Revenue Management with Forward-Looking Buyers*

Simon Board† Andrzej Skrzypacz‡

November 29, 2010

Abstract

We consider a seller who wishes to sell K goods by time T. Potential buyers enter IID over time and are forward-looking, so can strategically time their purchases. At any point in time, profit is maximized by awarding the good to the agent with the highest valuation exceeding a cutoff. These cutoffs are characterized by a one-period-look-ahead rule and are deterministic, depending only on the number of units and time remaining. The cutoffs decrease over time and in the inventory size, and are independent of the buyers' arrival times. In the continuous time limit, the seller's profits are maximized by posting anonymous prices, with an auction for the last unit at time T. Unlike the cutoffs, the optimal prices depend on the history of past sales.

1 Introduction

We derive the optimal strategy of a seller who wishes to sell K goods by time T to forward-looking buyers, and can commit to a dynamic mechanism. Potential buyers enter the market stochastically over time and possess privately known values and arrival times. Once they arrive, buyers prefer to obtain the good sooner rather than later, but can strategically time their purchases, incurring a costly delay and risking a stock-out in the hope of lower prices.

Dynamic pricing is a topic of a substantial literature on revenue management (see the book by Talluri and van Ryzin (2004)). It is estimated that these techniques have led to a substantial increase in profits for airlines (Davis (1994)), retailers (Friend and Walker (2001))

*We thank Jeremy Bulow, Yeon-Koo Che, Songzi Du, Drew Fudenberg, Willie Fuchs, Alex Frankel, Mike Harrison, Jon Levin, Yair Livne, Rob McMillan, Moritz Meyer-ter-Vehn and Benny Moldovanu for helpful comments. We also thank seminar audiences at Bologna, CSEF (Naples), EIEF (Rome), ES World Congress, EUI, Microsoft, Midwest Meetings (Northwestern), Stanford, SWET 2010, UCLA DOTM, WCU-Economics Conference (Yonsei), Yahoo! and Yale. A previous version of this paper went by "Optimal Dynamic Auctions for Durable Goods: Posted Prices and Fire-sales". JEL: D44, L12. Keywords: Dynamic Mechanism Design, Revenue Management, Durable Goods.

†Department of Economics, UCLA. http://www.econ.ucla.edu/sboard/

‡Graduate School of Business, Stanford University. http://www.stanford.edu/∼skrz/

and car manufacturers (Coy (2000)).[1] These models typically assume that buyers are impatient, exiting the market if they do not immediately buy. Our contribution is a full analysis of the problem with buyers who are forward-looking. This seems natural when considering markets such as airline tickets and cars, where buyers can easily time their purchases. It is also becoming more important as buyers use price prediction tools to aid such inter-temporal arbitrage (e.g. bing.com, where searches for flights results contain predictions about changes in prices).[2]

We derive our results in two stages. First, we consider the set of all dynamic selling mechanisms and use mechanism design to characterize the profit-maximizing allocations. Second, we show how to decentralize these allocations through posted prices, potentially with an auction for the last unit at time T. By tackling the problem in two stages, we significantly simplify this complex dynamic pricing problem. When the seller changes the price at time t, this affects both earlier and later sales. By using mechanism design these effects are built into the marginal revenues and the problem collapses to a standard single-agent dynamic programming problem.

Under the optimal allocation, the seller awards a good to the agent with the highest valuation, if their value exceeds a cutoff. The optimal cutoffs are *deterministic*, depending on the number of units and time remaining, but not on the number of buyers, their values, or when previous units were sold. This property is surprising: the presence of forward-looking buyers means that the seller must carry around a large state variable corresponding to the reservoir of potential buyers; however, this state does not affect the seller's optimal cutoff. Intuitively, the seller's decision to delay allocating a good does not affect when lower value agents buy, which only depends on their valuations and ranks. Hence changing these agents' values raises the profits from selling and delaying equally and does not affect the cutoff type. This property is robust to a number of different extensions, but is not automatic: it does not hold if agents exit the market or if there is a mixture of patient and impatient buyers.

In a general dynamic program, the seller would be indifferent between serving the cutoff type today and delaying. We show that the optimal cutoffs are determined by a *one-period-look-ahead policy*, whereby the seller is indifferent between serving the cutoff type today and waiting exactly one more period, potentially selling to a new entrant. The optimal cutoffs decrease over time and in the inventory size. Intuitively, if the seller delays awarding the k^{th} unit by one period then she can allocate it to the highest value entrant rather than the current leader. As k rises or we approach time T, the current leader is increasingly likely to be awarded the good eventually, decreasing the option value of delay and causing the cutoff to fall. Finally, the optimal cutoffs are independent of the agents' arrival times. Agents' arrival times are

[1] Robert Crandall, former-CEO of American Airlines, has said "I believe yield management is the single most important technical development in transportation management since we entered the era of airline deregulation in 1979 [...] We expect yield management to generate at least $500 million annually for the foreseeable future." (Smith, Leimkuhler, and Darrow (1992, p. 26))

[2] For evidence of forward looking behaviour see Zeithammer (2006) on eBay auctions, Hartmann (2006) on golf rounds, Chevalier and Goolsbee (2009) on textbooks or Gowrisankaran and Rysman (2009) on camcorders.

uncorrelated with their values so the incentive compatibility constraints on these arrival times are slack.

For the second set of results we ask how to achieve the optimal allocation using a natural mechanism. The optimal cutoffs are deterministic and, in the continuous time limit with buyers arriving via a Poisson process, do not jump down at any time $t < T$. As a result, these cutoffs can be implemented by posting anonymous prices. That is, we show how to compute prices such that optimal consumer responses induce the optimal allocations and, by revenue equivalence, the maximal expected profit to the seller. The deterministic cutoffs are crucial for the simplicity of this solution: if the decision to award agent 1 the good depended on agent 2's value then the seller would have to obtain reports about all agents' values upon entering the market, requiring indicative bidding of some sort.

The optimal prices depend on the number of units and time remaining, and, unlike the optimal cutoffs, also depend on the time of previous sales. Prices drift down over time if there are no sales and jump up with every unit sold. As time T approaches, there is a "fire-sale": prices of units $k \geq 2$ fall quickly; and if one unit remains there is an auction at T.

The pattern of posted prices and a "last-minute" auction is qualitatively consistent with internet sites selling plane and hotel reservations. Similarly, in the secondary market for baseball tickets, Sweeting (2010) shows that prices decline by 60% in the month before the game, with the price decline accelerating, the probability of sale increasing and auctions becoming more popular as the game day approaches. Sweeting also finds evidence that buyers strategically time their purchases: fans are more likely to delay if they live closer to the stadium or buy fewer tickets for other games.

The practical problems that we model have two key properties. First, the pricing problem is non-stationary. In our examples, the good may expire at a fixed date (e.g. plane tickets, online advertising slots), become much less valuable (e.g. seasonal clothing), or the number of interested buyers may decline over time (e.g. a house). Second, the total surplus is higher if the buyer and seller trade sooner. In the case of clothing or a house, the buyer has more days to enjoy the good; in case of online advertising slots, the advertiser has more time to plan complementary activities such as scheduling production, delivery and other advertising; in case of physical goods, the seller may incur inventory costs. Both these properties have important impact on the results. First, in stationary environments, the optimal prices never decline so buyers either buy immediately or never, while our buyers may delay on the equilibrium path. Second, without a cost of delay the optimal mechanism would involve waiting until the deadline and running an optimal auction, rather than having the market open continuously.

The paper proceeds as follows. Section 2 describes the model with proportionate discounting. Section 3 considers the one-unit case, and shows how to implement the optimal cutoffs in continuous time via posted prices and a fire-sale. Section 4 presents the main results, when the seller has K units. We then consider three basic extensions of the model. In Section 5 we extend our results to the case where time-preference comes from inventory costs rather than proportional discounting. In Section 6 we study the effect of inter-temporal changes in the distribution of the number of entrants, showing that our results extend when the number of entrants decreases over time, whereas the one-period-look-ahead policy fails if the number increases. Finally, in Section 7, we bridge the gap between the patient- and impatient-buyer models, and consider two models of partially patient agents.

1.1 Literature

There are a number of papers that examine how to sell to patient buyers entering over time. Our results are related to a classic result on "asset selling with recall". Bertsekas (1995, p. 177) derives the welfare-maximising policy with one good, when one agent enters each period and his value is publicly known. McAfee and McMillan (1988) analyze the profit-maximising mechanism when one agent enters each period and there is a fixed cost of delay. These results are special cases of Propositions 1 and 3 respectively. We derive the profit-maximising policy for many goods, when several agents enter each period and their values are privately known. We also show how to implement the optimal mechanism via posted prices.

Wang (1993) supposes that the seller has one object and that buyers arrive according to a Poisson distribution and experience a fixed per–period delay cost. Wang shows that with an infinite horizon, a profit-maximising mechanism is a constant posted-price. Gallien (2006) characterises the optimal sequence of prices when agents arrive according to a renewal process over an infinite time horizon. Assuming inter-arrival times have an increasing failure rate, Gallien proves that agents will buy when they enter the market (or not at all). In contrast to both Wang (1993) and Gallien (2006), we find that the optimal mechanism may induce delay of purchases on the equilibrium path.

Pai and Vohra (2008) consider a model without discounting where agents arrive and leave the market over time, and partially characterize the profit-maximising mechanism. Mierendorff (2009) considers a two-period version of a similar model and provides a complete characterisation of the optimal contract. Gershkov and Moldovanu (2010) show how to implement the efficient allocation with privately known values and arrival times, when buyers arrive according to a counting process.[3] In a separate line of work, Said (2009) characterises the optimal dominant strategy mechanism where agents are patient but goods are nonstorable, and describes a

[3]There are a number of papers on similar themes. Shen and Su (2007) summarize the operations research literature. For example, Aviv and Pazgal (2008) suppose a seller has many goods to sell to agents who arrive over time and are patient, but they restrict the seller to choosing two prices that are independent of the past sales. In economics, Board (2007) assumes a seller sells a single unit to agents whose values vary over time. Hörner and Samuelson (2008) consider a seller with no commitment power who sells a single unit to N agents by setting a sequence of prices.

dynamic open-auction implementation.

There is also a classic literature studying the sequential allocation of goods to impatient buyers. Karlin (1962) analyses the problem of allocating multiple goods to buyers who arrive sequentially but only remain in the market for one period. In the optimal policy, a buyer is awarded a unit if their valuation exceeds a cutoff. This cutoff is decreasing in the number of units available and increasing in the time remaining. These results have been extended in a number of ways. Derman, Lieberman, and Ross (1972) allow for heterogeneous goods. Albright (1974) allows for random arrivals with positive discount rates. More recently, a number of studies allow buyers' valuations to be private information. Gallego and van Ryzin (1994) considers profit-maximisation in continuous time, while Vulcano, van Ryzin, and Maglaras (2002) suppose N agents enter each period and allow the seller to hold an auction. Gershkov and Moldovanu (2009a) solve the profit-maximising policy for heterogenous goods. Gershkov and Moldovanu (2009b) allow the seller to learn about the distribution of valuations over time, introducing correlations in buyers' valuations.

Finally, the paper is related to the durable goods literature. Stokey (1979) characterises the optimal strategy for a seller with infinite supply who faces a fixed distribution of buyers. Conlisk, Gerstner, and Sobel (1984) suppose a homogenous set of buyers enters each period, while Board (2008) allows the entering generations to differ. As in this literature, we assume that the seller commits to a mechanism at the start of the problem. This is reasonable in many markets where yield management is used (e.g. airlines, hotels, cars), since the seller either designs a computer program to choose contingent prices or is a long-term player in a repeated game.[4]

2 Model

Basics. A seller has K goods to sell. Time is discrete and finite, $t \in \{1, \dots, T\}$. Time-preference comes from a common discount factor $\delta \in [0, 1)$.

