A Theory of Search with Deadlines and Uncertain Recall

S. Nuray Akin* and Brennan C. Platt†‡

January 12, 2012

Abstract

We analyze an equilibrium search model where buyers seek to purchase a good before a deadline and face uncertainty regarding the availability of past price quotes in the future. Sellers cannot observe a potential buyer's remaining time until deadline nor his quote history, and hence post prices that weigh the probability of sale versus the profit once sold. The model's equilibrium can take one of three forms. In a late equilibrium, buyers initially forgo purchases, preferring to wait until the deadline. In an early equilibrium, any equilibrium offer is accepted as soon as it is received. In a full equilibrium, higher prices are turned down until near the deadline, while lower prices are immediately accepted. Equilibrium price and sales dynamics are determined by the time remaining until the deadline and the quote history of the consumer.

Keywords: Equilibrium search, deadlines, uncertain recall, price posting, reservation prices

JEL Classification: D40, D83

---

*Department of Economics, University of Miami, 517-K Jenkins, Coral Gables, FL 33124, (305) 284-1627, nakin@miami.edu

†Department of Economics, Brigham Young University, 135 FOB, Provo, UT 84602, (801) 422-8904, brennan_platt@byu.edu

‡Dr. Akin thanks Patrick Kehoe, Tim Kehoe, Warren Weber, and the Research Department at the Minneapolis FED for their hospitality during the revision of this paper. We greatly benefited from comments by the editor, Christian Hellwig, and anonymous referees, as well as those of Manuel Santos, Val Lambson, and Joseph McMurray.

1 Introduction

Quite frequently, the timing of a purchase is of crucial importance to the buyer. When a worker relocates to a new city, he would ideally like to secure a new home before the start date; otherwise, he may have to use costly temporary housing, such as hotels or short-term leases. Similarly, a parent might buy a birthday gift for her child at any time before the birthday, but purchasing it thereafter would cause significant grief. For credit cards and home mortgages, consumers are often given an introductory rate that eventually expires; this may give the consumer incentive to search for a replacement provider as this deadline approaches.

These scenarios are non-trivial because finding the right home or gift requires search. If the right item at the best price were perfectly known, one could acquire it just before the deadline without worry. More likely, though, one will only come across an acceptable item infrequently. In addition, that item could be offered at a variety of prices, making decisions more difficult for the buyer. A high price quoted early in one's search may be unappealing, but the same offer may be unavailable later when desperation sets in.

Deadlines have been frequently studied in the bargaining literature. There, a buyer and a seller must agree on a price, effectively dividing their surplus (which may or may not be common knowledge). If they fail to reach an agreement before a commonly-known deadline, the available surplus is reduced. Here, we recast this deadline problem with the anonymity that is typical in a market: sellers are uncertain how close any particular buyer is to his deadline or what price quotes he has previously received.

In particular, buyers randomly encounter sellers of a homogeneous good. On meeting, the buyer learns the seller's asking price and must decide whether to purchase now, or defer the offer and continue search. Buyers face uncertain recall, meaning that with some probability, the offer deferred yesterday is still available for consideration today. During an initial grace period, buyers enjoy a positive flow of utility; after passing the deadline, however, they incur a penalty each period until the purchase is completed.[1] Thus, even though buyers identically value the good in question, they

[1] In housing search, the initial benefits reflect the buyer's consumer surplus in his current housing, while the penalty reflects the lower surplus of switching to short-term housing or long-distance commuting. In gift search, the analog is a stream of utility from a relationship, which falls if expectations of a gift are not met.

will be ex-post heterogeneous in their willingness to pay for it, depending on how close they are to the deadline and how lucky they have been in securing quotes in the past.

Sellers explicitly consider this in choosing their asking price. A higher price will generate greater profit if accepted, but also limit the pool of buyers who would accept the offer. In equilibrium, the higher markup exactly cancels with the lower probability of acceptance, which allows identical sellers to offer different prices and yet have equal expected profit. We study this endogenous price formation in our model of search on a deadline. The technical challenge of this model is that both the buyer's time until deadline and best prior offer enter as state variables; both play key roles in determining the endogenous distribution of offered prices.

The resulting equilibrium displays several interesting characteristics. First, the equilibrium takes one of three forms. In a late equilibrium, buyers initially forgo any purchases, preferring to wait until the deadline. In an early equilibrium, buyers accept any offer as soon as it is received. In a full equilibrium, buyers early in their search are only willing to accept some of the offers; higher prices are turned down until near the deadline. Typically only one of these equilibria will exist.

Second, in a full equilibrium, our model generates a continuum of offered prices. This occurs because some fortunate buyers are able to compare their best prior quote to their new offer. Thus sellers realize that a higher price has a greater chance of being undercut in such comparisons. This feature of our sequential model with recall mimics the behavior of simultaneous search models (Burdett and Judd, 1983) where individuals get multiple quotes in a single period and choose the lowest price. Unlike that setting, though, our price distribution also includes an atom at the lowest price, targeting those who have just entered the market. Any buyer offered this price accepts it immediately, so sellers offering the price have no risk of being undercut.

Third, the full equilibrium response to a change in parameters is often surprising. The anticipated direct effect is often overpowered by a larger indirect effect: changes in the number of buyers near their deadline. This *deadline concentration* strongly influences sellers' ability to extract consumer surplus. If more of the population is desperate, more sellers offer prices accepted only by the desperate. But this increases the average offered price; thus all buyers face worse prospects and are willing to pay more.

For example, if buyers enjoy higher grace-period utility, they are inclined to delay their purchases and enjoy more of this utility flow. This increases deadline concentration, and allows sellers to raise prices and capture most of this added utility. On the other hand, if the good becomes more valuable to the buyers (with no change in its cost of production), one typically expects sellers to extract more surplus through higher prices. Yet the buyers' eagerness to quickly obtain the good will reduce deadline concentration, causing prices to fall.

Perhaps most surprisingly, prices also fall when the ability to recall past offers is reduced. One would expect less recall to give more power to sellers, who have less chance of being undercut. Yet this also raises the incentive of buyers to accept offers earlier rather than defer them, so fewer of them reach their deadline. It is also noteworthy that equilibrium price dispersion can occur with any degree of recall. Even if quotes are never lost, the impending deadline can push fortunate buyers to accept a low-price offer early in their search, while higher-price offers are deferred in hopes of receiving a second quote before the deadline.

Our work contributes to several strands of the existing literature on price formation and sales in the presence of imperfect information, which we examine in greater detail in Section 5. One is the literature on bargaining before a known deadline, previously mentioned. Rather than having a single buyer and seller negotiate, such as in labor or other contract disputes, our environment examines private deadlines in a market setting that is more appropriate for consumer transactions. While both environments result in dispersed prices, the dynamics can differ significantly; for instance, in our model, the highest prices are paid by buyers near their deadline, while the earliest buyers do so in Fuchs and Skrzypacz (2011).

We also add to the literature on equilibrium search, which examines under what circumstances offering multiple prices is consistent with optimal search behavior of both buyers and sellers. This has been accomplished by assuming differences in buyers' valuations or search costs; we demonstrate that an impending deadline also generates equilibrium price dispersion. A similar point is made in a labor market search environment in Akin and Platt (2012), where the deadline arises due to limited unemployment benefits. That model assumed past offers could not be recalled, which greatly simplifies the analysis. Here, we demonstrate that recall does not collapse the price distribution; indeed, multiple prices can emerge even when perfect recall is allowed. This is because deadlines create ex-post differences among the buyers, making those closer to the deadline willing to pay a higher price.

This also makes an important contribution regarding search with uncertain recall. Most search models assume either no recall or perfect recall, but these extremes may not capture the uncertainty the buyer faces. In consumer search, it is often the case that the product is sold or the price has changed before the consumer returns to a previously-visited seller. Only a few articles have considered some form of uncertain recall, and these have assumed an exogenous distribution of prices (*e.g.* Landsberger and Peled, 1977). Our model brings uncertain recall into the equilibrium search literature; indeed, the comparative statics discussed above demonstrate the need to consider the endogenous response of sellers.

We proceed as follows: Section 2 presents the model and defines equilibrium. In Section 3, we present the equilibrium solution. Section 4 solves for and discusses equilibrium for two special cases: no recall and guaranteed sampling of a quote every period, as well as a numerical illustration of the general case. We defer full review of related work until Section 5, which compares our results to those in the search and bargaining literatures. We conclude in Section 6.

2 Model

Consider a dynamic discrete time environment, with infinitesimal buyers and sellers each entering the market at rate δ. All agents have a common discount factor β. The good being sold is homogeneous, which provides a utility x to any buyer, but because sellers may quote different prices for the good, buyers may find it worthwhile to search. Buyers and sellers encounter one another with probability λ; when this occurs, the buyer is quoted a price p, drawn from an endogenous distribution of offered prices, $F(p)$.

The buyer can either make the purchase immediately, or *defer* the offer to allow further search. Recall is uncertain; if an offer is deferred, with probability γ it will be unavailable tomorrow. If an offer is still available in the next period, we say a buyer has *retained* it. Only the best available offer can be deferred — all others are discarded. If retained, an offer can be deferred yet again with the same uncertainty of retaining it one more period.

2.1 The Buyer's Search Problem

On entering the market, a buyer has two periods in which he receives utility $b > 0$ if still searching; thereafter, his utility each period falls to $d < 0$ until a purchase is made. In each period, the buyer's timing proceeds as follows:

1. Utility b or d is received.
2. Buyers learn if they will receive a new offer, and if so, the quoted price q.
3. If a buyer has both a new and an old offer, the higher of these is discarded. Ties are broken by randomizing with equal probability on each offer. The lowest (or only) offer becomes the *current* offer.
4. Buyers with a current offer p decide whether to accept it, obtaining utility $x - p$ and exiting the market.
5. If the buyer defers offer p, he then learns whether the offer was retained for tomorrow.

This problem has two state variables. First, buyer choices can depend on the remaining time until the purchase deadline s, which can take three values: 2, 1, or 0. Note that after the deadline, the search problem is stationary in state 0. Second, buyer choices may also depend on whether a prior quote p is currently available.

