Revenue Management with Forward-Looking Buyers*

Simon Board[†] Andrzej Skrzypacz[‡]

June 22, 2013

Abstract

We consider a seller who wishes to sell K goods by time T. Potential buyers enter 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 buyer with the highest valuation exceeding a cutoff. These cutoffs are deterministic, depending only on the number of units and time remaining. Since the seller does not need to elicit all entrants' values, she can implement the optimal cutoffs in the continuous time limit by posting anonymous prices, with an auction for the last unit at time T. These prices depend on the number of units and time remaining and, unlike the optimal cutoffs, the timing of previous sales. When incoming demand is decreasing over time, the optimal cutoffs satisfy a one-period-look-ahead property and prices are defined by an intuitive differential equation.

1 Introduction

Revenue management considers the problem of selling a limited number of goods to customers entering a market over time. These techniques were first employed by the airline industry after deregulation, and have been successful in helping retailers sell apparel and hotels manage their inventory of rooms. More recently, with firms' increased ability to control sales in real time, revenue management is being used by online ticket distributors, websites pricing ad space and e-retailers in online marketplaces.

The purpose of this paper is to derive the optimal revenue management policy of a seller facing forward-looking buyers who can strategically time their purchases. The seller wishes

*We thank Jeremy Bulow, Yeon-Koo Che, Songzi Du, Drew Fudenberg, Willie Fuchs, Alex Frankel, Mike Harrison, Johanna He, Jon Levin, Yair Livne, Rob McMillan, Moritz Meyer-ter-Vehn, Benny Moldovanu and Dmitry Orlov 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 Operations, 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/

to sell K goods by time T and commits to a dynamic mechanism at the start of the game, analogous to a retailer designing an inventory management system. Potential buyers enter the market stochastically over time and possess privately known values and arrival times. Once they arrive, buyers can delay their purchases, incurring a costly delay and risking a stock-out in the hope of lower prices. Typical revenue management models assume that buyers are myopic, exiting the market if they do not immediately buy (see the book by Talluri and van Ryzin (2004)). Our contribution is to characterize the optimal mechanism when buyers are forward looking, and to show that it can be implemented with posted prices.

Our model applies to settings where a seller has a finite number of units to sell and (i) the selling time is finite, (ii) it is socially optimal for trade to occur sooner rather than later, and (iii) buyers are forward looking. There are a number of applications that have these features. In retail markets, sellers (e.g. Zara, H&M) manufacture a number of units of a new design that they seek to sell over the subsequent season. Buyers arrive over time, are impatient to buy the good to wear it more frequently, and are often forward looking. Soysal and Krishnamurthi (2012) structurally estimate a model of consumer choice for women's coats and find a substantial fraction of customers are forward looking. To illustrate, Figure 1 illustrates a typical sales pattern, showing large spikes in demand after price reductions that quickly fade away, suggesting these customers strategically wait for the discount. A second application is for dealers at the Fulton Fish Market who must sell their fish prior to the end of the week, or before they go bad (Graddy (2006)). Typical buyers are restaurants and wholesalers who are impatient because of the opportunity cost of their time, and are experienced enough to time their purchases to arbitrage any substantial price declines over the day. Third, in the market for display ads, companies such as the New York Times and Yahoo sell ads on their homepage prior to the broadcast date. Buyers (e.g. Netflix, Universal Studios) prefer to buy sooner to plan an advertising campaign, and are typically experienced players so it is natural to think of them as forward looking. Indeed, they sometimes purchase "contingent contracts" which reserve an ad slot at a discount if no other buyer subsequently pays the full price. This type of contract only makes sense with forward-looking buyers and is an alternative implementation of our optimal mechanism. Finally, airlines wish to sell tickets prior to the flight date to buyers who enter over time. Buyers are impatient to purchase so they can plan complementary activities and, as shown by Li, Granados, and Netessine (2011), are often forward looking. Indeed, the rise of price prediction sites such as Bing Travel and the availability of "price alerts" should make such behavior more common.

In the model, N_t buyers arrive in period t; N_t is independent of past arrivals, but the distribution may change over time. Motivated by online markets, we assume N_t is not observed by either buyers or the seller. Upon entering, each buyer has unit demand, draws a private value from a common distribution, and calculates whether to buy today or wait at the risk of

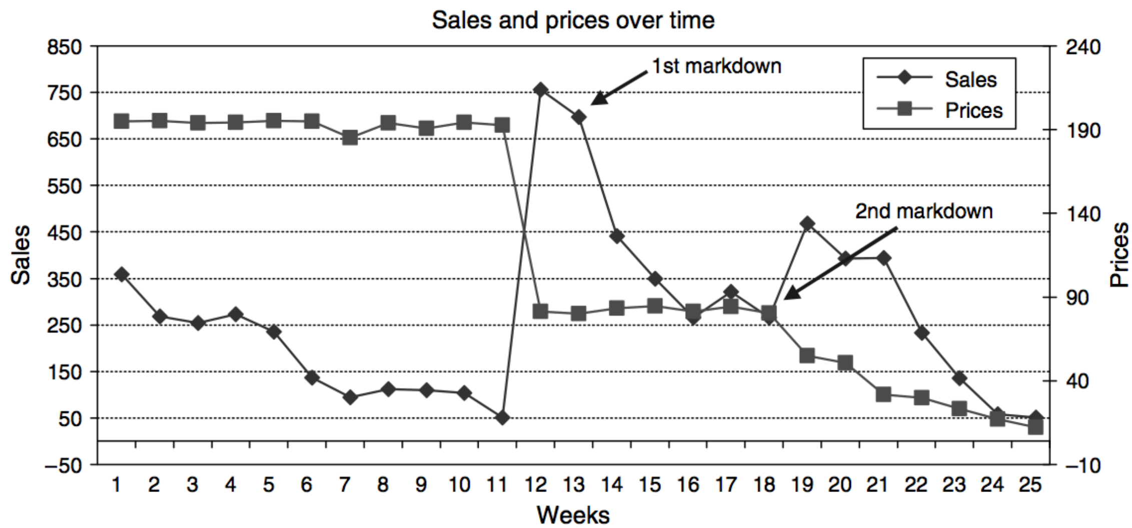

Figure 1: **Prices and Sales for a Sample Women's Coat.** Source: Soysal and Krishnamurthi (2012).

a stock-out. Finally, we model the impatience for purchasing a unit of the good by assuming proportional discounting or, alternatively, by having the seller incur inventory costs.

We have two sets of results. 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 implement these allocations through posted prices 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.

In Section 4, we characterize optimal allocations. We first show that the seller awards a good to the buyer with the highest valuation, if their value exceeds a cutoff. Multiple units may be allocated within a period if the highest value exceeds the cutoff when k goods remain, the second highest value exceeds the cutoff when $k-1$ goods remain, and so on. 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 past entrants; 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 buyers buy, which only depends on their valuations and ranks. Hence changing these buyers' values raises the profits from selling and delaying equally and does not affect the cutoff type. Since cutoffs are deterministic, the seller does not need to elicit the valuations of lower-value buyers when deciding whether or not to allocate it to the highest-value buyer.

The optimal cutoffs are decreasing in the inventory size, k. Intuitively, if the seller delays awarding the k^{th} unit then she can allocate it to an entrant rather than the current leader. As k rises the current leader is increasingly likely to be awarded the good eventually, decreasing the option value of delay and causing the cutoff to fall. When the number of entering buyers is weakly decreasing over time (in the usual stochastic order), the optimal cutoffs are also decreasing over time and satisfy a *one-period-look-ahead property* whereby the seller is indifferent between serving the cutoff type today and waiting exactly one more period before selling that unit. Analogous to the above intuition, as the seller gets closer to the deadline T, the option value of delaying awarding a unit falls, as does the cutoff.

In Section 5, we consider implementation in the continuous time limit, assuming buyers arrive according to a time-varying Poisson process. We show that the seller can implement the optimal mechanism with posted prices and an auction for the final good at time T. That is, the seller chooses a single price at each point in time and buyers only reveal their existence when they purchase a unit. Intuitively, given the optimal cutoffs, we can choose prices to make the cutoff type indifferent between buying and waiting. Prices are an imperfect way to implement optimal cutoffs for two reasons: the induced cutoffs cannot dynamically adjust within a period, and there may be rationing if demand outstrips supply. However, when periods are short, the optimal cutoffs do not jump down (much) at any time $t < T$, so the scope of inefficiency shrinks with the period length. Note that the deterministic cutoffs are crucial for the simplicity of this solution: if the decision to award buyer 1 the good depends on buyer 2's value then the seller has to obtain reports about all buyers' values upon entering the market, rather than using posted prices.

When entering demand is weakly decreasing, prices are determined by an intuitive differential equation. If the cutoff type waits a little then he gains from the price decrease, but he loses some utility from delay, and risks the good being bought by either a new entrant or another waiting buyer. As a result, the optimal prices depend on the number of units and time remaining and, unlike the optimal cutoffs, the timing 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. When compared to a model with myopic buyers, our model implies that profits and overall sales are higher, but that many more of these sales occur near the deadline T. This pattern of posted prices and a "last-minute" auction is qualitatively consistent with internet sites selling plane and hotel reservations.

In Section 6, we consider a number of applications. These serve the dual purpose of illustrating the applicability of the model while allowing us to explore the robustness of the main results. We first show that the spirit of the main results continues to apply if impatience comes from inventory costs (e.g. for a retailer where shelf space is costly), or if units arrive and expire over time (e.g. for a dealer at the Fulton Fish Market). Second, if there are different classes of buyers (e.g. buyers of rich media and static display ads) or if the distribution of entering buyers gets stronger over time (e.g. for an airline as the flight date approaches), then the cutoffs are defined in marginal revenue space, with the seller charging different prices for different types of buyers. Finally, if buyers disappear probabilistically (e.g. when selling a house) then optimal cutoffs are no longer deterministic. This helps explain why real estate sellers use an auction wherein all buyers bid and the seller makes a counteroffer to the highest, rather than using posted prices.

1.1 Literature

Gallien (2006) characterizes the optimal sequence of prices when a seller of multiple units faces buyers who arrive according to a renewal process over an infinite time horizon. Assuming inter-arrival times have an increasing failure rate, Gallien proves that buyers will buy when they enter the market (or not at all) and the solution thus corresponds to that without recall (e.g. Gallego and van Ryzin (1994)) with infinite time. In contrast, our finite horizon means that the optimal mechanism will induce buyers to delay their purchases on the equilibrium path.

Pai and Vohra (2013) consider a model where a seller has multiple units and wishes to sell them over finite time. This model is very rich, allowing buyers arrive and leave the market over time, and the distribution of entering buyers to vary over time. Mierendorff (2009) considers a two-period version of a similar model and provides a complete characterization of the optimal contract. Su (2007) considers a model with heterogeneous values and discount rates, examining how the interaction of these terms determines the optimal price paths. Aviv and Pazgal (2008) consider a model similar to ours, but restrict the seller to choose two prices that are independent of the past sales; this is extended to multiple markdowns by Elmaghraby, Gülcü, and Keskinocak (2008). In contrast to these papers, we characterize the optimal mechanism and show that the price path required to implement it will jump up when sales occur.

The single-unit version of our model is closely related to the classic "house selling" problem with recall, where an owner receives offers for his house and picks the largest (e.g. MacQueen and Miller (1960)). When there is a single house and valuations are IID, the cutoff value is constant, except for the last period (see Bertsekas (2005, p. 185)). McAfee and McMillan (1988) introduce private information into this model, changing values into marginal revenues. With regard to this literature, our price-posting implementation is new, as is our analysis of multiple units.