Entrants. At the start of period t, N_t agents/buyers arrive. We initially assume N_t are IID random variables. N_t is observed by the seller, but not by other agents (see below).

Preferences. After he has entered the market, an agent wishes to buy a single unit. An agent is thus endowed with type (v_i, t_i), where v_i denotes his valuation, and t_i his birth date. The agent's valuation, v_i, is private information and drawn IID with density $f(\cdot)$, distribution $F(\cdot)$ and support $[\underline{v}, \overline{v}]$. The agent's birthdate, t_i, is observed by the seller but not by other buyers.

[4]Analysis of our general problem without commitment is missing in the literature. A related paper is Fuchs and Skrzypacz (2010) who study a seller with one unit, bargaining with one buyer with unknown value and waiting for a stochastic arrival of one additional buyer.

If the agent buys at time s for price p_s, his utility is $(v - p_s)\delta^s$. Let v_s^k denote the k^{th} highest order statistic of the agents entering at time s. Similarly, let $v_{\leq s}^k$ denote the k^{th} highest order statistic of the agents who have entered by time s.

Mechanisms. Each agent makes report $\tilde{v}_i$ when he enters the market. A mechanism $\langle P_{i,s}, TR_i \rangle$ maps agents' reports into an allocation rule $P_{i,s}$ describing the probability agent i is awarded a good in period s, and a transfer TR_i expressed in time-0 prices. A mechanism is *feasible* if (a) $P_{i,s} = 0$ before the agent enters, (b) $\sum_s P_{i,s} \in [0,1]$; (c) $\sum_i \sum_s P_{i,s} \leq K$; and (d) $P_{i,s}$ is adapted to the seller's information, so $P_{i,s}$ can vary only with the reports of agents that have entered by s.[5]

Agent's Problem. Upon entering the market, agent i chooses his declaration $\tilde{v}_i$ to maximise his expected utility,

$$u_i(\tilde{v}_i, v_i) = E_0 \left[\sum_{s \geq t_i} v_i \delta^s P_{i,s}(\tilde{v}_i, v_{-i}) - TR_i(\tilde{v}_i, v_{-i}) \middle| v_i \right] \tag{2.1}$$

where E_s denotes the expectation at the start of period s, before agents have entered the market. A mechanism is incentive compatible if the agent wishes to tell the truth, and is individually rational if the agent obtains positive utility.

Seller's Problem. The seller chooses a feasible mechanism to maximise the net present value of profits

$$\Pi_0^K = E_0 \left[\sum_i TR_i(v_i, v_{-i}) \right] \tag{2.2}$$

subject to incentive compatibility and individual rationality.

Some remarks regarding interpretation are pertinent. First, time T can be viewed as the date at which the good expires (e.g. a plane ticket) or the last time agents enter the market, since no sales will occur after this point.

Second, we adopt a durable-goods utility specification, interpreting the discount rate as the rate of time preference. If instead the discount rate is the degree agents' values fall over time (e.g. values for summer clothes will be lower in July than in June), then utility is given by $v\delta^t - \tilde{p}_t$. Under this new specification, the analysis is unchanged with prices given by $\tilde{p}_t = \delta^t p_t$.

[5] This formulation ignores the correlation between allocations, for a fixed set of reports. We can model such correlation by considering allocation function $P_{i,s}(v, \omega) \in \{0, 1\}$ where $\omega \in \Omega$ is a random variable. Since the optimal mechanism is deterministic, the correlation plays no role.

Third, the assumption that the seller can observe agents' birth-dates is for definiteness: the optimal allocation and implementation are identical if the seller cannot observe t_i. The assumption that an agent cannot observe agents' birth-dates is motivated by anonymous markets, such as large retailers and online sellers. If t_i's are publicly observed, the optimal allocations are unaffected although, when implementing this allocation, the price at time t is a function of $\{N_1, \ldots, N_t\}$.

Fourth, in the mechanism, we assume that buyers do not know the number of units remaining (indeed, they simply make a report when the enter the market). However, when implementing the optimal allocation, we attain the same profits when agents know the number of units available, so the seller does not benefit from hiding his remaining inventory.

2.1 Preliminaries

Fix a cohort of agents, $\{i : t_i = t\}$. When an agent enters the market, he chooses his declaration $\tilde{v}_i$ to maximise his utility (2.1). As shown in Mas-Colell, Whinston, and Green (1995, Proposition 23.D.2), an allocation rule is incentive compatible if and only if the discounted allocation probability

$$E_0 \left[\sum_{s \geq t} \delta^s P_{i,s}(v_i, v_{-i}) \right] \tag{2.3}$$

is increasing in v_i. Using the envelope theorem and integrating by parts, expected utility is then

$$E_0[u_i(v_i, v_i)] = E_0 \left[\sum_{s \geq t} \delta^t P_{i,s} \frac{1 - F(v_i)}{f(v_i)} \right] \tag{2.4}$$

where we use the fact that an agent with value $\underline{v}$ earns zero utility in any profit-maximising mechanism. Profit (2.2) equals welfare minus agents' utilities. Summing utility (2.4) over each cohort, we obtain

$$\Pi_0^K = E_0 \left[\sum_i \sum_{s \geq 1} P_{i,s} \delta^s m(v_i) \right] \tag{2.5}$$

where the *marginal revenue* of agent i is given by $m(v_i) := v_i - (1 - F(v_i))/f(v_i)$. Throughout we assume $m(v)$ is continuously increasing in v, implying that the seller's optimal mechanism is characterised by cutoff rules, and allowing us to ignore the monotonicity constraint (2.3).

Suppose the seller has k goods at time t. Write *continuation profits* before the period-t

entrants have entered by[6]

$$\Pi_t^k := E_t\left[\sum_i \sum_{s\geq t} \hat{P}_{i,s}\delta^{s-t}m(v_i)\right]. \tag{2.6}$$

where $\hat{P}_{i,s}$ is the allocation function given the principal has k goods in period t. Let the expected continuation profits after period-t entrants have entered be denoted by $\tilde{\Pi}_t^k$. When $k=1$, we omit the superscript.

3 Single Unit

We first derive the optimal solution when the firm has one unit to sell. This is the first step in the inductive process for the K-unit analysis, develops intuition and has some special features not present when $K > 1$. By the principle of optimality, we can solve (2.6) by maximising continuation profits in every state. At time t, profit is

$$\begin{aligned}\Pi_t &= \max_{\hat{P}_{i,t}} E_t\left[\sum_i \hat{P}_{i,t}m(v_i) + \left(1-\sum_i \hat{P}_{i,t}\right)\delta\Pi_{t+1}\right] \\ &= \max_{\hat{P}_{i,t}} E_t\left[\sum_i \hat{P}_{i,t}(m(v_i)-\delta\Pi_{t+1})\right] + E_t[\delta\Pi_{t+1}] \end{aligned}\tag{3.1}$$

Equation (3.1) implies that the good is allocated to maximise the flow profit minus the opportunity cost of allocating the good, $\delta\Pi_{t+1}$. As a result, when the good is awarded, it will be given to the agent with the highest marginal revenue (and the highest valuation).

We can now think of the highest current valuation, v, as a state variable. Let $\Pi_t(v)$ be the profit just before entry in time t, so that

$$\begin{aligned}\Pi_t(v) &= E_t\big[\max\{m(v), m(v_t^1), \delta\Pi_{t+1}(\max\{v, v_t^1\})\}\big] \quad \text{for } t<T \\ \Pi_T(v) &= E_T\big[\max\{m(v), m(v_T^1), 0\}\big]\end{aligned}\tag{3.2}$$

The following result shows that the optimal cutoffs can be characterised by a simple one-period-look-ahead rule.

Proposition 1. *Suppose $K=1$ and N_t are IID. The optimal mechanism awards the good to the agent with the highest valuation exceeding a cutoff. The cutoffs $\{x_t\}$ are uniquely determined*

[6] While we call Π_t^k continuation profits, this includes the impact of time t decisions on the willingness to pay of agents who buy in earlier periods.

by:

$$\begin{aligned} m(x_t) &= \delta E_{t+1}[\max\{m(v^1_{t+1}), m(x_t)\}] \qquad \textit{for } t < T \\ m(x_T) &= 0 \end{aligned} \tag{3.3}$$

Consequently, the cutoffs are constant in periods $t < T$.

Proof. The proof is by induction. In period $t = T$, then $m(x_T) = 0$. In period $t = T - 1$, the seller should be indifferent between selling to agent x_{T-1} today and waiting one more period and getting a new set of buyers. Hence

$$m(x_{T-1}) = \delta E_T[\max\{m(v^1_T), m(x_{T-1})\}].$$

Continuing by induction, fix t and suppose x_s, as defined by (3.3), are optimal for $s > t$. If $v < x_t$ then

$$m(v) < \delta E_{t+1}[\max\{m(v^1_{t+1}), m(v)\}]$$

so the seller strictly prefers to wait one period rather than sell to type v today. Conversely, if $v > x_t$ then

$$m(v) > \delta E_{t+1}[\max\{m(v^1_{t+1}), m(v)\}]. \tag{3.4}$$

Since N_t is IID, (3.4) implies that $v > x_{t+1}$ so type v will buy tomorrow if he does not buy today. Hence

$$\Pi_{t+1}(v) = E_{t+1}[\max\{m(v^1_{t+1}), m(v)\}]$$

and (3.4) implies that the seller strictly prefers to sell to type v today rather than waiting. Putting this together, x_t is indeed the optimal cutoff. □

Proposition 1 uniquely characterises the optimal cutoffs, and shows they are constant in all periods prior to the last. The intuition is as follows. At the cutoff the seller is indifferent between selling to the agent today and delaying one period and receiving another draw. This indifference rule relies on the assumption that if type x_t does not buy today, then he will buy tomorrow. This is satisfied because the seller faces exactly the same tradeoff tomorrow and therefore is once again indifferent between selling and waiting.

The optimal cutoffs are deterministic, depending on the number of periods remaining, but not on the number of agents who have entered in the past and their valuations. While the value of the second highest agent may affect the seller's realised revenue, it does not alter the seller's expected revenue and hence the optimal cutoff. Since cutoffs are deterministic the seller can implement the optimal mechanism without observing the number of arrivals, as we show below.

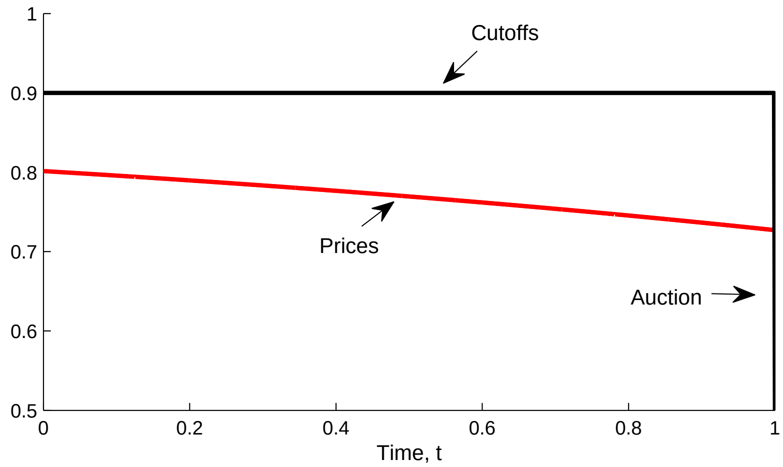

Figure 1: **Optimal Cutoffs and Prices with One Unit in Continuous Time.** When there is one unit, the optimal cutoffs are constant when $t < T$ and drop at time T. The price path is decreasing and concave, with an auction occurring at time T. In this figure, agents enter continuously with Poisson parameter $\lambda = 5$ have values $v \sim U[0, 1]$, so the static monopoly price is 0.5. Total time is $T = 1$ and the interest rate is $r = 1/16$.

