Online Retail." We thank Mark Armstrong, Lanier Benkard, Liran Einav, 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, and the Marketplace Innovation Conference for helpful comments. Coey was an employee of eBay Research Labs while working on this project. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Discounts and Deadlines in Consumer Search
Dominic Coey, Bradley Larsen, and Brennan Platt
NBER Working Paper No. 22038
February 2016, Revised September 2019
JEL No. C73,D44,D83

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 a good reveal consumers' reservation prices. We find robust evidence supporting the theory, and demonstrate that ignoring buyer deadlines can distort estimates of market welfare, consumer demand, and underlying causes of market shifts.

Dominic Coey
Facebook
Core Data Science
coey@fb.com

Bradley Larsen
Department of Economics
Stanford University
579 Serra Mall
Stanford, CA 94305
and NBER
bjlarsen@stanford.edu

Brennan Platt
Department of Economics
Brigham Young University
130 Faculty Office Building
Provo, UT 84602-2363
brennan_platt@byu.edu

1 Introduction

Searching for the best price on a product requires time,[1] 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, buyers face opportunities to acquire the good through some discount channel, but each opportunity has some chance of failure (*e.g.* not finding the item in stock, failing to win an auction, or failing to reach a deal while haggling). Meanwhile, the good is always available at a known higher-priced outlet. In equilibrium, a consumer will steadily increase her reservation price as her deadline increases, eventually turning to the full-price outlet. The model also allows sellers to offer their good through either the discount channel or at full price. In equilibrium, sellers use both options, and 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, in detecting the underlying causes of market shifts, and in the interpretation of demand data.

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.[2] We thus view

[1]Blake et al. (2016) follow the web searches of individual consumers over a month, finding that they query for the same item on average 36 times on 11 distinct (often non-consecutive) days before purchasing it.

[2]In a similar vein, Genesove (1995) exploits failed offers in used-auto auctions to study a stationary partial

eBay as an excellent laboratory for studying time-sensitive search, which is likely to apply in other settings that are harder to measure. We demonstrate that the theory extends to other discount sales mechanisms, including lotteries, bargaining, or discount posted prices.

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 these facts are reasonably close to the magnitudes observed in the data.

For example, past losers tend to bid more in subsequent auctions — 2.0% more on average in the data, compared to 3.6% more in the fitted model. In the data, 40% of the auctions are won by the bidder who has been searching the longest (compared to 45.7% in the fitted model). In contrast, this would only occur 7% of the time if search length and bids were uncorrelated, as implicitly assumed in traditional search models. We also find that as bidders search longer, they gravitate toward listings that offer faster shipping and that end sooner. 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 within homogenous goods. After controlling for seller and product fixed effects, auction prices have an interquartile range equal to 13% of the average price (compared to 6% in the fitted model). We also see price dispersion between selling mechanisms, with auction prices averaging 15% lower than posted-price sales, which the model replicates exactly. 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 5.2).

Second, consumer search with deadlines rationalizes the coexistence of discount and non-

equilibrium search model, but cannot observe repeat bids over time as we do. Panel data of unsuccessful 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.

[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.

discount mechanisms for identical products. In the data, we see that auctions represent 53% of sales (compared to 60% in the fitted model). The literature on competing mechanisms (reviewed in Section 5.3) 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.

If ignored, the presence of deadlines can skew the evaluation of welfare and market data. We compute the social value of the discount sales channel, finding that buyer welfare is 5.6 percentage points higher (as a fraction of the retail price) with the discount option available than without, but platform revenue is 7.9 percentage points lower. If the latter is pure profit, then the discount channel's existence reduces total welfare. This potential negative welfare effect of discount channels would be missed by prior theories of competing mechanisms: when mechanism coexistence is justified by exogenous differences rather than by deadlines, the discount channel only increases social welfare.

The model also provides a micro-foundation for the decline in auction transactions relative to posted prices, recently documented in Einav et al. (2018). We find that this shift is largely driven by changing supply fundamentals, as sellers' costs have shifted from the time of listing towards the time of sale, improving both the profitability and welfare benefits of the posted-price mechanism.

Finally, demand data generated by consumers with deadlines will be misinterpreted under standard methods. Using simulated data from our fitted model, we show that a static model would understate demand by 5.3 percentage points, while a stationary dynamic model would overstate demand by 2.5 percentage points. These errors reflect 16 to 33% of the potential markup, and thus could lead to sizable errors in entry, exit, or pricing decisions.

2 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.[4]

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]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.[5] 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. Each consumer identified a recent purchase for which they 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 instance, 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

[5] 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.* Board and Skrzypacz, 2016; Mierendorff, 2016; Dilme and Li, 2018), where a monopolist has multiple units of a good that expire at a known deadline.

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 homogenous good. These observed actions in specific eBay markets are strongly consistent with the motives reported in the broad market surveyed here.

3 Buyers with Deadlines

We begin by modeling buyers' choices when faced with deadlines; we present evidence for this predicted behavior in Section 4. Seller behavior is treated as exogenous until Section 5, where we consider the market implications of deadlines and document evidence of seller behavior.

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*).[6] The buyer receives βx dollars of utility at the time of purchase, while $(1-\beta)x$ dollars of utility are realized at the deadline,[7] 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 $\left(\beta + (1-\beta)e^{-\rho s}\right)x$ minus the purchase price.

A 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 from participation by other commitments. Consistent with our empirical application to eBay auctions, we refer to a successful attempt to acquire the good through the discount sales channel as *winning*. The outcome of the purchase opportunity is resolved immediately: the winner pays and exits, while losers continue their search.

Alternatively, a buyer can obtain the good at any time through a non-discount option with posted price z. We assume throughout that $x \geq z$, so that buyers weakly benefit from purchasing via the posted-price option, and $z > \beta x$, so that buyers prefer to try for discount opportunities until the deadline is reached. We also treat the consumption value

[6] 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.

[7] The extreme of $\beta = 0$ indicates that the good is literally of no use until the date of the deadline, while $\beta = 1$ indicates that it starts providing the same flow of value regardless of when it is purchased. The intermediate case seems reasonable for many deadlines: for instance, a gift is not needed until the birthday, but the giver may enjoy some peace 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.

x 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).

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. The cumulative distribution of buyer types is represented by $F(s)$, while the total stock of buyers in the market is denoted by H. Both of these are endogenously determined and are commonly known by market participants.

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

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

For now, we assume that $b(s)$ is strictly decreasing in s (*i.e.* the buyer is willing to pay more as the deadline approaches); we later verify this in Proposition 2. A buyer's expected utility in state s can then be expressed in the following Hamilton-Jacobi-Bellman (HJB) equation:

$$\rho V(s) = -V'(s) + \tau\alpha\Big(\Pr(win|s)\left[\left(\beta + (1-\beta)e^{-\rho s}\right) x - V(s)\right] - E(payment|s)\Big). \tag{2}$$

where $\Pr(win|s)$ is the probability that a consumer of type s wins the discount and $E(payment|s)$ represents the expected payment of the buyer. In this continuous-time formulation, the left-hand side of (2) represents the flow of expected utility that a buyer with s units of time remaining receives each instant. The right hand side depicts potential changes in (net) utility times the rate at which those changes occur. 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 $-V'(s)$. When a discount opportunity arises and the individual participates in it—which occurs at a rate of $\tau\alpha$ per unit of time—the buyer wins with probability $\Pr(win|s)$. When this occurs, the term $\left(\beta + (1-\beta)e^{-\rho s}\right) x - V(s)$ depicts the change in utility due to winning.

The exact form of $\Pr(win|s)$ and $E(payment|s)$ depends on the specifics of the discount sales mechanism; Section 6 offers several alternatives. Here, we focus on second-price sealed-bid auctions: each participant submits a bid and immediately learns the auction outcome, with the highest bidder winning and paying the second-highest bid. As in a static second-price auction, buyers find it weakly dominant to bid their reservation price $b(s)$.[8] The auctioneer is

[8] 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.

assumed to open the bidding at the lowest buyer reservation price, $b(T)$, which is only relevant in the event that just one bidder participates in the auction.

The number of participants n in each auction is drawn from a Poisson distribution with mean λ. Indeed, $\lambda \equiv \tau H$ because each of the H buyers participate in a given auction with probability τ. Thus, a buyer in state s will win an auction with probability

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

To win the auction, the buyer must have the highest bid, which means all n other participants must have more than s units of time remaining.[9] This occurs with probability $(1 - F(s))^n$. The expected payment of the buyer is the average second-highest bid times the probability of winning and thus paying it, and is stated as

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

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

Buyers also have the option to purchase from the non-discount posted-price option at any time, receiving utility $\left(\beta + (1 - \beta)e^{-\rho s}\right)x - z$. However, a buyer in state s can obtain a present expected utility of $(x - z)e^{-\rho s}$ by waiting until $s = 0$ to buy, which is strictly preferred for all $s > 0$ because $z > \beta x$. That is, buyers in the model prefer to exhaust all discount opportunities if the posted price exceeds the immediate utility. The posted-price option is exercised if and only if $s = 0$, so the expected utility of a buyer who reaches her deadline is

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

3.1 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 remains constant over time. As buyers exit the market, they are replaced by new consumers; as one group of buyers get closer to their deadline, a proportional group follows behind.

[9] 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.

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.[10] 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[11] on $(0, T]$. Indeed, the population of buyers changes according to

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

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 $\Pr(win|s)$), thereby exiting the market.

Equation (6) defines the law of motion for the interior of the state space $s \in (0, T)$. The end points are defined by requiring $F(s)$ to be a continuous distribution:

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

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

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

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

[10] 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.

[11] 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.

3.2 Buyer Equilibrium

The preceding optimization by buyers constitutes a dynamic game. We define a *steady-state equilibrium*[12] 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. Bids b^* satisfy equations (1) through (5), taking F^* and τH^* as given.
2. The distribution F^* satisfies the steady-state equations (6) through (8).
3. The average mass of buyers in the market H^* satisfies steady-state equation (9).

The first requirement requires buyers to bid optimally; the next two require that buyers correctly anticipate the distribution of competitors, consistent with steady state.

We now characterize the unique equilibrium of this game. Our equilibrium requirements can be translated into two second-order differential equations regarding $F(s)$ and $b(s)$. The 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 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 $\kappa \equiv \delta + \alpha e^{-\tau H^*}$.

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

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

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.*

[12] When necessary for clarity, we will refer to this equilibrium as a *buyer equilibrium*, and we will refer to the augmented equilibrium derived in Section 5, which takes into account sellers' decisions, as a *market equilibrium*.

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

Proposition 2. *In equilibrium,* $b'(s) < 0$ *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.[13] Yet the fractional term also increases as s falls, which reflects the increasing probability that the buyer will reach her deadline without winning an auction. This interpretation is discussed further in an alternative formulation of (12) reported in Technical Appendix B.

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 can be obtained by implicit differentiation, as derived and discussed in Technical Appendix C. These comparative statics contribute to understanding changes in online retail markets documented by Einav et al. (2018), which we address further in Section 5.5.

4 Evidence of Deadlines

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.[14] By considering new-in-box products within a single platform, we ensure product consistency across listings and across mechanisms.

[13] 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.

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

Table 1: Descriptive Statistics

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

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

Notes: Table displays descriptive statistics for 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 product-level standard deviations. Normalized revenue is computed by first dividing auction price by product-level average of posted-price sales. Panel B reports average and standard deviation, taken across all products, of the number transactions of a given product, the number of unique sellers selling a given product, and the number of unique buyers bidding in an auction for a given product or purchasing the product through a posted-price listing.

Our data consist of auctions and posted-price sales on eBay.com for the year from October 1st, 2013, to September 30th, 2014. 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 popular items, principally electronics, media, or health/beauty products.

Table 1 presents descriptive statistics for the listings that end in a sale. In all, there are 1,055,309 sales of 3,663 distinct products, split roughly evenly between auctions and posted prices. Means and standard deviations in panel A are computed by taking the mean and standard deviation over all transactions of a given product and then taking the mean of these product-level means and standard deviations. The average number of bidders per sold auction is 5.3. In our model, we treat the posted price for a given product as fixed (z), whereas in reality posted prices can vary from listing to listing just as in auctions. However, consistent with the model, the average selling price (calculated inclusive of shipping fees) is higher with posted prices than auctions ($107 versus $97). To adjust for differences across products, we follow Einav et al. (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.