There are a number of adjacent literatures. Buyers' values may vary over time (e.g. Board (2007)). The firm may be unable to commit (e.g. Hörner and Samuelson (2011)). There may be heterogenous goods (e.g. Gershkov and Moldovanu (2009a)) or learning about the distribution of valuations (e.g. Gershkov and Moldovanu (2009b)). The seller may pay the inventory cost until all units of the good have been sold (e.g. Bruss and Ferguson (1997)). There is also a large literature on selling durable goods without capacity constraints (e.g. Stokey (1979)).

2 Model

Basics. A seller (she) has K goods to sell to buyers (he) arriving over time. Time is discrete and finite, $t \in \{1, \ldots, T\}$. Time-preference comes from a common discount factor $\delta \in [0, 1)$.

Entrants. At the start of period t, N_t buyers arrive. N_t is independent of past arrivals; the distribution of arrivals may change over time, allowing us to talk of "increasing" and "decreasing" demand (in the usual stochastic order). For simplicity, we assume the number of arrivals N_t is observed by the seller, but not by other buyers. Our analysis is unchanged if N_t is also unobserved by the seller. In this latter case, one can think of us solving the "relaxed" problem, ignoring the (IC) constraints on birth-dates. In Section 4.1 we show that the optimal allocation is characterized by cutoffs that are independent of buyers' birth-dates, so the (IC) constraints are satisfied in the optimal mechanism.

Preferences. After a buyer has entered the market, he wishes to buy a single unit. A buyer is thus endowed with type (v_i, t_i), where v_i denotes his valuation, and t_i his birth-date. The buyer's valuation, v_i, is private information and drawn IID with continuous density $f(\cdot)$, distribution $F(\cdot)$ and support $[\underline{v}, \overline{v}]$. The buyer's birth-date, t_i, is observed by the seller but not by other buyers. If the buyer purchases at time s for price p_s, his utility is $(v - p_s)\delta^s$.[1] Let v_t^k denote the k^{th} highest order statistic of the buyers entering at time t.

Mechanisms. At time 0 the seller chooses a mechanism. Each buyer makes report $\tilde{v}_i$ when he enters the market. Without loss of generality we can restrict ourselves to mechanisms where the buyers do not observe the history, including the number of objects available (Myerson (1986)). A mechanism consists of an allocation rule and transfer $\langle \tau_i, TR_i \rangle$ that maps buyers' reports and birth-dates into a purchasing time τ_i for buyer i, and expected transfer TR_i expressed in time-0 prices. A mechanism is *feasible* if (a) $\tau_i \geq t_i$, (b) $\sum_i \mathbf{1}_{\tau_i < \infty} \leq K$, and (c) τ_i is adapted to the seller's information (i.e. the reports and birth-dates of entrants).[2]

[1] 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 buyers' values fall over time (e.g. values for summer clothes will be lower in July than in June), then utility is given by $v\delta^t - p_t$. Whether or not we discount money does not affect allocations, although prices must be rescaled under this new specification. For example, in (5.3) the last term becomes $-rx_t^k$ rather than $-r(x_t^k - p_t^k)$.

[2] We work with deterministic mechanisms, but one can allow for random allocation by letting the mechanism describe a probability space $\langle \Omega, \mathcal{F}, P \rangle$, and letting the purchasing time depend on $\omega \in \Omega$. Allocation is linear in

Buyer's Problem. Upon entering the market, buyer i chooses his report $\tilde{v}_i$ to maximize his expected utility,

$$u_i(\tilde{v}_i, v_i, t_i) = E_0\left[v_i \delta^{\tau_i(\tilde{v}_i, \mathbf{v_{-i}}, \mathbf{t})} - TR_i(\tilde{v}_i, \mathbf{v_{-i}}, \mathbf{t})\Big| v_i, t_i\right] \tag{2.1}$$

where E_s denotes the expectation over buyers' values at the start of period s, before buyers have entered the market, $\mathbf{v}$ is the vector of buyers' values, and $\mathbf{t}$ is the vector of their birth-dates. A mechanism is incentive compatible if the buyer wishes to tell the truth given all others are truthful, and is individually rational if the buyer obtains positive utility.

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

$$\text{Profit} = E_0\left[\sum_i TR_i(\mathbf{v}, \mathbf{t})\right] \tag{2.2}$$

subject to incentive compatibility and individual rationality.

Remarks. First, time T can be viewed as the date at which the good expires (e.g. a plane ticket), the last time the customers get utility from the item (e.g. seasonal clothing) or the last time buyers enter the market since, in the optimal mechanism, no sales will occur after this point.

Second, we assume that buyers do not know the number of units remaining in the mechanism (indeed, they only know their value and birth-date). However, when implementing the optimal allocation, the seller will share this information with the buyers, so the seller does not benefit from hiding her remaining inventory.

Third, we solve for the profit-maximizing allocation because of the practical relevance. If one replaces marginal revenues (defined below) with values, all our results apply to the welfare-maximizing allocation. Our results can therefore be viewed as characterizing a general K-unit optimal stopping problem and can be applied to a consumer who searches for K goods, or a firm that wishes to hire K employees.

Finally, we assume the seller can commit to a mechanism. We think this is reasonable with applications such as display ads, retailers and airlines where the seller automates the pricing scheme and uses it repeatedly. It is also appropriate when using the model from a normative perspective to design, say, the dynamic pricing strategy of an online ticket seller. In addition, one can view this as an upper bound on what a seller can obtain. If the seller cannot commit,

probabilities, and we assume marginal revenue is increasing in values, so a deterministic mechanism is optimal.

the problem is much harder to study, e.g. Fuchs and Skrzypacz (2010), Hörner and Samuelson (2011), Dilme and Li (2012).

2.1 Preliminaries

When a buyer enters the market, he chooses his report $\tilde{v}_i$ to maximize his utility (2.1). As shown by Mas-Colell, Whinston, and Green (1995, Proposition 23.D.2), an allocation rule is incentive compatible if and only if the discounted allocation probability

$$E_0\left[\delta^{\tau_i(\mathbf{v},\mathbf{t})}\middle|v_i,t_i\right] \tag{2.3}$$

is increasing in v_i. Using the envelope theorem, equilibrium utility is

$$u_i(v_i,v_i,t_i)=E_0\left[\int_{\underline{v}}^{v_i}\delta^{\tau_i(z,\mathbf{v_{-i}},\mathbf{t})}dz\middle|v_i,t_i\right]$$

where we use the fact that a buyer with value $\underline{v}$ earns zero utility in any profit-maximizing mechanism. Taking expectations over (v_i,t_i) and integrating by parts then yields,

$$E_0[u_i(v_i,v_i,t_i)]=E_0\left[\delta^{\tau_i(\mathbf{v},\mathbf{t})}\frac{1-F(v_i)}{f(v_i)}\right] \tag{2.4}$$

Profit (2.2) equals welfare minus buyers' utilities. Summing utility (2.4) over all buyers, we obtain

$$\text{Profit}=E_0\left[\sum_i\delta^{\tau_i(\mathbf{v},\mathbf{t})}m(v_i)\right] \tag{2.5}$$

where the *marginal revenue* of buyer i is given by $m(v_i):=v_i-(1-F(v_i))/f(v_i)$. Throughout we assume $m(v)$ is strictly increasing and continuously differentiable in v, implying that the seller's optimal allocations are monotone in valuations and allowing us to ignore the monotonicity constraint (2.3). We also assume that $m(\underline{v})<0$, so the optimal cutoff is interior.

3 Single-Unit Example

Before launching into the main analysis we develop some intuition by heuristically deriving the optimal allocation and prices for the case when $K=1$ and N_t is IID. To motivate this, suppose YouTube wishes to sell the main banner ad on its front page.

First, consider optimal allocations. Since marginal revenue is increasing, the seller will award the good to the buyer with the highest value if it exceeds a cutoff, x_t. As is well known (e.g. Bertsekas (2005, p. 185)), the cutoffs are constant $x_t=x^*$ in periods $t<T$ and are

uniquely given by[3]

$$m(x^*) = \delta E_{t+1}[\max\{m(v^1_{t+1}), m(x^*)\}] \tag{3.1}$$

In period T, the cutoff jumps down to the usual monopoly quantity, $m(x_T) = 0$. This last period is identical to a standard auction, so the seller wants to sell to the buyer with the highest value as long as their their marginal revenue is positive. In earlier periods the seller balances the benefit from selling today (the left-hand side) against the benefit of waiting one period, receiving a new draw but discounting the profit (the right-hand-side). The cutoff is constant because a seller who does not sell to x_t today faces exactly the same tradeoff tomorrow of waiting one more period, and is once again indifferent between selling and waiting.

The cutoffs do not depend on the number of buyers who have entered in the past and their valuations. This matters because the seller can implement the optimal mechanism without observing the number of arrivals. In addition, the cutoffs satisfy a one-period-look-ahead property, with the seller being indifferent between awarding the good to the cutoff type today and waiting exactly one period. These two properties also hold in the multi-unit case, as shown in Sections 4.1 and 4.2.

To consider the continuous time limit, suppose buyers enter with Poisson rate λ and the instantaneous discount rate is r. Then (3.1) becomes

$$rm(x^*) = \lambda E_v\big[\max\{m(v) - m(x^*), 0\}\big] \tag{3.2}$$

where E_v is the expectation over $v \sim F(\cdot)$. If the seller waits dt she loses the flow profit from the cutoff type (the left-hand-side) but gains the option value of waiting for a new entrant (the right-hand-side). At time T, the optimal cutoff is given by $m(x_T) = 0$.

The optimal allocation can be implemented by a deterministic sequence of prices with an auction at time T. In the last period, the seller uses a second-price auction with reserve $m^{-1}(0)$. At time $t < T$ the seller chooses a price p_t that 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,

$$p_T = E_0\left[\max\{y^2, m^{-1}(0)\}\middle|y^1 = x^*\right] \tag{3.3}$$

where y^j is the value of the j^{th} highest buyer in the market at time T. When $t < T$, the indifference equation for buyer x^* yields

$$\frac{dp_t}{dt} = -(x^* - p_t)\big(\lambda(1 - F(x^*)) + r\big). \tag{3.4}$$

[3]We omit the proof: This result is a special case of Theorem 2.

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 enters with a value above x^* (the right hand side). Even though the cutoffs are constant, prices decline since a delaying buyer loses one period's enjoyment of the good and risks a stockout. Price are also concave, falling more rapidly as $t \to T$.

While we focus on implementation via posted prices, the seller can also implement the optimal allocations via a *conditional contract* that awards a buyer the good at time $s > t$ if no-one offers a better price beforehand. Formally, suppose the price path p_t is given by (3.4) with boundary condition (3.3). In the game, a buyer bids b at any time after they enter. If $b \geq p_T$, then the buyer buys the good at time $\min\{s : p_s = b\}$ subject to no other buyer bidding more beforehand. If $b < p_T$, then this is treated as a bid in a first-price auction held at time T.[4] This implementation is related to the "red zone contracts" used by some firms (e.g. YouTube) to sell their front-page banner ad. Such a contract allows a firm to buy at a discount if no buyer is willing to pay the full price. It is also related to websites that allow buyers to set up a "price alert" to inform them of the availability of a particular good (for example, an airline ticket or a hotel room) in case the price drops below some pre-specified level.

4 Optimal Allocations

We now turn to the analysis of the multi-unit model. In Section 4.1 we consider general sequences of the demand process N_t, showing the optimal allocations are characterized by cutoffs that are deterministic. In Section 4.2 we specialize the model to assume N_t is weakly decreasing in the usual stochastic order, and show that the cutoffs satisfy the one-period-look-ahead property.

4.1 General Case

The seller's problem is to choose a feasible allocation rule $\langle \tau_i \rangle$ to maximize profits (2.5). Since this is a single-agent optimization problem, the principle of optimality means we can solve it by backward induction.