Proposition 1 is very different from the optimal mechanism when buyers are impatient (e.g. Vulcano, van Ryzin and Marglaras (2002)). In this case, the optimal cutoffs are fully forward-looking, and fall over time as the seller becomes increasingly keen to sell the good. In contrast, when agents are patient, the allocations are determined by a one-period-look-ahead rule.[7]

Finally, let us assess the welfare consequences of Proposition 1. Using an analogous proof, one can show that the welfare-maximising mechanism awards the good to the agent with the highest value exceeding a cutoff given by $x_t^W = E_{t+1}[\delta \max\{v_{t+1}^1, x_t^W\}]$ for $t < T$, and $x_T^W = 0$. If $(1 - F(v))/vf(v)$ is decreasing in v, then the profit-maximising cutoffs exceed the welfare-maximising cutoffs for all t, implying that a profit-maximising seller awards the good later than is efficient (and sometimes never at all).[8]

[7]This assumes T is finite. When $T = \infty$, the cutoffs are determined by (3.3) and are therefore constant in all periods. An agent therefore either buys immediately or never, and we can assume that buyers are impatient without loss of generality (Gallien (2006)).

[8]Proof: Since $(1 - F(v))/vf(v)$ is decreasing in v, $m(v)/v$ is increasing in v and $m(v)/m(x) \geq v/x$ for $v \geq x$ if $m(x) > 0$. If $x_t^W > x_t$, then

$$1 = E_{t+1}\left[\delta \max\left\{\frac{m(v_{t+1}^1)}{m(x_t)}, 1\right\}\right] \geq E_{t+1}\left[\delta \max\left\{\frac{v_{t+1}^1}{x_t}, 1\right\}\right] > E_{t+1}\left[\delta \max\left\{\frac{v_{t+1}^1}{x_t^W}, 1\right\}\right] = 1$$

yielding the required contradiction.

3.1 Implementation

The optimal mechanism allocates the good to the agent with the highest valuation exceeding the cutoff, x_t. This allocation becomes particularly simple to implement as time periods become very short.[9] Suppose agents enter the market according to a Poisson process with arrival rate λ,[10] and let r be the instantaneous discount rate. Taking the limit of equation (3.3), the optimal allocation at $t < T$ is given by

$$rm(x^*) = \lambda E\big[\max\{m(v) - m(x^*), 0\}\big] \tag{3.5}$$

where v is distributed according to $F(\cdot)$. Equation (3.5) says the seller equates the flow profit from the cutoff type (the left-hand-side) and the option value of waiting for a new entrant (the right-hand-side). At time T, the optimal cutoff is given by $m(x_T) = 0$. See Figure 1 for an illustration.

The optimal allocation can be implemented by a deterministic sequence of prices with a fire-sale at time T. In the last period, the seller uses a second-price auction with reserve $R_T = m^{-1}(0)$. At time $t < T$ the seller chooses a price p_t, which makes type x^* indifferent between buying and waiting. The final "buy-it-now" price, denoted by $p_T = \lim_{t \to T} p_t$, is chosen so type x^* is indifferent between buying at price p_T and entering the auction. That is,

$$p_T = E_0\left[\max\{v^2_{\leq T}, m^{-1}(0)\}\big|N_0 = 1, v^1_{\leq T} = x^*\right] \tag{3.6}$$

Note that the buyer conditions on his own existence; since arrivals are independent, we assume that the buyer arrives at time 0 without loss of generality.

When $t < T$, type x^* is indifferent between buying now and waiting dt. This yields

$$(x^* - p_t) = (1 - rdt - \lambda dt)(x^* - p_{t+dt}) + \lambda dt(x^* - p_{t+dt})F(x^*)$$

Rearranging and letting $dt \to 0$,

$$\frac{dp_t}{dt} = -(x^* - p_t)\big(\lambda(1 - F(x^*)) + r\big). \tag{3.7}$$

When a buyer waits a little, they gain from the falling prices (the left hand side), but lose the rental value of the good and risk a stock-out if a new buyer enter with a value above x^* (the right hand side).

[9] In discrete time, the optimal allocation can be implemented through a sequence of second-price auctions. In the auctions, agents do not bid if their value is below the cutoff, so the sale occurs with the first bid. The reserve prices are constructed so that the cutoff type is indifferent between buying today and delaying, similar to the continuous time prices below. (These results were derived in an earlier version of this working paper.)

[10] In discrete time, this means the number of agents arriving in any period, N_t, has a Poisson distribution with parameter λ.

The optimal prices are illustrated in Figure 1 and have several interesting features. First, while cutoffs are constant in periods $t < T$, the optimal prices decline. When the agent delays he forgoes one period's enjoyment of the good, so the price has to drop at least at quickly as the interest rate, but since he is also risking the arrival of new competition, the price has to fall faster. Second, the price path is concave. As time progresses, prices fall and the agent loses more by delaying, so prices need to fall even faster to keep the cutoff type indifferent. Third, agents below x^* refrain from buying, even though their valuations may exceed the reserve price. Such agents delay in order to take advantage of the fire-sale in period T.

4 Multiple Units

In this section we suppose the seller has K goods to sell. Using the principle of optimality, the seller maximises continuation profits at each point in time. Consider period t and suppose the seller has k units.

Lemma 1. *The seller allocates goods to high value agents before low value agents.*

Proof. Suppose in period t the seller sells to agent j but does not sell to agent i, where $v^i \geq v^j$. To be concrete, suppose the seller eventually sells to agent i in period $\tau > t$, where we allow $\tau = \infty$. Now suppose the seller leaves all allocations unchanged but switches i and j. This increases profit by $(1-\delta^{\tau-t})(m(v^i) - m(v^j))$, contradicting the optimality of the original allocation. $\square$

Using Lemma 1, we need only keep track of the k highest remaining valuations. At the start of time t suppose the seller has agents with valuations $\{y^1, \ldots, y^k\}$, where $y^i \geq y^{i+1}$. Profit is described by the Bellman equation[11]

$$\tilde{\Pi}_t^k(y^1, \ldots, y^k) = \max_{j \in \{0,\ldots,k\}} \left[\sum_{i=1}^{j} m(y^i) + \delta \Pi_{t+1}^{k-j}(y^{j+1}, \ldots, y^k) \right]$$

where $\Pi_{t+1}^k := E_{t+1}[\tilde{\Pi}_{t+1}^k]$. The Bellman equation says the seller receives the marginal revenue from the units she sells today plus the continuation profits from the remaining units. The seller's optimal strategy is thus to sell the first object to the highest value agent, subject to his value exceeding cutoff x_t^k. She then sells the second object to the second highest value agent, subject to his value exceeding cutoff x_t^{k-1}, and so forth. We can thus think of the items being awarded sequentially within a period.[12]

[11] When $j = 0$ the first term in the summation is zero.

[12] Formally, a *cutoff* x_t^k is defined as the value of y^1 such that the seller is indifferent between selling today and waiting.

The following Lemma shows that when $\{x_t^k\}$ are decreasing in k we can treat each unit separately, comparing the j^{th} cutoff to the corresponding agent's valuation.

Lemma 2. *Fix t and suppose $\{x_t^k\}$ are decreasing in k. Then unit j is allocated to agent i at time t if and only if*
(a) v_i exceeds the cutoff x_t^j.
(b) v_i has the $(k-j+1)^{th}$ highest valuation of the currently present agents.

Proof. Suppose agent i is allocated good j, then (a) and (b) are satisfied.

Suppose (a) and (b) are satisfied. Then there are $(k-j)$ agents with higher valuations than i. Since the cutoffs are decreasing in k, these valuations exceed their respective cutoffs. Hence object j is allocated to agent i. □

Proposition 2 is our main result: it explicitly solves for the optimal cutoffs and shows they fall over time and decrease in the number of units remaining. In the period $t = T$, the seller wishes to allocate the goods to the k highest value buyers, subject to these values exceeding the static monopoly price. Hence,

$$m(x_T^k) = 0 \qquad \text{for all } k. \tag{4.1}$$

Next, consider period $t = T-1$. If she allocates the k^{th} good she gets $m(x_{T-1}^k)$. The opportunity cost is to wait one period and award the good either to agent x_{T-1}^k or the k^{th} highest new entrant. Hence,[13]

$$m(x_{T-1}^k) = \delta E_{T-1}\left[\max\{m(x_{T-1}^k), m(v_T^k)\}\right] \tag{4.2}$$

In periods $t \le T-1$, the seller is indifferent between selling to the cutoff type today and waiting one more period. If she sells today, she only sells one unit since $\{x_t^k\}$ are decreasing in k. If she waits, she sells at least one unit tomorrow by the one-period-look-ahead policy. We thus have:

$$\begin{aligned} m(x_t^k) &+ \delta E_{t+1}\left[\tilde{\Pi}_{t+1}^{k-1}(v_{t+1}^1, \dots, v_{t+1}^{k-1})\right] \\ &= \delta E_{t+1}\left[\max\{m(x_t^k), m(v_{t+1}^1)\}\right] + \delta E_{t+1}\left[\tilde{\Pi}_{t+1}^{k-1}(\{x_t^k, v_{t+1}^1, \dots, v_{t+1}^k\}_k^2)\right]. \end{aligned} \tag{4.3}$$

where the notation $\{x_t^k, v_{t+1}^1, \dots, v_{t+1}^k\}_k^2$ represents the ordered vector of the 2^{nd} to k^{th} highest choices from $\{x_t^k, v_{t+1}^1, \dots, v_{t+1}^k\}$. Notably, equation (4.3) is independent of the state $\{y^2, \dots, y^k\}$ for reasons explained below.

[13] To be more formal, if $y^1 > v_T^k$, the seller loses $y^1(1-\delta)$ by delaying. If $y^1 < v_T^k$, the seller loses $y^1 - \delta v^k$ by delaying. The seller is thus indifferent if y^1 satisfies (4.2).

Proposition 2. *Suppose the seller has K units to sell and N_t are IID. The optimal allocation awards unit k at time t to the agent with the highest value exceeding a cutoff x_t^k. The cutoffs are characterised by equations (4.1), (4.2) and (4.3). These cutoffs are deterministic, and decreasing in t and k.*

Proof. See Appendix A.1 and A.2. □

Proposition 2 has a number of important consequences. First, the cutoffs are uniquely determined. Intuitively, the sooner an agent buys a good the more his value affects overall profit. Hence the left hand sides of (4.2) and (4.3) have a steeper slope than the right hand sides.

Second, the cutoffs are independent of the current state $(y^2, \ldots, y^k)$. Intuitively, at the cutoff, the seller is indifferent between selling to y^1 and waiting. In either case the allocation to $(y^2, \ldots, y^k)$ is unaffected since, in any future state, this decision does not affect their rank in the distribution of agents available to the seller. This fact is used in equation (4.3), where we set $y^j = 0$ for $j \geq 2$. Moreover, since cutoffs are deterministic, we do not have to elicit the values of agents and, in continuous time, can implement the optimal allocations through posted prices (see below).

Third, the cutoffs increase when there are fewer units available (see Figure 2). Intuitively, if the seller delays awarding the k^{th} unit by one period then she can allocate it to the highest value entrant, rather than agent y^1. When there are more goods remaining, agent y^1 is more likely to be awarded the good eventually, reducing the option value of delay and decreasing the cutoff.