Let V_s denote the expected utility of a consumer entering period s with no past offers available, and $W_s(p)$ denote the same for a consumer who retained offer p. Let R_0 denote the endogenously-determined reservation price of a buyer whose deadline has passed. Note that this will be the highest price offered in equilibrium, since no buyer in any state is willing to pay more. Thus, the expected utility of a buyer who has exceeded his deadline is:

$$V_0 = d + \lambda \int_{-\infty}^{R_0} (x - q)dF(q) + \beta(1 - \lambda)V_0. \tag{1}$$

Each period, he receives utility d. He also encounters purchase opportunities with probability λ, and will accept any price at or below R_0. If a transaction occurs at price q, he also receives $x - q$.

Buyers who have passed their deadline are willing to purchase at any price offered in equilibrium; but to determine their reservation price, one must determine the consequences of declining an offer, even if this does not occur in equilibrium. We consider a one-period deviation, in which the buyer defers offer p to perform one more search, but then accepts any price thereafter. For instance, if p is retained, the next period utility would be:

$$W_0(p) = d + \lambda \int_{-\infty}^{p} (x - q) dF(q) + (1 - \lambda F(p))(x - p). \tag{2}$$

If the buyer is fortunate enough to receive a price $q < p$, it is accepted; but if the new quote is higher or none is encountered, price p is exercised instead.

Of course, the buyer is not guaranteed that p will be retained tomorrow. Thus, if offer p is deferred, his expected utility tomorrow is $(1 - \gamma)W_0(p) + \gamma V_0$. This is weighed against the immediate utility $x - p$ received from making the purchase now. The reservation price R_0 is defined such that the buyer is indifferent between the current purchase and further search:

$$x - R_0 = \beta((1 - \gamma)W_0(R_0) + \gamma V_0). \tag{3}$$

Next, we move to those who have one period remaining until expiration. Their problem is nearly identical. In particular, if they defer an offer p today, their expected utility tomorrow is also $(1 - \gamma)W_0(p) + \gamma V_0$. Thus, the reservation price in period 1 is also R_0. Indeed, the only difference in expected utility is that the current benefit is b rather than d:

$$V_1 = b + \lambda \int_{-\infty}^{R_0} (x - q) dF(q) + \beta(1 - \lambda)V_0. \tag{4}$$

Finally, consider buyers who have just entered the market, having two periods until expiration. Let R_2 denote the reservation price of this group. Their expected utility is:

$$V_2 = b + \lambda \int_{-\infty}^{R_2} (x - q) dF(q) + \lambda\beta \int_{R_2}^{R_0} ((1 - \gamma)W_1(q) + \gamma V_1)\, dF(q) + (1 - \lambda)\beta V_1. \tag{5}$$

Note that any offer greater than R_2 is not accepted, but will be deferred for potential use tomorrow. Here, deferrals *can* occur on the equilibrium path. If offer p is retained in state 1, his expected utility would be:

$$W_1(p) = b + \lambda \int_{-\infty}^{p} (x - q)dF(q) + (1 - \lambda F(p))(x - p) \tag{6}$$

With offer p in hand, he will gladly accept anything better, which is reflected in the integral. If he receives no second offer, or the offer is higher than his first offer (which occurs with probability $1 - \lambda F(p)$), he has the option to either purchase at p or defer it to state 0. However, deferral provides expected utility $(1-\gamma)W_0(p)+\gamma V_0$ tomorrow, just as it does when an offer is deferred in state 1 with no retained offers from state 2. Thus, the same reservation price R_0 applies. Since that is also the highest price offered in equilibrium, $R_0 \geq p$ and the buyer will purchase rather than defer.

Of course, it is uncertain whether the offer deferred in state 2 will be retained in state 1; if offer p is deferred, the expected utility is $(1 - \gamma)W_1(p) + \gamma V_1$. The reservation price R_2 must make buyers in period 2 indifferent between purchasing at R_2 or deferring until tomorrow:

$$x - R_2 = \beta \left((1 - \gamma)W_1(R_2) + \gamma V_1\right). \tag{7}$$

When buyers are indifferent, equilibrium will direct them to make the purchase rather than defer. Indeed, this is necessary. For instance, suppose some fraction of state 2 buyers deferred offers of R_2. By offering $R_2 - \epsilon$, with $\epsilon > 0$ arbitrarily small, a seller could break this indifference, earning strictly more profit thereby.

2.2 The Seller's Problem

Sellers produce their good at cost $c < x$, at the time of the transaction. They are unable to observe the state of the buyer with whom they have been paired. Thus, quoting a higher price bears the risk of a lower likelihood of being accepted. At the same time, it results in higher realized profits if accepted. If a set of prices produce the same maximal expected profit, sellers can randomize over this set.

Seller profit is defined as:

$$\pi(p) \equiv (p - c)S(p), \tag{8}$$

where $S(p)$ is the expected sales from charging price p. Informed by the preceding buyer's problem, we can make several observations regarding the equilibrium price distribution, $F(p)$. Let P denote the support of $F(p)$.

Lemma 1. *1.* $P \cap (R_0, +\infty) = \emptyset$.

2. $P \cap (-\infty, R_2) = \emptyset$.

3. If $\gamma < 1$, *no atoms*[2] *can occur in* $F(p)$ *for any* $p \in (R_2, R_0]$.

The first two claims are straightforward. For the first, no seller will charge more than R_0, because even the most desperate of buyers (those in state 0) always defer such an offer, which results in zero profit. For the second, no seller will charge less than R_2 because every buyer in every state is willing to immediately purchase the good at price R_2. Charging less will decrease p without any compensating increase in $S(p)$.

The third claim is slightly more subtle, and is a consequence of buyers comparing retained offers to new offers. In essence, if an atom did occur at a given price, sellers would prefer to undercut that price to avoid a tie, much like Bertrand competition. The proof is provided in the appendix. Note that this does not apply when $\gamma = 1$, since no offers can be retained. Also, this does not rule out an atom at R_2, since no one defers a purchase when offered R_2.

Thus, let a be the weight of the atom at R_2 (while still allowing that a may be 0). The remaining distribution is continuous. Let $\underline{p}$ denote the lower bound of $P \backslash \{R_2\}$; that is, the lowest price in the continuous portion of the the distribution. This notation is needed because typically $\underline{p} > R_2$.

2.3 Steady State Conditions

For sellers to choose a pricing strategy, it will be critical to know how many buyers there are at each state of the search process. We consider a steady state equilibrium, where the measure of buyers in each state stays constant over time. Let h_s denote the measure of buyers with s periods remaining who have no offer at the start of the period, and $g_s(p)$ denote the measure of buyers with s periods remaining who have retained offer p.

[2] An *atom* with weight $a > 0$ occurs in F at p when $\lim_{q \searrow p} F(q) - \lim_{q \nearrow p} F(q) = a$.

Buyers enter the market (in state 2) at rate δ. All of them either move to state 1 or exit the market. Thus,

$$h_2 = \delta. \tag{9}$$

Next, consider buyers with one period remaining and no retained offers. Buyers arrive in this state either because they received no offer in state 2, or because they deferred an offer and did not retain it. The latter occurs with probability $\gamma\lambda(1-a)$, since state 2 buyers defer any offer except R_2. All of the buyers in state 1 either move to state 0 or exit. Thus, the steady state requires:

$$h_1 = ((1-\lambda) + \gamma\lambda(1-a))h_2. \tag{10}$$

On the other hand, those who deferred and retained an offer arrive in state $(1, p)$. All of these buyers exit the market. Their steady state condition is:

$$g_1(p) = (1-\gamma)\lambda F'(p)h_2, \tag{11}$$

for $\underline{p} \leq p \leq R_0$, with $g_1(p) = 0$ for all other prices.

Last, consider those who have reached their deadline and have not retained an offer. Note that buyers only reach this state if they received no offers in state 1. On the other hand, only λh_0 buyers exit this state each period, since the others received no offer. Hence, in steady state:

$$\lambda h_0 = (1-\lambda)h_1. \tag{12}$$

These steady state conditions allow us to compute the expected sales $S(p)$ from quoting price p. First, if a seller offers R_2, any buyer immediately accepts it. Thus,

$$S(R_2) = \lambda, \tag{13}$$

since λ is the probability of being paired with a buyer in a given period.

On the other hand, if a seller offers price $p \in [\underline{p}, R_0]$, there are three possibilities. First, buyers with no retained offer in state 1 or 0 will immediately accept the offer. Second, buyers in state 1 with retained offer q will only accept if $q > p$. Finally, buyers in state 2 will defer the offer, and if it is retained, will exercise it only if they receive no second quote or a higher second quote. Thus:

$$S(p) = \lambda \frac{h_0 + h_1 + \int_p^{R_0} g_1(q)dq + \beta(1-\gamma)(1-\lambda F(p))h_2}{h_0 + h_1 + \int_{\underline{p}}^{R_0} g_1(q)dq + h_2} \tag{14}$$

The denominator is the total population of buyers.

2.4 Equilibrium Definition

A steady state search equilibrium consists of seller profit π^*, the distribution of sellers' offered prices $F^*(p)$, reservation prices R_0^* and R_2^*, and the measures of buyers in each state h_0^*, h_1^*, h_2^*, and $g_1^*(p)$, such that:

1. R_0^* and R_2^* are consistent with Bellman Eqs. 1 through 7, given $F^*(p)$.
2. All prices in the support of F^* produce the same profit π^*, while all other prices produce no more than π^*, using Eqs. 8, 13 and 14.
3. h_0^*, h_1^*, h_2^*, and $g_1^*(p)$ satisfy the steady state conditions in Eqs. 9 through 12.

An equilibrium can take one of three forms, as we demonstrate and discuss in further detail in the next section. In the first, R_2 is the only price offered in equilibrium; charging a higher price (say, R_0) is unprofitable because too few buyers reach states 1 or 0. We refer to this as an *early equilibrium*, since all buyers conclude their search as soon as they receive their first offer.

In the other extreme, sellers only offer prices between $\underline{p}$ and R_0, effectively ignoring buyers in state 2. This occurs because state 2 buyers are a relatively small portion of the market and the price drop between $\underline{p}$ and R_2 is too large. Yet the price distribution does not become degenerate because of competition over those buyers who obtain quotes in both state 2 and 1 (having retained the former). We refer to this as a *late equilibrium*.