In what follows, we document striking patterns in the data and discuss their connection to the qualitative predictions of our deadlines model. We also use the data to fit the parameters of the model in order to compare the magnitude of these patterns in the data versus the model. This procedure is relatively straightforward, as each parameter corresponds directly to a transformation of sample moments. For example, the average number of bidders per auction identifies the model parameter λ; the average number of auction sales per month identifies α; and the number of auctions the average bidder engages in each month identifies τ. Details of these transformations and those used to identify each of the other model parameters are described in Technical Appendix E. We normalize a unit of time to equal one month. We compute these data moments by taking averages across product-level averages; thus, this exercise should be interpreted as fitting the model for the average product.

Figure 1: Bids Over Search Duration — Data

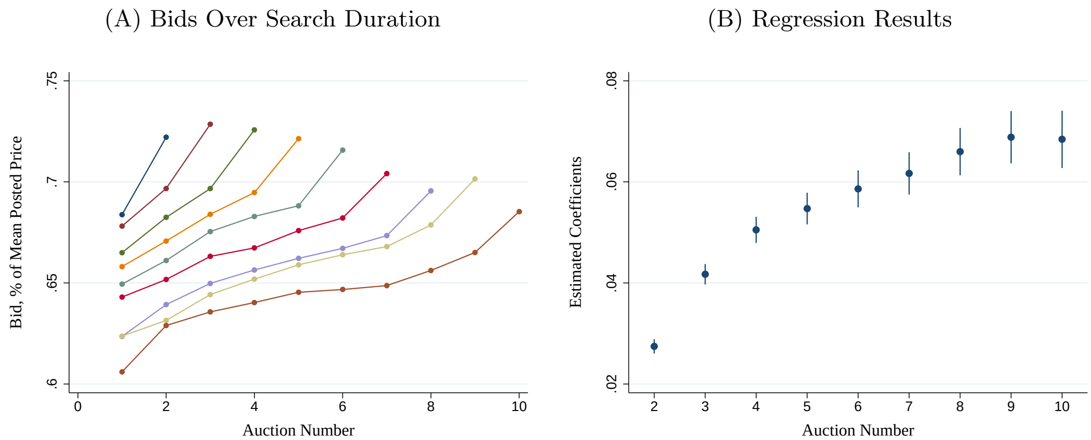

Notes: In Panel (A), a given line with n points corresponds to bidders who bid in n auctions total for a given product without winning in the first $n-1$ auctions. Horizontal axis represents auction number within the sequence (from 1 to n) and vertical axis represents the average normalized bid. Panel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

This estimation procedure explicitly takes into account the selection introduced through using listings that end in a sale and the unobserved-bidders problem raised in Song (2004). The estimation approach is also robust to certain types of model misspecification, such as some bidders preferring to bid in only one auction. The sample moments we exploit in estimation do not correspond to the key data patterns we seek to explain, 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, as we discuss below.

4.2 Bids Over Duration of Search

In our data, we follow each bidder across multiple auctions of the same product. We find that the average willingness to pay tends to increase from one auction to the next. To compute this, we keep each bidder's highest bid in each auction in which she participates, yielding a sample of 4,077,410 bids. We order these bids in a chronological sequence for each bidder and product pair, ending when the bidder either wins an auction or does not participate in any more auctions in our sample. This yields 2,728,258 unique bidding-sequence and product pairs. We then compute the average of the normalized bids, separately for each sequence length and each step within the sequence.

Figure 1 displays the resulting trend across repeated auction participation. 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 ten—the average bid steadily rises over time from the first to the last auction in the sequence.[15] Panel (B) frames the same trend in terms of a regression result. Averaging across all sequence lengths and auction numbers, the bid increases by a statistically significant amount of 2.0 percentage points in each successive auction.

By way of comparison, Figure 2.A provides the same analysis on data simulated from the model under fitted parameters. Again, we see the average bid steadily rises within each sequence. The underlying bidding strategy (as a function of time remaining) is depicted in the solid line of Figure 2.B. Initially (for s near T), the price path is more or less linear, but as the deadline approaches, greater curvature is introduced. On average, a buyer increases her bid at a rate of 4.3 percentage points per month. Since the average bidder participates in 1.18 auctions per month, this translates to an increase of 3.6 percentage points between each auction of a given product — a reasonable fit relative the 2.0 percentage point gain seen in the data (Figure 1.B), considering that our parameter estimation does not exploit any details about bids over time.

Notably, in both Figures 1.A and 2.A, the line for one sequence length never crosses the line of another sequence length. This occurs in the model (Figure 2.A) because 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 (Figure 1.A).[16]

The dotted line in Figure 2.B indicates the utility that the buyer gets by purchasing at time s (simulated from the model at 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. This also highlights the fact that the increasing bids pattern predicted by the model is not solely due to impatience, but also reflects the reduced option value of future auction opportunities.

It is worth noting that canonical search models (*e.g.* Stigler, 1961; Diamond, 1987; Stahl,

[15] 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, however. To see this, note that Panel (A) of Figure 1 shows that even before the final bid in a sequence, bids tend to increase, and Panel (D) of Figure 5 presents the same trends while excluding winning bids. Nor is it due to selection in the product mix across the auction number variable, as the sequences are constructed at the bidder-by-product level, so conditional on sequence length the product mix is constant across auction number.

[16] Note also in Figure 2.A that the average final bid in each sequence is increasing in the auction number, according to the model. Figure 1.A displays this same pattern for shorter sequences, but the final bid appears to decrease for longer search sequences. In a regression, we find that this apparent difference between the data and model is actually minor: any decrease in final bids across sequence lengths is not statistically significant when comparing across sequences of six or fewer auctions, a sample which represents 96% of the observations underlying Figure 1.A.

Figure 2: Bids Over Search Duration — Model

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

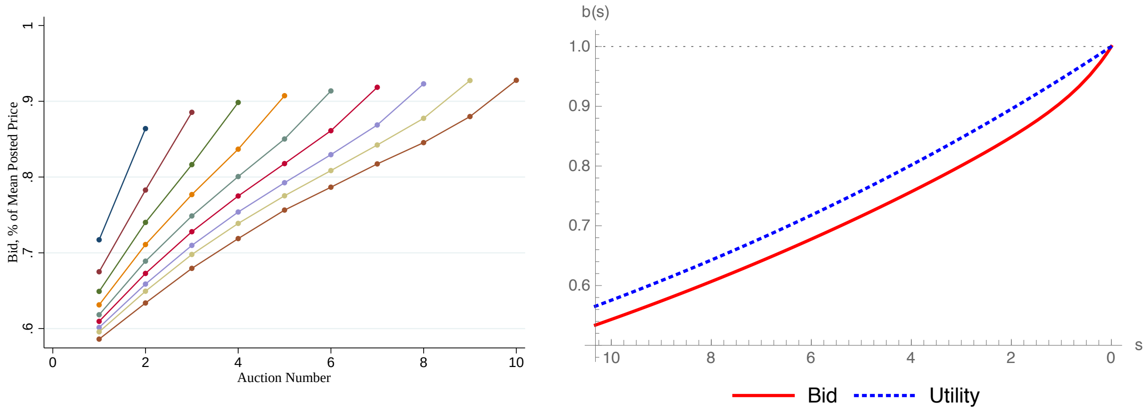

Notes: Panel (A) reproduces Figure 1.A from simulated data under the fitted parameters. Likewise, Panel (B) reports bids (solid line) and utility (dotted line) as a function of time remaining s. Since $z = 1$, these may be read as percentages relative to the retail price.

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.[17] 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 repeated sequential auctions (Zeithammer, 2006; Said, 2011; Hendricks et al., 2012; Backus and Lewis, 2016; Bodoh-Creed et al., 2016; 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.

As highlighted in Section 2, the deadlines we model are inherently idiosyncratic and unobservable—such deadlines cannot be directly observed in any data of which we are aware (other than the consumer survey we introduced in Section 2). Deadlines common to large groups of people (holidays) are indeed observable, but in eBay data we found no clear evidence that consumers treat these as a common deadline: aggregate price trend leading up

[17] 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). The spatial search model of taxi drivers in Buchholz (2018) also yields non-stationarity, and includes a deadline element with drivers approaching the end of their shift.

to such holidays (*e.g.* Xbox or other children's toys for Christmas, jewelry for Valentines, or costumes for Halloween) exhibit no clear pattern, but *individual-specific* price trends and other patterns consistent with the model jump out clearly from the data, as documented here.[18] In our consumer survey, where we see a self-reported measure of consumers' individual deadlines, we do find evidence of increasing reservation prices as the deadline approaches; these results are discussed in Technical Appendix A.

It is tempting to attempt to isolate specific types of products in the data for which deadlines play a role and contrast them to other products for which this is not the case; however, we were unable to find any product for which idiosyncratic deadlines may not play a role: everything from pet food to beauty products to consumer electronics to clothing could potentially be a product for which a consumer searched under a deadline, and in our survey data consumers report having deadlines for a wide range of product categories. 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-deadlines-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 since the bidder's first observed bid) is frequently the winner, occurring in 40% of auctions. Under the fitted parameters, the model predicts a very similar frequency, 45.7% of the time, and this moment was not exploited in fitting the model's parameters.[19] 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 would be drastically lower, given by $\frac{1}{\lambda} = 7.7\%$, since such orthogonality would make each bidder equally likely to win regardless of her time spent searching so far.[20] Moreover, elapsed time and likelihood of winning would be *inversely* correlated if valuations were constant over time, since high valuation buyers would win shortly after entering the market while low valuations bidders would require many repeated attempts to get lucky.

Second, patterns among auction losers who switch to a posted-price purchase are also consistent with the model. We do not observe all posted-price purchases of these auction losers (many could take place on other platforms, such as Amazon), but among auction losers

[18]As highlighted in Section 2, even common holidays can generate idiosyncratic deadlines given individual preferences and anxieties for having the good by a certain date. Our data contains a full year of transactions, and thus it includes each of these holidays, and exhibits clear evidence in favor of the idiosyncratic deadlines model.

[19]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, we compare the winning frequency for bidders with the longest *observed* time since first bid.

[20]Our estimated value for λ, shown in Table A6, is $\lambda = 13.10$, and thus $\frac{1}{\lambda} = 7.7$.

Figure 3: Sales Rates and Switching Rates

(A) Cumulative Fraction Sold by Days Since Listing

(B) Time To Posted-Price Purchase Since Last Losing Auction

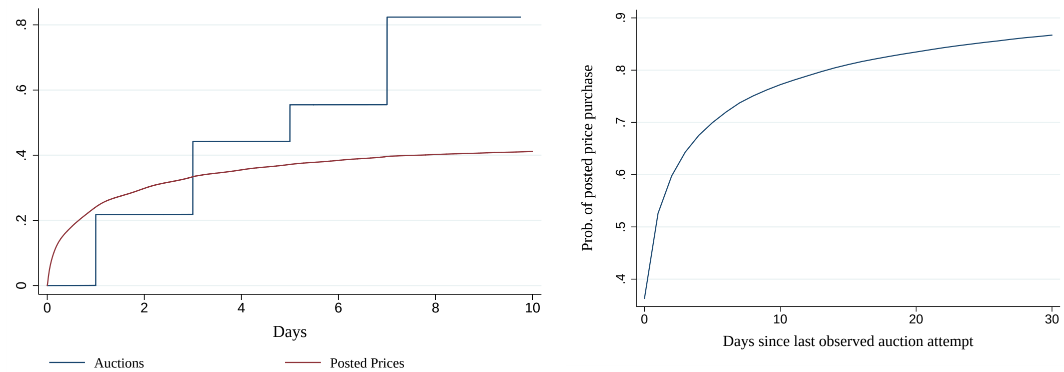

Notes: Panel (A) displays the cumulative fraction of listings sold (vertical axis) against the number of days since the listing was posted (horizontal axis) for auctions and posted-price listings. Panel (B) displays cumulative density of the time difference between the last observed auction attempt and 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.

who eventually purchase from a posted-price listing observed in our eBay sample, nearly all such buyers do so within a very short time of their last observed auction attempt: 52% do so within one day of their last losing attempt, 70% within 5 days, 77% 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 3.B. 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.

4.4 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. The overall fraction of buyers bidding in auctions with fast shipping available is 45%, and this fraction increases on average by 0.6 percentage points per auction attempt. This is shown in Figure 4.A.