Fourth, the cutoffs for the last unit are identical to the one unit case and are therefore constant in periods $t \leq T-1$. The other cutoffs are decreasing over time (see Figure 2). The intuition, as above, is based on the fact that if the seller delays awarding the k^{th} unit by one period then she can allocate it to the highest value entrant, rather than agent y^1. As the game progresses, agent y^1 is more likely to be awarded the good eventually, reducing the option value of delay and decreasing the cutoff. Figure 2 shows that the cutoffs decrease rapidly as $t \to T$. Figure 3 shows the corresponding hazard rate of sale. The hazard rate with one unit remaining stays low until $t = T$, at which point it jumps to infinity (because of the fire sale). The hazard rate with two units remaining is qualitatively similar: it is low initially and rapidly rises as we approach T, and therefore still resembles a fire-sale.[14]

[14] We can also bound the k^{th} unit cutoff from above and below in periods $t < T$ by $\overline{x}^k$ and $\underline{x}^k$ as determined by $m(\overline{x}^k) = \delta E_{t+1}[\max\{m(\overline{x}^k), m(v_{t+1}^1)\}]$ and $m(\underline{x}^k) = \delta E_{t+1}[\max\{m(\underline{x}^k), m(v_{t+1}^k)\}]$.

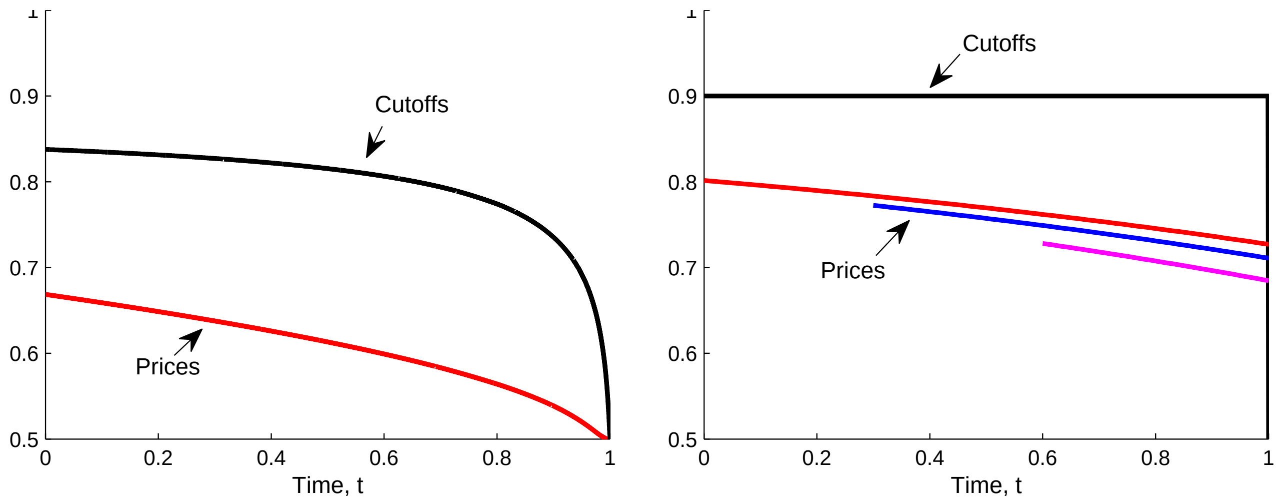

Figure 2: **Optimal Cutoffs and Prices with Two Units.** The **left** panel shows the optimal cutoffs/prices when the seller has two units remaining. The **right** panel shows the optimal cutoffs/prices when the the seller has one unit remaining. The three price lines illustrate the seller's strategy when it sells the first unit at times $t = 0$, $t = 0.3$ and $t = 0.6$. In this figure, agents enter continuously with $\lambda = 5$ and have values $v \sim U[0, 1]$. Total time is $T = 1$ and the interest rate is $r = 1/16$.

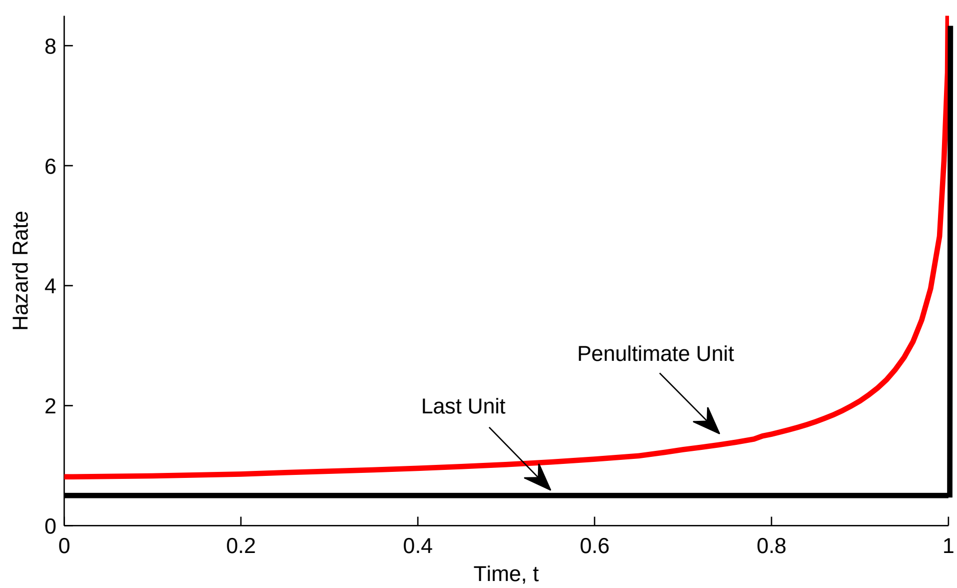

Figure 3: **Hazard Rates with Two Units** This figure shows the probability the last/penultimate unit is sold at time $t + dt$, conditional on there being one/two units remaining at time t. We assume agents enter continuously with Poisson parameter $\lambda = 5$ and have values $v \sim U[0, 1]$. Total time is $T = 1$ and the interest rate is $r = 1/16$.

4.1 Implementation

Suppose agents enter according to a Poisson process with parameter λ.[15] In period T, the optimal cutoffs are given by $m(x_T^k) = 0$. In period $t < T$, equation (4.3) becomes

$$rm(x_t^k) = \lambda E\Big[\max\{m(v) - m(x_t^k), 0\} + \Pi_t^{k-1}\big(\min\{v, x_t^k\}\big) - \Pi_t^{k-1}(v)\Big] \tag{4.4}$$

where v is drawn from $F(\cdot)$. Equation (4.4) states the seller is indifferent between selling today and delaying a little. The cost of delay is the forgone interest (the left-hand side); the benefit is the option value of a new buyer entering the market (the right-hand side). When compared to the single unit case (3.5), we see that delay leads to a higher marginal revenue tomorrow, if a new agent enters, and a lower state variable in the continuation game. While the continuation value depends on the values of the highest $k-1$ agents, the difference in continuation values only depends on the highest value (Lemma 4). This enables us to write Π_t^{k-1} as a function of one variable and, when computing the cutoffs, assume there is only one buyer present. Using Lemmas 5–7, equation (4.4) implies that x_t^k is uniquely determined, and decreasing in k and t. When $k = 1$, x_t^1 is constant for all $t < T$, and jumps down discontinuously at $t = T$. For $k \geq 2$, $\Pi_t^{k-1}(v) \to m(v)$ as $t \to T$, so (4.4) implies that $x_t^k \to m^{-1}(0)$, as shown in Figure 2.

We can implement the optimal allocations with prices $\{p_t^k\}$ and a fire-sale at time T for the last unit. We first wish to understand the limit of prices as $t \to T$, giving us a boundary point. For $k \geq 2$, $m(x_t^k) \to 0$ and hence the prices converge to $m^{-1}(0)$. For $k = 1$, the seller can use a second-price auction with reserve $m^{-1}(0)$ at time T. When $t < T$, the price converges to

$$p_T = E_0\left[\max\{v_{\leq T}^2, m^{-1}(0)\}\middle| v_{\leq T}^1 = x^*, s_T(x^*)\right]$$

where x^* is the constant cutoff with one unit remaining, and $s_T(x)$ denotes the last time the cutoff went below x. Note that p_T depends on when other agents purchased units. In particular, the earlier those units were purchased, the more competition agent with type x^* expects at time T, and the higher is p_T (see Figure 2).

In earlier periods, the prices are such that the cutoff type is indifferent between buying now and waiting a little. If he waits, the price is a little lower; however the agent forgoes some utility, and the good may be taken by a new buyer or a buyer with a slightly lower valuation. Equating these terms yields[16]

$$\frac{dp_t^k}{dt} = \left[\frac{dx_t^k}{dt}(t - s_t(x_t^k))\lambda f(x_t^k) - \lambda(1 - F(x_t^k))\right]\left[x_t^k - p_t^k - U_t^{k-1}(x_t^k)\right] - r\left(x_t^k - p_t^k\right) \tag{4.5}$$

[15] Li (2009) extends our results by providing an implementation in discrete time.

[16] For a derivation see Appendix A.4.

where $U_t^{k-1}(x_t^k)$ is the buyer's utility at time t when there are $k-1$ goods left, conditional on $v_{\le t}^1 = x_t^k$ and the history of the price path.[17] When a buyer waits a little they gain from the falling prices (the left hand side), but lose the rental value of the good and risk a stock-out if either a new buyer enters with a value above x_t^k or if an old buyer has a value just below the today's cutoff (the right hand side). This last feature means that prices drop faster if buyers think they have more competition from existing buyers. Overall, the price path falls smoothly over time, but jumps up with every sale.

5 Extension: Inventory Costs

In some applications the cost of delay is likely to be a function only of time, rather than proportional to values. (e.g. floor space in a shop). Suppose these costs are given by a convex function c_t for $t \in \{1, \dots, T+1\}$ and let $\Delta c_t := c_{t+1} - c_t$ be the cost of a one period delay. A buyer's utility is given by (2.1), where we set $\delta = 1$. Adapting (2.5), the firm's profits are given by

$$\Pi_0^K = E_0\left[\sum_i \sum_{s\ge 1} P_{i,s}[m(v_i) - c_t] + \left(K - \sum_i \sum_{s\ge 1} P_{i,s}\right)(-c_{T+1})\right]$$

We can now state the analogue of Proposition 1.

Proposition 3. *Suppose $K = 1$ and c_t is convex. The optimal cutoffs x_t are uniquely determined by*

$$\begin{aligned} m(x_t) &= E_{t+1}[\max\{m(v_{t+1}^1), m(x_t)\}] - \Delta c_t \qquad \text{for } t < T \\ m(x_T) &= -\Delta c_T \end{aligned} \tag{5.1}$$

These cutoffs are decreasing over time.

Proof. Since Δc_t is increasing in t the cutoffs, as defined by (5.1), are decreasing in t. The rest of the proof is the same as Proposition 1. □

In continuous time, these allocations can be implemented by a deterministic price sequence p_t and a fire-sale at date T. Suppose buyers enter with Poisson arrival rate λ and the inventory cost function $c(t)$ is differentiable, increasing and weakly convex.[18] The optimal cutoff at time $t < T$ is given by

$$c'(t) = \lambda E\big[\max\{m(v) - m(x_t), 0\}\big]$$

[17]Note: Utilities can be expressed in terms of future allocations via the envelope theorem.
[18]If $c(t)$ has kinks then x_t will sometimes jump down, requiring the use of an auction.

where v is drawn according to $F(\cdot)$. At time T, the optimal cutoff is given by $m(x_T) = -c'(T)$.

The agent's utility (2.1) is not affected by the inventory costs, so the implementation is the same as before. At time T, the seller can use a second-price auction with reserve $R_T = m^{-1}(-c'(T))$. The final "buy-it-now" price is given by

$$p_T = E_0 \left[\max\{v^2_{\leq T}, m^{-1}(-c'(T))\} \middle| N_0 = 1, v^1_{\leq T} = \overline{x}_T \right]$$

where $\overline{x}_T := \lim_{t \to T} x_t$. In earlier periods, prices are determined by

$$\frac{dp_t}{dt} = -(x_t - p_t)\left(-\frac{dx_t}{dt} \lambda t f(x_t) + \lambda_t (1 - F(x_t)) \right)$$

which is similar to equation (4.5).