Suppose the seller has k units at the start of period t. First, observe that the seller does not discriminate on the basis of birth-dates. That is, if buyer (v_i, t_i) is in the market at time t then their allocation (and the allocation of all other buyers) is independent of t_i. This follows because

[4] The revenue equivalence theorem applies to the auction at time T since we have assumed symmetric bidders with independent private values. Under the assumption of Poisson arrivals, a bidder makes no inference from his time of arrival and hence all bidders below x^* have the same beliefs about their competition and there exists a symmetric equilibrium in the first-price auction. Note that it is important that buyers are not informed about the existing contingent contracts in the system since this would affect their incentives to wait for a lower price and their optimal bids in the auction.

the birth-date only enters through the feasibility requirement that $\tau_i \geq t_i$, and is therefore not payoff relevant at time t. Since a buyer's birth-date does not affect their allocation, the (IC) constraint on the truthful reporting of the buyer's birth-date is slack and the seller need not see when buyers arrive in order to choose the optimal allocations. Intuitively, this follows because the birth-date provides the seller no information about a buyer's valuation.

Second, observe that buyers with high values are allocated goods prior to buyers with low values. That is, if buyers $v_i > v_j$ are in the market at time t then the seller awards a unit to buyer i before buyer j. To see this suppose, by contradiction, that a unit is allocated to buyer v_j at period t', whereas buyer v_i is not allocated a good until period $t'' > t'$. By swapping these two buyers, but leaving everything else unchanged, profits are increased by $(1 - \delta^{t''-t'})(m(v_i) - m(v_j))$, which is strictly positive since $m(\cdot)$ is strictly increasing.

These two observations imply that we can solve the seller's problem via backward induction, only keeping track of the highest k remaining buyers. Denote the ordered vector of the k highest buyers' values in the market at time t by $\mathbf{y} := \{y^1, \ldots, y^k\}$, where $y^i \geq y^{i+1}$. In the final period the seller sells to the buyers with the highest marginal revenues, subject to their marginal revenue being positive, as in Myerson (1981). In earlier periods, the *continuation profit* at the start of period t is[5]

$$\Pi_t^k(\mathbf{y}) := \max_{\tau_i \geq t} E_{t+1} \left[\sum_i \delta^{\tau_i(\mathbf{y})-t} m(v_i) \right]. \tag{4.1}$$

where the E_{t+1} reflects the fact that the period-t entrants have entered, and are included in $\mathbf{y}$. If the seller makes j sales in period t, then we denote the period-$(t+1)$ continuation profit before the period $t+1$ entrants have entered as

$$\tilde{\Pi}_{t+1}^{k-j}(\mathbf{y^{-j}}) := \max_{\tau_i \geq t+1} E_{t+1} \left[\sum_i \delta^{\tau_i(\mathbf{y^{-j}}, \mathbf{v_{t+1}})-(t+1)} m(v_i) \right]. \tag{4.2}$$

for $\mathbf{y^{-j}} := \{y^{j+1}, \ldots, y^k\}$. These equations are related via the Bellman equation

$$\Pi_t^k(\mathbf{y}) = \max_{j \in \{0,\ldots,k\}} \left[\sum_{i=1}^{j} m(y^i) + \delta \tilde{\Pi}_{t+1}^{k-j}(\mathbf{y^{-j}}) \right] \tag{4.3}$$

The following lemma shows that allocation is monotone in buyers' values $\mathbf{y}$. As a result, a mechanism can be characterized by *cutoffs* $x_t^j(\mathbf{y^{-(k-j+1)}})$, $j \leq k$, which is the lowest type that can be awarded the j^{th} unit, assuming the previous $k-j$ units have been sold. Within a period,

[5]While we call Π_t^k continuation profit, this includes the impact of time t decisions on the willingness to pay of buyers who purchase in earlier periods. That is, if we allocate a unit to type v in period t then the seller only receives $m(v)$ since all higher types gain rents, even those that purchase prior to period t.

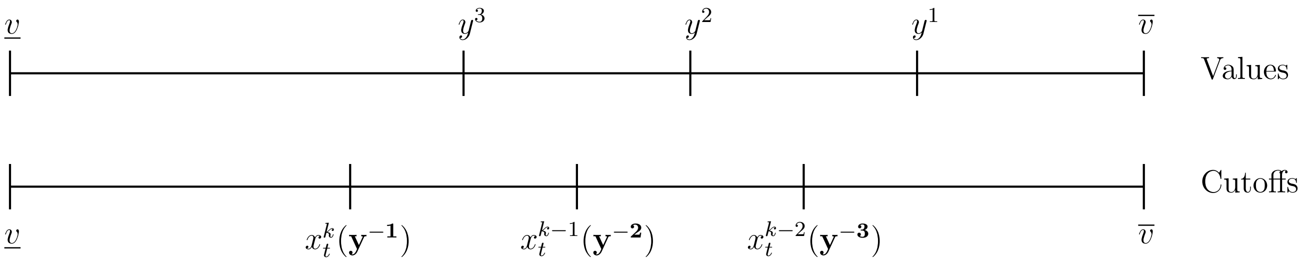

Figure 2: **Allocation within a Period.** This figure shows the top three value buyers and the cutoffs for three units. Unit k is allocated to y^1 since $y^1 \geq x_t^k(\mathbf{y^{-1}})$. Similarly, unit $k-1$ is allocated to y^2 since unit k was sold and $y^2 \geq x_t^{k-1}(\mathbf{y^{-2}})$. However, unit $k-2$ remains unsold as $y^3 < x_t^{k-2}(\mathbf{y^{-3}})$.

several units may be allocated. We allocate unit k to the highest buyer if $y^1 \geq x_t^k(\mathbf{y^{-1}})$, unit $k-1$ to the second highest buyer if $y^2 \geq x_t^{k-1}(\mathbf{y^{-2}})$, and so on. In general, we allocate the j^{th} unit if $y^{k-\ell+1} \geq x_t^\ell(\mathbf{y}^{-(\mathbf{k}-\ell+\mathbf{1})})$ for all $\ell \in \{j, \ldots, k\}$. This is illustrated in Figure 2.

Lemma 1. *Suppose the seller starts period t with k units and buyers with values $\mathbf{y}$. The optimal mechanism can be characterized by cutoffs $x_t^j(\mathbf{y^{-(k-j+1)}})$.*

Proof. Suppose we have sold $k-j$ units in period t. By contradiction, suppose the seller awards unit j to the buyer when their value is y^{k-j+1}, but not when it is $\hat{y}^{k-j+1} > y^{k-j+1}$. By revealed preference,

$$m(y^{k-j+1}) + \Pi_t^{j-1}(\mathbf{y^{-(k-j+1)}}) \geq \delta\tilde{\Pi}_{t+1}^j(y^{k-j+1}, \mathbf{y^{-(k-j+1)}})$$
$$m(\hat{y}^{k-j+1}) + \Pi_t^{j-1}(\mathbf{y^{-(k-j+1)}}) \leq \delta\tilde{\Pi}_{t+1}^j(\hat{y}^{k-j+1}, \mathbf{y^{-(k-j+1)}})$$

Subtracting the first equation from the second,

$$m(\hat{y}^{k-j+1}) - m(y^{k-j+1}) \leq \delta[\tilde{\Pi}_{t+1}^j(\hat{y}^{k-j+1}, \mathbf{y^{-(k-j+1)}}) - \tilde{\Pi}_{t+1}^j(y^{k-j+1}, \mathbf{y^{-(k-j+1)}})] \leq \delta[m(\hat{y}^{k-j+1}) - m(y^{k-j+1})]$$

where the second inequality uses the fact the y^{k-j+1} seller can mimic the strategy of the $\hat{y}^{k-j+1}$ seller from period $t+1$. This yields a contradiction, implying the allocation of unit j is monotone in y^{k-j+1}. As a result, there is a lowest value that purchases, denoted $x_t^j(\mathbf{y^{-(k-j+1)}})$. □

If the seller starts period t with k units, she sells the j^{th} unit if $y^{k-\ell+1} \geq x_t^\ell(\mathbf{y}^{-(\mathbf{k}-\ell+\mathbf{1})})$ for all $\ell \in \{j, \ldots, k\}$, so in order to know if we can sell the j^{th} unit we also need to check all the previous units. That is, the seller may be willing to sell unit $k-1$ once she has sold unit k, but refrains from doing so because she is not willing to sell unit k. The following lemma shows that if cutoffs are decreasing in k, this problem does not arise and we can treat each unit separately, simply comparing the j^{th} cutoff to the corresponding buyer's valuation. Indeed, this is the case in Figure 2.

Lemma 2. *Suppose period-t cutoffs* $x_t^j(\mathbf{y}^{-(\mathbf{k}-\mathbf{j}+\mathbf{1})})$ *are decreasing in* j. *Then unit* j *is allocated iff* $y^{k-j+1} \geq x_t^j(\mathbf{y}^{-(\mathbf{k}-\mathbf{j}+\mathbf{1})})$.

Proof. If unit j is allocated, then the corresponding buyer's value must exceed the cutoff. Conversely, if $y^{k-j+1} \geq x_t^j(\mathbf{y}^{-(\mathbf{k}-\mathbf{j}+\mathbf{1})})$ then $y^{k-\ell+1} \geq y^{k-j+1} \geq x_t^j(\mathbf{y}^{-(\mathbf{k}-\mathbf{j}+\mathbf{1})}) \geq x_t^\ell(\mathbf{y}^{-(\mathbf{k}-\ell+\mathbf{1})})$ for all $\ell \geq j$, since the cutoffs are decreasing. Hence all units $\ell \geq j$ are allocated to their corresponding buyer. □

So far we have been concerned with the cutoff for unit j, assuming that the seller starts the period with k units. Since we solve by backward induction, this introduces redundant notation, and it is without loss to suppose that unit j is the first unit sold in the period. Henceforth, we characterize the cutoffs by considering the sale of unit k to buyer y^1, taking into account that the seller may wish to sell further units.

If cutoffs are decreasing in k then the analysis is further simplified: if the seller is indifferent between selling to buyer y^1 today and waiting, then the seller weakly prefers not to sell a second unit today. Correspondingly, it will be useful to define the profit if the seller sells 0 units today, or if she sells only one.

$$\Pi_t^k(\text{sell 0 today}) = \delta\tilde{\Pi}_{t+1}^k(y^1, \mathbf{y}^{-\mathbf{1}}) \tag{4.4}$$
$$\Pi_t^k(\text{sell 1 today}) = m(y^1) + \delta\tilde{\Pi}_{t+1}^{k-1}(\mathbf{y}^{-\mathbf{1}}) \tag{4.5}$$

Denote the difference function by

$$\Delta\Pi_t^k(y^1, \mathbf{y}^{-\mathbf{1}}) = \Pi_t^k(\text{sell 1 today}) - \Pi_t^k(\text{sell 0 today})$$

which reflects the incentives to sell today rather than wait. Lemma 3 establishes some basic properties of $\Delta\Pi_t^k(y^1, \mathbf{y}^{-\mathbf{1}})$. We say the cutoffs $x_t^k(\mathbf{y}^{-\mathbf{1}})$ are *deterministic* if they are independent of $\mathbf{y}^{-\mathbf{1}}$.

Lemma 3. *Suppose future cutoffs* $\{x_s^j\}_{s\geq t+1}$ *are deterministic and decreasing in* $j \leq k$. *Then*
(a) $\Delta\Pi_t^k(y^1, \mathbf{y}^{-\mathbf{1}})$ *is independent of* $\mathbf{y}^{-\mathbf{1}}$.
(b) $\Delta\Pi_t^k(y^1)$ *is continuous and strictly increasing in* y^1.
(c) $\Delta\Pi_t^k(y^1)$ *is increasing in* k.