Second, we find that as bidders move further along in their search process they are increasingly likely to participate in auctions that are ending soon. A buyer's highest bid in a given auction is, on average, placed when there are 1.33 days remaining. Figure 4.B demonstrates that this number decreases steadily and significantly across auction attempts (with the aver-

Figure 4: Shipping Speed and Closing Time

(A) Shipping Speed Regression (B) Days Left in Auction Regression

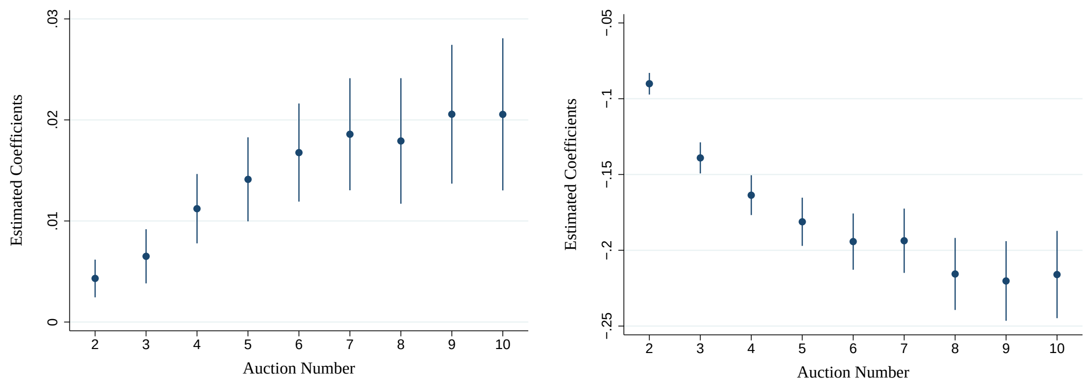

Notes: Each panel displays estimated coefficients for dummy variables for each auction number (*i.e.* where the auction appears in the sequence) from a regression of a dependent variable on these auction number dummies and on dummies for the length of auction sequence. In Panel (A), the dependent variable is a fast-shipping dummy (an indicator for whether the listing offered a shipping option guaranteed to arrive within 96 hours) and in Panel (B), 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. The displayed coefficients are relative to a regression constant of 0.43 in panel (A) and 1.34 in panel (B).

age 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.

4.5 Bidder Learning and Alternative Explanations

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

[21] We note that learning does not necessarily imply increases in bids across auctions. In Jeitschko (1998), bidders can learn their opponents types from their bids in the first auction, but in equilibrium, they reach the same expected price in the second auction. The model in Iyer et al. (2014) generates bids that rise on average, but there learning occurs only for the auction winner, who needs to experience the good to refine her information about its value. Learning stories can also generate *decreasing* reservation prices, as in De los Santos et al. (2017).

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, 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. Yet the same pattern of increasing bids appears among experienced bidders, as shown in panel (A) of Figure 5 for bidders who participated in at least 50 auctions prior to the current auction and in panel (B) for bidders who, over the past year, participated at least 10 auctions in the same product grouping as the reference auction.[22] Third, learning by participation is more costly with expensive products, due to the danger of bidding too high and winning when initially uninformed. We would expect buyers to be more cautious with such products and use alternative methods of learning. Yet panel (C) of Figure 5 shows the same increasing bid pattern for products with a median price over \$100.[23] The shipping speed and closing time patterns documented in Section 4.4 serve as further evidence that is consistent with the increasing impatience we model and that has no clear connection to a bidder learning story.

We emphasize here that we do not attempt to entirely rule out the possibility of bidder learning or other alternative explanations for some of the patterns we observe. However, the appeal of our model of time-sensitive buyers is that it provides a single, unified explanation of a number of facts together. For example, one alternative explanation for the increase in bids at the end of bidding sequences in Figure 1 is that, from one auction to the next, bidders receive random shocks to their valuations and that the increase at the end of the sequence is caused by bidders winning and exiting after a positive shock. However, a story of random

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

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

valuation shocks would fail to explain the pattern of increasing bids *prior* to the final auction in the sequence (as shown in panel (D) of Figure 5). While these alternative explanations could play a role, the bulk of the evidence also seems to indicate a role for time sensitivity.

5 Market Implications of 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 sketch additional implications for the market as a whole, and provide evidence of this behavior in our data. We fully develop the seller's side of the model and the full market equilibrium in Technical Appendix D. Here we summarize the model's key insights, all of which naturally follow from buyer deadlines and are relevant in all plausible variations of our seller assumptions.

First, sellers face a trade off between selling quickly (via the discount mechanism) or selling at full price (via posted prices). Discount sellers transact immediately due to the large number of potential participants (H) awaiting discount opportunities. Effectively, sellers are on the short side of the discount market, able to quickly find trading partners.

The full-price market is nearly the opposite: a large number of potential sellers hold their inventory, awaiting the arrival of buyers who have reached their deadline. These buyers therefore have no problem finding the item at full price, but the sellers may wait a considerable time before successfully transacting. This lowers the present value of the eventual profit from the sale. Moreover, sellers may have to incur some costs before the sale, such as *insertion costs* (production or other fees paid at the time of listing) and *holding costs* (expenses required throughout a listing).[24]

We allow for free entry by homogenous sellers in either market, which drives expected profits in both markets to zero. If profits in the auction segment were positive, more sellers would flow into the auction market and auctions would occur more frequently. Buyers would then be more likely to win a discount, so they would bid less and auction revenues would fall. Similarly, if profits in the posted-price market were positive, more sellers would flow into that market, driving up the time required for any given seller to find a buyer, thus reducing expected profits. These forces balance in equilibrium, as demonstrated in Technical Appendix D.

This market equilibrium 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 we should see considerable dispersion in prices both within the discount mechanism and between the mechanisms. These results are even more stark due to our focus on homogenous products (both in the model and in the data):

[24] Examples include warehouse rental space to store inventory, fees paid to any intermediating platform (such as the fees paid to eBay.com in the case of our empirical application), and costs of monitoring the listing, answering buyer questions, or paying employees to maintain a showroom.

despite being identical new-in-box products, transaction prices vary widely and sellers use multiple sales mechanisms. The market equilibrium also enables us to understand market evolution, identifying the underlying forces driving a recent trend toward more posted-price sales.

5.1 Selling Time

As predicted in the model, transactions in our data are typically completed faster through the discount (auction) mechanism relative to 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 3.A plots the cumulative fraction of listings sold against the number of days after listing the item for sale, for auctions and posted-price listings. Auctions are approximately equally likely to sell as posted-price listings within a day (21% vs 24%), and are over twice as likely to sell within 10 days (83% vs 41%). The model predicts that 88% of auctions and 45% of posted-price listings will sell in 10 days, and the latter takes 11 times longer to close than auctions, on average.

5.2 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 interquartile range of the normalized second-highest bid across auctions is 32 percentage points.[25] 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 resulting residuals have an interquartile range of 13 percentage points, or 5 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 discussed in Section 4.2.

Figure 6.A highlights the fitted model's predicted dispersion within the discount mechanism. The dashed line reports the relative likelihood of a given quantity being placed as a bid.[26] This density only gradually declines (especially at lower prices) because the probability of winning and exiting is small early on. This creates wide dispersion in bids across a range

[25]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.

[26]This plots $F'\left(b^{-1}(p)\right) \cdot (b^{-1})'(p)$, obtained via a parametric plot since $b^{-1}(p)$ cannot be analytically derived.

equal to 47% of the posted price. The dotted line in Figure 6.A reports the distribution of highest bids in each auction, which is heavily concentrated towards higher prices, even though bidders with those valuations are somewhat scarce. We find an average of $\lambda = 13$ participants per auction, and thus the highest bids come from very near the top of the bid distribution. Even so, the transaction occurs at the second-highest price, whose density is depicted with the solid line in Figure 6.A. Despite the uniform nature of the auctioned goods, closing prices are significantly dispersed, with the model predicting an interquartile range of 5.9 percentage points — a quantity closely aligned with what we measure in the data after controlling for seller and product fixed effects.

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 (Salop and Stiglitz, 1977; Stahl, 1989) or valuations (Diamond, 1987) or the frequency with which the item will be needed in the future (Sorensen, 2000). 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.[27] It is only after randomly arriving to the market that buyers differ ex-post, leading to a continuum of dispersed prices. 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.

5.3 Coexistence of Auctions and Posted-Price Sales

Our data indicates that a mixture of both mechanisms are simultaneously used for frequently-listed homogeneous goods. In our sample, 53% of transactions occurred through auctions. In our fitted model, sellers use a mixed strategy, transacting through auctions at a similar, though somewhat higher, rate of 60%. This is also true from the buyers' perspective. In a cohort of δ buyers who enter together, $H \cdot F'(0)$ buyers will never win an auction; under the estimated parameters, this means that $(\delta - H \cdot F'(0))/\delta = 60\%$ of buyers successfully purchase through auctions.

Of course, products are only included in our sample if at least 25 transactions occurred under both mechanisms. Thus, to document this coexistence more broadly, we turn to a larger sample that includes all products sold at least 50 times in our sample period regardless of listing method, yielding 12,368 products. For each product, we compute the fraction of

[27] Directed search models (surveyed in Wright et al., 2017) 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.

listings sold via auction. The histogram in Panel (B) of Figure 6 indicates the distribution of this fraction across products. Under 1% of products are sold exclusively by auction or by posted price. It is far more common to see a nontrivial fraction of sales through both formats: over 90% of products have between 10% and 90% of sales by auction. Averaging across products in this broader sample, the average rate of auction use is 47%.

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.

5.4 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. We illustrate this with a 10% increase in eBay's listing fee. Evaluating the fitted model with this increased fee yields a decrease in the number of auctions held per month by 0.9%, a decrease in consumers' value from searching, and an increase in their reservation prices (though expected revenue rises a mere 0.2%). This decreases the stock of posted-price sellers by 7.8%, shortening the average wait for the remaining posted-price sellers to find a buyer. Surprisingly, the net effect is that a given buyer is 1.4% *more* likely to make their purchase through posted prices. Sellers collectively end up paying 3.1% more, not the full 10% increase, due to the reduction of waiting time in the posted-price market.

This response illustrates a potential hazard of ignoring deadlines in market design: if buyer valuations 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 valuations and changes buyer behavior.

5.5 Market Evolution Over Time

Work by Einav et al. (2018) documents a recent decline in the popularity of online auctions relative to posted prices. Our model provides a fully micro-founded framework for studying the underlying structural changes associated with this decline. We begin by compiling datasets for each year from September 2010 through October 2015 using the same sample restrictions

Table 2: Welfare Estimates Over Time

	2010	2011	2012	2013	2014
Zero-Profit Platform Case					
Auc+PP	0.063	0.056	0.050	0.056	0.061
PP only	0	0	0	0	0
Costless-Operation Platform Case					
Auc+PP	0.112	0.105	0.107	0.121	0.120
PP only	0.106	0.106	0.124	0.144	0.141

Notes: Table displays estimates, using data from 2010–2014, of welfare under the dispersed (auctions and posted prices) equilibrium versus a degenerate (posted prices only) equilibrium, depending on whether platform fees exactly cover their costs or are pure profit. Data was constructed using same sample restrictions as in main sample. Each year label corresponds to one year of data beginning in Oct. of that year and continuing through Sep. of the following year (*i.e.* the 2013 column corresponds to the primary data sample of the paper). Units for welfare measures are fraction of retail (posted) price. Standard errors, from 200 bootstrap replications at the product level, are omitted to save space; they are all 0.001 or less.

as in our main data sample and fit the model separately in each year.[28] Over this time period, we find that the flow of buyers into the market (δ) and the rate at which auctions are offered (α) have decreased by nearly half, although average auction revenue has held fairly steady at approximately 85% of the retail price.

Einav et al. (2018) suggest that, from 2003–2009, the primary cause for the decline of auctions was a shift in buyer preferences, with buyers experiencing more hassle costs from auctions. Our model is more explicit about the underlying cause of buyers' time sensitivity, and is thus well suited to evaluate this potential cause of the shift. In the period we study (2010–2015), we find that the decline is instead explained by differences on the *supply* side of the market: over time, eBay sellers are incurring more of their costs at the time of purchase (perhaps due to the rise of just-in-time inventory practices) rather than at the time of the listing. We find that auction participants in the post-2010 period have not become more time sensitive, as the search time T has grown longer and buyers are more patient (smaller ρ). Indeed, the remaining buyers in the market appear to be paying greater attention (τ) to each auction.[29] Even so, this could be driven by a changing composition of buyers: it may be the case that eBay users who remain on the auction platform over time are the less-time-sensitive

[28]Descriptive statistics for each year are reported in Table A8 of the Technical Appendix, while Table A9 summarizes the data moments used to estimate our parameters.