For simplicity, we have assumed there is only one good. When $K \geq 1$ and the per-unit inventory cost c_t is convex, the proof of Proposition 2 can be adapted to show the optimal cutoffs x^k_t are characterised by the one-period look ahead rule:[19]

$$\begin{aligned} m(x^k_t) &+ E_{t+1}\left[\tilde{\Pi}^{k-1}_{t+1}(v^1_{t+1}, \ldots, v^{k-1}_{t+1}) \right] \\ &= E_{t+1}\left[\max\{m(x^k_t), m(v^1_{t+1})\} \right] + E_{t+1}\left[\tilde{\Pi}^{k-1}_{t+1}(\{x^k_t, v^1_{t+1}, \ldots, v^k_{t+1}\}^2_k) \right] - \Delta c_t \end{aligned}$$

for $t < T$, with $m(x^k_T) = -\Delta c_T$. These cutoffs are deterministic and decreasing in t and k. In continuous time, the optimal allocation can be implemented by posted prices plus an auction for the last unit in period T. The continuous time cutoffs are determined by

$$c'(t) = \lambda E\Big[\max\{m(v) - m(x^k_t), 0\} + \Pi^{k-1}_t\big(\min\{v, x^k_t\}\big) - \Pi^{k-1}_t(v) \Big]$$

where v is drawn from $F(\cdot)$. Prices are then determined by (4.5) with $r = 0$ and the auction for the last unit as above.

6 Extension: Varying Entry

This section analyses the optimal mechanism when the expected number of entrants varies over time. In Section 6.1 we suppose fewer agents enter over time, as the stock of potential entrants is used up. In Section 6.2 we suppose more agents enter over time, as word of the market's existence spreads. This analysis forms a bridge between models with no entry (e.g. Harris and Raviv (1981)) and the constant entry model in Section 3.

[19] The proof of Proposition 2 has to be slightly adjusted. First, equations (A.1) and (A.2) have to be adjusted to include inventory costs; similarly equation (A.13) in Lemma 7. Second, in Lemmas 5-7, δ^τ should be replaced by $\mathbf{1}_{\tau \leq T}$.

6.1 Decreasing Entry

We first show that, when entry is decreasing, the IID allocations and prices are easily generalized. In particular, the cutoffs are determined by a one-period-look-ahead policy and are deterministic.

Proposition 4. *Suppose $K = 1$ and N_t is decreasing in the usual stochastic order. Then the optimal cutoffs are characterised by (3.3). These cutoffs are decreasing over time.*

Proof. Since N_t is decreasing in the usual stochastic order, v_t^1 is decreasing in the usual stochastic order and x_t, as defined by (3.3), is decreasing in t. The rest of the proof is the same as Proposition 1. $\square$

In continuous time, these allocations can be implemented by a deterministic price sequence p_t and a fire-sale at date T. Suppose buyers enter with Poisson arrival rate λ_t, which is continuously decreasing in t. The optimal cutoff at time $t < T$ is given by

$$rm(x_t) = \lambda_t E\big[\max\{m(v) - m(x_t), 0\}\big]$$

where v is drawn according to $F(\cdot)$. At time T, the optimal cutoff is given by $m(x_T) = 0$.

At time T, the seller can implement the optimal allocation through a second-price auction with reserve $R_T = m^{-1}(0)$. The final "buy-it-now" price is given by

$$p_T = E_0\left[\max\{v^2_{\leq T}, m^{-1}(0)\}\middle|N_0 = 1, v^1_{\leq T} = \overline{x}_T\right]$$

where $\overline{x}_T := \lim_{t\to T} x_t$. In earlier periods, prices are determined by

$$\frac{dp_t}{dt} = -(x_t - p_t)\left(-\frac{dx_t}{dt}\left(\int_0^t \lambda_\tau \, d\tau\right) f(x_t) + \lambda_t(1 - F(x_t)) + r\right) \tag{6.1}$$

which is similar to equation (4.5).

To illustrate, suppose a seller puts her house on the market. There is an initial stock of buyers who have a high probability of seeing the newly listed house, plus a constant inflow of new buyers (where $T = \infty$). In the optimal mechanism, there is an introductory period where cutoffs and price fall quickly, with some buyers strategically waiting. In the limit, where existing buyers see the new house immediately, the seller reduces prices instantly in the form of a Dutch auction. After this introductory period, prices coincide with the cutoffs, and are constant over time, so that no buyer ever delays.

For simplicity, we have assumed there is only one good. When $K \geq 1$, Proposition 2 applies to the decreasing entry case, and optimal cutoffs are characterised by equations (4.1), (4.2) and

(4.3).[20] As before, these cutoffs are deterministic, and decreasing in t and k. In continuous time, the optimal allocation can be implemented by posted prices plus an auction for the last unit in period T. The continuous time cutoffs are determined by (4.4), replacing λ with λ_t. Similarly, the price path is determined by (4.5), again replacing λ with λ_t.

6.2 Increasing Entry

When the number of entrants increases over time, the one-period-look-ahead policy fails. Intuitively, because the number of entrants is rising, the seller wishes to increase the cutoff. If the seller does not serve a cutoff type x_t in period t, she will therefore not return to that agent until period T. As a result the optimal allocations depend on the number of entrants in all future periods, not just the adjacent period.

Recursively define the following functions:

$$\begin{aligned} \pi_t(v) &= E_t\big[\max\{m(v_t^1), \delta\pi_{t+1}(\max\{v, v_t^1\})\}\big] \qquad \text{for } t < T \\ \pi_T(v) &= E_T\big[\max\{m(v), m(v_T^1), 0\}\big] \end{aligned} \tag{6.2}$$

This looks similar to equation (3.2), but is simpler because an agent who does not receive the good at time t need not be considered again until period T.

Proposition 5. *Suppose $K = 1$ and N_t is increasing in the usual stochastic order. Then the optimal cutoffs are given by*

$$\begin{aligned} m(x_t) &= \delta\pi_{t+1}(x_t) \qquad \textit{for } t < T \\ m(x_T) &= 0. \end{aligned} \tag{6.3}$$

These cutoffs are increasing over time, for $t < T$.

Proof. See Appendix A.5. □

When the number of entrants increases over time, the optimal cutoffs (6.3) also increase. As a result, an agent either buys when he enters the market or waits until the final period. This means that, unlike the one-period-look-ahead policies in Propositions 1–4, the optimal cutoffs depend on the future of the game. Consequently, today's cutoff increases if either the game becomes longer, or the future number of entrants rises.

In continuous time, these allocations can be implemented by a deterministic price sequence p_t and a fire-sale at date T. Suppose buyers arrive with Poisson arrival rate λ_t, which is continuously increasing in t. We can define functions corresponding to (6.2) using the end

[20] Proof: Replace Lemma 7 with Lemma 7′ in Appendix A.3.

point $\pi_T(v) = v$ and the differential equation

$$r\pi_t(v) = \frac{d\pi_t(v)}{dt} + \lambda_t E\Big[\max\big\{m(v'), \pi_t(\max\{v, v'\})\big\} - \pi_t(v)\Big] \tag{6.4}$$

where v' is the value of the new entrant and is drawn from $F(\cdot)$. Equation (6.4) says that asset value of profits are determined by the increase in their value and the option value from new entrants arriving. We can now define the optimal cutoffs. At time T, the optimal cutoff is given by $m(x_T) = 0$. At time $t < T$, the optimal cutoff is given by the smallest x_t such that $m(x_t) = \pi_t(x_t)$.

At time T, the seller can implement the optimal allocation through a second-price auction with reserve $R_T = m^{-1}(0)$. For $t < T$, the prices are determined so that buyer x_t is indifferent between buying in period t and waiting until the fire-sale. That is,

$$(x_t - p_t) = e^{-r(T-t)} \Pr(v^1_{\geq t} \leq x_t) E\big[x_t - \max\{v^2_{\leq T}, m^{-1}(0)\}\big|N_0 = 1, v^1_{\leq T} = x_t\big] \tag{6.5}$$

where $v^1_{\geq t}$ is the highest order statistic of the buyers who have entered after time t. Let

$$\psi_t := e^{-r(T-t)} \Pr(v^1_{\geq t} \leq x_t) = e^{-r(T-t)} e^{-(\int_t^T \lambda_\tau \, d\tau)(1-F(x_t))}.$$

Note that ψ_t increases in t, and that $\psi_T = 1$. Prices are then given by

$$p_t = (1 - \psi_t)x_t + \psi_t E[\max\{v^2_{\leq T}, m^{-1}(0)\}|N_0 = 1, v^1_{\leq T} = x_t], \tag{6.6}$$

Over time, the optimal posted prices will tend to rise and then fall. Intuitively, as t grows so the cutoff increases, increasing the first term in (6.6). However, as $t \to T$, the fire-sale at T comes closer, decreasing agents willingness to delay and increasing the weight on the second term in (6.6). If we take $T \to \infty$, then the right hand side of (6.5) converges to zero and $p_t \to x_t$ for all t. This follows from the fact that a buyer who delays at time t must wait until period T to have another opportunity to buy.

7 Extension: Partially Patient Agents

One limitation of our analysis is that we do not allow for heterogeneity in the timing of buyers' demands. That is, a type-v agent who enters in period 1 has the same valuation in period t as a type-v agent who enters in period t. This is a problematic assumption for some applications, since buyers may exit the market (for example, a customer may buy another airline ticket), or buyers' valuations may decline relative to the entrants (for example, a customer's value for a seasonal piece of clothing declines after his vacation). In this section we consider these two perturbations of the model: In Section 7.1 we suppose buyers' values decline deterministically relative to those of new entrants; In Section 7.2 we assume that buyers exit stochastically. These results highlight the difficulties these considerations create and bridge our results with the analysis of impatient agents (e.g. Vulcano, van Ryzin, and Maglaras (2002)).

7.1 Declining Values

For the first model, assume that an agent with value v who enters in period t and buys in period s receives utility

$$\delta^s \beta^{s-t} v \tag{7.1}$$

where $\beta \in [0,1]$. When $\beta = 1$ this coincides with the model in Section 3; when $\beta = 0$ this coincides with the model of impatient agents. Following the derivation in Section 2, profits are given by

$$\Pi_0 = E_0 \left[\sum_i \sum_{s \geq 1} P_{i,s} \delta^s \beta^{s-t_i} m(v_i) \right]$$

It will be convenient to think of the state variable as the highest marginal revenue, $\hat{m}$, rather than the highest valuation. Recursively define the following functions:

$$\begin{aligned} \pi_t(\hat{m}) &= E_t\big[\max\{m(v_t^1), \delta\pi_{t+1}(\max\{\beta\hat{m}, m(v_t^1)\})\}\big] \qquad \text{for } t < T \\ \pi_T(\hat{m}) &= E_T\big[\max\{\beta\hat{m}, m(v_T^1), 0\}\big] \end{aligned} \tag{7.2}$$

This looks similar to equation (3.2), but is simpler because, if the seller delays at time t then she does not return to that buyer until period T.

Proposition 6. *Suppose $K = 1$, N_t is IID, and agents have declining values (7.1). At time $t < T$, the optimal mechanism awards the good to the highest value agent who enters in time t, if this value exceeds a cutoff x_t defined by*

$$m(x_t) = \delta\pi_{t+1}^t(m(x_t)). \tag{7.3}$$

These cutoffs have the property that $m(x_t) \geq m(x_{t+1}) \geq \beta m(x_t)$ for $t < T - 1$. At time T, the good is awarded to the agent with the highest discounted marginal revenue, $\beta^{T-t_i} m(v_i)$, providing it is positive.