Proof. See Appendix A.1. □

Part (a) says that $\Delta\Pi_t^k(y^1, \mathbf{y}^{-\mathbf{1}})$ is independent of $\mathbf{y}^{-\mathbf{1}}$. Intuitively, the decision of whether or not to allocate one object to buyer y^1 does not affect buyer y^2's rank and therefore when they are allocated a good. Hence the value of y^2 does not affect the decision of whether or not to allocate a unit today. Part (b) says that a higher value of y^1 increases $\Delta\Pi_t^k(y^1)$ since the cost of waiting is higher. Part (c) says that more units raise $\Delta\Pi_t^k(y^1)$ reflecting the idea that such a seller is more eager to allocate goods.

We now have our first main result:

Theorem 1. *Suppose the seller has k goods in period t. The optimal allocation rule awards a unit to the highest remaining buyer if their value exceeds a deterministic cutoff x_t^k. The cutoffs x_t^k are decreasing in k and are uniquely determined by $\Delta\Pi_t^k(x_t^k) = 0$.*

Proof. We wish to show that cutoffs x_t^k are decreasing in k and deterministic. We do this by backward induction. In period T, cutoffs are defined by $m(x_T^k) = 0$ and therefore are deterministic and (weakly) decreasing in k. Now fix t and suppose future cutoffs $\{x_s^k\}_{s\geq t+1}$ are deterministic and decreasing in k.

Let $k = 1$, so $\mathbf{y} = y^1$. Lemma 3(b) states that $\Delta\Pi_t^k(y^1)$ is continuously strictly increasing in y^1, so the cutoff is uniquely defined by $\Delta\Pi_t^k(x_t^k) = 0$, and so is (trivially) deterministic.[6]

Continuing by induction, fix $k > 1$ and suppose $x_t^j \leq x_t^{j-1}$ for $j < k$, and that these cutoffs are deterministic. Lemma 3(a) implies that $\Delta\Pi_t^{k-1}(\mathbf{y})$ is independent of $\mathbf{y^{-1}}$, and can thus be written as $\Delta\Pi_t^{k-1}(y^1)$. Lemma 3(b) states that $\Delta\Pi_t^{k-1}(y^1)$ is continuously strictly increasing in y^1. Since $x_t^{k-1} \leq x_t^{k-2}$, the cutoff x_t^{k-1} is uniquely defined by $\Delta\Pi_t^{k-1}(x_t^{k-1}) = 0$.

Now suppose, by contradiction, that $x_t^k(\mathbf{y^{-1}}) > x_t^{k-1}$ for some $\mathbf{y^{-1}}$. By the envelope theorem, profits are continuous in y^1, so the cutoff is defined by the indifference condition

$$\Pi_t^k(\text{sell 0 today}) = \Pi_t^k(\text{sell} \geq 1 \text{ today}) \geq \Pi_t^k(\text{sell 1 today}),$$

where the inequality uses revealed preference. We thus have

$$0 \geq \Delta\Pi_t^k(x_t^k(\mathbf{y^{-1}})) > \Delta\Pi_t^k(x_t^{k-1}) \geq \Delta\Pi_t^{k-1}(x_t^{k-1}) = 0,$$

where the second inequality comes from Lemma 3(b), and the third inequality follows from Lemma 3(c). We thus have a contradiction.

We thus know that $x_t^k(\mathbf{y^{-1}}) \leq x_t^{k-1}$, so Lemma 3(a),(b) imply that the cutoff is defined by $\Delta\Pi_t^k(x_t^k(\mathbf{y^{-1}})) = 0$. The cutoff x_t^k is thus independent of $\mathbf{y^{-1}}$. □

The optimal cutoffs have two important properties. First, they are deterministic in that they are independent of the values of lower buyers, $\mathbf{y^{-1}}$. Economically, this means the decision to allocate the good to the highest-value buyer depends on the number of periods and units remaining, but not on the number of buyers, their valuations, or when previous units were sold.

[6] We can apply the intermediate value theorem since the $\Delta\Pi_t^k(\underline{v}) \leq m(\underline{v}) < 0$, while $\Delta\Pi_t^k(\overline{v}) = (1-\delta)m(\overline{v}) > 0$.

While the value of $\mathbf{y^{-1}}$ does affect the seller's realized revenue, it does not alter the seller's allocation decision. As a result, the seller does not need to elicit values from buyers as they arrive; this property will become crucial for implementation.

Second, the cutoffs increase when there are fewer units available. Intuitively, if the seller delays awarding a unit by one period then she can allocate it to an entrant, rather than buyer y^1. When there are more goods remaining, buyer y^1 is more likely to be awarded one of the units eventually, reducing the option value of delay and decreasing the cutoff.

4.2 Weakly Decreasing Demand

In this section we suppose incoming demand is weakly decreasing over time. This includes the canonical case of IID arrivals. It captures the idea that the pool of potential new customers falls over time (e.g. if Zara launches a line of coats).

In Theorem 1, a seller is indifferent between selling to buyer x_t^k today and waiting for future entry. We say an allocation satisfies the *one-period-look-ahead property* if the seller is indifferent between selling to buyer x_t^k today and waiting one period, and allocating that unit tomorrow. To analyze this, it will be useful to change the definition of $\Delta\Pi_t^k$ to force a waiting seller to sell at least one unit at time $t+1$. Write the vector of entering buyers as $\mathbf{v_t} := \{v_t^1, \dots, v_t^k\}$, and let $\{y^1, \mathbf{v_{t+1}}\}_k^2$ represent the ordered vector of the 2^{nd} to k^{th} highest choices from $\{y^1, \mathbf{v_{t+1}}\}$. Define

$$\Pi_t^k(\text{sell} \geq 1 \text{ tomorrow}) = \delta E_{t+1}\left[\max\{m(y^1), m(v_{t+1}^1)\} + \Pi_{t+1}^{k-1}(\{y^1, \mathbf{v_{t+1}}\}_k^2)\right]$$

and

$$D\Pi_t^k(y^1) = \Pi_t^k(\text{sell 1 today}) - \Pi_t^k(\text{sell} \geq 1 \text{ tomorrow})$$

which we can write as a function of y^1 alone using the reasoning of Lemma 3(a). Observe that $D\Pi_t^k(y^1) \geq \Delta\Pi_t^k(y^1)$ by revealed preference, with $D\Pi_t^k(y^1) = \Delta\Pi_t^k(y^1)$ if $x_t^k \geq x_{t+1}^k$.

Lemma 4. *Suppose N_t is weakly decreasing in the usual stochastic order, and future cutoffs are decreasing in time, $x_s^j \geq x_{s+1}^j$ for $s \in \{t+1, \dots, T-1\}$ and $j \leq k$. Then $D\Pi_{t+1}^k(y^1) \geq D\Pi_t^k(y^1)$.*

Proof. See Appendix A.2. □

Theorem 2. *Suppose N_t is weakly decreasing in the usual stochastic order. Then the optimal cutoffs x_t^k are decreasing in t. As a result, allocations satisfy the one-period-look-ahead property and are uniquely characterized by $D\Pi_t^k(x_t^k) = 0$.*

Proof. We now show that cutoffs x_t^k are decreasing in t by induction. When $k = 1$, $x_{T-1}^1 \geq x_T^1 = m^{-1}(0)$. Now, consider x_t^k and suppose that $x_s^j \geq x_{s+1}^j$ for all $j < k$ for all s, and for $j = k$ and $s \geq t+1$. Since $x_{t+1}^k \geq x_{t+2}^k$, $D\Pi_{t+1}^k(x_{t+1}^k) = \Delta\Pi_{t+1}^k(x_{t+1}^k) = 0$, where the second equality uses

Theorem 1. Now suppose, by contradiction, that $x_t^k < x_{t+1}^k$, so that $D\Pi_t^k(x_t^k) \geq \Delta\Pi_t^k(x_t^k) = 0$. We then have,

$$0 \leq D\Pi_t^k(x_t^k) < D\Pi_t^k(x_{t+1}^k) \leq D\Pi_{t+1}^k(x_{t+1}^k) = 0.$$

where the second inequality follows from the proof of Lemma 3(b), and the third inequality uses Lemma 4. We thus have a contradiction, implying that $x_t^k \geq x_{t+1}^k$, as required. Given that x_t^k are decreasing in t, the optimal cutoffs are uniquely defined by $D\Pi_t^k(x_t^k) = 0$. □

Intuitively, if the seller delays awarding the k^{th} unit by one period then she can allocate it to an entrant, rather than buyer y^1. As the game progresses, buyer y^1 is more likely to be awarded the good eventually, reducing the option value of delay and decreasing the cutoff.

The one period look ahead property means that cutoffs can be characterized by a series of local indifference conditions. In 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,

$$m(x_T^k) = 0. \tag{4.6}$$

In period $t = T - 1$, the seller balances the revenue from allocating the k^{th} good, $m(x_{T-1}^k)$, against the opportunity cost derived from the possibility of denying the good to the k^{th} highest new entrant. Hence,

$$m(x_{T-1}^k) = \delta E_T \left[\max\{m(x_{T-1}^k), m(v_T^k)\} \right], \tag{4.7}$$

In periods $t \leq T-1$, 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 property. Hence,

$$m(x_t^k) + \delta E_{t+1} \left[\Pi_{t+1}^{k-1}(\mathbf{v_{t+1}}) \right] = \delta E_{t+1} \left[\max\{m(x_t^k), m(v_{t+1}^1)\} \right] + \delta E_{t+1} \left[\Pi_{t+1}^{k-1}(\{x_t^k, \mathbf{v_{t+1}}\}_k^2) \right] \tag{4.8}$$

In equation (4.8), we have set $\mathbf{y^{-1}} = 0$ because cutoffs are deterministic.[7]

5 Implementation

In this section we show that the optimal cutoffs can be implemented with posted prices as periods become short. Under a *price mechanism* the seller announces how many goods are remaining and charges a single price in each period. The buyers only reveal their existence when they purchase a unit. The entire price path is public information; if there is excess demand in a given period the units are rationed randomly.

[7] Since we know future cutoffs, the value functions in (4.8) can be calculated via the sequence problem (4.1) or the bellman equation (4.3).

Suppose time is continuous and, motivated by the law of rare events, buyers enter the market continuously according to a Poisson process with non-homogeneous arrival rate λ_t. Let r be the instantaneous discount rate and assume sales occur at discrete intervals of length h. In Section 5.1, we first consider general sequences of the demand process λ_t, showing the optimal allocations can be implemented by prices. In Section 5.2 we assume λ_t is weakly decreasing and show that the prices are given by an intuitive differential equation.

5.1 General Case

We say that the optimal allocations can be *implemented with prices* if the lost profit from using a price mechanism is of order $O(h)$.

Theorem 3. *Suppose that λ_t is Lipschitz continuous in t. Then the optimal allocations can be implemented with prices, with a second-price auction for the last unit at time T.*

Proof. See Appendix A.3 □

Theorem 3 is based on the fact that the cutoffs are deterministic, which means the seller does not have to elicit values $\mathbf{y^{-1}}$ in order to decide whether or not to allocate to buyer y^1. The proof consists of two parts. First, if λ_t is Lipschitz continuous in t, then the optimal allocations x_t^k cannot jump down more than $O(h)$, except for the last unit at time T. Second, by backward induction, we can pick prices to make the cutoff types indifferent. The prices imperfectly implement the cutoffs for two reasons: (i) the cutoffs cannot dynamically adjust within a given period; and (ii) when buyers are rationed, the good may be allocated to the wrong buyer. However, the probability of two sales within any given period is $O(h^2)$, so the lost profit is $O(h)$.