[29]An additional distinction between our data and that of Einav et al. (2018) is that their data contains a unknown mixture of new and used items, whereas we restrict our data to new items only; eBay's internal categorization of new vs. used items is only available from 2010 onward.

buyers.

Beyond simply understanding why these market shifts have occurred, our model can also address whether the changes are socially optimal. On their face, discount channels duplicate full-price retailers and could encourage consumers to purchase items far in advance of when they are needed. On the other hand, posted-price channels require sellers to incur insertion and holding costs during a long wait for customers. We evaluate these competing inefficiencies in Technical Appendix F.1, finding that the answer depends on parameter values and on the level of platform (eBay) profits. Table 2 displays structural estimates of total welfare under two polar scenarios: eBay makes zero profits (*i.e.* fees paid to eBay only cover operating costs), or eBay makes all profits (*i.e.* operating the platform is costless). For these results, we treat $x = z$, so consumer surplus in a posted-price-only platform is normalized to zero. The units for welfare estimates in Table 2 are fractions of the retail price.

In the case where the platform exactly covers its costs, sellers and eBay earn zero profits, so total welfare equals consumer surplus. In this setting, the discount channel always benefits consumers by offering a chance at a lower price. Table 2 indicates that welfare has stayed roughly constant over the years at about 6% of the retail price. In the other extreme of a costless platform, profits are included in total welfare, which has improved over our time period. Welfare in the full market (with auctions and posted prices available) has increased by 7% (from 0.112 to 0.120 of the retail price) as sellers have been able to delay more of their costs. However, if the market had been restricted to posted-prices, total welfare would have grown even more, by 33% (from 0.106 to 0.141). Thus, if the platform is costless to operate, auctions have shifted over the past several years from efficient (increasing welfare by 6% over posted prices alone) to inefficient (reducing welfare by 15% relative to posted prices alone). Of course, the actual welfare effects over time likely lie somewhere in between these two extremes cases. Further details are provided in Technical Appendix F.1.

These welfare consequences stand out compared to the models reviewed in Section 5.3. The models where only one mechanism is used in equilibrium cannot consider such a counterfactual, but among those can sustain multiple mechanisms by relying on exogenous differences, Hammond (2013) and Bauner (2015) evaluate a similar counterfactual of shutting down the auction market. As with the costless platform case in our model, these studies find that sellers prefer to only have a fixed price market. However, assuming risk neutral buyers as we do, these studies find that total welfare is doubled when both markets are allowed to operate. This is sensible in a model where buyers and sellers sort themselves into different markets based entirely on ex-ante heterogeneity in preferences (valuations and costs): in such a model, depriving agents of their preferred market is likely to harm efficiency. In contrast, when it is the presence of deadlines that sustains coexisting mechanisms, this effect will be dampened or even reversed because buyers change their *type* over time. In our deadlines model, shutting down the auction mechanism leads buyers to delay their purchases until reaching their deadline, but by then, the fixed price channel *is* their preferred mechanism, thus reducing the harm of losing the discount channel relative to the prior models.

6 Extensions

6.1 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 buyer's search problems could be formulated, changing the discount mechanism in (2) while maintaining the deadlines embedded in the $-V'(s)$ term and the full price option z. Corresponding seller's problems are presented in Technical Appendix D.7.

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:

$$\Pr(win|s) = 1 - G(s), \tag{13}$$

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

$$E[payment|s] = \int_s^T b(t)dG(t). \tag{14}$$

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 $G(s)$). The equilibrium in this environment is closely related to the labor market model of Akın and Platt (2012).

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:

$$\Pr(win|s) = 1 - G(s). \tag{15}$$

The expected payment would be:

$$E[payment|s] = \int_s^T \left(\omega b(s) + (1-\omega)b(s')\right) dG(s'). \tag{16}$$

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:

$$\Pr(win|s) = \int_0^T \frac{k(s)}{k(s)+k'} dG(k'). \tag{17}$$

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

$$E[payment|s] = pk(s). \tag{18}$$

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

6.2 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, since 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.

The model also assumes that auction sellers always set their reserve price equal to $b(T)$, the lowest bid any buyer might make in equilibrium. This too can be relaxed, allowing the seller to chose a reserve price. In this augmented model, the unique optimal reserve price will be $b(T)$ so long as $b(T)$ exceeds the seller's costs, and is equal to the seller's cost otherwise. We prove this result in Technical Appendix D.8.

6.3 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).[30] 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.[31]

Another dimension of heterogeneity that we see in the data, but that is not accommodated for in our baseline model, is that half the bids come from buyers who only participate in one auction. A simple extension of the model can accommodate this by introducing a type of buyer who is unwilling to bid in more than one auction, along with an additional parameter (to be estimated) that represents the mass of such buyers. One could perfectly fit the distribution of the number of bidders in such a model, but with little value added to the model's insights. This is because our measurements of data patterns and our estimation of model parameters are conditioned on repeat bidding rather than relying on one-time bidders. It is precisely the repeat bidding behavior that the model seeks to explain, and our measurements and estimation steps are robust to model misspecification in terms of the number of one-time bidders. We discuss this further in Technical Appendix E.1.

6.4 Endogenous participation

A final group of extensions endogenize 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.

[30] 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.

[31] 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.

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 in the initial 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 setting, buyers closer to their deadline have higher reservation prices; thus, 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 Technical Appendix E.

6.5 Demand Estimation

Finally, we consider a potential use of our model in estimating consumer demand. While auction data is frequently used to this end, it has wildly different interpretations depending on the model in which it is analyzed. Under the traditional static auction model, bids precisely reveal the underlying utility of the bidder. In our setting, however, bids are shaded down from the buyer's true utility. As a consequence, if bids were generated by deadline-motivated buyers but interpreted using the static model, it would understate demand — by 5.3% of the retail price under our fitted parameters. This is a non-trivial distortion when compared to the potential markup of $z - c = 16\%$; firms could be off by one third of their potential markup!

Of course, other dynamic models (Zeithammer, 2006; Said, 2011; Backus and Lewis, 2016; Bodoh-Creed et al., 2016; Hendricks and Sorensen, 2018) can make a similar critique, since 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 2.5% of the retail price). A detailed comparison across models is found in Technical Appendix F.2.

Incorrect estimates of the demand curve could easily 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. While these policies are beyond the scope of the current paper, our model suggests that willingness to pay will systematically change over the duration of search, and that taking into account a consumer's current cumulative search duration could allow sellers to better target and tailor their prices for the individual consumer.[32]

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, since 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, deadlines, valuations of goods, seller costs, etc.

In our empirical application we document a variety of reduced-form findings consistent with the time sensitivity we model. In particular, buyers offer more in each successive attempt to win a discount, are more likely to win, and are more likely to pursue items with fast shipping and imminent closing times. 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 ignoring buyer deadlines in consumer search can

[32]For example, firms engaged in personalized pricing based on Big Data (*e.g.* Kehoe et al. 2018) could benefit by simply including in their models a measure of a given consumer's observed search duration.

lead to miscalculations of welfare and pricing, and can miss important insights in the design of platform markets and the analysis of how these markets evolve over time.

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.

Argyle, Bronson, Taylor Nadauld, and Christopher Palmer, "Real effects of search frictions in consumer credit markets," 2017. MIT Sloan Working Paper 5242-17.

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

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.

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?," 2016. Working paper, UC Berkeley.

Buchholz, Nicholas, "Spatial Equilibrium, Search Frictions and Efficient Regulation in the Taxi Industry," 2018. Working paper, Princeton University.

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.

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.

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," 2018. Forthcoming, *Review of Economic Studies*.

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.

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.

Salop, Steven and Joseph Stiglitz, "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion," *Review of Economic Studies*, 1977, *44*, 493–510.

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

Sorensen, Alan T, "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs," *Journal of Political Economy*, 2000, *108* (4), 833–850.

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: A Guided Tour," September 2017. NBER Working Paper 23884.

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

Appendix: Proofs

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

$$F''(s) = \alpha\tau F'(s)e^{-\tau H F(s)}. \tag{19}$$

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

$$F(s) = \frac{1}{\tau H}\ln\left(\frac{\alpha\tau - e^{\tau Hk(s+m)}}{\tau Hk}\right). \tag{20}$$

The constants are determined by our two boundary conditions. Applying (8), we obtain $m = \frac{1}{\tau Hk}\ln\left(\alpha\tau - \tau Hke^{\tau H}\right) - T$. The other boundary condition, (7), requires that k satisfy:

$$\alpha\tau\left(1 - e^{-\tau HTk}\right) - \tau Hk\left(1 - e^{\tau H(1-Tk)}\right) = 0. \tag{21}$$

From (9), we know that $HF'(T) = \delta$, and using the solution for F in (20), this yields:

$$k = \frac{\delta + \alpha e^{-\tau H}}{H}. \tag{22}$$

When we substitute for m and k in (20), we obtain the equilibrium solution for F^* depicted in (11). Also, (22) is used to replace k in the boundary condition in (21), 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, $\phi(0) = \delta\left(1 - e^{-\tau(\alpha+\delta)T}\right) > 0$. Since 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 positive:

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

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 (3) and (4). The first sum is similar to that in (6). For the second, we first change the order of operation, to evaluate the sum inside the integral. This is permissible by the monotone convergence theorem, because $F(s)$ is monotone and $\sum \frac{e^{-\lambda}\lambda^n}{n!}b(t)n(1-F(t))^{n-1}$ converges uniformly on $t \in [0, T]$. After evaluating both sums, we obtain:

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

Next, by taking the derivative of $b(s) = \left(\beta + (1-\beta)e^{-\rho s}\right)x - V(s)$ in (1), we obtain

$b'(s) = -\rho(1-\beta)xe^{-\rho s} - V'(s)$. We use these two equations to substitute for $V(s)$ and $V'(s)$, obtaining:

$$(\rho + \alpha\tau e^{-\lambda F(s)})b(s) + b'(s) = \rho\beta x + \alpha\tau\left(e^{-\lambda}b(T) + \int_s^T \lambda e^{-\lambda F(t)}b(t)F'(t)dt\right). \tag{23}$$

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

$$(\rho + \alpha\tau e^{-\lambda F(s)})b'(s) + b''(s) = 0. \tag{24}$$

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:

$$b(s) = a_1 \cdot \left(\frac{\delta e^{\tau H^* - \tau\kappa T}}{\rho} + \frac{\alpha e^{-\tau\kappa s}}{\rho + \tau\kappa}\right) e^{-s\rho} + a_2. \tag{25}$$

One constant of integration is pinned down by (23). We substitute for $b(s)$ in (23) using (25), and solve for a_2. This can be done at any $s \in [0, T]$ with equivalent results, but is least complicated at $s = T$ since the integral disappears: $(\rho + \alpha\tau e^{-\lambda F(T)})b(T) + b'(T) = \rho\beta x + \alpha\tau e^{-\lambda}b(T)$. After substituting $b(T)$, $b'(T)$, and $F(T)$, solving for a_2 yields:

$$a_2 = \beta x + a_1 \frac{\alpha\tau\kappa}{\rho(\rho + \tau\kappa)} e^{-(\rho+\tau\kappa)T}. \tag{26}$$

The other constant of integration is determined by boundary condition (5). If we translate this in terms of $b(s)$ as we did for the interior of the HJB equation, we get $b(0) = z$. We then substitute for $b(0)$ using (25) evaluated at 0, and substitute for a_2 using (26), then solve for a_1:

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

If the solutions for a_1 and a_2 are both substituted into (25), one obtains (12) with minor simplification. □