The intermediate outcome we refer to as a *full equilibrium*. In it, sellers earn the same expected profit whether offering R_2 or any price between $\underline{p}$ and R_0, and all of these are in the support of F.

3 Equilibrium Characterization

In the case of early or late equilibria, we can provide a closed form solution. For the full equilibrium, we must implicitly solve for a, the size of the atom on R_2; all other equilibrium objects are expressed in terms of a. A numerical illustration of this equilibrium is provided in Section 4.3. Conveniently, all three solutions can be expressed in the same equations presented below. The early equilibrium is found by setting $a = 1$, while the late equilibrium uses $a = 0$.

We begin by providing the equilibrium prices. We note that even in a late equilibrium, R_2 is still well defined (though not offered in equilibrium). In essence, this provides the hypothetical price at which state 2 buyers would be willing to immediately accept an offer. The same applies for R_0 in an early equilibrium. In equilibrium, these prices (stated in terms of a^*) are:

$$R_2^* = x - \frac{\beta b + \gamma(x - c - \beta d - \kappa(a^*))}{1 - \beta + \beta\gamma} \tag{15}$$

$$R_0^* = \kappa(a^*) + c, \tag{16}$$

where

$$\kappa(a) \equiv \frac{((x-c)(1-\beta) - \beta d)(1 - \beta + \beta\gamma) + a\beta\lambda(1-\gamma) - a\beta^2\lambda(b-d)}{1 + \beta(\gamma + \lambda - 2 - a\gamma\lambda + \beta(1-\gamma)(1-\lambda) + (1 - \beta + \beta\gamma)\phi(a))},$$

and

$$\phi(a) \equiv \frac{1 - a\lambda\gamma - \lambda(1-\gamma)(1-\beta+\beta\lambda)}{\lambda(1-\gamma)(1+\beta)} \ln\left(1 - \frac{\lambda^2(1-\gamma)(1-a)(1+\beta)}{1 - a\lambda\gamma - \lambda(1-\gamma)(1-\beta-\lambda(1-a(1+\beta)))}\right).$$

We note that $\phi(1) = 0$.

Next, we turn to the equilibrium price distribution. In the continuous portion of the distribution, the lowest price is:

$$\underline{p}^* = \frac{(1-\gamma)\lambda^2(1-a^*)(1+\beta)c + (1 - \gamma\lambda a^* - \lambda(1-\gamma)(1-\beta+\beta\lambda))(\kappa(a^*)+c)}{1 - \gamma\lambda a^* - \lambda(1-\gamma)(1-\beta-\lambda(1-a^*(1+\beta)))},$$

which approaches R_0^* as $a^* \to 1$.

For prices $p \in [\underline{p}^*, R_0^*]$, the cumulative distribution is:

$$F^*(p) = 1 - \frac{(\kappa(a^*) + c - p)(1 - \gamma\lambda a^* - \lambda(1-\gamma)(1-\beta+\beta\lambda))}{(p-c)(1-\gamma)\lambda^2(1+\beta)}, \tag{17}$$

while $F^*(p) = a^*$ for $p \in [R_2^*, \underline{p}^*)$, and $F^*(p) = 0$ for $p < R_2^*$. An example of this distribution, in our numerical illustration, is plotted in Figure 1.

The steady state distribution of buyers is simply:

$$\begin{aligned}
h_2^* &= \delta \quad (18)\\
h_1^* &= \delta(1 - \lambda(1 - \gamma(1 - a^*))) \quad (19)\\
h_0^* &= \delta\frac{1-\lambda}{\lambda}(1 - \lambda(1 - \gamma(1 - a^*))) \quad (20)\\
g_1^*(p) &= \delta\lambda(1-\gamma)F'(p) \quad (21)
\end{aligned}$$

The comparison of profits plays a pivotal role in equilibrium. Let $\Pi_s(a)$ denote the expected profit of a seller offering price R_s, given the distribution of prices and the steady state population of buyers implied by a. Thus, the expected profit from offering R_2 is:

$$\Pi_2(a^*) \equiv \pi(R_2^*) = \lambda\left(x - c - \frac{\beta b + \gamma(x - c - \beta d - \kappa(a^*))}{1 - \beta + \beta\gamma}\right) \tag{22}$$

while offering any price $p \in [\underline{p}, R_0]$ yields:

$$\Pi_0(a^*) \equiv \pi(p) = \lambda\frac{1 - (1-\gamma)(\beta + a^*\lambda) - (1-\gamma)(1-\beta)\lambda}{1 + \lambda(\gamma + \lambda - \gamma\lambda)(1 - a^*)}\kappa(a^*). \tag{23}$$

Proposition 1 establishes that any equilibrium must take this form, with the pivotal question being whether sellers exclusively target early or late buyers, or target both in a mixed strategy. The following proposition assures us that at least one of these will occur; which one depends on parameter values. In particular, expected profits must be consistent with equilibrium; *e.g.*, if all sellers offer R_2^*, this must be weakly more profitable than offering R_0^*.

Proposition 1. *An equilibrium exists, and must satisfy Eqs. 15 through 23, where:*

1. (Early) $a^* = 1$ *only if* $\Pi_2(1) \geq \Pi_0(1)$.

2. (Late) $a^* = 0$ *only if* $\Pi_2(0) \leq \Pi_0(0)$.

3. (Full) $a^* \in (0,1)$ *only if* $\Pi_2(a^*) = \Pi_0(a^*)$.

The proof (provided in the appendix) proceeds by construction, and illuminates the genesis of this solution. For instance, $\phi(a)$ arises while calculating the mean offered price by integrating over the required price distribution. Also, $\kappa(a)$ is a by-product of solving for the reservation prices R_2 and R_0 from Eqs. 3 and 7.

Next we consider the uniqueness of equilibrium. First we note that when $a = 1$ is imposed, Equations 15 through 23 provide a single solution; thus, for any given parameter values, there is at most one early equilibrium. Similarly, there can be at most one late equilibrium.

Concerning the full equilibrium, note that the equal profit condition is highly non-linear in a and includes a large number of parameters. Therefore (aside from the special cases we discuss in Section 4), it is not possible to prove uniqueness of full equilibrium. However, $\Pi_2(a) - \Pi_0(a)$ is continuous in a; if the derivative of this function is also negative at a^*, then then it can only cross 0 once and uniqueness is guaranteed.[3]

Monotone Property. *If there exists an* $a^* \in [0,1]$ *such that* $\Pi_2(a^*) = \Pi_0(a^*)$*, then* $\Pi_2'(a^*) - \Pi_0'(a^*) < 0$.

For given parameter values, $\Pi_2'(a) - \Pi_0'(a)$ is easy to compute. Indeed, we have evaluated this derivative across a fine grid of parameter values,[4] and have consistently found it to be negative when $\beta \in \left(\frac{8}{9}, 1\right)$, $\gamma \in \left[\frac{1}{12}, 1\right]$, and $\lambda \in \left[\frac{1}{8}, 1\right]$, with $a^* \in [0,1]$.

Now we consider whether early, full, or late equilibria might simultaneously exist. Possible outcomes can be expressed in terms of Π_2 and Π_0; once evaluated at 0 and 1, these are simply conditions on the exogenous parameters of the model.

Proposition 2. *Suppose the Monotone Property holds.*

1. If $\Pi_2(1) < \Pi_0(1)$ *and* $\Pi_2(0) > \Pi_0(0)$*, then only the full equilibrium exists.*

2. If $\Pi_2(1) \geq \Pi_0(1)$*, then only the early equilibrium exists.*

---

[3] Intuitively, this condition makes the full equilibrium stable with respect to local deviations. For instance, suppose ϵ additional sellers offered R_2 instead of higher prices. This increases a, and under the Monotone Property, Π_2 falls below Π_0. This would prompt sellers to return to the equilibrium a^*; hence, the market self-corrects after small deviations.

[4] In the technical appendix available at `http://economics.byu.edu/Documents/BrennanPlatt/DeadlinesTA.nb`, we provide details on the expression of the derivative as well as an interactive tool allowing the user to experiment with all possible parameter values.

3. If $\Pi_2(0) \le \Pi_0(0)$, then only the late equilibrium exists.

If the monotone property did not hold, it would be possible to have multiple equilibria. In particular, if $\Pi_2(1) - \Pi_0(1) > 0 > \Pi_2(0) - \Pi_0(0)$, then an early, a late, and full equilibria would simultaneously exist. In practice, this requires rather extreme parameters, such as very low probability of receiving offers (*e.g.* $\lambda < 0.05$) or remarkably impatient buyers and sellers (*e.g.* $\beta < 0.5$).

Before examining more detailed examples, we also offer a useful measure of total welfare. For each transaction, buyer surplus and seller profit combine for a net surplus of $x - c$. In addition, we must include the benefits b or penalties d that buyers accrue during their search. With discounting, the timing of purchase also matters. We compute total expected welfare from the perspective of a new market entrant. That is, we determine the probability of each possible search spell, multiplied by the discounted flow of surplus under that spell. Though we do this from a buyer's perspective, sellers exit at the same rate, and their realized profits are included in the net surplus.

Search ends in state 2 only if the buyer is offered the lowest price, so with probability λa^*. Search ends in state 0 only if the buyer received no state 1 offer, and either in state 2 either received no offer or lost it. This occurs with probability $(1-\lambda)\,(1-\lambda+\gamma\lambda(1-a^*))$. Otherwise search ends in state 1. For those who reach state 0, we also must consider how many periods elapse before an offer is received, to account for penalties and discounting. This results in the following:

$$\begin{aligned}
TW \;&\equiv\; \lambda a^*(b+x-c) + (1-\lambda a^* - (1-\lambda)\,(1-\lambda+\gamma\lambda(1-a^*)))\,(b+\beta(b+x-c)) \\
&\quad +(1-\lambda)\,(1-\lambda+\gamma\lambda(1-a^*))\left(b+\beta b+\beta^2\sum_{t=0}^{\infty}(1-\lambda)^t\beta^t(d+\lambda(x-c))\right) \\
&= (1+\beta-\beta\lambda a^*)b + \frac{1}{1-\beta+\beta\lambda}\Big(\beta^2(1-\lambda)(1-\lambda(1-\gamma(1-a^*)))d \\
&\quad +\lambda(a^*+(1-a^*)\beta(1+(1-\gamma)(1-\beta-\lambda+\beta\lambda)))(x-c).
\end{aligned}$$

4 Equilibrium Behavior

To illustrate the equilibrium behavior, we provide two special cases as well as a numerical example of the general model. The special cases offer valuable insight applicable in the general case. Moreover, they help relate our model to other strands of search literature. In the numerical example, we depict a typical full equilibrium and provide details as to when an early or late equilibrium would emerge.