Prices are chosen to make the cutoff type indifferent between buying immediately and waiting. They therefore depend on the inventory and time remaining via the cutoff type. In addition, prices depend on the timing of past sales since this affects a buyer's belief about other buyers in the market, and hence his continuation utility. It is worth stressing that prices do not depend on the reports of the buyers (else it would not be a price mechanism), and they do not depend on the number of arrivals to the market.[8] In addition, it is notable that the seller publicly announces her inventory, so she does not gain from keeping k private.[9]

Theorem 3 assumes λ_t is Lipschitz continuous. If λ_t jumps down, then multiple sales may occur at one point in time, so one would need an auction to allocate efficiently. Saying this, one can approximate any single jump by quickly declining prices, analogous to a Dutch auction. That is, if the cutoffs x_t^k jump down then one can define a second sequence of cutoffs that are

[8]Prices would depend on N_t if buyers observed other buyers' entry into the market.

[9]The seller need not actually announce her inventory since the cutoff is a decreasing function of her inventory so the buyers could infer k from the price.

Lipschitz continuous and therefore can be implemented in prices. One can also do this for the last unit as $t \to T$, so the requirement of a "final auction" should be interpreted liberally.

The assumption of Poisson entry is more important since it implies that a buyer's entry time tells them nothing about the arrival rate of other buyers. As a result, all buyers share the same expectations over the evolution of future cutoffs. If this were not the case then a buyer's information about their entry time would give them information about other entrants' existence and even their values. For example, if buyers enter in pairs then knowing he entered earlier and no-sale had occurred implies that a buyer's "partner" has a lower valuation.

Finally, we assume the seller uses a second-price auction, but a first-price auction with reserve $m^{-1}(0)$ will also suffice. Since entry is Poisson, all buyers have the same information about others values deduced from observing the path of prices and there will be a symmetric equilibrium in increasing bidding strategies.

5.2 Weakly Decreasing Demand

When entering demand is decreasing over time, Theorem 2 says that cutoffs are decreasing and satisfy the one-period-look-ahead property. This allows us to heuristically derive the allocations and prices in the continuous time limit via local indifference conditions.[10]

First, consider optimal allocations. In period T, the optimal cutoffs are given by $m(x_T^k) = 0$. In period $t < T$, equation (4.8) becomes

$$rm(x_t^k) = \lambda_t E_v \Big[\max\{m(v) - m(x_t^k), 0\} + \Pi_t^{k-1}\big(\min\{v, x_t^k\}\big) - \Pi_t^{k-1}(v) \Big] \tag{5.1}$$

Equation (5.1) 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). Such delay leads to a higher marginal revenue tomorrow, if a new buyer enters, and a lower state variable in the continuation game. As $t \to T$, the cutoff jumps down discontinuously to $m^{-1}(0)$ if $k = 1$. However, if $k \geq 2$, then $\Pi_t^{k-1}(v) \to \max\{m(v), 0\}$, the right-hand side converges to zero and the cutoffs converge continuously, $x_t^k \to m^{-1}(0)$. Intuitively, in the last instant, there is an option value from the possibility of a single entrant arriving with a value higher than y^1; however, the probability of two or more entrants is zero.

Figure 3 illustrates the optimal cutoffs when the seller starts with two goods and buyers enter with constant hazard rate λ. When there is one unit remaining (the right panel), the cutoffs are constant in periods $t < T$ and drop down at time T (see Section 3). When there are two units remaining (the left panel), the option value of waiting falls over time since the seller needs two entrants to make it worthwhile to delay allocation. As a result, the cutoffs decrease

[10] See Ross (1971) for a formal treatment of the one unit case directly in continuous time.

over time.

The optimal cutoffs can be implemented by a sequence of decreasing prices p_t^k with an auction for the final good in period T. These prices can be derived backwards, starting at time T. When $k = 1$, the seller can use a second-price auction with reserve $m^{-1}(0)$ at time T. As $t \to T$, the price must be set so that the cutoff type x_{T-h}^1 is indifferent between taking the "buy it now" price and entering the auction at time T. This yields a price

$$p_T^1 = E_0 \left[\max\{y^2, m^{-1}(0)\} \middle| y^1 = \lim_{h \to 0} x_{T-h}^1, \{s_T(x)\}_{x \leq y^1} \right] \tag{5.2}$$

where $s_T(x)$ denotes the last time the cutoff went below x when looking back from time T. To understand this last term, note that buyer y^1 uses the sequence of past cutoffs to update about the presence of lower value buyers in the market at time T; since he only cares about the buyers remaining, a sufficient statistic is the last time the cutoff went below x. As a result, p_T^1 depends on when other buyers purchased units; in particular, the more time that has passed since those units were sold, the more competition buyer y^1 expects at time T, and the higher is p_T^1. When $k \geq 2$, the allocation converges to the static monopoly price, $x_t^k \to m^{-1}(0)$, as does the price, $p_t^k \to m^{-1}(0)$.

At time $t < T$, the cutoffs x_t^k are decreasing over time, so the prices are such that the cutoff type is indifferent between buying now and waiting a little. This becomes,

$$\frac{dp_t^k}{dt} = \left[\frac{dx_t^k}{dt} f(x_t^k) \int_{s_t(x_t^k)}^{t} \lambda_s ds - \lambda_t (1 - F(x_t^k)) \right] \left[x_t^k - p_t^k - U_t^{k-1}(x_t^k) \right] - r \left(x_t^k - p_t^k \right) \tag{5.3}$$

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 x_t^k being the highest-value buyer at time t. Intuitively, 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 good k is bought by either a new buyer with a value above x_{t+dt}^k, or an old buyer with value between x_t^k and x_{t+dt}^k (the right hand side). The possibility of a stockout 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.

Figure 3 illustrates the optimal prices for a seller with two goods. When there is one unit remaining (the right panel), the prices fall even though the cutoff stays constant. Intuitively, when the buyer 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. While cutoffs are deterministic, only depending on the number of units and time remaining, prices also depend on when the penultimate unit was sold. We illustrate this with three price lines. Intuitively, if the penultimate unit is sold early on, then buyers think there may be many other buyers in the market waiting for the price to drop,

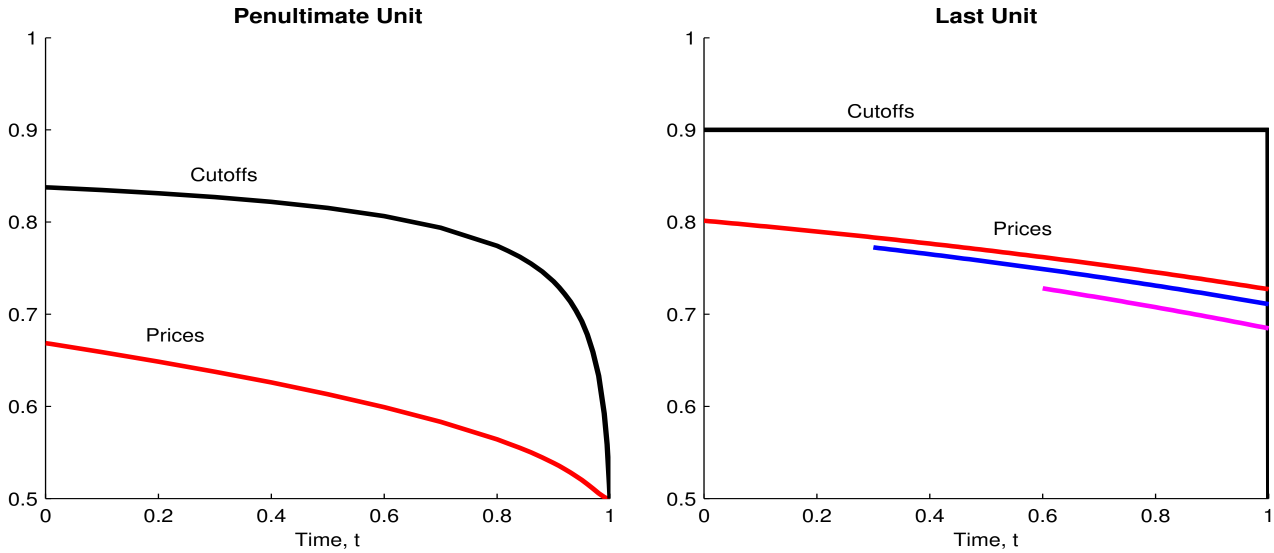

Figure 3: **Optimal Cutoffs and Prices with Two Units.** The **left** panel shows the optimal cutoffs/prices when the seller has two units remaining. The **right** panel shows the optimal cutoffs/prices when the the seller has one unit remaining. The three price lines illustrate the seller's strategy when she sells the penultimate unit at times $t = 0$, $t = 0.3$ and $t = 0.6$. In this figure, buyers 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 $r = 1/16$.

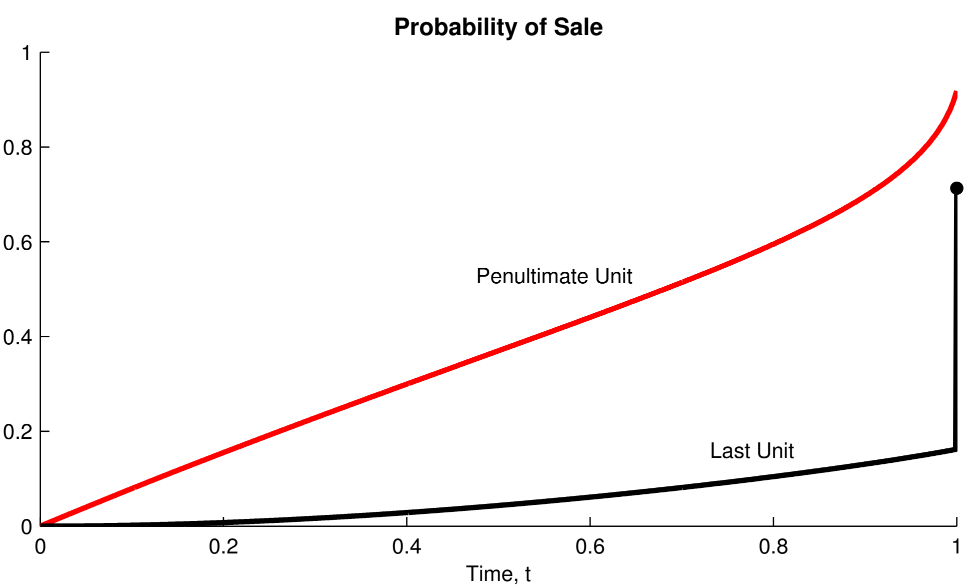

Figure 4: **Probability of Sale with Two Units** This figure shows the unconditional probability the last/penultimate unit is sold by time t. The parameters are the same as Figure 3.

meaning the seller can charge a higher price to implement the same cutoff. When $k = 2$, the prices fall over time reflecting the declining cutoffs, buyers' impatience of getting the good early, and buyers' concern of another buyer poaching the good.

Figure 4 illustrates the unconditional probability of both units being sold as a function of time. With both units, the probability of the sale increases rapidly as $t \to T$. When $k = 1$, there is an atom at time T; when $k = 2$, the probability rises rapidly and resembles a fire-sale (although there is no atom at T). Intuitively, the existence of the fire-sale comes from the concavity of the path of cutoffs (see Figure 3) and the stock of buyers building up, waiting to buy. This 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 (2012) 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.

While we focus on continuous time, one can also implement the optimal cutoffs in discrete time. With a single good, $K = 1$, this can be done via a sequence of second-price auctions (Board and Skrzypacz (2010)). With more goods, Li (2011) shows the seller can use a sequence of ascending auctions in which buyers compete against a robot who acts like the cutoff type. The basic problem in the discrete time game is that more is known about older buyers' values, implying a new and old buyer with the same valuation calculate continuation utilities differently and therefore bid differently. To overcome this, Li follows Said (2012) in using an ascending auction; this reveals all buyers' values each period, allowing buyers to use memoryless strategies.