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

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

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

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

Again, each parenthetical term is strictly positive. □

Figure 5: Experienced Bidders, Expensive Products, and Removing Winning Bids

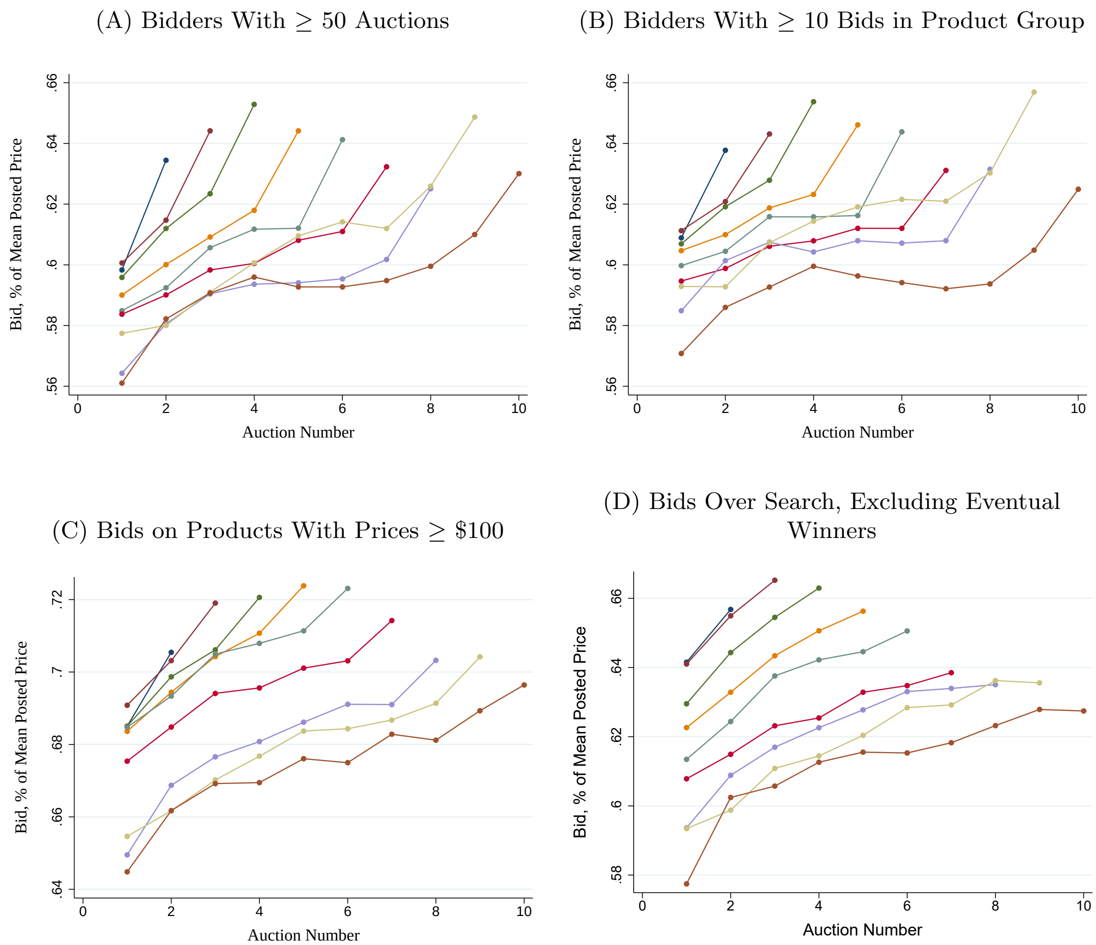

Notes: Each panel reports average bids over duration of search as in Panel (A) of Figure 1. Panel (A) limits to bidders who have bid in at least 50 auctions; Panel (B) limits to bidders who have bid in at least 10 auctions for products in the same product grouping in the past year prior to the current auction; Panel (C) limits to products with average transaction price ≥ \$100; Panel (D) uses the full sample but removes bidders who eventually win.

Figure 6: Price and Mechanism Dispersion

(A) Distribution of Bids and Prices — Model

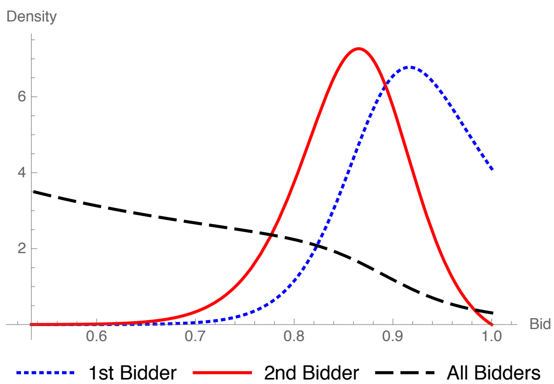

(B) Distribution of % of Sales by Auction — Data

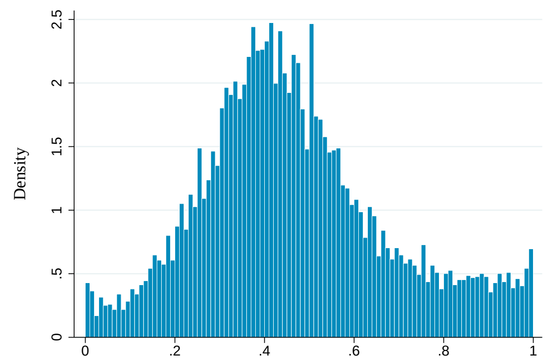

Notes: Panel (A) reports the equilibrium density of the highest bid in an auction (dotted), the second highest bid (solid), and all bids (dashed), derived from model under fitted parameters. Panel (B) shows a histogram at the product level of the fraction of listings sold by auction (rather than posted price) for a given product. The sample used to generate this figure is a superset of our main sample, containing the 12,368 distinct products with at least 50 transactions observed in the sample period without regards to listing method.

For Online Publication Only

Technical Appendix to "Discounts and Deadlines in Consumer Search"

Dominic Coey   Bradley Larsen   Brennan C. Platt[33]

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. 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.

[33]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

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?
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

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
—	Span > 20 days	(%)	69	72	78	73
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
Q4b	Searched eBay	(%)	31	28	25	28
Q4c	Searched Google	(%)	24	25	27	25
Q4a	Searched Amazon	(%)	74	73	59	69
Q7a-b	For a special event	(%)	36	38	23	32
Q7a-c	Needed more over time	(%)	65	66	64	65
Q7d	Knew a high-price option	(%)	43	47	50	47
Q8a	Needed as soon as possible	(%)	40	44	53	46
Q8b	Needed for specific use (not urgent)	(%)	45	42	38	42
Q9b	Willing to pay more in future	(%)	66	63	64	64
Q12a	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
Purchase Price	(mean)	15,613	398	56	375	63	92	884
	(sd)	14,213	476	98	6194	109	455	4,299
% saved	(mean)	14	27	35	29	33	32	30
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
# of Physical Searches	(mean)	2.7	1.8	1.3	1.7	4.1	1.6	2.1
# of Online Searches	(mean)	6.5	4.5	3.4	4.2	3.4	3.8	4.1
Searched eBay	(%)	13	29	37	21	25	33	28
Searched Google	(%)	7.7	29	27	23	26	25	25
Searched Amazon	(%)	17	73	84	66	65	71	69
For a special event	(%)	1.9	30	47	25	43	34	32
Needed more over time	(%)	42	66	65	63	65	67	65
Knew a high-price option	(%)	56	46	45	45	51	45	47
Needed as soon as possible	(%)	54	52	34	44	38	48	46
Needed for specific use (not urgent)	(%)	35	36	50	40	50	42	42
Willing to pay more in future	(%)	62	64	57	66	64	66	64
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).

greater savings. The estimate is quite noisy initially, but we find both the point estimate and its precision increase as we narrow the sample to those whose reported motives most closely fit 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: 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. ** indicates significance at the 95% level.

B Winning Probabilities and the Bid Function

For easier interpretation, the equilibrium bidding function from (12) can be expressed as follows:

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

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

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

C Comparative Statics

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 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

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	$b^*(T)$	–	**	–	+	–

Note: $**$ 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.

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 since 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 further away; there, the additional participants override the lower initial bid, driving up expected revenue.

Section D.6 reports comparative statics in the market equilibrium, which is particularly interesting for questions of market design such as adjusting the timing or size of costs to sellers.

C.1 Derivation of Comparative Statics Results

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:

$$\frac{\partial \phi}{\partial \lambda} = -\alpha e^{-\lambda} - (\tau T \alpha + e^{\lambda})\delta e^{-\tau T \kappa} < 0, \tag{27}$$

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:

$$b(T) = z e^{-\rho T} \cdot \frac{\kappa\left(\tau\kappa + \rho\right) e^{\lambda^*}}{\tau\kappa\left(\delta e^{\lambda^*} + \alpha e^{-\rho T}\right) + \rho\left(\delta e^{\lambda^*} + \alpha e^{\tau T \kappa}\right)}. \tag{28}$$

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:

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

C.1.1 Auction Rate, α

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

$$\begin{aligned} \frac{\partial \phi}{\partial \alpha} &= -1 + e^{-\lambda} + \tau T \delta e^{-\tau T \kappa} &(30)\\ &= -1 + e^{-\lambda}\left(1 + \left(\frac{\delta + \alpha e^{-\lambda} - \alpha}{\delta + \alpha e^{-\lambda}}\right)\ln\left(\frac{\delta e^{\lambda}}{\delta + \alpha e^{-\lambda} - \alpha}\right)\right). &(31) \end{aligned}$$

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:

$$\ln\left(\frac{\delta e^{\lambda}}{\delta + \alpha e^{-\lambda} - \alpha}\right) - \left(e^{\lambda} - 1\right)\frac{\delta + \alpha e^{-\lambda}}{\delta + \alpha e^{-\lambda} - \alpha} \leq 0 \tag{32}$$

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:

$$-\frac{\left(e^{\lambda} - 1\right)\left(\alpha + \delta e^{\lambda}\right)\left(2\alpha + e^{\lambda}(\delta - \alpha)\right)}{\left(\alpha + (\delta - \alpha)e^{\lambda}\right)^2} \tag{33}$$

Each parenthetical term is strictly positive for all $\lambda > 0$, so the left-hand side of (32) is strictly decreasing in λ. Thus, (32) 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,

$$\frac{\partial \lambda}{\partial \alpha} = -\frac{1 - (1 + \tau T(\kappa - \alpha))e^{-\lambda}}{\kappa - \alpha + (1 + \tau T(\kappa - \alpha))\alpha e^{-\lambda}}. \tag{34}$$

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

$$\frac{\partial F'(0)}{\partial \alpha} = e^{-\lambda} - 1 + \alpha \cdot \frac{\partial \lambda}{\partial \alpha}. \tag{35}$$

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.

C.1.2 Attention, τ

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

$$\frac{\partial \phi}{\partial \tau} = \delta\kappa T e^{-\tau T\kappa} > 0. \tag{36}$$

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

$$\frac{\partial \lambda}{\partial \tau} = \frac{\delta \kappa T e^{\lambda}}{\alpha e^{\tau T \kappa} + \delta e^{\lambda}\left(e^{\lambda} + \alpha \tau T\right)}. \tag{37}$$

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:

$$\frac{\partial F'(0)}{\partial \tau} = \frac{\partial F'(0)}{\partial \lambda} \cdot \frac{\partial \lambda}{\partial \tau} = -\alpha e^{-\lambda} \cdot \frac{\partial \lambda}{\partial \tau} \tag{38}$$

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

$$\frac{\partial b(T)}{\partial \tau} = \frac{z \alpha e^{\lambda} \psi}{(\kappa - \alpha)\left(\tau \alpha + (\tau(\kappa - \alpha) + \rho) e^{T(\rho + \tau \kappa)}\right)^{2}\left(\alpha + (\kappa - \alpha)\left(\tau \alpha T + e^{\lambda}\right)\right)}. \tag{39}$$

where

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

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

To verify the sufficient conditions listed under Table 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, since $\kappa > \alpha$ and $1 > e^{-\lambda}$, then $\delta \kappa \tau^2 T > \alpha \tau e^{-\lambda}$. Thus, the sufficient condition $\tau(\kappa - \alpha) > \rho$ ensures that $\delta \kappa \tau^2 T > \frac{\alpha e^{-\lambda} \rho}{\kappa - \alpha}$.

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

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

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

C.1.3 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:

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

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

C.1.4 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

$$\frac{\partial b(T)}{\partial \beta} = x \cdot \frac{(1 - e^{-\rho T})\delta e^{\lambda^*}\left(\tau\kappa + \rho\right) + \rho\alpha\left(e^{\tau\kappa T} - e^{-\rho T}\right)}{\tau\kappa\left(\delta e^{\lambda^*} + \alpha e^{-\rho T}\right) + \rho\left(\delta e^{\lambda^*} + \alpha e^{\tau\kappa T}\right)} > 0. \tag{40}$$

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

C.1.5 Deadline, T

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

$$\frac{\partial \phi}{\partial T} = \delta\kappa\tau e^{\lambda^*} e^{-\tau T \kappa}, \tag{41}$$

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,

$$\frac{\partial \lambda}{\partial T} = \frac{\delta\tau\kappa}{\delta\left(1 + \tau T \alpha e^{-\lambda^*}\right) + \alpha e^{\tau T \kappa - 2\lambda^*}} \tag{42}$$

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:

$$\frac{\partial F'(0)}{\partial T} = \frac{\partial F'(0)}{\partial \lambda} \cdot \frac{\partial \lambda}{\partial T} = -\alpha e^{-\lambda} \cdot \frac{\partial \lambda}{\partial T} \tag{43}$$

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:

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

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

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

and

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

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

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

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

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

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

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

D Market Equilibrium Model

This section provides the details of optimization by sellers in an environment in which buyers face deadlines. We consider a continuum of sellers producing an identical good,[34] and allow free entry to offer their product via either mechanism. 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.