4.1 Special Case: No Recall

One parameter of particular interest is γ, the likelihood that deferred offers are lost. If taken to the extreme of $\gamma = 1$, this becomes search without recall. We briefly report the solution under this assumption, which has one notable difference compared to the general case: the continuous portion of the distribution collapses to an atom on R_0. This dramatically simplifies the solution, as κ reduces to:

$$\kappa(a) = x - c - \frac{\beta}{1-\beta}\left(d + a\lambda(b-d)\right).$$

The distribution collapses because no one reaches state 1 with a retained offer, and everyone who does reach state 1 has the same reservation price R_0. Thus, if a seller offered slightly less than R_0, he will (still) sell to buyers in states 0 and 1 without picking up additional sales from buyers in state 2.

Again, we let a^* denote the atom on R_2^*; and hence R_0^* is offered with probability $1 - a^*$. Equilibrium prices become:

$$\begin{aligned} R_2^* &= x - \frac{\beta\left(b - \beta(b-d)(1-\lambda a^*)\right)}{1-\beta} && (24) \\ R_0^* &= x - \frac{\beta\left(d + \beta(b-d)\lambda a^*\right)}{1-\beta}. && (25) \end{aligned}$$

The steady state population has the same solution as in Eqs. 18 through 21, setting $\gamma = 1$. The expected profit from offering R_2 is:

$$\Pi_2(a^*) = \lambda\left(x - c - \frac{\beta\left(d - (b-d)(1-\beta+\beta\lambda a^*)\right)}{1-\beta}\right), \tag{26}$$

while offering R_0 yields:

$$\Pi_0(a^*) = \lambda\frac{1-\lambda a^*}{1+\lambda-\lambda a^*}\left(x - c - \frac{\beta\left(d + \beta(b-d)\lambda a^*\right)}{1-\beta}\right). \tag{27}$$

Here, the comparison of profits $\Pi_2(a^*) \gtreqless \Pi_0(a^*)$ simplifies to:

$$x - c \gtreqless \frac{\beta}{(1-\beta)\lambda} \left(\lambda b - (\beta\lambda + \beta - 1)(1 - \lambda a^*)(b - d)\right). \tag{28}$$

If a full equilibrium exists, the condition holds with equality and we obtain a unique value for a^*:

$$a^* = \frac{1}{\lambda} - \frac{(x-c)(1-\beta) - \beta b}{(b-d)\beta(1-\beta-\beta\lambda)}.$$

This constitutes a full equilibrium if and only if $0 < a^* < 1$. If so, the equilibrium prices simplify further to:

$$R_2^* = \frac{(1-\beta)x - \beta(b + \lambda c)}{1 - \beta - \beta\lambda} \tag{29}$$

$$R_0^* = \frac{(1-\beta)x - \beta(d + \lambda c + (b-d)\beta(1+\lambda))}{1 - \beta - \beta\lambda}. \tag{30}$$

The uniqueness of the full equilibrium is obvious in this case, since the condition is linear in a^*, with a derivative $-\frac{\beta(b-d)(\beta\lambda+\beta-1)}{1-\beta}$. Note that the Monotone Property holds if and only if $\lambda > \frac{1-\beta}{\beta}$; that is, the probability of obtaining a quote each period must be greater than the effective interest rate.[5]

The no recall solution (assuming $\lambda > \frac{1-\beta}{\beta}$) also makes the uniqueness claims in Proposition 2 more transparent, since Equation 28 readily indicates when each equilibrium emerges. For instance, if $x - c > \frac{\beta(b\lambda-(b-d)(1-\lambda)(\beta\lambda+\beta-1))}{(1-\beta)\lambda}$, we only obtain the early equilibrium ($a^* = 1$). This is likely to occur if the good is highly valuable relative to its production cost, or the post-deadline utility gap, $b - d$, is fairly large. On the other hand, we only obtain the late equilibrium ($a^* = 0$) if $x - c < \frac{\beta(b\lambda-(b-d)(\beta\lambda+\beta-1))}{(1-\beta)\lambda}$. This is likely to occur in the opposite extreme, when x is close to c, or when $b - d$ is small.

For moderate parameter values, the following will hold:

$$\frac{\beta\left(b\lambda - (b-d)(1-\lambda)(\beta\lambda+\beta-1)\right)}{(1-\beta)\lambda} > x - c > \frac{\beta\left(b\lambda - (b-d)(\beta\lambda+\beta-1)\right)}{(1-\beta)\lambda},$$

in which case, the full equilibrium exists. Note that the left-most term is strictly

[5] If this property fails to hold, there is still at most one full equilibrium, but when the full equilibrium exists, an early and a late equilibrium also exist. This requires either a low λ or a low β. Since this implies massive discounting or very little chance of obtaining the good before the deadline, such a situation is not likely to be of much economic interest.

greater than the right-most term iff $\lambda > \frac{1-\beta}{\beta}$.

One surprising result is that when the buyers' valuation x of the good increases, the average price increases and the range shrinks. In our environment, $R_0 - R_2$ is independent of x; moreover, in the full equilibrium when $\beta\lambda > 1-\beta$, an increase in x leads to a *decrease* in prices.

The latter result is quite surprising, but is driven by reduced concentration of buyers near the deadline. As the good provides more utility, buyers want to secure its utility earlier. This increases their willingness to pay, which would allow sellers to charge more. In a late or early equilibrium, this is the only effect. But in a full equilibrium, this also changes the steady state population; more buyers exit the market before reaching state 1 or 0, so sellers are more inclined to target buyers in state 2 (a^* increases). But this also weakens the desperation of those in state 1 or 0 (they are more likely to be offered R_2), so the sellers cannot extract as much from them. When $\beta\lambda > 1-\beta$, the effect of fewer desperate buyers dominates the effect of higher willingness to pay.

4.2 Special Case: Guaranteed Sampling

Another useful special case is when buyers are guaranteed to observe a quote each period ($\lambda = 1$). If we also eliminate discounting ($\beta = 1$), this bears strong resemblance to many sequential search problems. This special case still permits us to examine the importance of recall in a more transparent setting.

The same equilibrium structure holds, though its expression is greatly simplified. In particular,

$$\kappa(a) \equiv -\frac{2(1-\gamma)(ab+(1-a)d\gamma)}{(1-a)\gamma\left(2(1-\gamma)+\gamma\ln\left(\frac{\gamma}{2-\gamma}\right)\right)}, \tag{31}$$

while prices simplify to

$$\begin{aligned} R_2^* &= d - \frac{b}{\gamma} + \kappa(a^*) + c & (32) \\ R_0^* &= \kappa(a^*) + c & (33) \\ \underline{p}^* &= c + \frac{\gamma}{2-\gamma}\kappa(a^*). & (34) \end{aligned}$$

For prices $p \in [\underline{p}^*, R_0^*]$, the cumulative distribution is:

$$F^*(p) = 1 - \frac{\gamma(1-a^*)}{2(1-\gamma)}\left(\frac{\kappa(a^*)}{p-c} - 1\right), \tag{35}$$

while $F^*(p) = a^*$ for $p \in [R_2^*, \underline{p}^*)$, and $F^*(p) = 0$ for $p < R_2^*$.

The expected profit from offering R_2 is:

$$\Pi_2(a^*) = \left(d - \frac{b}{\gamma} + \kappa(a^*)\right), \tag{36}$$

while offering R_0 yields:

$$\Pi_0(a^*) = \frac{(1-a^*)\gamma\kappa(a^*)}{2-a^*}. \tag{37}$$

Again, the comparison of profits $\Pi_2(a^*) \gtreqqless \Pi_0(a^*)$ simplifies greatly to:

$$\gamma d - b \gtreqqless \frac{(2-a^*-(1-a^*)\gamma)(1-\gamma)(ba^*+(1-a^*)d\gamma)}{(1-a^*)(2-a^*)\left(2(1-\gamma)+\gamma\ln\left(\frac{\gamma}{2-\gamma}\right)\right)}. \tag{38}$$

In this special case, we can also prove that the Monotone Property holds, as the derivative of $\Pi_2(a^*) - \Pi_0(a^*)$ is always negative. Thus, in addition to a^* being unique (when a full equilibrium exists), Proposition 2 applies, and only one type of equilibrium will exist for any given parameters.

Proposition 3. *Let $\beta = \lambda = 1$ and $\gamma < 1$. If there exists an $a^* \in (0,1)$ such that $\Pi_2(a^*) = \Pi_0(a^*)$, then $\Pi_2'(a^*) - \Pi_0'(a^*) < 0$. Thus a^* is unique.*

When a late equilibrium emerges ($a^* = 0$), the resulting price distribution exactly coincides with that in Burdett and Judd (1983) (B&J). There, buyers choose the number of quotes to simultaneously request then choose the best offer received, with no option for additional search. In equilibrium, the buyers are indifferent between requesting one or two quotes, and employ a mixed strategy with fraction q requesting a single quote.

In our setting, a late equilibrium mimics a simultaneous search environment. All buyers defer on the first draw, hoping to choose the best of two; but γ of them lose this draw and are stuck with their second (and only available) quote. While γ plays the role of their q here, the difference is that γ is exogenously given. On the other hand, B&J take the maximum price as exogenous, while R_0 is endogenously determined in our model by the (off-equilibrium) possibility for an additional round of search in state 0.

The full equilibrium shares this flavor of simultaneous search; yet it also permits an atom in the distribution, which could not occur in B&J. Despite having ex-ante identical buyers, the impending deadline makes them differ ex-post: state 2 buyers have a different willingness to pay than state 1 buyers. If a firm decides to target state 2 buyers by offering R_2, they do not concern themselves with the likelihood of being undercut, since the offer will be accepted before a second draw is possible. This sidesteps the tie-breaking which leads to an atomless distribution in simultaneous search (and the upper price range in our model).