5.3 Myopic vs. Forward-Looking Buyers

Typical revenue management models assume that buyers are myopic, leaving the market if they do not buy immediately (e.g. Gallego and van Ryzin (1994)). In this case, the state variable is time and the number of units remaining (k, t), so the cutoffs are automatically deterministic. If V_t^k is the continuation value, then the optimal cutoffs are given by $m(x_t^k) = \delta(V_{t+1}^k - V_{t+1}^{k-1})$. These optimal allocations can be implemented with auctions in discrete time, or prices in continuous time, with the (reserve) price being set equal to the corresponding cutoff. In contrast, with forward-looking buyers, cutoffs are deterministic while prices depend on the timing of past sales.

Figure 5 illustrates the optimal cutoffs/prices and the probability of sale when buyers are myopic under the same parameters as Figures 3–4. A first observation is that profits are higher when buyers are forward looking.[11] This initially might seem surprising since forward-looking buyers can time their purchase to lower their payments. For example, fixing the retail

[11] Proof: Since the arrival time is observable, the seller could replicate the myopic allocation. Yet Lemma 1 shows that it is optimal to treat all generations of buyers equally.

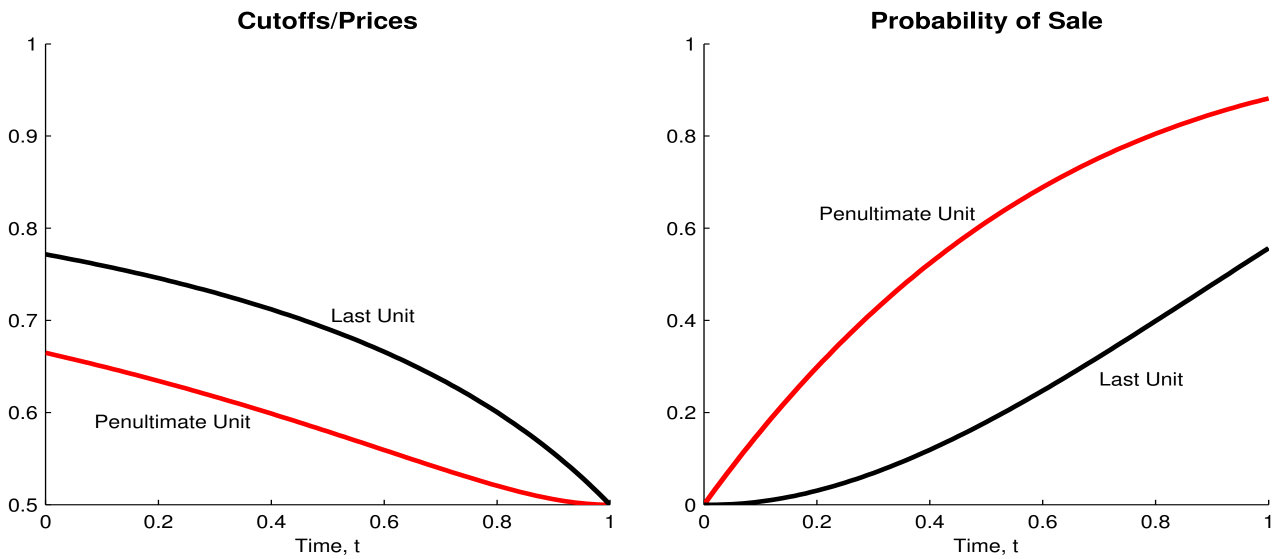

Figure 5: **Myopic Buyers with Two Units.** The **left** panel shows the optimal cutoffs/prices when the seller has two units remaining and buyers are myopic. The **right** panel shows unconditional probability of sale of one/both units. The parameters are the same as Figure 3.

prices, Soysal and Krishnamurthi (2012) found that profits for women's coats are 9% higher when their forward-looking customers are eliminated. However, when the seller is choosing the optimal mechanism, the ability to delay means the seller can pool different cohorts of buyers together, raising the efficiency of allocation and revenue.

Second, the total number of sales is higher when buyers are forward looking. With forward looking customers, the seller sells k goods if there are at least k entrants with values above $m^{-1}(0)$. With myopic customers, the seller might refuse to sell to a buyer with value above $m^{-1}(0)$ early in the game, and be unable to return to them later.

Third, sales occur later when buyers are forward looking. When buyers are myopic, sales occur fairly evenly throughout $[0, T]$, as seen in Figure 5. When buyers are forward looking, the combination of concave cutoffs (Figure 3 vs. 5) and waiting buyers produces a fire-sale, as illustrated by Figure 4. Indeed, while total sales are higher with forward-looking buyers, sales in the first period are higher with myopic buyers.[12] This is easily seen in the limit as $\delta \to 1$, where a seller facing forward-looking buyers simply holds an auction at time T.

[12] Proof: The forward-looking cutoffs exceed the myopic cutoffs since delaying sale has a higher option value. In the first period, there is no backlog of buyers in either case, so the probability of sale is higher under myopia.

6 Applications

In this section we consider a number of specific applications. This has the dual purpose of demonstrating the usefulness of the model in a wide variety of positive and normative contexts. We also use the applications to explore the robustness of the model to changes in the underlying assumptions. This includes the introduction of inventory costs (Section 6.1), arriving and expiring supply (Section 6.2), third-degree price discrimination (Section 6.3), changing distributions of incoming values (Section 6.4), and disappearing buyers (Section 6.5).

6.1 Retail Markets

Retail markets are a frequent application of revenue management models. A retailer (e.g. Zara, H&M) orders a new design that essentially has a fixed supply because of production lags. They then prefer to sell these units sooner rather than later because they take up valuable shelf space in the store. In order to model this form of impatience, suppose buyers do not discount over the shopping season, $\delta = 1$, but the seller pays a per-unit inventory costs, c_t, so a firm selling in period t gets profit $p_t - c_t$, where c_t increases in t.[13] We can adapt (2.5) to obtain the firm's profits,

$$\text{Profit} = E_0 \left[\sum_i [m(v_i) - c_{\tau_i}] \mathbf{1}_{\tau_i \le T} - \left(K - \sum_i \mathbf{1}_{\tau_i \le T} \right) c_{T+1} \right]$$

Much of the previous analysis carries over to this new setting. The optimal cutoffs are deterministic (Theorem 1). If the number of arriving buyers N_t falls over time and costs c_t are convex in t, then the cutoffs decline and the one-period-look-ahead property holds (Theorem 2). In the continuous time limit, if the arrival rate λ_t and marginal cost $\Delta c_t := c_{t+1} - c_t$ are Lipschitz continuous then the optimal cutoffs can be implemented with prices (Theorem 3).[14]

To see the effect of inventory costs, suppose entry is decreasing and costs c_t are convex, so the marginal cost Δc_t weakly increases in t. Adapting (4.8), the one-period-look-ahead property implies that cutoffs are determined by

$$m(x_t^k) + E_{t+1}\left[\Pi_{t+1}^{k-1}(\mathbf{v_{t+1}})\right] = E_{t+1}\left[\max\{m(x_t^k), m(v_{t+1}^1)\}\right] + E_{t+1}\left[\Pi_{t+1}^{k-1}(\{x_t^k, \mathbf{v_{t+1}}\}_k^2)\right] - \Delta c_t$$

for $t < T$, with $m(x_T^k) = -\Delta c_T$. The resulting cutoffs are decreasing over time because the

[13] We eliminate discounting for simplicity; the nature of the results are unchanged if one has inventory costs and discounting.

[14] To prove Theorem 1 one must change the difference formula $\Delta\Pi_t^k$ to account for the cost of delay. One should also interpret discounted stopping time δ^τ in Lemma 3(b),(c) as $\mathbf{1}_{\tau \le T}$; this is still strictly less than one because of the possibility of stocking out. For Theorem 2, the first step of Lemma 4 should be changed so that $\hat{D}\Pi_{t+1}^k$ assumes that there are N_{t+1} entrants (rather than N_{t+2}) and the delay cost is Δc_t (rather than Δc_{t+1}). If the number of entrants is weakly decreasing and cutoffs are convex then $D\Pi_{t+1}^k \ge \hat{D}\Pi_{t+1}^k$. For Theorem 3, the $D\Pi_t^k$ terms have to be adjusted to account for the changing marginal costs, but this new term is also Lipschitz continuous by assumption.

option value of delay falls, while the cost of delay Δc_t rises. In the continuous time limit, assuming c_t is differentiable, we can adapt (5.1) to obtain

$$\frac{dc_t}{dt} = \lambda_t E\Big[\max\{m(v) - m(x_t^k), 0\} + \Pi_t^{k-1}\big(\min\{v, x_t^k\}\big) - \Pi_t^{k-1}(v)\Big].$$

Adapting (5.3), prices are then determined by the differential equation

$$\frac{dp_t^k}{dt} = \left[\frac{dx_t^k}{dt} f(x_t^k) \int_{s_t(x_t^k)}^{t} \lambda_s ds - \lambda_t(1 - F(x_t^k))\right] \left[x_t^k - p_t^k - U_t^{k-1}(x_t^k)\right]$$

with boundary condition (5.2). Note that although buyers do not discount, they are still impatient because delay may lead the seller to sell the good to another buyer and stock out.

Our model can also be used as a positive theory of pricing. The model predicts that (i) price reductions lead to large numbers of sales as the waiting customers purchase, (ii) this burst of sales quickly dies down, (iii) that prices fall rapidly near the end of the season, with a large fraction of sales occurring during the second half of the sales season, and (iv) prices start off steady and become more volatile throughout the season. All of these predictions are consistent with Soysal and Krishnamurthi (2012), as illustrated by Figure 1, while (i)-(iii) are hard to rationalize using a model with myopic buyers. The one place that our model runs into problems is the prediction that prices should rise after a unit is sold, which seems relatively rare for retail goods (although perhaps happens with car dealers). One possibility is that, with a fashion good, the buyers' discount rate for the good is much higher than the discount rate of money implying that our pictures understate the optimal price reductions (see footnote 1). Alternatively, the falling prices may come from commitment problems (Bulow (1982)) or learning about demand (Lazear (1986)).

6.2 Dealers in the Fulton Fish Market

A second application for revenue management is to study how dealers in the Fulton Fish Market manage their inventory over time. This is analogous to the problem of an intermediary in a platform market (e.g. Amazon Marketplace, eBay) who buy in bulk and try to subsequently sell these units online. The Fulton Fish Market has been studied extensively by Graddy (1995, 2006), with Graddy and Hall (2011) formally estimating a revenue management model with myopic customers. The market is open Monday-Friday, so the dealer wishes to sell their fish by Friday, or earlier if the fish go bad. There is a small number of dealers who sell a particular type of fish; due to repeated interaction between the sellers, Graddy argues that they act fairly monopolistically. Buyers come from restaurants and wholesalers and are typically experienced; they are also impatient because of the opportunity cost of their time.

Relative to the baseline model, a major difference is the arrival and departure of units over time.[15] During the week new stock arrives; this is largely determined by weather conditions and can be taken as an exogenous process. In addition, fish expire after a few days, so there is an exogenous death process. We model arrivals by supposing that there is an arrival process $(a_1, \ldots, a_T)$ such that a_t units have arrived in the market by date t. Similarly, there is an exogenous death process $(b_1, \ldots, b_T)$ such that b_t goods must be sold by date t, else they disappear. Finally let ζ_t be the number of goods disposed by time t. The seller's problem is then to maximize profit (2.5) subject to the constraint that the number of sales plus disposals satisfy

$$a_t \geq \sum_i \mathbf{1}_{\tau_i \leq t} + \zeta_t \geq b_t.$$

One can then view the baseline model as a special case where there are K goods at time $t = 1$ that expire at time T. If we let $K = a_T$ be the total number of units available, and k be the number that have yet to be sold/destroyed, Lemma 1 implies that we can characterize the optimal allocations by a sequence of cutoffs $\{x_t^k\}_{k \in \{1,\ldots,K\}}$, where the seller must sell/destroy between a_t and b_t units by time t.