Proof. See Appendix A.6. □

Proposition 6 tells us that, when agents are only partially patient, the one-period-look-ahead policy fails to hold. In particular, an agent either buys when he enters the market or waits until period T. As in models with impatient agents, the cutoffs fall over time as the seller's options shrink. However, since the seller can always return to an old agent, the rate of decline is bounded below by β.

From equation (7.2) one can see that the seller's profits are increasing in β. That is, the seller prefers agents to be forward looking. While it may seem counterintuitive that allowing inter-temporal arbitrage benefits the seller, delay allows the seller to merge buyers from different cohorts and obtain a more efficient allocation. This suggests that if the seller can run an optimal mechanism, she should embrace price prediction sites such as bing.com, rather than viewing them as a threat to inter-temporal price discrimination.

We now turn to implementation. In period $t < T$, the optimal mechanism awards the good to the entrant with the highest value exceeding the cutoff. If the generalized failure rate $vf(v)/(1-F(v))$ is increasing in v then $m(x_{t+1}) \geq \beta m(x_t)$ implies that $x_{t+1} \geq \beta x_t$ so the time-t cutoff type will not wish to buy at time $t+1$. As a result, we can implement the optimal mechanism via simple second-price auctions with appropriate reserve prices, despite the new and old entrants being asymmetric. Similarly, in continuous time the optimal mechanism can be implemented via posted prices.

At time T, the optimal mechanism allocates the good to the agent with the highest $\beta^{T-t_i} m(v_i)$ while a second-price auction would allocate it to the one with the highest $\beta^{T-t_i} v_i$. If the generalized failure rate is increasing in v, $\beta^{T-t_1} v_1 = \beta^{T-t_2} v_2$ implies $\beta^{T-t_1} m(v_1) \geq \beta^{T-t_2} m(v_2)$ for $t_2 > t_1$. As a result, allocation is biased towards agents who enter the market earlier. Intuitively, allocating the object to an older buyer gives away fewer information rents because they have a higher valuation relative to their cohort. We can thus implement the optimal allocation by having agents register with the seller when they arrive in the market. The seller can then give a "discount voucher" to an agent who arrives early. For example, if $v \sim U[0,1]$ then the seller should give an agent who registers in period t a discount of $(1-\beta^{T-t})/2$.[21]

7.2 Disappearing Buyers

Another natural way to model the heterogeneity in agents' timing decisions is to allow buyers to exit probabilistically over time. If entry and exit times are private information of the buyers, the optimal mechanisms are very complicated, as discussed by Pai and Vohra (2008) and Mierendorff (2009). Even if we simplify the model to assume that agents have no private information about their exit times and assume each agent exits the game with probability ρ, the optimal allocations become much more complicated. In particular, the following example illustrates that the striking feature of our model — that the optimal cutoffs are deterministic — does not hold in a general model with random exits.

Suppose time is discrete, $T = 2$ and $K = 1$. Suppose there are two entrants in period

[21] Proof: The seller wishes award the good to the agent who maximises $\beta^{T-t}(2v-1)$, or equivalently $v\beta^{T-t} + (1-\beta^{T-t})/2$.

$t=1$ and one more entering at $T=2$, and values are distributed uniformly over $[0,1]$, so that $m(v)=2v-1$. Solving for optimal cutoffs, $x_2=\frac{1}{2}$ and x_1 is given by:

$$m(x_1)=\delta E_{v_3}\left[\max\{m(x_1),m(v_3)\}\right] \tag{7.4}$$

Next, suppose agents independently exit with probability ρ.[22] How does the optimal mechanism change? Without loss, suppose that $v_1 \geq v_2$. Then it is optimal to sell the good to agent 1 if and only if

$$m(v_1)\geq \delta E_{v_3}\left[(1-\rho)\max\{m(v_1),m(v_3)\}+\rho(1-\rho)\max\{m(v_2),m(v_3),0\}+\rho^2\max\{0,m(v_3)\}\right]$$

This expression is much more complicated than (7.4) because we need to take into account the risk of losing either of the two agents. Importantly, the possibility that agent 1 will exit and agent 2 will stay, makes the decision of whether to award the good to agent 1 today depend on the value of agent 2! For $\delta=1$ and $\rho=\frac{1}{9}$ the optimal cutoff for agent 1 as a function of agent 2 value is:

$$x_1(v_2)\begin{cases}\approx 0.91 & \text{for } v_2>0.91\\ =\frac{9}{8}-\frac{1}{24}\sqrt{79-64v_2^2} & \text{for } v_2\in[0.5,0.91]\\ \approx 0.79 & \text{for } v_2<0.5\end{cases}$$

Hence the optimal cutoffs depend on the values of all players that have entered, and are not deterministic. While we can implement such an allocation through a direct revelation mechanism, it seems unlikely that any natural indirect mechanism, such as auctions or posted prices, will work.

8 Conclusion

We have characterized the optimal mechanism for a seller who wishes to sell K goods to buyers who enter the market over time and are patient. We have also shown that the optimal mechanism is deterministic and can be implemented by a sequence of prices with an auction for the final good at time T.

A major motivation for the paper was to introduce forward-looking buyers into a standard revenue management model. When compared to a model with myopic buyers, this change has three major consequences: first, cutoffs are determined by a simple one-period-look-ahead rule rather than a fully forward-looking optimal stopping problem; second, prices depend on the history of sales, since this affects the competition faced by current buyers; third, the seller should hold an auction at the end (or reduce prices rapidly) to harvest delaying buyers. The

[22]If exits are perfectly correlated (e.g. the good expires) then the expiration probability can be incorporated into the discount rate, and the analysis of Sections 3–4 is unchanged.

importance of these differences depends on the environment. The assumption of myopic buyers is without loss if entry is IID and either there are infinite periods (Gallien (2006)), or markets are large (Segal (2003)), since a constant price is optimal under either scenario. This means that properly modelling patient buyers is most important where the seller's options decline over time, or where the market is small. It suggests, for instance, that airline companies should be more concerned with forward-looking customers on their small flights than on their large ones.

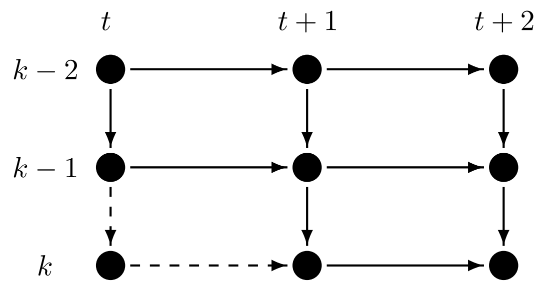

Figure 4: **Proof of Proposition 2.** This figure shows the order of cutoffs in induction step, where an arrow from $(k-1,t)$ to $(k-1,t+1)$ indicates that $x_t^{k-1} \geq x_{t+1}^{k-1}$. We know the relations indicated by dark arrows hold, and wish to prove the relations indicated by dashed arrows. Cases (a) and (b) correspond to the two ways these inequalities may not hold.

A Omitted Proofs

A.1 Proof of Proposition 2

At time $t = T$ and $t = T-1$ the cutoffs are given by (4.1) and (4.2), as argued in the text. We now claim that $\{x_t^k\}$ are deterministic and decreasing in t and k. This is true for the last two periods. We now continue by induction.

Definitions. At time t, suppose the state is $(y^1, y^2, \ldots, y^k)$. If the seller sells one unit today then continuation profit is

$$\begin{aligned}\Pi_t^k(\text{sell 1 today}) &= m(y^1) + \delta \Pi_{t+1}^{k-1}(y^2, \ldots, y^k) \\ &= m(y^1) + \delta E_{t+1}\left[\tilde{\Pi}_{t+1}^{k-1}(\{y^2, \ldots, y^k, v_{t+1}^1, \ldots, v_{t+1}^k\}_{k-1}^1)\right] \qquad \text{(A.1)}\end{aligned}$$

If the seller sells one or more units tomorrow then she will obtain

$$\Pi_t^k(\text{sell tomorrow}) = \delta E_{t+1}\left[\max\{m(y^1), m(v_{t+1}^1)\}\right] + \delta E_{t+1}\left[\tilde{\Pi}_{t+1}^{k-1}(\{y^1, y^2, \ldots, y^k, v_{t+1}^1, \ldots, v_{t+1}^k\}_k^2)\right] \qquad \text{(A.2)}$$

Denote the difference function by

$$\Delta\Pi_t^k(y^1, \ldots, y^k) = \Pi_t^k(\text{sell 1 today}) - \Pi_t^k(\text{sell tomorrow}).$$

As shown in Lemma 4 in Appendix A.2, $\Delta\Pi_t^k$ is independent of $\{y^2, \ldots, y^k\}$, so we can write it as a function of y^1 only.

Monotonicity in k and t. Since selling the last unit is identical to selling a single unit, the cutoffs are determined by (3.3) and obey $x_t^1 \geq x_{t+1}^1$. We now proceed by induction (see Figure 4). By contradiction, let $k \geq 2$ be the smallest number that either (a) $x_t^k > x_t^{k-1}$ or (b)

$x_t^k < x_{t+1}^k$. If there are multiple t's that satisfy either condition, pick the higher number. [23]

Case (a). Consider good k in period t. We have

$$x_t^k > x_t^{k-1} \geq x_{t+1}^{k-1} \geq x_{t+1}^k, \tag{A.3}$$

so the k^{th} cutoff is decreasing in t. At the cutoff the seller is indifferent between selling today and waiting. If she sells today she earns $\Pi_t^k(\text{sell today}) \geq \Pi_t^k(\text{sell 1 today})$. If she waits then (A.3) implies that she chooses to sell good k tomorrow and $\Pi_t^k(\text{wait}) = \Pi_t^k(\text{sell tomorrow})$. The indifference condition therefore implies that, at the cutoff,

$$\Delta\Pi_t^k(x_t^k) \leq 0. \tag{A.4}$$

Consider the cutoff for good $k-1$ in period t. Since $\{x_t^j\}_{j<k}$ are decreasing in k, $\Pi_t^{k-1}(\text{sell today}) = \Pi_t^{k-1}(\text{sell 1 today})$. Since $x_t^{k-1} \geq x_{t+1}^{k-1}$, $\Pi_t^{k-1}(\text{wait}) = \Pi_t^{k-1}(\text{sell tomorrow})$. As a result,

$$\Delta\Pi_t^{k-1}(x_t^{k-1}) = 0. \tag{A.5}$$

We therefore conclude that

$$0 \geq \Delta\Pi_t^k(x_t^k) > \Delta\Pi_t^k(x_t^{k-1}) \geq \Delta\Pi_t^{k-1}(x_t^{k-1}) = 0. \tag{A.6}$$

yielding the required contradiction. In equation (A.6), the first inequality comes from (A.4). The second comes from $x_t^k > x_t^{k-1}$ and Lemma 5, which says that $\Delta\Pi_t^k(x)$ is strictly increasing in x. The third inequality comes from Lemma 6, which says that $\Delta\Pi_t^k(x)$ is increasing in k. The final equality comes from (A.5).

Case (b). Consider good k in period t. We have

$$x_t^k < x_{t+1}^k \leq x_{t+1}^{k-1} \leq x_t^{k-1}$$

so $\{x_t^j\}_{j\leq k}$ are decreasing in k. At the cutoff the seller is indifferent between selling today and waiting. If the seller sells today she earns $\Pi_t^k(\text{sell today}) = \Pi_t^k(\text{sell 1 today})$, since $\{x_t^j\}_{j\leq k}$ are decreasing in k. If she waits then she obtains $\Pi_t^k(\text{wait}) \geq \Pi_t^k(\text{sell tomorrow})$. The indifference condition implies that, at the cutoff,

$$\Delta\Pi_t^k(x_t^k) \geq 0 \tag{A.7}$$

Consider the cutoff for good k in period $t+1$. Since $\{x_t^j\}_{j\leq k}$ are decreasing in k, $\Pi_{t+1}^k(\text{sell today}) =$

[23] Since $x_t^{k-1} \geq x_{t+1}^{k-1} \geq x_{t+1}^k$, these two cases are mutually exclusive.