Another curiosity in this special case is that an early equilibrium can never exist. This is seen in that $\Pi_0(a) - \Pi_2(a) \to +\infty$ as $a \to 1$. Intuitively, if nearly all sellers offer R_2, a state 2 buyer takes practically no risk in deferring the offer, since he will almost certainly receive the same offer tomorrow (with no discounting between periods). Moreover, by deferring, he will enjoy utility b next period. Thus, since $b > 0$, it is impossible to choose a price R_2 at which state 2 buyers will want to forgo future search.

In a similar vein, if we also add perfect recall ($\gamma = 0$) here, the only possible equilibrium is for all sellers to charge $p = c$. It is literally costless to defer in state 2, since there is no discounting and no chance of losing the first offer, but it does provide additional utility b in state 1. Since all buyers receive two quotes, competition among sellers drives prices to marginal cost, as in B&J. In the general case, however, it is possible to maintain a full equilibrium if enough buyers do not receive a quote in state 2; we offer an example in Section 4.3.

Uncertain recall ($\gamma \in (0, 1)$) has a surprising effect on the full equilibrium. As γ increases, buyers are less likely to retain a deferred offer. Standard intuition suggests that this should allow sellers to extract more rent from the purchase. Indeed, in the late equilibrium, this is exactly what happens. Sellers are less likely to be undercut, so they are able to offer higher prices and earn greater profits.

However, in the full equilibrium, an increase in γ causes a secondary effect of decreasing deadline concentration. The lower probability of retaining deferred offers would, all else equal, encourage state 2 buyers to accept higher offers; this makes R_2 more profitable, drawing more sellers to offer that price. But the ironic consequence of making more R_2 offers is that fewer buyers reach state 1 or 0, and they have better chances of still getting R_2. Thus, this smaller pool of desperate buyers is less desperate than before.

Surprisingly, this secondary effect dominates the first; R_2^* and R_0^* both fall as quotes are *less likely* to be retained. At the same time, $\underline{p}^*$ rises; but the weight that is lost on this low end of the continuous distribution is transferred into the atom at an even lower price, R_2^*. Thus, average prices unambiguously fall as well. These comparative statics can be analytically derived and retain these signs in any full equilibrium.

4.3 Numerical Illustration

To demonstrate the properties of the general model, we present a numerical example with its equilibrium price distribution and comparative statics, letting $x = 1500$, $c = 1000$, $b = 100$, $d = -100$, $\lambda = 2/3$, $\gamma = 1/2$, and $\beta = 0.99$. Under these parameters, $a^* = 0.219$. This set of parameters was chosen such that a unique full equilibrium exists (*i.e.* the Monotone Property applies), and is representative of the many parameterizations we have computed satisfying this property.[6]

First, consider the equilibrium price distribution, depicted in Figure 1. First, we note the atom at $R_2^* = 1,284$; just over a fifth of sellers offer this price. The remaining four-fifths are distributed between $\underline{p}^* = 1,388$ and $R_0^* = 1,579$, with slightly higher density on lower prices ($F''(p) < 0$). Sellers offering $\underline{p}^*$ make the sale roughly 50% more often than those offering R_0^*, since the latter are almost surely undercut whenever a second quote is obtained, while the former are only undercut if the second quote is R_2^*.

The gap between R_2^* and $\underline{p}^*$ is also noteworthy. Sellers offering the former enjoy two benefits. First, they secure a sale on all offers, rather than risking a positive probability of being undercut. Second, they receive the proceeds immediately, while those offering higher prices will have to wait a period before some offers (those to type 2 buyers) are accepted.

In this example, the support of F includes some prices greater than x. It is also possible for the whole support to lie above or below this value. This is determined by the magnitude of penalty d and the likelihood of incurring it. Buyers pay more than x whenever penalties loom large.

[6] Note that a^* is invariant to any proportional change applied to x, c, b, and d together; hence, one may normalize c without loss of generality.

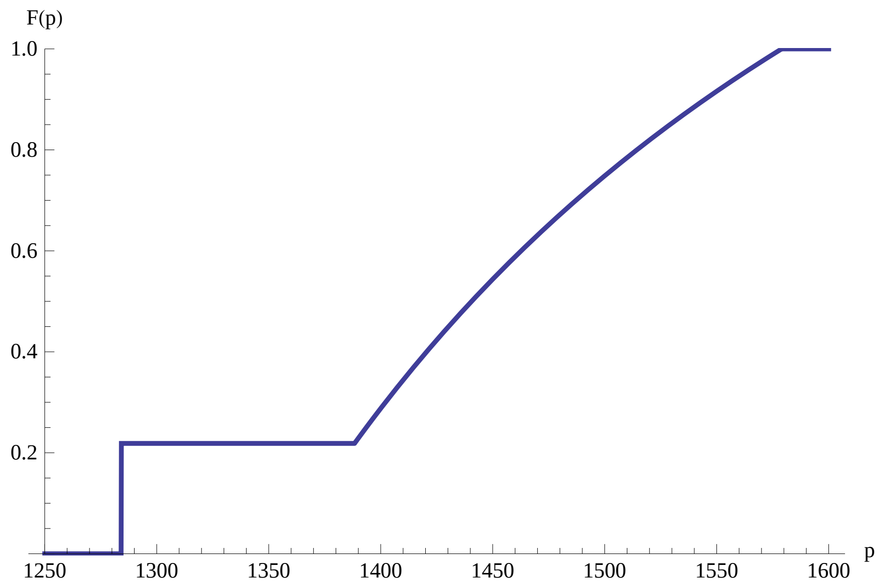

Figure 1: Cumulative distribution of prices offered in equilibrium.

The steady state distribution of buyers is $h_2 = \delta$, $h_1 = 0.59\delta$, $h_0 = 0.30\delta$ and $g_1(p) = 0.33\delta F'(p)$. Note that δ only increases the total population, but not the relative distribution, and thus has no effect on seller incentives.

Table 1 provides a different perspective on the distribution of buyers, indicating the probability of each outcome for a buyer newly entering the market. The majority of buyers accept an offer in state 1, with 15% accepting an offer in state 2 and another 20% in state 0. Roughly 17% of buyers have the opportunity to compare two offers, since others either lost or did not receive an offer in state 2 (or accepted it); this is remarkable, since this small group of comparison shoppers are driving the continuous portion of the price distribution. Also noteworthy is that 9% of buyers are penalized in state 0 despite having received a state 2 offer; this was deferred but lost, and the unfortunate buyer received no state 1 offer either.

Such misfortune reduces ex-ante total welfare as well, which is $TW^* = 650$, while ex-ante expected utility is $V_2^* = 247$. To gauge how the search frictions affects total welfare, contrast this to the equilibrium outcome if buyers had perfect recall of past offers.[7] This results in a late equilibrium with $TW^*_{\gamma=0} = 677$. The efficiency gains are twofold: first, fewer buyers reach state 0, as no one loses a deferred state 2 offer.

[7]In all our computed examples, total welfare strictly increases as γ decreases, even after accounting for the equilibrium response of a^*. The same applies as λ increases.

Table 1: Buyer Search Outcomes

State 2 offer	State 1 offer	Outcome Probability	
Accepted	—	λa^*	0.15
Retained	Received	$\lambda^2(1-\gamma)(1-a^*)$	0.17
Retained	None	$\lambda(1-\lambda)(1-\gamma)(1-a^*)$	0.09
Lost	Received	$\lambda^2\gamma(1-a^*)$	0.17
Lost	None	$\lambda(1-\lambda)\gamma(1-a^*)$	0.09
None	Received	$\lambda(1-\lambda)$	0.22
None	None	$(1-\lambda)^2$	0.11

Second, all buyers now enjoy the benefits in states 2 and 1, rather than exiting early. In sum, partial recall causes inefficiency in the timing of purchases, pushing some to be too early and other to be too late.

On the other hand, if buyers were guaranteed a quote each period (all else equal), the same benefits arise but also no one reaches state 0 or incurs its penalties. In particular, $\lambda = 0$ also results in a late equilibrium, with all buyers purchasing in state 1; this produces $TW^*_{\lambda=0} = 694$.

Next, we consider how smaller changes in parameters impact reservation prices and the price distribution, with comparative statics reported in Table 2. Though unreported, average prices consistently move the same direction as R^*_0 and R^*_2, even when $\underline{p}^*$ moves opposite these. The comparative static on a^* also informs us regarding when the early or the late equilibrium will emerge. Anything causing an increase in a^* will bring us closer to an early equilibrium, which occurs once $a^* = 1$; the reverse leads to a late equilibrium, reached once $a^* = 0$.

In many cases, the comparative statics are quite intuitive. For instance, if the pre-deadline benefits increase, buyers are less inclined to make a purchase in state 2, preferring to receive the added benefits in state 1 as well. Thus, fewer R_2 offers are made, and more buyers end up in state 1 or 0. This allows sellers to offer higher prices, effectively capturing some of the increased benefits by taking advantage of buyer desperation.

In a similar vein, an increase in the post-deadline utility (*e.g.* a smaller penalty for exceeding the deadline) also reduces the haste of buyers to make state 2 purchases. While this also piles more buyers into state 1 and 0, these buyers are also less desperate

Table 2: Comparative Statics for the Full Equilibrium

	$\partial/\partial x$	$\partial/\partial b$	$\partial/\partial d$	$\partial/\partial \gamma$	$\partial/\partial \lambda$
a^*	0.0002	-0.008	-0.008	0.026	-0.026
R_2^*	-0.02	2.34	-0.59	-3.34	-8.32
$\underline{p}^*$	-0.002	2.04	-1.85	0.906	-13.4
R_0^*	-0.03	4.34	-1.59	-7.17	-8.40
V_2^*	1.03	-2.39	0.30	6.81	7.89
TW^*	0.98	2.23	0.64	-1.44	2.67

Notes: Reported values are the resulting change after a \$1 (or for the last two, 1%) increase in the parameter.

than before. The net effect is that sellers cannot extract as much from them, and prices fall.

The surprising results come from increases in x or γ. These mimic the analytically-derived results in the No Recall and Guaranteed Sampling cases, respectively, and are driven by the same underlying forces discussed there. In both cases, the counter-intuitive effect on price arises because of the parameter's indirect effect on deadline concentration. Note that in Equations 18 through 21, the population in state 1 and state 0 falls relative to state 2 as a^* increases. As buyers are discouraged from reaching a desperate state, sellers are unable to extract as much from them, which is reflected in the prices.