Much of the previous analysis carries over to this new setting. The optimal cutoffs are deterministic (Theorem 1). The cutoffs x_t^k fall to $m^{-1}(0)$ in the period when unit k expires. More precisely, the cutoff for the last unit to expire in a given period jumps to $m^{-1}(0)$, while the cutoffs for previous units that expire in the same period continuously converge to $m^{-1}(0)$. Intuitively, if a single unit expires there may be an entrant at the last moment with a value higher than y^1; the probability of two or more entrants is zero. Prior to expiring, if N_t is decreasing over time and a_t, b_t are deterministic, then cutoffs fall over time and the one-period-look-ahead property holds (Theorem 2). This result fails, however, if entry or departure is stochastic; for example, if the market expects a large delivery of fish but few turn up, then the cutoff will rise. Turning to prices, the cutoffs can jump down for the last unit to expire in each period, so one then needs an auction. Although, as discussed in Section 5.1, one can approximate these auctions with posted prices analogous to a Dutch Auction.[16]

6.3 Display Ad Sales

A third application of the model is to the sale of display ads by companies such as the New York Times and Yahoo. These firms have a limited number of ad slots (e.g. the front page banner on Yahoo) that they need to sell prior to the broadcast date. The buyers are often experienced

[15] This is also an issue in airlines or display ads, where the seller sells a sequence of goods with different broadcast/flight times.

[16] Theorem 1 derives from the fact that the decision to sell to y^1 does not affect when y^2 gets a unit, so the new constraints on the supply side have no impact. For Theorem 2, they key step is to observe that $\tau_{t+1}^{k-1}(z) - (t+1) \geq \tau_{t+2}^{k-1}(z) - (t+2)$ in Lemma 4, meaning the option value of delay falls over time, along with the cutoff. For Theorem 3, it remains true that cutoffs don't jump down (much) except for the last good expiring in any period.

and forward looking, as illustrated by the existence of "conditional contracts" (see Section 3). Buyers are also impatient because they wish to coordinate advertising campaigns and because paying attention to the market is costly.

Relative to the benchmark model, a major consideration for such firms is how to divide ads between rich media (e.g. video, flash) and static display ads. The buyers of rich media (e.g. movie studios) are very different from those buying static ads (e.g. insurance companies) so this is a case of third degree price discrimination. To model this, suppose buyers are drawn from two distributions f_R and f_S, inducing marginal revenues m_R and m_S, and the seller knows which distribution a given buyer comes from. The seller has K units that they can allocate to either type of buyer. The seller's problem is thus to maximize profit

$$\text{Profit} = E_0 \left[\sum_i \delta^{\tau_i} m_i(v_i) \right],$$

subject to the constraint $\sum_i \mathbf{1}_{\tau_i \leq T} \leq K$, where $m_i \in \{m_R, m_S\}$.

The major difference relative to the benchmark model is that the ranking of buyers' values no longer corresponds with the ranking of the marginal revenues. If f_R hazard-rate dominates f_S, then a static buyer with the same value as a rich-media buyer will have a higher marginal revenue, and the seller will bias allocation in favor of the static buyer.

To solve the problem, the seller should now treat the k highest marginal revenues $\{m^1, \ldots, m^k\}$ as the state variable (rather than the underlying values). The optimal cutoffs (in marginal revenue space) are deterministic (Theorem 1). If the number of both types of entrant are weakly decreasing over time then the order statistics of the entrants marginal values fall over time and the one-period-look-ahead property holds (Theorem 2). To implement the optimal cutoffs, the seller needs two different price paths for the two types of buyer (Theorem 3). These are related through the inventory, so a sale of a static ad raises the prices for both types of buyer.

6.4 Airlines

One of the most natural applications for revenue management is for airlines. Robert Crandall, former-CEO of American Airlines, 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." (Smith, Leimkuhler, and Darrow (1992)). In the airline model, the seller has a fixed number of seats which must be sold prior to the flight date. A typical buyer is happy to wait to buy if they think the price will be reduced, but is impatient because they must regularly check prices and may wish to make complementary plans.

One feature of the data that seems inconsistent with the model is that prices often tend to rise in the last few days before a flight (McAfee and Te Velde (2006)). A natural reason for this is that buyers who enter nearer date T tend to have higher values. If we suppose buyers who enter in period t have distribution F_t, each generation is associated with a different marginal revenue and equation (2.5) can be adjusted to yield

$$\text{Profit} = E_0\left[\sum_i \delta^{\tau_i} m_{t_i}(v_i)\right]$$

As in Section 6.3, the seller would thus like to bias allocation towards generations with higher marginal revenues for a given value, which broadly corresponds to generations with weaker distributions (in the hazard rate order). As above, if the seller can discriminate between different generations then she can use the highest marginal revenues as state variables and implement these with cohort-specific price paths $p_t^k(t_i)$. This may take a relatively simple form: for example, if the distributions are exponential $F_t(v) = 1 - e^{-v/\mu_t}$ then $m_t(v) = v - \mu_t$ and the seller can use nondiscriminatory posted prices p_t^k with a cohort-specific fee of μ_{t_i}. Even if the seller cannot discriminate between different cohorts, then this mechanism is still incentive compatible if demand gets stronger over time, as with airlines. In this case, the seller would like to bias allocation towards earlier generations, and the inter-generational (IC) constraints will be slack since a generation t buyer would not wish to pretend to be born in $t+1$ (and cannot pretend to be born in $t-1$). For example, in the above exponential example, the seller could issue a coupon worth μ_{t_i} to a buyer born in period t_i.[17]

In many applications, it is natural to consider a nondiscriminatory price scheme, p_t^k, which would lead buyers in the market with values exceeding some cutoff x_t^k to buy at time t, independent of their birth-date. Since buyers from different generations will merge, one must consider the average marginal revenue from selling to a particular type (see Board (2008) for a related construction). While we will not solve this problem, it is interesting to note that the introduction of forward-looking buyers may now lower the firm's profits. When buyers become forward looking, the seller gains from the option value to delay selling a unit, but may lose if inter-generational price discrimination becomes harder. If demand gets stronger over time, profits are still higher with forward-looking buyers because a delaying buyer has a higher marginal revenue than a younger buyer.[18] However, if demand grows weaker over time, then profits may be lower with forward-looking customers as stronger generations delay to merge with weaker generations.

[17] If demand weakens over time then the intertemporal (IC) constraints will bind. This is an interesting problem but is beyond the scope of this paper.

[18] Proof: Fix some optimal cutoffs with myopic buyers x_t^k. Now, suppose a buyer v who enters in period $s < t$ gives rise to a doppelgänger in period t with marginal revenue $m_t(v)$. Since we've just increased the number of buyers in each period, this raises the seller's expected profit. Next, suppose this buyer contributed marginal revenue $m_s(v) > m_t(v)$ if they buy in period t. This second step increases the seller's profits in every state; it also corresponds to the profit with forward-looking buyers and cutoffs x_t^k.

6.5 House Selling

For the final application, consider the problem of an owner wishing to sell their house, with buyers entering over time. If $T = \infty$ then the seller should set a single posted price forever, and buyers either buy immediately or never (Gallien (2006)). However, if the seller must sell the house by date T (e.g. they wish to move to another city) or if the number of entering buyers falls over time, then optimal prices decline and buyers delay on the equilibrium path.

Relative to the benchmark model, it would seem natural to allow buyers to exit probabilistically over time (perhaps because they buy another house). However, this extension considerably complicates the analysis. First, it means that the seller must keep track of all remaining buyers, rather than just the k highest, since any buyer may disappear at any time. Second, it means that Theorem 1 fails and optimal cutoffs are no longer deterministic. To understand why suppose there are two buyers with values $v_H > v_L$ and one good. The seller's decision to award the good to buyer v_H will depend on the level of v_L because if the seller delays, buyer v_H may disappear, forcing the seller to award the good to v_L. It immediately follows that posted prices are not optimal: The seller would like to elicit the value v_L before deciding whether or not to award the good to v_H.

The general problem with this example is that the seller's ranking of buyers can change over time. In the above example she initially prefers v_H to v_L, but may prefer v_L in period 2 if v_H disappears (since disappearing is isomorphic to having one's value jump to zero). This problem is also seen if different types of buyers have different discount rates, $\delta \in \{\delta_L, \delta_H\}$, since the seller's ranking of a more patient buyer can rise above that of a less patient buyer over time. Hence the optimal mechanism is again not deterministic: the seller should elicit the value of patient buyers before awarding a unit to an impatient buyer.

When thinking about the housing application, this analysis helps explain the use of auction mechanisms in real estate pricing. If the house seller thinks that a buyer may disappear at any stage then the optimal mechanism will have buyers first submit indicative bids before the seller makes a counteroffer to the highest bidder.

7 Conclusion

We have considered a seller who wishes to sell K goods within T periods to buyers who enter the market over time and are forward looking. The optimal mechanism consists of a sequence of cutoffs that are deterministic and, in continuous time, can be implemented with posted prices. If the number of entrants decreases over time, the cutoffs are also decreasing and satisfy the one-period-look-ahead property.

This paper provides a benchmark for the analysis of revenue management with forward-looking buyers, but specific applications raise a number of issues that are not covered by our analysis. First, one would like to allow the seller to learn about the demand curve implying the N_t variables are correlated over time. In this case, cutoffs are still deterministic, but they will depend on the number of past entrants (although not their values) since they are indicative of future entry. Having a buyer report their arrival would be incentive compatible if it lowers future cutoffs, but not if it raises the cutoffs.[19] Second, while we have modeled impatience in a reduced-form manner (e.g. discounting, inventory costs), it would be interesting to model "attention costs" or "coordination costs" in a more sophisticated way. Finally, since ad slots on a website differ by position and size, one would like to allow for different qualities of goods.

[19] See Gershkov, Moldovanu, and Strack (2013) for a related model.

A Appendix: Omitted Proofs

A.1 Proof of Lemma 3

Proof. Part (a). We wish to prove that

$$\Delta\Pi_t^k(\mathbf{y}) = m(y^1) + \delta\tilde{\Pi}_{t+1}^{k-1}(\mathbf{y^{-1}}) - \delta\tilde{\Pi}_{t+1}^k(\mathbf{y})$$

is independent of $\mathbf{y^{-1}}$. Consider buyer y^j at time t for $j \geq 2$, and let $r_s(j)$ denote his rank in the distribution of buyers, including $\mathbf{y}$ and all subsequent entrants, at time $s > t$. Since future cutoffs are deterministic and do not depend on the seller's choice to sell at time t, Lemma 2 implies that, whether or not the seller sells at time t, buyer y^j is allocated a good at the first time τ such that $y^j \geq x_\tau^{k-r_\tau(j)+1}$. Since the allocation of y^j is independent of the decision of whether or not to sell a unit at time t, so y^j makes the same contribution to profits (2.5) in both cases, and $\Delta\Pi_t^k(\mathbf{y})$ is independent of y^j.

Part (b). Continuity follows from the envelope theorem. Using equation (4.5),

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

Using equation (4.4) and the envelope theorem,

$$\frac{d}{dy^1}\Pi_t^k(\text{sell 0 today}) = m'(y^1)E_{t+1}[\delta^{\tau_1^k(y^1)-t}]$$

where $\tau_1^k(y^1)$ 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 $\tau_1^k(y^1) > t$ and $\delta < 1$.