Π^k_{t+1}(sell 1 today). Since $x^k_{t+1} \geq x^k_{t+2}$, $\Pi^k_{t+1}(\text{wait}) = \Pi^k_{t+1}$(sell tomorrow). As a result,

$$\Delta\Pi^k_{t+1}(x^k_{t+1}) = 0. \tag{A.8}$$

We therefore conclude that

$$0 \leq \Delta\Pi^k_t(x^k_t) < \Delta\Pi^k_t(x^k_{t+1}) \leq \Delta\Pi^k_{t+1}(x^k_{t+1}) = 0. \tag{A.9}$$

yielding the required contradiction. In equation (A.9), the first inequality comes from (A.7). The second comes from $x^k_t < x^k_{t+1}$ and Lemma 5, which says that $\Delta\Pi^k_t(x)$ is strictly increasing in x. The third inequality comes from Lemma 7, which says that $\Delta\Pi^k_t(x)$ is increasing in t. The final equality comes from (A.8).

Summary. Given that $\{x^k_t\}$ are decreasing in k and t the optimal cutoffs are given by $\Delta\Pi^k_t(x^k_t) = 0$. Using Lemma 4 we can assume $y^j = 0$ for $j \geq 2$ and write this as (4.3).

A.2 Lemmas for Proof of Proposition 2

Lemma 3. *Fix t and suppose $\{x^k_s\}_{s\geq t+1}$ are decreasing in k. Suppose $y^1 \geq y^2 \geq \ldots \geq y^k$, and let $y^{j-1} \geq \tilde{y}^1 \geq y^j$. Then the difference*

$$\tilde{\Pi}^k_{t+1}(y^1, y^2, \ldots, y^k) - \tilde{\Pi}^k_{t+1}(\tilde{y}^1, y^2, \ldots, y^k)$$

is independent of $\{y^j, \ldots, y^k\}$.

Proof. Suppose the state is $(y^1, y^2, \ldots, y^k)$ and pick $i \geq j$. Since cutoffs are decreasing in k, Lemma 2 says the good is allocated to value y^i if and only if (a) given previous allocations (including those within the period), y^i has the highest value; and (b) y^i exceeds the current cutoff. Since this rule only depends on the rank of y^i, the allocation rule is the same as in state $(\tilde{y}^1, y^2, \ldots, y^k)$. Hence the difference in continuation profits is independent of y^i, as required. □

Lemma 4. *Fix t and suppose $\{x^k_s\}_{s\geq t+1}$ are decreasing in k. Then $\Delta\Pi^k_t(y^1, \ldots, y^k)$ is independent of $\{y^2, \ldots, y^k\}$.*

Proof. We can write (A.1) and (A.2) as

$$\Pi^k_t(\text{sell 1 today}) = m(y^1) + \delta E_{t+1}\left[\tilde{\Pi}^{k-1}_{t+1}(\{v^1_{t+1}, y^2, \ldots, y^k, v^2_{t+1}, \ldots, v^k_{t+1}\}^1_{k-1})\right] \tag{A.10}$$

$$\Pi^k_t(\text{sell tomorrow}) = \delta E_{t+1}\left[m(\overline{z})\right] + \delta E_{t+1}\left[\tilde{\Pi}^{k-1}_{t+1}(\{\underline{z}, y^2, \ldots, y^k, v^2_{t+1} \ldots, v^k_{t+1}\}^1_{k-1})\right]$$

where $\overline{z} := \max\{y^1, v^1_{t+1}\}$ and $\underline{z} := \min\{y^1, v^1_{t+1}\}$. The difference between these two equations only depends on $(y^2, \dots, y^k)$ through the difference

$$\tilde{\Pi}^{k-1}_{t+1}(\{v^1_{t+1}, y^2, \dots, y^k, v^2_{t+1}, \dots, v^k_{t+1}\}^1_{k-1}) - \tilde{\Pi}^{k-1}_{t+1}(\{\underline{z}, y^2, \dots, y^k, v^2_{t+1} \dots, v^k_{t+1}\}^1_{k-1}). \quad \text{(A.11)}$$

If $y^1 \geq v^1_{t+1}$, then (A.11) equals zero. If $v^1_{t+1} > y^1$ then, since $y^1 \geq y^2 \geq \dots \geq y^k$ and $\{x^k_s\}_{s \geq t+1}$ are decreasing in k, we can apply Lemma 3, implying that (A.11) is independent of $(y^2, \dots, y^k)$. Hence $\Delta\Pi^k_t$ is also independent of $(y^2, \dots, y^k)$. □

Lemma 5. *$\Delta\Pi^k_t(y^1)$ is strictly increasing in y^1.*

Proof. Using equation (A.1),

$$\frac{d}{dy^1}\Pi^k_t(\text{sell 1 today}) = m'(y^1)$$

Using equation (A.2) and the envelope theorem,

$$\frac{d}{dy^1}\Pi^k_t(\text{sell tomorrow}) = m'(y^1)E_{t+1}[\delta^{\tau^k_1(y^1)-t}]$$

where $\tau^k_1(y^1)$ is the time y^1 buys when he's in first position at time t and there are k goods to sell. The result follows from the fact that $\tau^k_1(y^1) > t$ and $\delta < 1$. □

Lemma 6. *Fix t and suppose $\{x^k_s\}_{s \geq t+1}$ are decreasing in k. Then $\Delta\Pi^k_t(y^1)$ is increasing in k.*

Proof. Let $\{y^1, \dots, y^k\}$ and $\{\tilde{y}^1, \dots, \tilde{y}^k\}$ be arbitrary vectors, where $y^j \geq \tilde{y}^j$ for each j. Using equation (2.6),

$$\begin{aligned}
\tilde{\Pi}^k_{t+1}(y^1, \dots, y^k) - \tilde{\Pi}^k_{t+1}(\tilde{y}^1, \dots, \tilde{y}^k) &\geq \tilde{\Pi}^k_{t+1}(y^1, \dots, y^{k-1}, \tilde{y}^k) - \tilde{\Pi}^k_{t+1}(\tilde{y}^1, \dots, \tilde{y}^{k-1}, \tilde{y}^k) \\
&= E_{t+2}\left[\delta^{-(t+1)} \int_{\{\tilde{y}^1, \dots, \tilde{y}^{k-1}\}}^{\{y^1, \dots, y^{k-1}\}} (m'(z^1)\delta^{\tau^k_1(z^1)}, \dots, m'(z^{k-1})\delta^{\tau^k_{k-1}(z^{k-1})}) d(z^1, \dots, z^{k-1})\right] \\
&\geq E_{t+2}\left[\delta^{-(t+1)} \int_{\{\tilde{y}^1, \dots, \tilde{y}^{k-1}\}}^{\{y^1, \dots, y^{k-1}\}} (m'(z^1)\delta^{\tau^{k-1}_1(z^1)}, \dots, m'(z^{k-1})\delta^{\tau^{k-1}_{k-1}(z^{k-1})}) d(z^1, \dots, z^{k-1})\right] \\
&= \tilde{\Pi}^{k-1}_{t+1}(y^1, \dots, y^{k-1}) - \tilde{\Pi}^{k-1}_{t+1}(\tilde{y}^1, \dots, \tilde{y}^{k-1}) \qquad \text{(A.12)}
\end{aligned}$$

The first line comes from the fact that $y^k \geq \tilde{y}^k$. The second line use the envelope theorem, where τ^k_j is the stopping time of the agent in the j^{th} position when there are k objects for sale. The third line follows from the fact that stopping times increase when the seller has one less object since $\{x^k_s\}_{s \geq t+1}$ are decreasing in k. The final line again uses the envelope theorem.

Looking at equations (A.1) and (A.2), observe that the vector

$$\{y^2, \ldots, y^k, v^1_{t+1}, \ldots, v^{k-1}_{t+1}\}^1_{k-1}$$

is pointwise larger than the vector

$$\{y^1, y^2, \ldots, y^k, v^1_{t+1}, \ldots, v^{k-1}_{t+1}\}^2_k.$$

The result follows from equation (A.12). □

Lemma 7. *Fix t and suppose $\{x^k_s\}_{s \geq t+1}$ are decreasing in s and k. Then $\Delta\Pi^k_{t+1}(y^1) \geq \Delta\Pi^k_t(y^1)$.*

Proof. Using Lemma 4 we can set $y^j = 0$ for $j \geq 2$. For shorthand, write

$$\tilde{\Pi}^{k-1}_{t+1}(z) := \tilde{\Pi}^{k-1}_{t+1}(z, v^2_{t+1} \ldots, v^{k-1}_{t+1}).$$

Using (A.10) we can write

$$\Delta\Pi^k_t(y^1) = m(y^1) - \delta E_{t+1}[m(\overline{z})] + \delta E_{t+1}[\tilde{\Pi}^{k-1}_{t+1}\left(v^1_{t+1}\right) - \tilde{\Pi}^{k-1}_{t+1}\left(\underline{\overline{z}}\right)] \tag{A.13}$$

where $\overline{z} := \max\{y^1, v^1_{t+1}\}$, $\underline{z} := \min\{y^1, v^1_{t+1}\}$ and $\underline{\overline{z}} := \max\{\underline{z}, v^k_{t+1}\}$. Using the envelope theorem,

$$\tilde{\Pi}^{k-1}_{t+1}(v^1) - \tilde{\Pi}^{k-1}_{t+1}(\underline{\overline{z}}) = \delta^{-1}E_{t+2}\left[\int_{\underline{\overline{z}}}^{v^1_{t+1}} m'(z)\delta^{\overline{\tau}^{k-1}(z)-t}\,dz\right]$$

where $\overline{\tau}^k(z)$ is the time the object is allocated to type z, holding $\{v^2_{t+1}, \ldots, v^{k-1}_{t+1}\}$ constant. As t increases the cutoff x^k_t decreases and $\overline{\tau}^k(z) - t$ falls. Hence $\delta^{\overline{\tau}^k(z)-t}$ and $\tilde{\Pi}^{k-1}_{t+1}(v^1_{t+1}) - \tilde{\Pi}^{k-1}_{t+1}(\underline{\overline{z}})$ increases. Since N_t is IID, $\Delta\Pi^k_t$ increases, as required. □

A.3 Extending Proposition 2 to Decreasing Entry

Lemma 7′. *Fix t and suppose $\{x^k_s\}_{s \geq t+1}$ are decreasing in s and k. Then $\Delta\Pi^k_{t+1}(y^1) \geq \Delta\Pi^k_t(y^1)$.*

Proof. Let $\hat{v}^j_{t+2}$ be the order statistics at time $t+2$ if the number of bidders N_{t+2} were drawn from the distribution of entrants at time $t+1$. Define

$$\begin{aligned}\Delta\hat{\Pi}^k_{t+1}(y^1) = m(y^1) &+ \delta E_{t+2}\left[\tilde{\Pi}^{k-1}_{t+2}(\{y^2, \ldots, y^k, \hat{v}^1_{t+2}, \ldots, \hat{v}^k_{t+2}\}^1_{k-1})\right] \\ &- \delta E_{t+2}\left[\max\{m(y^1), m(\hat{v}^1_{t+2})\}\right] + \delta E_{t+2}\left[\tilde{\Pi}^{k-1}_{t+2}(\{y^1, y^2, \ldots, y^k, \hat{v}^1_{t+2}, \ldots, \hat{v}^k_{t+2}\}^2_k)\right]\end{aligned}$$

where we have replaced v^j_{t+2} with $\hat{v}^j_{t+2}$, for $j \in \{1,\dots,k\}$. Since N_t is decreasing in the usual stochastic order, $\hat{v}^j_{t+2}$ exceeds v^j_{t+2} in the usual stochastic order. Since each entrant buys earlier under "sell tomorrow", this change increases $\Pi^k_t(\text{sell tomorrow})$ more than $\Pi^k_t(\text{sell 1 today})$. Hence $\Delta\Pi^k_{t+1}(y^1) \geq \Delta\hat{\Pi}^k_{t+1}(y^1)$.[24]