These results also highlight the importance of endogenously-determined prices. The standard intuition, insufficient here, has evolved primarily from models with a single price or perhaps an exogenous price distribution. But the equilibrium response can in fact overpower the direct effect. This would be particularly important in policy interventions, such as in the housing market. Tax subsidies for home ownership may increase x, but sellers may not capture any of these subsidies if it draws buyers to purchase earlier, so fewer buyers reach their deadline.

We also note that it is possible to sustain a full equilibrium even with full recall. For instance, if we set $\gamma = 0$ and $b = 3.33$ (leaving all other parameters as above), then $a^* = 0.51$, with $R_2^* = 1,170$, $\underline{p}^* = 1,210$ and $R_0^* = 1,374$. Indeed, with $\gamma = 0$, the full equilibrium occurs in an open set of the remaining parameters. Hence, price dispersion is not solely a product of uncertain recall; rather, it depends on the severity of the deadline balanced against the concentration of buyers near the deadline.

4.4 Empirical Relevance

Aside from our theoretical contributions discussed in Section 5, our model is also relevant to the empirical literature on pricing and sales. To explain the complex dynamics of pricing a commodity, one would likely need to draw on several theories; our model of deadlines offers a particular channel that helps explain which buyers are targeted in particular discounts.

For instance, in Varian (1980), sellers who offer a good at a lower price than competitors do so to target buyers with superior information or lower search costs. Warner and Barsky (1995) test this assertion by examining the price of eight goods at seventeen retail stores in Ann Arbor, Michigan. They find that retail markdowns follow predictable patterns (occurring on weekends, near holiday periods, and in many stores simultaneously); they argue that this is inconsistent with Varian's information channel, unless the fraction of highly-informed buyers ebbs and flows in the same pattern.

Lazear (1986) proposes that because sellers are initially uncertain how well a product will sell, they should introduce the product at a high price and lower it if demand appears to be weak. This is similar to the bargaining solution in Fuchs and Skrzypacz (2011). However, this does not seem to occur in holiday shopping; Pashigian and Bowen (1991) document that the highest markdowns are more likely to be found at the beginning of the season.

Our model offers an alternative explanation for these patterns. Sellers offering low prices are targeting new entrants who have time to spare, while higher prices only snare those who are closest to their deadline. Thus, if many buyers initiate their search over the weekend or on holidays, we would expect sellers to shift towards attracting such customers. At the same time, some remaining buyers will be closer to their deadline; thus, some sellers can maintain higher prices in hopes of capturing this desperate crowd.

In the case of holiday shopping, the deadline is more apparent to sellers than our model assumes, though other buyers could still have earlier or later private deadlines (*e.g.*, for winter birthday or anniversary gifts). Our model still suggests that sellers will increasingly concentrate on higher prices as the season proceeds, since buyers become more anxious as their Christmas deadline approaches. This is consistent with Miyazaki (1993), who finds that consumers' feelings of time pressure increase as Christmas approaches, particularly for shoppers with self-imposed deadlines (*e.g.*, needing to purchase before leaving town on December 22nd).

In addition, several studies confirm our counterintuitive comparative static on x: equilibrium prices fall when a good is more valuable to a buyer. In the DVD market, retailers are more likely to give heavy discounts on the most popular titles (Tang and Xing, 2001). This also holds for hardcover and paperback bestsellers, as noted by Lee and Png (2004). One typically would not expect lower prices on high-demand products, but our model rationalizes this pricing behavior; in equilibrium, fewer buyers delay purchase and more sellers offer a deep discount to target early buyers.

5 Comparison to prior literature

We now examine the related theoretical work to identify our contributions as well as commonalities. These broadly fall into three strands: equilibrium search theory, search with uncertain recall, and bargaining on a deadline.

5.1 Equilibrium Search Theory

Stigler (1961) initiated the formal modeling of search behavior, but only modeled buyer behavior, taking as given the distribution of prices offered by sellers. The goal of equilibrium search theory has been to complete the model, so that sellers also behave optimally, given the search strategy of buyers. Diamond (1971) highlights the difficulty in sustaining more than one price in equilibrium when both buyers and sellers behave optimally. Subsequent work has shown that equilibrium price dispersion can arise if buyers are heterogeneous in their cost of search (Salop and Stiglitz, 1976; Wilde and Schwartz, 1979; Rob, 1985; Butters, 1977) or their valuations of the good (Diamond, 1987). Burdett and Judd (1983) demonstrate this can also arise with simultaneous search.

In our model, all buyers are ex-ante identical. Of course, heterogeneity is necessary to justify sellers offering different prices, but it arises ex-post as some buyers experience longer search spells than others. Luck plays a role in creating these differences, as a buyer may not encounter an opportunity to buy; but the buyer is responsible as well, as he chooses the reservation price strategy that may lead him to defer (and potentially lose) high prices offered in state 2. These differences in buyer types are endogenously created by adding a very natural feature of deadlines to the search environment.

We also note that a large parallel literature has applied equilibrium search environments to labor markets[8] with the intent to study wage formation and the effects of unemployment insurance. The mechanisms mentioned above also can sustain wage dispersion; Burdett and Mortensen (1998) add on-the-job search as an additional mechanism. That is, workers continue to receive job offers while employed, but can only be enticed to switch jobs if the new wage exceeds their current wage, so it shares a flavor of the price comparisons occurring in our model.

The job search process is also influenced by deadlines, which are naturally imposed by most unemployment insurance systems by cutting off benefits after a fixed period of time. This institutional detail alone can create wage dispersion among ex-ante identical workers, as shown in Akin and Platt (2012). That model has similar behavior to the No Recall case presented in Section 4.1. For instance, an equilibrium can be characterized by whether firms target workers who are early or late in their search spell, and responds similarly to changes in analogous parameters.

We extend that work here by adding the possibility of recall. While comparative statics mostly behave the same as in the no recall setting, partial recall also allows for the comparison of active offers, causing sellers to offer a continuum of prices aimed at those near their deadline. The other key difference is methodological, in that labor search was modeled in continuous time. This allowed the equilibrium conditions to be translated into a solvable system of differential equations. When any degree of recall is incorporated, however, this system becomes intractable.[9] The discrete environment here uses a distinct though parallel solution method. It is also possible to incorporate more periods before the deadline, though at greater computational complexity than in the no recall labor environment.[10]

[8] This literature is well surveyed in Rogerson, *et al* (2005).

[9] An obvious complication is that retained offers must be included as a state variable. In addition, the steady state conditions become much more intricate, as buyers exit due to both current and past offers.

[10] If our discrete model in the no recall case is expanded by a period, a similar result emerges. In equilibrium, the price distribution F only places weight on one of the following sets: $\{R_3\}$, $\{R_3, R_2\}$, $\{R_3, R_2, R_0\}$, $\{R_2, R_0\}$, or $\{R_0\}$. Of course, this extra period also multiplies the number of conditions one must check to determine which equilibrium occurs.

5.2 Search with Recall

Another important contribution of our paper is to fill the gap between no recall and perfect recall in search models. While these two extremes are regularly studied, the intermediate case has received less attention, and even then, taking the distribution of prices as exogenously given.

The closest analog to our environment is found in Landsberger and Peled (1977); indeed, they depict uncertain recall as a constant probability that the best offer so far will be available in the next period.[11] They assume no discounting, and that an offer is received each period, as in our Guaranteed Sampling case.

If one can receive new offers indefinitely, the reservation price remains constant. In our equilibrium search model, sellers would only offer that single price; thus, deadlines are essential to price dispersion in our model. Landsberger and Peled impose an analog to deadlines by capping the number of offers one may receive; this leads to an increasing reservation price. Even so, their assumption of a fixed price distribution leads to very different comparative statics. For instance, they conclude that reservation prices will fall as the probability of recall increases; but this neglects the indirect effect on the concentration of buyers near their deadline, which produces the opposite result in our model.

In Janseen and Parakhonyak (2008), buyers are always able to recall a past price, but incur a cost to do so. Thus, recall is not always exercised. Even though no deadline is imposed, this results in a non-stationary problem, with reservation price depending on the full history.

Another strand of literature endogenizes recall in an environment with two sellers that decide how long to honor their own quotes. In Armstrong and Zhou (2010), a seller can identify buyers who return to their store (after seeking a second quote at the other seller). The seller must decide whether to honor the original price, or make the initial offer contingent on immediate purchase. This typically results in no recall, in that new customers are offered a discount that is unavailable to returning customers.

In Daughety and Reinganum (1992), sellers have no knowledge of buyers' search history, and are not guaranteed to have the item in stock. If he does, the seller offers

[11] In Karni and Schwartz (1977), the probability of retaining an offer is more flexible, allowing it to diminish if more time has elapsed since receiving the quote. Also, they consider the case in which buyers learn about the distribution of prices over time.

a price and states how many periods he will hold the item. The buyer can then check the other seller (at a cost). If both sellers have the item, they engage in Bertrand competition, driving price to marginal cost. In equilibrium, sellers sometimes offer recall for a finite but positive number of periods. Only two prices emerge: the price everyone quotes and the marginal cost if a second quote is obtained.

In contrast, our model produces a richer distribution of prices, making search more meaningful. Also, in our full equilibrium, buyers accept some offers immediately, rather than always deferring in hopes of sparking a price war. If γ in our model were set to maximize seller profits, this would typically occur by lowering γ until the equilibrium transitions from from full to early, since the comparative static on R_2^* also indicates the effect on expected profits and changes sign at that transition. Thus, firms would favor some degree of recall, though not necessarily perfect recall.

5.3 Bargaining on a Deadline

In the literature on bilateral bargaining with a deadline, the closest analog to our environment is Fuchs and Skrzypacz (2011).[12] There, a seller makes repeated offers until the buyer accepts. If the deadline is reached, both parties receive an exogenous fraction of the buyer's valuation. In addition to the possible loss of surplus after the deadline, discounting motivates both parties to complete the transaction earlier.

The equilibrium strategy has the seller make a decreasing sequence of price offers; only a buyer with the highest valuations would be willing to accept initial offers. If these early offers are rejected, the buyer is revealed to have a lower valuation; furthermore, the deadline is closer. Thus the seller offers successively lower prices. In spite of the lost surplus, a strictly positive fraction of trades occur after the deadline.