Each seller can produce one unit of the good at a marginal cost of $c < z$, with fraction γ of this cost incurred at the time the good is sold (the *completion cost*), and $1-\gamma$ incurred when the seller first enters the market (the *insertion cost*).[35] For either selling format, sellers also pay a *holding* fee of ℓ each unit of time from when the seller enters the market to when the good is sold. If the intermediating platform charges insertion fees (incurred at the start of listing), these could be included in $(1-\gamma)c$. Similarly, if an intermediary charges a commission fee (a percentage at the time of sale), this can be included in γc. Upon entry, each seller must decide whether to join the discount or the posted-price market.

[34]While we refer to each seller as producing a single unit, one could also think of a seller offering multiple units so long as the production and holding fees scale proportionately. Also, when sellers employ mixed strategies, they can be interpreted literally as each seller randomizing which mechanism to use, or as dividing sellers into two groups playing distinct pure strategies in the proportion dictated by equilibrium.

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

D.1 Discount Sellers

The advantage of the discount sector to a seller is that the sale occurs more quickly. Let θ denote the average revenue at the time of sale, and let η represent the Poisson rate of closing, meaning that the average time between listing and closing is $1/\eta$. As in Section 3, we focus on auctions as the discount channel. Section D.7 suggests θ and η for other discount environments (discount posted prices, haggling, and lotteries). In our auction mechanism we treat the listing length as exogenous. From the seller's perspective, a Poisson number of buyers (with mean λ) will participate in her auction at its conclusion, producing an expected revenue θ which is computed as follows:

$$\theta \equiv \frac{1}{1-e^{-\lambda}} \left(\lambda e^{-\lambda} b(T) + \sum_{n=2}^{\infty} \frac{e^{-\lambda}\lambda^n}{n!} \int_0^T b(s) n(n-1) F(s)(1-F(s))^{n-2} F'(s) ds \right). \tag{44}$$

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 $b(T)$. 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.

To determine the expected profit of the auction seller, we must also account for the costs of production and the expected time delay. Let Π_a denote the expected profit from the vantage of a seller who has already incurred the initial production cost $(1-\gamma)c$.

$$\rho \Pi_a = -\ell + \eta \left(1 - e^{-\lambda}\right) (\theta - \gamma c - \Pi_a). \tag{45}$$

The right-hand side of (45) indicates that the seller incurs the holding fee per unit of time. The listing closes at Poisson rate η, but if no bidders arrive then the seller re-lists the item (without incurring additional production costs but continuing to pay the holding fee ℓ) and waits for the new auction to close. If at least one bidder participates, the seller's net gain is the expected revenue minus the completion cost, relative to the expected profit.

From the perspective of a potential discount entrant, expected profits are net of the initial production cost: $\Pi_a - (1-\gamma)c$.

D.2 Non-discount Sellers

Sellers who charge the full retail price will obtain a higher price than discount sellers, since $b(s) < z$ for all $s \in (0, T]$. The disadvantage of this format is that sellers may wait a considerable time before being chosen by a buyer. Let ζ denote the 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.

For posted-price sellers in the market, expected profit Π_p is thus:

$$\rho\Pi_p = -\ell + \zeta\left(z - \gamma c - \Pi_p\right). \tag{46}$$

Like auction sellers, posted-price sellers incur the holding fee of ℓ each unit of time that they await a buyer; for simplicity, we assume the same holding fees, but little changes if holding fees differ across the two mechanisms. When sellers encounter a buyer (at rate ζ), the purchase always occurs, with a net gain of $z - \gamma c$ relative to Π_p. For sellers contemplating entry into the posted-price market, their expected profit is $\Pi_p - (1-\gamma)c$.

D.3 Steady-State Conditions

As with the population of buyers, the stock and flow of sellers are assumed to remain stable over time. We now derive these steady steady conditions as equations (47)–(50) below.

In aggregate, recall that δ buyers enter (and exit) the market over a unit of time; thus, we need an identical flow of δ sellers entering per unit of time so as to replenish the δ units sold. In addition, the mass of sellers in each market must remain steady. Let σ be the fraction of newly entered sellers joining the auction market, so that $\sigma\delta$ choose to list an auction over a unit of time. This must equal the mass of auctions that close with at least one bidder over the same unit of time:

$$\sigma\delta = \alpha\left(1 - e^{-\lambda}\right). \tag{47}$$

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:

$$(1-\sigma)\delta = HF'(0). \tag{48}$$

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. Each seller's 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 to ensure that buyers correctly anticipate the aggregate flow of auctions:

$$\eta A = \alpha. \tag{49}$$

A similar condition applies to posted-price sellers. In aggregate, $HF'(0)$ posted-price purchases occur over a unit of time (sold to buyers who reach their deadlines). Each seller expects to sell one unit every $1/\zeta$ units of time; so collectively, the P sellers transact ζP units in one unit of time. These must equate in equilibrium, ensuring that posted-price sellers correctly anticipate the aggregate rate at which buyers purchase from them:

$$\zeta P \quad = \quad HF'(0). \tag{50}$$

D.4 Market Equilibrium

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

1. Expected revenue θ^* is computed from equation (44) using the bidding function $b^*(s)$ and distribution $F^*(s)$ derived from the buyer equilibrium, given α^*.
2. Prospective posted-price entrants earn zero expected profits: $\Pi_p^* = (1-\gamma)c$, given ζ^*.
3. Prospective discount entrants earn zero expected profits: $\Pi_a^* = (1-\gamma)c$ if $\alpha^* > 0$, or $\Pi_a^* \leq (1-\gamma)c$ if $\alpha^* = 0$.
4. α^*, ζ^*, σ^*, A^*, and P^* satisfy the Steady-State equations (47) through (50).

The first requirement imposes that buyers behave optimally as developed in Section 3. The fourth imposes the steady-state conditions. The second and third requirements impose zero expected profits for both markets due to free entry. If either market offered positive profits, additional sellers would be attracted there and profits would fall: more posted-price sellers P would reduce the rate of selling ζ, and more auction sellers A would increase the auction arrival rate and decrease expected revenue θ. Together, these two requirements also ensure that sellers are indifferent about which market they enter, allowing them to randomize according to the mixed strategy σ.

The third requirement allows auction profits to be negative if no auctions are offered. A similar possibility could be added to the posted-price market, but that market would never shut down in equilibrium. Due to the search friction—that is, the uncertainty of winning in the discount market—a fraction of buyers will inevitably reach their deadlines; thus, a sufficiently small stock of posted-price sellers can always break even.

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. The third market equilibrium requirement can be written:

$$\theta^* = c + \frac{\ell + \rho c(1-\gamma)}{\eta\left(1 - e^{-\tau H^*}\right)}. \tag{51}$$

This ensures that the expected revenue from each auction precisely covers the expected cost of listing and producing the good. Equilibrium is attained when both (10) and (51) simultaneously hold. To compute θ^*, (44) 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:

$$\begin{aligned}
\Pi_a^* &= (1-\gamma)c && (52)\\
\Pi_p^* &= (1-\gamma)c && (53)\\
A^* &= \frac{\alpha^*}{\eta} && (54)\\
P^* &= \frac{(z-c)\left(\delta-\alpha^*\left(1-e^{-\tau H^*}\right)\right)}{\ell+\rho(1-\gamma)c} && (55)\\
\zeta^* &= \frac{\ell+\rho(1-\gamma)c}{z-c} && (56)\\
\sigma^* &= \frac{\alpha^*\left(1-e^{-\tau H^*}\right)}{\delta}. && (57)
\end{aligned}$$

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 $\phi(H^*) = 0$, equations (11) through (44), and equations (51) through (57).*

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.[36]

Proposition 4. *The degenerate market equilibrium, described by equations (11) through (12) and equations (53) through (57) with $\alpha^* = 0$ and $H^* = \delta T$, exists if and only if*

$$\beta x + \frac{(z-\beta x)\tau\delta}{1-e^{-\tau\delta T}} \cdot \frac{\tau\delta + \left(\rho T(\rho+\tau\delta) - \tau\delta\right)e^{-(\rho+\tau\delta)T}}{(\rho+\tau\delta)^2} \leq c + \frac{\ell+\rho c(1-\gamma)}{\eta\left(1-e^{-\tau\delta T}\right)}. \tag{58}$$

[36]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 (58) does not hold.

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

The left side of (58) 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 (58)). 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 (58) indicates that auctions are not viable when expected costs are high, such as high production costs or holding 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 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 can occur, and that a dispersed equilibrium cannot occur when (58) 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.

D.5 Proofs

Proof of Proposition 3. By Proposition 1, equations (11) through (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 (49) and (47), 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 (52) and (53) are required by the third and second equilibrium conditions, respectively. From (46), profit solves as: $\Pi_p = \frac{\eta(z-\gamma c)-\ell}{\rho+\zeta}$, so for this to equal $(1-\gamma)c$, we require $\zeta^* = \frac{\ell+\rho c(1-\gamma)}{z-c}$ as in equation (56). With this, (48) readily yields P^* as listed in (55).

The only remaining element regards expected auction profit. Equation (45) solves as: $\Pi_a = \frac{\eta\left(1-e^{-\tau H}\right)(\theta-c\gamma)-\ell}{\eta(1-e^{-\tau H})+\rho}$. By setting this equal to $(1-\gamma)c$ and solving for θ, we obtain (51).

To evaluate the integrals in (44), we first note that by interchanging the sum and integral and evaluating the sum, expected revenue simplifies to:

$$\theta = \frac{\lambda}{1 - e^{-\lambda}} \left(e^{-\lambda} b(T) + \lambda \int_0^T b(s) F(s) F'(s) e^{-\lambda F(s)} ds \right). \tag{59}$$

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

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

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 $[0, T]$, since none of them exit early; so $F^*(s) = s/T$ and $H^* = \delta T$. Moreover, the buyer's willingness to bid (if an auction unexpectedly occurred) reduces to: $b(s) = \beta x + (z - \beta x)e^{-\rho s}$.

For $\alpha^* = 0$ to be a market equilibrium, we need $\Pi_a^* \leq \Pi_p^*$. To prevent further enter, $\Pi_p^* = (1 - \gamma)c$ is still required. From (45), if an auction were unexpectedly offered, the seller would generate $\Pi_a^* = \frac{\eta(\theta - \gamma c) - \ell}{\rho + \eta\left(1 - e^{\tau\delta T}\right)}$. Thus, the expected profit comparison simplifies to: $\theta \leq c + \frac{\ell + \rho c(1-\gamma)}{\eta\left(1 - e^{-\tau\delta T}\right)}$. This is equivalent to (58), where the left-hand side is evaluated from (59):

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

Thus, if (58) 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 (58) fails to hold, then $\alpha^* = 0$ cannot be an equilibrium, since some firms will earn greater profit by deviating and offering an auction.

To prove the last claim, first note that in a buyer equilibrium, $H \to 0$ as $\alpha \to \infty$. In addition, $b(s) \to 0$ for all $s > 0$, because auctions occur every instant, in which the buyer faces

Table A5: Comparative Statics on Key Statistics: Market Equilibrium

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

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

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 (58) is equivalent to $\Pi_a > 0$ for $\alpha = 0$. Since expected revenue is continuous in α, by the intermediate value theorem there must exist an $\alpha^* > 0$ such that $\Pi_a(\alpha^*) = 0$, which will constitute a dispersed equilibrium. □

D.6 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 fewer sellers join the discount market when buyers are less attentive (τ), less patient (ρ), or have more time (T), all of which lead to lower bids (with a steeper bid profile). Higher seller costs (whether for holding, ℓ, or production, c) or more backloaded costs (γ) 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 and expected revenue.

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.

D.7 Alternative Mechanisms

Here, we indicate how the profit of discount sellers can be formulated under the alternative mechanisms listed in Section 6.1.

First, consider physical search from the seller's perspective. Deeper discounts result in lower revenue but a higher likelihood of sale, since 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:

$$\rho\Pi_a = -\ell + \eta F(s)\left(b(s) - \gamma c - \Pi_a\right). \tag{60}$$

To obtain price dispersion, each targeted price $b(s)$ must yield the same expected profit Π_a. A similar labor market environment is characterized in Akın and Platt (2012).