We now prove that $\Delta\hat{\Pi}^k_{t+1}(y^1) \geq \Delta\Pi^k_t(y^1)$. Since $\hat{v}^j_{t+2}$ and v^j_{t+1} have the same distribution, we can assume that $\hat{v}^j_{t+2} = v^j_{t+1}$ for each j. As in the proof of Lemma 7, we then have,

$$\Delta\Pi^k_t(y^1) = m(y^1) - \delta E_{t+1}[m(\underline{z})] + \delta E_{t+1}[\tilde{\Pi}^{k-1}_{t+1}(v^1_{t+1}) - \tilde{\Pi}^{k-1}_{t+1}(\overline{z})].$$

Using the envelope theorem,

$$\tilde{\Pi}^{k-1}_{t+1}(v^1) - \tilde{\Pi}^{k-1}_{t+1}(\overline{z}) = \delta^{-1} E_{t+2}\left[\int_{\overline{z}}^{v^1_{t+1}} m'(z)\delta^{\overline{\tau}^{k-1}(z)-t}\, dz\right]$$

where $\overline{\tau}^k(z)$ is the time the object is allocated to type z, holding $\{v^2_{t+1},\dots,v^{k-1}_{t+1}\}$ constant. Since the cutoff types are decreasing in k, agent z buys the first time (a) he has the highest valuation, and (b) his type exceeds the cutoff. Since (a) future order statistics are decreasing in t, and (b) future cutoffs decrease in t, $\overline{\tau}^k(z) - t$ decreases in t. Hence $\delta^{\overline{\tau}^k(z)-t}$ and $\tilde{\Pi}^{k-1}_{t+1}(v^1_{t+1}) - \tilde{\Pi}^{k-1}_{t+1}(\tilde{y}^1)$ increases in t. Since $\hat{v}^1_{t+2}$ and v^1_{t+1} have the same distribution, $\Delta\hat{\Pi}^k_{t+1}(y^1) \geq \Delta\Pi^k_t(y^1)$, as required. $\square$

A.4 Derivation of Equation (4.5)

Suppose the highest agent at time t, y^1, has value equal to the cutoff, x^k_t. Let y^2 be the value of the second highest agent, and denote its distribution function conditional on the entire history of cutoffs by H_t. Cutoffs are decreasing over time so if y^1 delays, agent y^2 may buy. Given that arrivals are independent, this occurs with probability

$$1 - \Pr(y^2 \leq x^k_{t+dt} | y^1 = x^k_t, N_t = 1) = 1 - H_t(x^k_{t+dt}).$$

Note that $H_t(x^k_t) = 1$ and the density is

$$h_t(x^k_t) = \Pr(y^2 = x^k_t | \text{ past cutoffs}) = \lambda(t - s_t(x^k_t))f(x^k_t)$$

[24]This is analogous to Lemma 6.

where $s_t(x)$ is the last time the cutoff went below x. Prices are determined by the cutoff type's indifference condition,

$$(x_t^k - p_t^k) = (1 - rdt - \lambda dt)\left[x_t^k - p_{t+dt}^k\right] H_t(x_{t+dt}^k) + (1 - rdt - \lambda dt)U_t^{k-1}(x_t^k)\left[1 - H_t(x_{t+dt}^k)\right]$$
$$+ (\lambda dt)(x_t^k - p_{t+dt}^k)F(x_t^k) + (\lambda dt)U_t^{k-1}(x_t^k)(1 - F(x_t^k))$$

Rearranging and letting $dt \to 0$ yields (4.5).

A.5 Proof of Proposition 5

In period T, the seller awards the good to the agent with the highest value, subject to his marginal revenue exceeding zero, implying that $m(x_T) = 0$. We next claim that x_t are weakly increasing for $t < T$. Suppose, by contradiction, that there exists $t < T-1$ such that $x_t > x_{t+1}$. Then the cutoff x_t is given by

$$m(x_t) = \delta E_{t+1}[\max\{m(x_t), m(v_{t+1}^1)\}] \tag{A.14}$$

This follows from the fact that type x_t will buy in period $t+1$ if he does not buy in period t. Now consider period $t+1$ and suppose the seller faces a buyer of value x_{t+1}. If the seller delays she obtains at least $\delta E_{t+2}[\max\{m(x_{t+1}), m(v_{t+2}^1)\}]$. Indifference therefore implies that

$$m(x_{t+1}) \geq \delta E_{t+2}[\max\{m(x_{t+1}), m(v_{t+2}^1)\}] \tag{A.15}$$

Since N_t is increasing in the usual stochastic order, v_{t+2}^1 is larger than v_{t+1}^1 in the usual stochastic order, so (A.14) and (A.15) imply $x_{t+1} \geq x_t$, yielding a contradiction.

Fix $t < T$. If the seller sells to type x_t, she obtains $m(x_t)$. If the seller delays, she obtains $\delta\pi_{t+1}(x_t)$, as defined by (6.2), where we use the fact that x_t will not buy in period $t+1$ because the cutoffs are increasing. The seller is indifferent between selling to type x_t and delaying, yielding (6.3), as required.

A.6 Proof of Proposition 6

We first show that for $t < T-1$, $m(x_{t+1}) \geq \beta m(x_t)$. By contradiction, let t be the last time this inequality is not satisfied, so $m(x_{t+1}) < \beta m(x_t)$. It follows that, if the seller chooses not to sell to type x_t at time t then she will sell at time $t+1$. Hence the time-t cutoff is determined by

$$m(x_t) = \delta E_{t+1}[\max\{\beta m(x_t), m(v_{t+1}^1)\}] \tag{A.16}$$

where the left-hand side is the payoff today, and the right-hand side is the payoff from delaying using the fact that $m(x_{t+1}) < \beta m(x_t)$. At time $t+1$ the cutoff satisfies

$$m(x_{t+1}) \geq \delta E_{t+2}[\max\{\beta m(x_{t+1}), m(v^1_{t+2})\}] \tag{A.17}$$

where the right-hand side is lower bound on the value from delaying. From (A.16) and (A.17), $m(x_{t+1}) \geq m(x_t)$, which contradicts the assumption that $m(x_{t+1}) < \beta m(x_t)$.

Since $m(x_{t+1}) \geq \beta m(x_t)$, we know that if type x_t does not obtain a good in period t then he will not obtain one until period T. The resulting indifference equation yields equation (7.3).

Finally, since the seller can always choose allocate the good by period $T-1$, the profit function obeys $\pi_t(v) \geq \pi_{t+1}(v)$. Equation (7.2) thus implies that x_t decreases over time.

References

Albright, S. C. (1974): "Optimal Sequential Assignments with Random Arrival Times Optimal Sequential Assignments with Random Arrival Times," *Management Science*, 21(1), 60–67.

Aviv, Y., and A. Pazgal (2008): "Optimal Pricing of Seasonal Products in the Presence of Forward-Looking Consumers," *Manufacturing & Service Operations Management*, 10(3), 339–359.

Bertsekas, D. P. (1995): *Dynamic Programming and Optimal Control*, vol. II. Athena Scientific, second edn.

Board, S. (2007): "Selling Options," *Journal of Economic Theory*, 136(1), 324–340.

——— (2008): "Durable-Goods Monopoly with Varying Demand," *Review of Economic Studies*, 75(2), 391–413.

Chevalier, J., and A. Goolsbee (2009): "Are Durable Goods Consumers Forward Looking? Evidence from the College Textbook Market," *Quarterly Journal of Economics*, 124(4), 1853–1884.

Conlisk, J., E. Gerstner, and J. Sobel (1984): "Cyclic Pricing by a Durable Goods Monopolist," *Quarterly Journal of Economics*, 99(3), 489–505.

Coy, P. (2000): "The Power of Smart Pricing," *Business Week*, pp. 160–164.

Davis, P. (1994): "Airline Ties Profitability Yield to Management," *SIAM News*, 27(5).

Derman, C., G. J. Lieberman, and S. M. Ross (1972): "A Sequential Stochastic Assignment Problem," *Management Science*, 18(7), 349–355.

Friend, S., and P. Walker (2001): "Welcome to the New World of Merchandising," *Harvard Business Review*, 79, 133–141.

Fuchs, W., and A. Skrzypacz (2010): "Bargaining with Arrival of New Traders," *American Economic Review*, 100(3), 802–836.

Gallego, G., and G. van Ryzin (1994): "Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons," *Management Science*, 40(8), 999–1020.

Gallien, J. (2006): "Dynamic Mechanism Design for Online Commerce," *Operations Research*, 54(2), 291–310.

Gershkov, A., and B. Moldovanu (2009a): "Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach," *American Economic Journal: Microeconomics*.

——— (2009b): "Learning about the Future and Dynamic Efficience," *American Economic Review*.

——— (2010): "Incentive Compatible Optimal Stopping," Working paper, Bonn.

Gowrisankaran, G., and M. Rysman (2009): "Dynamics of Consumer Demand for New Durable Goods," Working paper, Boston University.

Harris, M., and A. Raviv (1981): "A Theory of Monopoly Pricing Schemes with Demand Uncertainty," *American Economic Review*, 71(3), 347–65.

Hartmann, W. (2006): "Intertemporal effects of consumption and their implications for demand elasticity estimates," *Quantitative Marketing and Economics*, 4(4), 325–349.

Hörner, J., and L. Samuelson (2008): "Managing Strategic Buyers," Working paper, Yale University.

Li, A. (2009): "Selling Storable Goods to a Dynamic Population of Buyers: A Mechanism Design Approach," Working paper, Stanford University.

Mas-Colell, A., M. D. Whinston, and J. R. Green (1995): *Microeconomic Theory*. Oxford University Press.

McAfee, P., and J. McMillan (1988): "Search Mechanisms," *Journal of Economic Theory*, 44(1), 99–123.

Mierendorff, K. (2009): "Optimal Dynamic Mechanism Design with Deadlines," Working paper, University of Bonn.

Pai, M., and R. Vohra (2008): "Optimal Dynamic Auctions," Working paper, Northwestern University.

Said, M. (2009): "Auctions with Dynamic Populations: Efficiency and Revenue Maximization," Working paper, Northwestern University.

Segal, I. (2003): "Awarding Monopoly Franchises," *American Economic Review*, 93(3), 509–529.

Shen, Z.-J. M., and X. Su (2007): "Customer Behavior Modeling in Revenue Management and Auctions: A Review and New Research Opportunities," *Production & Operations Management*, 16(6), 713–728.

Smith, B. C., J. F. Leimkuhler, and R. M. Darrow (1992): "Yield Management at American Airlines," *Interfaces*, 22(1), 8–31.

Stokey, N. L. (1979): "Intertemporal Price Discrimination," *Quarterly Journal of Economics*, 93(3), 355–371.

Sweeting, A. (2010): "Dynamic Pricing Behavior in Perishable Goods Markets: Evidence from Secondary Markets for Major League Baseball Tickets," Working paper, Duke University.

Talluri, K., and G. van Ryzin (2004): *The Theory and Practice of Revenue Management*. Springer.

Vulcano, G., G. van Ryzin, and C. Maglaras (2002): "Optimal Dynamic Auctions for Revenue Management," *Management Science*, 48(11), 1388–1407.

Wang, R. (1993): "Auctions versus Posted-Price Selling," *American Economic Review*, 83(4), 838–851.

Zeithammer, R. (2006): "Forward-looking bidding in online auctions," *Journal of Marketing Research*, 43(3), 462–476.