There are three important differences from our environment: first, the buyer's valuation is not known but the deadline is. We reverse these, so that the seller cannot observe the time until deadline, but knows the buyer's valuation of the good. Second, one seller and one buyer are engaged in repeated interaction, rather than a one-shot encounter. In our environment, a seller cannot learn anything about the buyer's state until it is too late; we assume there is no opportunity for that seller to make a second offer, though the seller does hope the buyer retains and uses the original offer. Third, our seller does not directly face a deadline, and only cares about

[12] This in turn builds on Sobel and Takahashi (1983), which assumed that if the deadline passes without agreement, both parties receive a zero payoff.

the (unknown) deadline of her current customer because it affects his willingness to pay.

In both their bargaining and our search environments, some buyers will refuse offers in equilibrium. Indeed, with positive probability, some buyers will not make a purchase before the deadline. However, in the bargaining model, buyers who accept an early offer pay the highest prices; in ours, they pay the lowest prices. Also, as the delay between offers goes to zero, the limit equilibrium has a continuum of transaction prices, including an atom on the lowest price. Our price distribution has similar features, but is caused by the comparison of retained and new offers rather than learning the buyer's type.

In Ma and Manove (1993), the agents make alternating offers on how to split a commonly known surplus, which evaporates after a known deadline. When an agent makes an offer, he can strategically delay how quickly the offer is received (giving the other agent less time to reject and make a counter-offer); at the same time, nature adds a random delay on top of this, creating the risk that the deadline will be crossed. In equilibrium, early offers are rejected, agreements are reached late in the game, and the deadline is missed with positive probability.

Our model serves as a bridge between the literature on bargaining on a deadline and equilibrium search, but admittedly other assumptions could be used. The only existing alternative is found in Albrecht, *et al* (2007), which depicts consumer search where both buyers and sellers become desperate. Their model differs from ours in two key aspects. First, the deadline arrives randomly, at a given Poisson rate, rather than deterministically after a fixed period of time. Second, they assume that buyers and sellers can observe each others' states on meeting, and then determine prices via Nash bargaining.[13] In our incomplete information environment, a seller must make a probabilistic judgment as to the likelihood of its product being purchased when it posts a particular price.

As a consequence, their model responds differently to parameter changes. For instance, when the value x of the good to buyers increases, average price and the range rises; in our full equilibrium, prices fall and the range is unchanged. These differences are mostly driven by the deterministic versus random transition to the desperate state.

[13]If a pair of buyers and sellers cannot reach a mutually beneficial trade, both continue search with no ability to recall the meeting once their state changes.

6 Conclusion

We analyze price formation in a market where buyers face a deadline to complete their transaction and have limited ability to recall their best prior offer. Sellers are uninformed of the state of a given buyer; therefore, posting a given price results in a tradeoff between the probability of acceptance and the realized profit if accepted. In this incomplete information environment, we solve for the endogenous price distribution that prevails in the steady state.

We show that the equilibrium takes one of three forms. In a late equilibrium, buyers initially forgo any purchases, preferring to wait until the deadline. In an early equilibrium, buyers accept any offer as soon as it is received. In a full equilibrium, buyers early in their search are willing to accept some of the offers, but higher prices are turned down until near the deadline. This carries the risk of losing such offers; in equilibrium, some buyers will cross their deadline having lost an offer that is now strictly preferred.

In a full or late equilibrium, our model generates a continuum of offered prices. This occurs because fortunate buyers are able to compare their best prior quote to their new offer. This feature of our sequential model with recall mimics the behavior of simultaneous search models where individuals get multiple quotes in a single period and choose the lowest price.

We identify a deadline concentration effect that gives rise to interesting comparative statics in equilibrium. For example, if the good becomes more valuable to the buyers, one typically expects sellers to extract more surplus through higher prices. Yet the buyers' eagerness to quickly obtain the good will reduce deadline concentration, causing prices to fall. Prices also fall when the ability to recall past offers is reduced. One would expect less recall to give more power to sellers, who have less chance of being undercut. Yet this also raises the incentive of buyers to accept offers earlier in their search spell, so fewer reach their deadline.

It is relatively easy to incorporate several features in the model. For instance, we could assume delayed consumption, so that buyers in state 2 or 1 who make a purchase do not enjoy utility x until the date of the deadline. This might apply to housing search, where often the new home is not needed until a certain date; an early purchase results just leave the new home vacant in the interim. This feature results in minor changes to the equilibrium conditions, but still produces similar overall behavior.

In addition to markets previously mentioned, our model could potentially be adapted to study other interesting deadlines. For instance, sellers may face a deadline in pricing a commodity over time. Copeland, *et al* (2011) present evidence that in the automobile market, manufacturers reduce the price of an automobile model over the year, as they feel the pressure to eliminate inventory before next year's model is introduced to the market. The model release cycle induces a deadline, after which the prior model will drop in value.

An important biological deadline occurs in marriage markets, as female fertility declines with age. In contrast, male fertility remains roughly constant. Both face uncertainty about future prospects and how long current suitors will wait for a decision. Thus, the marriage search process could be characterized by a reservation quality (of spouse), which would decline as the deadline approaches. In Akin, Butler, and Platt (2011), we show that a dynamic search model with deadlines can explain interesting facts found in the U.S. census regarding variations in spouse quality by a woman's age at first marriage.

A Proofs

A.1 Proof of Lemma 1, Claim 3

Proof. Suppose an atom of weight $a > 0$ existed at price $p \in (R_2, R_0]$. A seller offering p would find that a positive measure of their offers are tied with another offer of p. In particular, this occurs with probability $(1-\gamma)\lambda a > 0$, since λa is the probability that another offer is made at price p, and $1-\gamma$ is the probability that the buyer retains the first offer made.

If this happens, the buyer randomizes with equal weight between the two offers. If the seller had offered $p-\epsilon$, where $\epsilon > 0$ is arbitrarily small, he would have garnered $\frac{(1-\gamma)\lambda a}{2}$ more sales among buyers in state 2, without losing any of the other $S(p)$ sales from offering p. This discrete jump in expected sales (with a negligible reduction in profit per sale) means offering p is strictly dominated. Thus no firm would offer p and thus cannot be part of the support.

While we have assumed equal probability in randomizing over the two offers, other tie-breaking rules have similar effect, as in in Bertrand competition. $\square$

A.2 Proof of Proposition 1

Proof. We begin by restating the buyer's value functions, using the fact that F has at most one atom of size a at R_2, and a continuous distribution between $\underline{p}$ and R_0:

$$\begin{aligned}
V_0 &= d + \lambda a(x - R_2) + \lambda \int_{\underline{p}}^{R_0} (x-q)F'(q)dq + \beta(1-\lambda)V_0 \\
W_0(R_0) &= d + \lambda a(x - R_2) + \lambda \int_{\underline{p}}^{R_0} (x-q)F'(q)dq + (1-\lambda)(x-R_0) \\
V_1 &= b + \lambda a(x - R_2) + \lambda \int_{\underline{p}}^{R_0} (x-q)F'(q)dq + \beta(1-\lambda)V_0 \\
V_2 &= b + \lambda a(x - R_2) + \beta\lambda \int_{\underline{p}}^{R_0} \left((1-\gamma)W_1(q) + \gamma V_1\right) F'(q)dq + \beta(1-\lambda)V_1 \\
W_1(R_2) &= b + x - R_2.
\end{aligned}$$

Next, note that $W_0(R_0) = (1-\beta+\beta\lambda)V_0 + (1-\lambda)(x-R_0)$. By substituting this into the indifference condition, Eq. 3, we get:

$$R_0 = x - \beta V_0. \tag{39}$$

Also, a comparison of V_0 and V_1 reveals that $V_1 = b - d + V_0$. Substituting for V_1 and $W_1(R_2)$ in the indifference condition, Eq. 7, we obtain:

$$R_2 = x - \frac{\beta(b - \gamma(d - V_0))}{1 - \beta + \beta\gamma}. \tag{40}$$

Also, the first value function above allows us to solve for V_0:

$$V_0 = \frac{d + \lambda\left(x - aR_2 - \int_{\underline{p}}^{R_0} qF'(q)dq\right)}{1 - \beta + \beta\lambda}. \tag{41}$$

Next, we turn to the steady state population. Note that Eqs. 9 though 12 are linear in terms of h_2, h_1, h_0, and $g_1(p)$, and easily solve to obtain Eqs. 18 through 21.

To obtain the price distribution, we turn to the equal profit conditions. Of course, offering R_2 results in profit $\lambda(R_2 - c)$. After substituting for the steady state populations, for prices between $\underline{p}$ and R_0, profit (in Eq. 14) is:

$$\pi(p) = \lambda\left(1 - \lambda\frac{1 - (1-\gamma)\left(\beta + a\lambda - \lambda(1+\beta)F(p)\right)}{1 + \lambda(\gamma + \lambda - \gamma\lambda)(1-a)}\right)(p - c). \tag{42}$$

We then solve for $F(p)$ in this price range by equating $\pi(p) = \pi(R_0)$. Note that because this portion of the distribution is atomless and R_0 is the highest offered price, $F(R_0) = 1$. This yields:

$$F(p) = 1 - \frac{(R_0 - p)(1 - \gamma\lambda a - \lambda(1-\gamma)(1 - \beta + \beta\lambda))}{(p - c)(1-\gamma)\lambda^2(1+\beta)}, \tag{43}$$

and taking the first derivative w.r.t. p provides:

$$F'(p) = \frac{(R_0 - c)(1 - \gamma\lambda a - \lambda(1-\gamma)(1 - \beta + \beta\lambda))}{(p - c)^2(1-\gamma)\lambda^2(1+\beta)}. \tag{44}$$

In addition, we must determine the value of $\underline{p}$. This is pinned down by the fact that

$F(\underline{p}) = a$, since no prices between R_2 and $\underline{p}$ are offered. This provides the solution:

$$\underline{p} = \frac{(1-\gamma)\lambda^2(1-a)(1+\beta)c + (1-\gamma\lambda a - \lambda(1-\gamma)(1-\beta+\beta\lambda))R_0}{1-\gamma\lambda a - \lambda(1-\gamma)(1-\beta-\lambda\,(1-a(1+\beta)))}. \tag{45}$$

With the price distribution pinned down, we can now compute expected prices in Eq. 41 to express V_0 in terms of a, R_0 and R_2. The evaluation of the integral introduces the term $\phi(a)$, defined in the text.

$$V_0 = \frac{d + \lambda\,(x - c - a(R_2 - c)) + (R_0 - c)\phi(a)}{1-\beta+\beta\lambda}. \tag{46}$$

Notice that Equations 39, 40, and 46 provide three linear equations in terms of R_0, R_2, and V_0; of course, only the former two are of explicit interest. These jointly solve to obtain the proposed R_0^* and R_2^*; the $\kappa(a)$ term arises in solving this system.