Second, 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.

$$\rho\Pi_a(s') = -\ell + \eta\left(\int_0^{s'} \left(\omega b(s) + (1-\omega)b(s') - \gamma c\right) dF(s) - F(s')\Pi_a(s')\right). \tag{61}$$

Finally, 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:

$$\rho\Pi_a = -\ell + \eta\left(\int_0^{T} pk' dG(k') - \Pi_a\right). \tag{62}$$

D.8 Endogenizing Reserve Prices

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

Now consider a seller who contemplates raising the reserve price to $R > b(T)$, taking the behavior of all others in the market as given. This will only affect the seller 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. Of course, the seller would give up the immediate revenue and completion cost, which is no greater than $R - \gamma c$.

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

If $b^*(T) < c$, then the seller would prefer to set a reserve price of c. One can still analyze this optimal reserve price in our model by endogenizing the buyer deadline, T. For instance, suppose that buyers who enter six months before their deadline are only willing to bid below the cost of production. By raising the reserve price, these bidders are effectively excluded from all auctions; it is as if they do not exist. They only begin to participate once they reach time S such that $b^*(S) = c$. In other words, it is as if all buyers enter the market with S units of time until their deadline. To express this in terms of our model, we 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 choice in this case would be to shut down the auction market, forcing all buyers to purchase at the highest price z. When there are numerous independent sellers, however, they cannot sustain this degenerate equilibrium (at least when Proposition 4 does not apply). There is always an advantage to offering an auction if all other sellers offer posted-price listings: the product sells faster through auctions, even if at a slightly lower price.

E Parameter Estimation

The data facts reported in Section 4 provide evidence consistent with our model's qualitative predictions. To proceed further, we estimate parameter values to illustrate equilibrium behavior. The fitted model closely matches the magnitudes of the data facts (which are not used in the estimation). Moreover, the parameter values provide insight into the forces shaping this market behavior, allowing us to identify underlying changes in these online retail markets over time and their welfare consequences.

Our data is particularly well suited for estimation of our model, since eBay users are likely to be familiar with both auction and posted-price mechanisms, and encounter both options when searching for a given product. While buyers (and sellers) in the data potentially had other platform options, our estimation procedure only requires data moments of their decisions while they participated at eBay, rather than their full history across all platforms.

E.1 Repeat versus Single Bidders

Our model sheds light on the dynamic search strategy of buyers, who are willing to pay more over time and eventually abandon the discount mechanism. Thus, empirical testing of the model relies on observing a buyer across multiple purchase attempts — a requirement that our eBay data uniquely delivers. Of course, not all bidders make repeat attempts after failing to win, but such bidders do constitute an economically significant fraction of all bidding activity: 20% of bidders in our data participate in more than one auction of some product, and collectively these repeat bidders place 47% of all bids. However, the auctions in our data also contain a large number of bidders who bid only once.[37] While our model does not attempt to explain this behavior, we consider here any systematic differences between bidders who bid in only one auction and those who bid in multiple auctions (whom we will refer to as *single* and *repeat* bidders, respectively), and examine what effect (if any) the presence of these single bidders has on our data facts and estimated parameters.

In most observable aspects, single bidders closely resemble repeat bidders. Most notably, the bids of repeat bidders are only 0.17% higher than those of single bidders, and this difference is not statistically significant. Due to her persistence, the average repeat bidder eventually wins 27.3% of the products on which she bids, which is modestly larger than the 22.7% won by single bidders. Repeat bidders are also moderately more experienced: 85% have received feedback over 100 times, compared to 76% among single bidders.

Fortunately, single bidders are mostly irrelevant to the data trends we study. For example, Figures 1 and 4 exclusively consider choices of repeat bidders. Moreover, since single and repeat bids are similarly distributed, the resulting price dispersion would be unaffected if restricted to repeat bidders (holding the number of participants fixed). Similarly, in our estimation procedure below, the only parameter that could potentially be sensitive to the inclusion of single bidders is δ, the flow of new bidders. However, we avoid this sensitivity by identifying δ from the flow of new *repeat* bidders in the data, as detailed below.

[37]If all bidders were to play according to our model, bidding in just one auction would still occur in equilibrium, but only 13% of the time.

E.2 Estimation Procedures

Fitting our model to the data is relatively straightforward, as each parameter either corresponds directly to a sample moment or to a known transformation of sample moments. The moments we match and the corresponding model parameter estimates implied by these moments are shown in Table A6.

When linking data to theory in an auction environment, it is crucial to realize that some participation may not be observable. In our model, auctions proceed by sealed bid, yet eBay conducts ascending-like auctions. As pointed out by Song (2004), this can prevent a willing participant from submitting a bid if the auction's standing price passes his valuation before his arrival. While this does not affect the eventual winner or final price, it will cause the observed number of bidders to underestimate the true number of participants. We adjust for this bias using the approach in Platt (2017), which computes the probability that a participant places a bid as:

$$P(\lambda) \equiv \frac{1}{\lambda}\left(2\left(\ln(\lambda) - \Gamma'(1) - \Gamma(0,\lambda)\right) - 1 + e^{-\lambda}\right), \tag{63}$$

where $-\Gamma'(1) \approx 0.57721$ is Euler's constant and $\Gamma(0,\lambda) \equiv \int_{\lambda}^{\infty} \frac{e^{-t}}{t} dt$ is the incomplete gamma function.[38] Thus, the expected number of bidders per auction is a strictly increasing function $\lambda \cdot P(\lambda)$, allowing us to estimate the average number of would-be participants based on the average number of observed bidders. Similarly, while the model predicts buyers will participate at a rate of $\tau\alpha$ auctions per month, they will only be observed in $\tau\alpha P(\lambda)$ auctions per month. These adjustments allow for identification and estimation of λ, α, and τ, as shown in Table A6.

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

$$D \equiv \tau\alpha \int_0^T \left(\frac{2\left(1 - e^{-\lambda F(s)}\right)}{\lambda F(s)} - e^{-\lambda F(s)}\right) ds. \tag{64}$$

The integrand is the probability that a participant in state s can place a bid (Platt, 2017). The function D can then be used to estimate δ as shown in Table A6.

We set $\beta = 0$ throughout this estimation, since it is not separately identifiable from x — the two parameters always appear multiplied together in any of our equilibrium conditions. One could exogenously set $x = z$, and then estimate β from the median (or some other moment) of auction revenue. We performed this exercise and found that the corner solution of $\beta = 0$

[38] Platt (2017) assumes that participants in an auction arrive in random order, which is compatible with our model since the payoff of losing participants is the same whether or not they were able to actually place a losing bid. The assumption implies that 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).

consistently yields the closest fit possible.[39]

In computing these moments, we normalize the posted price as $z = 1$. This has no effect on the distribution $F(s)$, while bids $b(s)$ will scale proportionally. Therefore we similarly transform bidding data: for each item, we compute the average price (including shipping costs) among all sold posted-price listings (the analog of z) of that product, then divide all bids (including shipping costs) for that item by this product-level average. This rescaling is equivalent to "homogenizing" bids (Balat et al., 2016). All monetary values (θ, ℓ, c) are therefore in units of fractions of the average posted price. We consider one unit of time to be a month, which is also a normalization and it merely adjusts the interpretation of T and ρ. The statistics in the second column of Table A6 are computed from the data by taking averages across product-level averages; thus, this exercise should be interpreted as fitting the model for the average product.

E.3 Estimation Moments

We now discuss the moments used to identify the parameters. The buyer equilibrium depends on one endogenous variable (λ) and five parameters (α, δ, τ, T, and ρ). The market equilibrium (developed in Technical Appendix D) depends on the additional parameters ℓ, η, γ, and c. These are estimated one at a time, in the following sequence. In particular, all the seller parameters are determined after (and hence cannot affect) the buyer parameters.

The first moment in Table A6 is the number of observed bidders per completed (sold) auction. The auction ends in a sale as long as at least one bidder arrives, so the theoretical equivalent of this moment is $\frac{\lambda P(\lambda)}{1-e^{-\lambda}}$, because λ is the mean number of participants per auction, $P(\lambda)$ is the probability that a participant is able to place a bid and be observed, and $1 - e^{-\lambda}$ is the probability of at least one bidder arriving.

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

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

To compute the market size δ, we measure the number of newly-observed repeat bidders per month who never win; this is matched to the theoretic equivalent, which is given by $(\delta - \alpha)(1 - (1 + D)e^{-D})$. As described above, D is the Poisson rate of bids per buyer over the

[39] When $\beta = 0$, the model predicts a median revenue that is 85.4% of the posted price, while the data reports a median revenue of 83.2% . However, any increase in β would increase the predicted median.

duration of his search. D is computed for a complete search spell (*i.e.* all T periods); thus, multiplying by $\delta - \alpha$ rather than δ yields the average flow per month of newly observed repeat bidders who never win.

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

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

Moving to the parameters in the full market equilibrium, we first treat inventory holding fees as equaling the listing fees paid to the platform. These fees are directly observable, yielding an estimate of ℓ.[40] The average time for which an auction is listed is also observable in the data, yielding an estimate of $1/\eta$.

We also observe the fraction of posted-price listing that sell within the 30 day window of their listing. Since the model predicts that posted-price sellers exit at Poisson rate ζ, $1 - e^{-\zeta}$ will have exited within one month; the table simply inserts the equilibrium value for ζ. We also note that α, which was computed in the second row of Table A6, is endogenous. In the market equilibrium, we can use the equilibrium condition for α in (51) to determine the underlying cost of production. The last two conditions are solved jointly to recover γ and c.

We now summarize what sources of variation in the data leads to the most change in the estimate of each parameter. Variation in many of the parameters (λ, α, τ, δ, ℓ, and η) maps directly back to variation in the data moments found in the corresponding rows of Table A6. The effect is less obvious for those involving the equilibrium conditions (T, ρ, γ, and c), which we now consider in turn.

The parameter T relies on the requirement for a buyer equilibrium found in (10), which effectively balances buyer entry and exit in steady state. It is most sensitive to increases in the number of bidders per auction (the first data moment in Table A6). If the number of bidders per auction increases, buyers are more likely to hit their deadline without winning, causing too much exit relative to entry; but the estimation restores balance by keeping buyers in the market longer, giving them more chances to win. While α, τ, and δ also appear in (10), the effect of their corresponding moments on the estimated T are somewhat muted.

[40] In practice, the structure of eBay fees varies across item categories and has changed over time. During the early part of our sample period, for some categories, sellers were allowed an allotment of listings that were only charged "final value fees" (a commission) and not "insertion fees" (a per-listing fee), and beyond that allotment insertion fees were required. In May, 2014, the insertion fee exemption allotment was abolished for posted-price listings in many categories. For simplicity of modeling, we treat all fees as insertion fees. For further details, see http://community.ebay.com/t5/Announcements/Sellers-Updates-to-fees-in-select-categories-could-mean-more/ba-p/26163313

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

The parameters γ and c are jointly determined in the last two rows of Table A6. The market equilibrium requirement in (51) determines the endogenous number of auctions so that both mechanisms will be equally profitable. As a consequence, c is closely tied to average auction revenue, keeping ex-post profits low. Thus, γ is left to reconcile the time it takes to sell a posted-price listing and the number of auctions listings. The latter has a particularly large impact over time, as seen in Table A7. While the auction revenue and speed of selling have hardly changed, the number of auctions declined precipitously; this can be justified if γ increased, because posted-price sellers will have fewer up-front costs before the long wait for a buyer.

E.4 Parameter Estimates

We conclude by discussing insights from the estimated parameter values. In Section 4.2 of the paper, we reported theoretical analogs of key empirical facts. Here, we discuss each parameter's contribution to this match between data and model.

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

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

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

Average auction revenue (θ) is 0.853. A sale in the posted-price market therefore generates $z - \theta = 14.6\%$ more revenue than a sale in the auction market, but this is offset in that the sale occurs after 1.8 months on average, which is 11 times longer than the average auction (0.156 months, implying an η of 6.39 auctions ending per month). While informative, these parameters are essentially direct reflections of the data.

The model uncovers the cost structure that sellers must face to optimally split between auctions in this way. With a production cost of $c = 0.840$, discount sellers are only paid a 1.3% markup on average. Moreover, nearly all the costs of production are incurred at the time of sale, with $\gamma = 0.984$, suggesting that a large portion of costs are shipping fees or that sellers may source their inventory at the time of the sale rather than at the time of listing.