Note that we obtain the proposed profit functions and price distribution by substituting R_0^* and R_2^* into Eqs. 42, 43, and 45.

If the respective profit condition holds, this solution is an equilibrium by construction. Existence is assured because $\Pi_2(a) - \Pi_0(a)$ is continuous in a. Thus, if $\Pi_2(1) - \Pi_0(1) < 0$ and $\Pi_2(0) - \Pi_0(0) > 0$, there must exist an $a^* \in (0,1)$ such that $\Pi_2(a^*) = \Pi_0(a^*)$. Of course, if $\Pi_2(1) - \Pi_0(1) \geq 0$ or $\Pi_2(0) - \Pi_0(0) \leq 0$, then an early or late equilibrium exists, respectively.

Finally, this proof by construction also demonstrates that no other equilibrium (beyond these three) can exist, as only the constructed solution will satisfy all equilibrium requirements. □

A.3 Proof of Proposition 2

Proof. Let $D(a) = \Pi_2(a) - \Pi_0(a)$, which is continuous on $a \in [0,1]$.

For claim 1, consider that $D(1) < 0$ and $D(0) > 0$ rule out an early and a late equilibrium, respectively; thus, existence of a full equilibrium follows from Proposition 1. This does not depend on the Monotone Property except to give uniqueness of the full equilibrium.

For claim 2, by Proposition 1 the early equilibrium exists. Suppose the late equilibrium also existed. This can only occur if $D(0) \leq 0$. But then $D(0) \leq 0 \leq D(1)$. By the intermediate value theorem, there must be an $a^* \in (0,1)$ such that $D(a^*) = 0$.

But since $D(0) \leq D(1)$, we would need $D'(a^*) = \Pi_2'(a^*) - \Pi_0'(a^*) \geq 0$ for at least one such a^*. This contradicts the Monotone Property.

Suppose instead that no late equilibrium existed ($D(0) > 0$), but a full equilibrium did for some $a^* \in (0, 1)$. By the Monotone Property, this a^* is unique (which rules out $D(1) = 0$) and $D'(a^*) < 0$. Since $D(0)$ and $D(1)$ are strictly positive and $D(a)$ is continuous, either $D'(a^*) = 0$ (reaches its minimum at a^*) or there exists multiple solutions to $D(a) = 0$, at least one of which has $D'(a^*) > 0$. Either contradicts the Monotone Property.

Claim 3 follows similar logic to preclude the existence of an early or a full equilibrium when a late equilibrium exists. □

A.4 Proof of Proposition 3

Proof. Suppose $\Pi_2(a^*) = \Pi_0(a^*)$ for some $a^* \in (0, 1)$. If we solve Eq. 38 for b, this yields:

$$b = \frac{(1-a^*)d\gamma^2 \left(2(1-a^*)(1-\gamma) + (2-a^*)\ln\left(\frac{\gamma}{2-\gamma}\right)\right)}{2(1-\gamma)(2-a^*-(1-a^*)a\gamma) + (2-a^*)(1-a^*)\gamma\ln\left(\frac{\gamma}{2-\gamma}\right)}.$$

We then take the first derivative of $D(a) \equiv \Pi_2(a) - \Pi_0(a)$ w.r.t. a, and evaluate it at a^*. This allows us to substitute for b using the equation above. This results in

$$\begin{aligned} D'(a^*) \quad = \quad & 2(1-\gamma)\gamma^2 d \cdot \\ & \left(2(1-\gamma)(1-a)(3-a(1-\gamma)-\gamma) - \left(\gamma\,(1-a)^2 - (2-a)^2\right)\ln\left(\frac{\gamma}{2-\gamma}\right)\right) \Big/ \\ & \left((2-a)(1-a)\left(2(1-\gamma) + \gamma\ln\left(\frac{\gamma}{2-\gamma}\right)\right) \cdot \right. \\ & \left.\left(2(1-\gamma)\left((1-a)a\gamma - 2 + a\right) - (2-a)(1-a)\gamma\ln\left(\frac{\gamma}{2-\gamma}\right)\right)\right). \end{aligned}$$

The term $2(1-\gamma)\gamma^2$ is always positive. The next term in the numerator is strictly increasing[14] in γ, with a value of zero when $\gamma = 1$. Thus, since $d < 0$ by assumption, the numerator is positive for any $a^* \in (0, 1)$.

[14]This and the following assertions are demonstrated in the technical appendix.

In the denominator, $(2-a)(1-a)$ is always positive. The next term is strictly decreasing in γ, with a value of 0 at $\gamma = 1$; hence, it is always positive. Finally, the last term is always negative, as it is strictly increasing in γ, with a maximum of 0 at $\gamma = 1$.

Thus, with a positive numerator and negative denominator, $D'(a^*) < 0$ so long as $\gamma < 1$. $\square$

References

Akin, S. Nuray, M. Butler, and B. Platt (2011) "Racing the Biological Clock: Accounting for Trends in Spousal Choice in First Marriage," Working Paper, University of Miami.

Akin, S. Nuray, and B. Platt. (2011) "Searching on a Deadline," Working Paper, University of Miami.

Akin, S. Nuray, and B. Platt. (2012) "Running Out of Time: Limited Unemployment Benefits and Reservation Wages," *Review of Economic Dynamics*, forthcoming.

Albrecht, James, A. Anderson, E. Smith, and S. Vroman. (2007) "Opportunistic Matching in the Housing Market," *International Economic Review* **48**, 641-664.

Armstrong, Mark and J. Zhou. (2010) "Conditioning Prices on Search Behaviour," ELSE Working Paper 351.

Burdett, Kenneth, and K. Judd. (1983) "Equilibrium Price Dispersion," *Econometrica* **51**, 955-970.

Burdett, Kenneth, and D. Mortensen (1998) "Wage Differentials, Employer Size, and Unemployment," *International Economic Review* **39**,257-273.

Butters, Gerard. (1977) "Equilibrium Distribution of Sales and Advertising Prices," *Review of Economics Studies* **44**, 465-491.

Copeland, Adam, G. Hall, and W. Dunn. (2011) "Inventories and the Automobile Market," *Rand Journal of Economics*, forthcoming.

Daughety, Andrew F. and J. F. Reinganum. (1992) "Search Equilibrium with Endogenous Recall," *The RAND Journal of Economics* **23**, 184-202.

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

Diamond, Peter. (1987) "Consumer Differences and Prices in a Search Model," *Quarterly Journal of Economics* **102**, 429-436.

Fuchs, William, and A. Skrzypacz. (2011) "Bargaining with Deadlines and Private Information," Mimeo, University of California Berkley.

Janseen, Maarten C. W. and A. Parakhonyak. (2008) "Consumer Search with Costly Recall," Tinbergen Institute Discussion Paper, TI 2008-002/1.

Karni, Edi, and A. Schwartz. (1977) "Search Theory: The Case of Search with Uncertain Recall," *Journal of European Theory* **16**, 38-52.

Landsberger, Michael, and D. Peled. (1977) "Duration of Offers, Price Structure, and The Gain From Search," *Journal of European Theory* **16**, 17-37.

Lazear, Edward. (1986) "Retail Pricing and Clearance Sales," *American Economic Review* **76**, 14-32.

Lee, Sang-Yong, and I. P. L. Png. (2004) "Buyer Shopping Costs and Retail Pricing: An Indirect Empirical Test," *Review of Marketing Science* **2**, 1-20.

Ma, Ching-to Albert, and M. Manove. (1993) "Bargaining with Deadlines and Imperfect Player Control," *Econometrica* **61**, 1313-1339.

Miyazaki, Anthony. (1993) "How many shopping days until Christmas? A Preliminary Investigation of Time Pressures, Deadlines, and Planning Levels on Holiday Gift Purchases," *Advances in Consumer Research*, Vol. 20. Association for Consumer Research:Provo, UT. 331-335.

Pashigian, Peter, and B. Bowen. (1991) "Why are Products Sold on Sale? Explanations on Pricing Regularities," *Quarterly Journal of Economics* **106**, 1015-1038.

Rob, Rafael. (1985) "Equilibrium Price Distributions," *Review of Economic Studies* **52**, 487-504.

Rogerson, Richard, R. Shimer, and R. Wright. (2005) "Search-Theoretic Models of the Labor Market: A Survey," *Journal of Economic Literature* **43**, 959-988.

Roth, Alvin, J. Murnighan, and F. Schoumaker. (1988) "Deadline Effect in Bargaining: Some Experimental Evidence," *American Economic Review* **78**, 806-823.

Salop, Steven, and J. Stiglitz. (1976) "Bargains and Ripoffs: A Model of Monopolistically Competitive Prices," *Review of Economic Studies* **44**, 493-510.

Sobel, Joel, and I. Takahashi. (1983) “A Multistage Model of Bargaining,” *Review of Economic Studies* **50**, 411-426.

Stigler, George. (1961) “The Economics of Information,” *Journal of Political Economy* **69**, 213-225.

Tang, Fang-Fang and X. Xing. (2001) “Will the Growth of Multi-Channel Retailing Diminish the Pricing Efficiency of the Web?” *Journal of Retailing* **77**, 319-331.

Varian, Hal. (1980) “A Model of Sales,” *American Economic Review* **70**, 651-659.

Warner, Elizabeth, and R. J. Barsky. (1995) “The Timing and Magnitude of Retail Store Markdowns,” *Quarterly Journal of Economics* **110**, 321-352.

Wilde, Louis, and A. Schwartz. (1979) “Equilibrium Comparison Shopping,” *Review of Economic Studies* **46**, 543-554.