F Welfare and Consumer Surplus

Our model provides important insight into who benefits from the existence of a discount market, which we examine in this Appendix. We also consider how demand estimation differs in the context of deadlines compared to a model of static valuations or stationary search.

F.1 Welfare Consequences of the Existence of a Discount Channel

Our model rationalizes the seemingly redundant coexistence of discount and posted-price channels when goods, sellers, and buyers are ex-ante homogenous—an uncommon result, as noted in Section 5.3 of the paper. A question that naturally follows is what social value does the discount mechanism offer? Does this mechanism improve efficiency, or would welfare instead be higher if only posted prices were available? The structure of our model allows us to investigate this question.

Total expected welfare is derived from three sources in our environment: buyers, sellers, and any recipient of holding fees, ℓ. We will refer to this third category as *intermediaries*; in our empirical application, this corresponds to the eBay platform. Here we consider who benefits from the existence of a discount channel, and whether those gains outweigh any losses accruing to others. Although buyers and sellers in our model find it individually rational to participate in both mechanisms, it is possible that, given the nature of buyers' non-stationary search process, overall welfare would increase if the discount channel were removed.[41]

For sellers, free entry ensures that newly-entering sellers earn zero expected profit under any mechanism. Any potential benefit (or harm) to a seller from the discount channel is dissipated through competition with other sellers.

Buyers, in contrast, are always better off with the existence of the discount channel, since it offers a chance at lowering the cost of purchase. Formally, a newly entering buyer expects

[41] For this counterfactual exercise, we compare the dispersed equilibrium (auctions and posted prices) to an environment where only the posted-price channel can operate — effectively, a degenerate equilibrium but without imposing the constraint in (58). Reported welfare calculations are evaluated at $x = z$.

utility of $V(T)$ by participating in the discount market (measured in terms of dollars, and net of any payments to sellers). If only the posted-price channel is available, the buyer's expected utility is simply $(x-z)e^{-\rho T}$. The latter is always smaller than the former, as we demonstrate in Proposition 5.

Proposition 5. *A buyer's utility is $V(T)$, which strictly exceeds his utility $(x-z)e^{-\rho T}$ in a degenerate equilibrium.*

Proof of Proposition 5. The expected utility $V(T)$ is greater than $(x-z)e^{-\rho T}$ if:

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

The second line holds because $z > \beta x$ is assumed throughout. If this were not the case, buyers would not use the discount mechanism, opting to purchase immediately from the posted-price mechanism on entering the market. The l.h.s. of the last line is strictly increasing in T, with derivative: $\tau\kappa\rho\left(e^{(\tau\kappa+\rho)T}-1\right)e^{-\rho T}$. Moreover, at $T = 0$, it evaluates to 0. Therefore, the expression is always greater than 0 for $T > 0$, and expected welfare is strictly greater with auctions than without. □

For intermediaries, the impact of the discount channel depends on the cost of intermediation. To analyze the welfare benefits of the existence of the discount channel, we examine the two extreme cases of intermediary costs in which these benefits are the highest and lowest. In one extreme, suppose that holding fees merely cover the real resources that are consumed in the process of connecting buyers and sellers, such as maintaining warehouse space for inventory, building the intermediating platform, staffing showrooms, etc. If so, intermediaries earn zero profit and are indifferent about the existence of the discount channel. This implies that social welfare is equal to buyer utility, so auctions would always be socially beneficial. Under our estimated parameters in this scenario, the existence of auctions increases welfare by 5.6% of the product's retail price (relative to a market with only posted prices).

To appreciate this result, one should note that there are two avenues through which sellers are being driven to zero profit: one is to receive less revenue through auctions, and the other is to incur more costs by waiting longer for buyers in the posted-price market. But when intermediation is costly, the latter is inefficient because it burns up real resources. Auctions also create some inefficiency as buyers pay for the product earlier than they need it, but most of the dissipated seller profit is transferred to buyers rather than destroyed.

Next, consider when holding fees are pure profit for intermediaries. Also, let $x = z$ and $\beta = 0$. Together, these conditions will yield the least favorable case for examining the discount channel. If a product is listed for t periods, then the present value of those holding fees is $\left(1-e^{-\rho t}\right)\ell/\rho$. With a posted-price listing, the duration is exponentially distributed with parameter ζ, whereas with an auction, the parameter is $\eta\left(1-e^{-\lambda}\right)$. Recall that fraction σ of products are listed as auctions, so the present value of expected profit to intermediaries is:

$$\sigma\int_0^\infty \eta\left(1-e^{-\lambda}\right)e^{-\eta\left(1-e^{-\lambda}\right)t}\frac{\left(1-e^{-\rho t}\right)\ell}{\rho}dt+(1-\sigma)\int_0^\infty \zeta e^{-\zeta t}\frac{\left(1-e^{-\rho t}\right)\ell}{\rho}dt. \tag{65}$$

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

$$\frac{\ell}{\delta}\left(\frac{(z-c)\,(\kappa-\alpha)}{\rho(z-c\gamma)+\ell}+\frac{\alpha\left(1-e^{-\lambda}\right)}{\eta\left(1-e^{-\lambda}\right)+\rho}\right) \tag{66}$$

In the degenerate equilibrium, $\alpha=0$, so the expected intermediary profit becomes: $\frac{(z-c)\ell}{\rho(z-c\gamma)+\ell}$. The change in expected intermediary profit from shutting down auctions is their difference, which, after substituting for η using (51), becomes:

$$\frac{\ell\alpha\left(1-e^{-\lambda}\right)}{\delta}\left(\frac{z-c}{\rho(z-c\gamma)+\ell}-\frac{\theta-c}{\rho(\theta-c\gamma)+\ell}\right). \tag{67}$$

Since $z>\theta$, the parenthetical term is positive, so intermediaries always collect more revenue when the discount channel is not available. Thus, the welfare impact of discount mechanisms (when intermediation is costless) depends on whether the benefit to buyers exceeds the lost revenue to intermediaries. This must be numerically evaluated.

When holding fees are all profit, the net effect of the discount channel on total welfare depends on whether the benefits to consumers outweigh the lower profits of intermediaries, and this comparison depends on parameter values. Under our estimated parameters, when auctions are available, consumer surplus equals 5.6% of the retail price and intermediaries earn a profit of 6.5% of the retail price (for total welfare of 12.1%). If auctions were not available, intermediaries would capture all of the surplus with a profit (and total welfare) equal to 14.4% of the retail price. Thus, the presence of the discount mechanism reduces total welfare by 2.3% of the retail price.

This comparison is essentially a competition between two forms of inefficiency: in the discount sales channel, the item is produced earlier than needed, but in the posted-price channel, sellers are entering the market earlier than needed. The latter inefficiencies tend to be small. As a result, the discount mechanism is welfare improving only when holding fees ℓ are small, the flow of buyers δ is large, or production occurs mostly at time of entry (γ is small).

In sum, the discount channel is always welfare-improving in a setting where holding fees merely cover costs; but if these fees are pure profit for the intermediaries, then typically the discount channel is socially wasteful. In reality, one would expect that a platform such as eBay can sustain some profits (for instance, network externalities make it beneficial for sellers and buyers to rely on a single platform), but it is also unlikely to be costless to facilitate the listings. A moderate level of platform profits could result in the discount channel being welfare neutral.

Finally, we highlight that a first-best outcome in this environment would be a case where buyers were able to produce their own units at a cost of c, in which case total welfare would be $x - c$, which is 16.0% of the retail price under our estimated parameters (compared to 12.1% in the dispersed equilibrium). To see this formally, we compare this first-best welfare to expected welfare in the pure transfer degenerate equilibrium:

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

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

F.2 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, ignoring this non-stationarity would lead to mis-measurement of demand and consumer surplus. 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 continuation value of a bidder is:

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

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

Figure A1: Derived Demand Curve

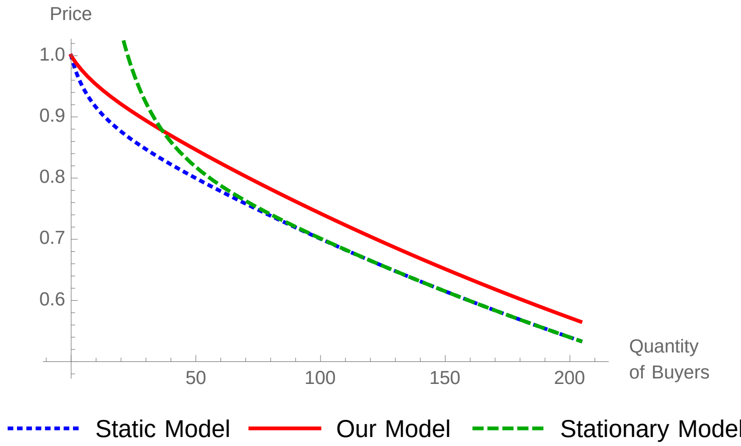

Notes: The figure reports the demand curve inferred from bids reported in Panel (C) of Figure 2 using our deadlines model (solid) vs. treating the data as though it came from a static model (dotted) or a stationary dynamic model (dashed). The dashed line is truncated, but would intersect the vertical axis at a price of 2.5.

simplifies to:

$$x(s) \equiv b(s) + \frac{\tau\alpha}{\rho}\left(e^{-\lambda F(s)}b(s) - e^{-\lambda}\left(b(T) + e^{\lambda}\int_s^T b(t)\lambda e^{-\lambda F(t)}F'(t)\,dt\right)\right). \tag{68}$$

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 A1, 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,[42] it generates the true demand curve, depicted as the solid line in Figure A1. The static interpretation of data generated from a dynamic process will *underestimate* demand — on average by 5.3% of the retail price. On the flip side, if the data were generated by in our non-stationary environment, but then interpreted using the stationary dynamic model, demand would be *overstated* by an average of 2.5% of the retail price (the dashed line in Figure A1). In the stationary model, low-valuation buyers are unlikely to win in the current auction and also and expect very little benefit from 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

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

auctions, so they shade their bids aggressively (by as much as 60%). In our non-stationary model, however, high-valuation buyers are closer to their deadline and are shading *less* than low valuation bidders.

Table A6: Key Data Moments and Matching Parameter Values

	Observed Value in Data	Theoretical Equivalent	Fitted Parameter
Bidders per completed auction	5.30	$\frac{\lambda \cdot P(\lambda)}{1-e^{-\lambda}}$	$\lambda = 13.10$ (0.243)
Completed auctions per month	12.76	$\alpha\left(1-e^{-\lambda}\right)$	$\alpha = 12.76$ (0.525)
Auctions a bidder tries per month	1.18	$\frac{\tau\alpha P(\lambda)}{1-e^{-\tau\alpha P(\lambda)}}$	$\tau = 0.064$ (0.002)
New repeat bidders per month who never win	7.41	$(\delta-\alpha)\left(1-(1+D)e^{-D}\right)$	$\delta = 21.13$ (1.163)
—	—	Eq. (10)	$T = 10.30$ (0.255)
Average revenue per completed auction	0.853	θ	$\rho = 0.055$ (0.003)
Average listing fee paid	0.087	ℓ	$\ell = 0.087$ (0.0003)
Average duration of an auction listing (months)	0.156	$1/\eta$	$\eta = 6.39$ (0.028)
Average % of posted-price listing sold in 30 days	48.1%	$1-e^{-\frac{\ell+\rho(1-\gamma)c}{z-c}}$	$\gamma = 0.985$ (0.039)
—	—	Eq. (51)	$c = 0.840$ (0.003)

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. The theoretical objects are described in Technical Appendix E. 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.

Table A7: Parameter Values and Welfare Estimates in 2010–2014

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

Notes: Table displays estimates of parameter values using data from 2010–2014. Data was constructed using same sample restrictions as in main sample. Each year label corresponds to one year of data beginning in Oct. of that year and continuing through Sep. of the following year (*i.e.* the 2013 column corresponds to the primary data sample of the paper). Units for parameters are as in Table A6. Standard errors, from 200 bootstrap replications at the product level, are contained in parentheses.

Table A8: Descriptive Statistics From 2010–2014 Data Samples

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

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

Table A9: Data Moments Used in Estimation From 2010–2014 Data Samples

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

Notes: Table displays observed data used for estimation (as in Table A6) for data sampless from 2010–2014. Data was constructed using same sample restrictions as in main sample. Each year label corresponds to one year of data beginning in Oct. of that year and continuing through Sep. of the following year (*i.e.* the 2013 column corresponds to the main data sample used in the body of the paper). Displayed moments were used to obtain parameter estimates displayed in Table A7.