Part (c). Suppose $\{x_s^k\}_{s\geq t+1}$ are deterministic and decreasing in k. We first prove a preliminary result. Let $\mathbf{y} = \{y^1, \ldots, y^k\}$ and $\tilde{\mathbf{y}} = \{\tilde{y}^1, \ldots, \tilde{y}^k\}$ be arbitrary vectors, where $y^j \geq \tilde{y}^j$ for each j. We claim that for time $\sigma \geq t+1$,

$$\begin{aligned}
\Pi_\sigma^k(\mathbf{y}) - \Pi_\sigma^k(\tilde{\mathbf{y}}) &= E_{\sigma+1}\left[\delta^{-\sigma}\int_{\{\tilde{y}^1,\ldots,\tilde{y}^k\}}^{\{y^1,\ldots,y^k\}}(m'(z^1)\delta^{\tau_1^k(z^1)},\ldots,m'(z^k)\delta^{\tau_k^k(z^k)})d(z^1,\ldots,z^k)\right] \\
&\geq E_{\sigma+1}\left[\delta^{-\sigma}\int_{\{\tilde{y}^1,\ldots,\tilde{y}^k\}}^{\{y^1,\ldots,y^k\}}(m'(z^1)\delta^{\tau_1^{k-1}(z^1)},\ldots,m'(z^k)\delta^{\tau_k^{k-1}(z^k)})d(z^1,\ldots,z^k)\right] \\
&= \Pi_\sigma^{k-1}(\mathbf{y}) - \Pi_\sigma^{k-1}(\tilde{\mathbf{y}}) \qquad \text{(A.1)}
\end{aligned}$$

The first line applies the envelope theorem to equation (4.2), where τ_j^k is the purchasing time of the buyer in the j^{th} position at time σ when there are k objects for sale. The second line follows from the fact that $\tau_j^k(z^j) \leq \tau_j^{k-1}(z^j)$ since $\{x_s^k\}_{s\geq\sigma+1}$ are decreasing in k. Note that

$\tau_k^{k-1} = \infty$ since a seller with $k-1$ goods cannot allocate a k^{th} good. The final line again uses the envelope theorem.

Suppose the seller has k units at time t. In periods $s \geq t$, the seller follows the optimal strategy as dictated by the deterministic, decreasing cutoffs $\{x_s^k\}_{s \geq t+1}$. By part (a), $\Delta\Pi_t^k(y^1)$ is independent of the other buyers, so we can set $\mathbf{y^{-1}} = \varnothing$.

Letting $v_{(t,s]}^j$ be the j^{th} highest value of a buyer who has entered over $\{t+1, \ldots, s\}$, define $\sigma = \min\{s \geq t+1 : \max\{y^1, v_{(t,s]}^1\} \geq x_s^{k-1}\}$ as the (random) time the seller with k units at time $t+1$ next makes a sale. Define $\mathbf{v_{(t,\sigma]}} := \{v_{(t,\sigma]}^1, \ldots, v_{(t,\sigma]}^k\}$ and let $\{y^1, \mathbf{v_{(t,\sigma]}}\}_k^2$ be the ordered vector of the 2^{nd} to k^{th} highest choices from $\{y^1, \mathbf{v_{(t,\sigma]}}\}$. We claim that

$$
\begin{aligned}
\Delta\Pi_t^k(y^1) &= m(y^1) + \delta\tilde{\Pi}_{t+1}^{k-1}(\varnothing) - \delta\tilde{\Pi}_{t+1}^k(y^1) \\
&= m(y^1) + E_{t+1}\left[\delta^{\sigma-t}\left[\Pi_\sigma^{k-1}(\{\mathbf{v_{(t,\sigma]}}\}_{k-1}^1) - \max\{m(y^1), m(v_{(t,\sigma]}^1)\} - \Pi_\sigma^{k-1}(\{y^1, \mathbf{v_{(t,\sigma]}}\}_k^2)\right]\right] \\
&\geq m(y^1) + E_{t+1}\left[\delta^{\sigma-t}\left[\Pi_\sigma^{k-2}(\{\mathbf{v_{(t,\sigma]}}\}_{k-1}^1) - \max\{m(y^1), m(v_{(t,\sigma]}^1)\} - \Pi_\sigma^{k-2}(\{y^1, \mathbf{v_{(t,\sigma]}}\}_k^2)\right]\right] \\
&\geq m(y^1) + \delta\tilde{\Pi}_{t+1}^{k-2}(\varnothing) - \delta\tilde{\Pi}_{t+1}^{k-1}(y^1) \\
&= \Delta\Pi_t^{k-1}(y^1)
\end{aligned}
$$

The first line is the definition of $\Delta\Pi_t^k(y^1)$. The second line uses the fact that a seller with k units makes a sale weakly before a seller with $k-1$ units since future cutoffs are decreasing in k. The third line comes from (A.1) and the fact that $\{y^1, \mathbf{v_{(t,\sigma]}}\}_{k-1}^1$ is pointwise larger than $\{y^1, \mathbf{v_{(t,\sigma]}}\}_k^2$. The fourth line uses the fact that a seller with $k-1$ goods stops at a weakly later time than a seller with k units, so $\delta^{t+1}\tilde{\Pi}_{t+1}^{k-2}(\varnothing) = E_{t+1}[\delta^\sigma \Pi_\sigma^{k-2}(\{\mathbf{v_{(t,\sigma]}}\}_{k-1}^1)]$, and $\delta^{t+1}\tilde{\Pi}_{t+1}^{k-1}(y^1) \geq E_{t+1}[\max\{m(y^1), m(v_{(t,\sigma]}^1)\} + \Pi_\sigma^{k-2}(\{y^1, \mathbf{v_{(t,\sigma]}}\}_k^2)]$. □

A.2 Proof of Lemma 4

The proof is in two steps. First, we wish to nullify the effect of the decreasing demand so we can compare like-with-like. Writing out the value of selling immediately, we have

$$
D\Pi_{t+1}^k(y^1) = m(y^1) + \delta E_{t+2}\left[\Pi_{t+2}^{k-1}(\{\mathbf{v_{t+2}}\}_{k-1}^1)\right] - \delta E_{t+2}\left[\max\{m(y^1), m(v_{t+2}^1)\} + \Pi_{t+2}^{k-1}(\{y^1, \mathbf{v_{t+2}}\}_k^2)\right],
$$

where we use the analogue of Lemma 3(a) to ignore $\mathbf{y^{-1}}$. We now show that the option value of waiting is higher if the entrants have higher values. If we use the envelope theorem to differentiate

$$
\Pi_{t+2}^{k-1}(\{\mathbf{v_{t+2}}\}_{k-1}^1) - \max\{m(y^1), m(v_{t+2}^1)\} - \Pi_{t+2}^{k-1}(\{y^1, \mathbf{v_{t+2}}\}_k^2) \qquad \text{(A.2)}
$$

with respect to v_{t+2}^j, we obtain

$$m'(v_{t+2}^j)[\delta^{\hat{\tau}_j^1(v_{t+2}^j)} - \delta^{\hat{\tau}_j^0(v_{t+2}^j)}]\delta^{-(t+2)}, \tag{A.3}$$

where $\hat{\tau}_j^1$ is the purchasing time of v_{t+2}^j under "sell 1 today", and $\hat{\tau}_j^0$ is the purchasing time under "sell 0 today and ≥ 1 tomorrow". In the former case v_{t+2}^j has rank j at time $t+2$; in the latter case v_{t+2}^j may have rank j or $j-1$. Given future cutoffs are deterministic, $\hat{\tau}_j^0(v_{t+2}^j) \leq \hat{\tau}_j^1(v_{t+2}^j)$ and (A.3) is negative. Hence (A.2) is decreasing in $\mathbf{v_{t+2}}$.

Now, let $\hat{\mathbf{v}}_{\mathbf{t+2}}$ be order statistics at time $t+2$ drawn from the same distribution as N_{t+1}. Replacing $\mathbf{v_{t+2}}$ with $\hat{\mathbf{v}}_{\mathbf{t+2}}$ in $D\Pi_{t+1}^k(y^1)$, define

$$\hat{D}\Pi_{t+1}^k(y^1) = m(y^1) + \delta E_{t+2}\left[\Pi_{t+2}^{k-1}(\{\hat{\mathbf{v}}_{\mathbf{t+2}}\}_{k-1}^1)\right] - \delta E_{t+2}\left[\max\{m(y^1), m(\hat{v}_{t+2}^1)\} + \Pi_{t+2}^{k-1}(\{y^1, \hat{\mathbf{v}}_{\mathbf{t+2}}\}_k^2)\right],$$

Since N_t is decreasing in the usual stochastic order, $\hat{\mathbf{v}}_{\mathbf{t+2}}$ exceeds in the $\mathbf{v_{t+2}}$ usual stochastic order and, since (A.2) is decreasing in $\mathbf{v_{t+2}}$, $D\Pi_{t+1}^k(y^1) \geq \hat{D}\Pi_{t+1}^k(y^1)$. Intuitively, the seller has more to gain from selling today if there are fewer entrants tomorrow.

For the second step, we prove that $\hat{D}\Pi_{t+1}^k(y^1) \geq D\Pi_t^k(y^1)$. To do this, we can write the Π_{t+1}^{k-1} terms in $D\Pi_t^k(y^1)$ in terms of a single variable and then apply the envelope theorem to obtain

$$\begin{aligned}\Pi_{t+1}^{k-1}(\{\mathbf{v_{t+1}}\}_{k-1}^1) - \Pi_{t+1}^{k-1}(\{y^1, \mathbf{v_{t+1}}\}_k^2) &= \Pi_{t+1}^{k-1}(\{v_{t+1}^1, \mathbf{v_{t+1}^{-1}}\}_{k-1}^1) - \Pi_{t+1}^{k-1}(\{\max\{v_{t+1}^k, \min\{y^1, v_{t+1}^1\}\}, \mathbf{v_{t+1}^{-1}}\}_{k-1}^1) \\ &= E_{t+2}\left[\int_{\max\{v_{t+1}^k, \min\{y^1, v_{t+1}^1\}\}}^{v_{t+1}^1} m'(z)\delta^{\tau_{t+1}^{k-1}(z)-(t+1)}\,dz\right],\end{aligned}$$

where $\tau_{t+1}^{k-1}(z)$ is the time the object is allocated to type z looking forward from time $t+1$, holding $\mathbf{v_{t+1}^{-1}}$ constant. The same term in $\hat{D}\Pi_{t+1}^k(y^1)$ is defined the same way, but advanced one period. That is,

$$\Pi_{t+2}^{k-1}(\{\hat{\mathbf{v}}_{\mathbf{t+2}}\}_{k-1}^1) - \Pi_{t+1}^{k-1}(\{y^1, \hat{\mathbf{v}}_{\mathbf{t+2}}\}_k^2) = E_{t+3}\left[\int_{\max\{\hat{v}_{t+2}^k, \min\{y^1, \hat{v}_{t+2}^1\}\}}^{\hat{v}_{t+2}^1} m'(z)\delta^{\tau_{t+2}^{k-1}(z)-(t+2)}\,dz\right].$$

Recall buyer z buys a unit at time s if he has the highest value, and his value is above the corresponding cutoff. Since $\hat{\mathbf{v}}_{\mathbf{t+2}}$ and $\mathbf{v_{t+1}}$ have the same distribution we can suppose $\hat{\mathbf{v}}_{\mathbf{t+2}} = \mathbf{v_{t+1}}$. If $\tau_{t+1}^{k-1}(z) = s$ for $s < T$ then $\tau_{t+2}^{k-1}(z) \leq s+1$ since future cutoffs decrease in t and N_t falls over time.[20] In addition, if $\tau_{t+1}^{k-1}(z) = T$ then $\tau_{t+2}^{k-1}(z) \leq T$ since more entrants enter over time. Putting this together, $\tau_{t+1}^{k-1}(z) - (t+1) \geq \tau_{t+2}^{k-1}(z) - (t+2)$ for all z. Taking expectations over the distribution of entrants, the integral equations then imply that $\hat{D}\Pi_{t+1}^k(y^1) \geq D\Pi_t^k(y^1)$.

[20] If $N_s \geq N_{s+1}$ in the usual stochastic order, then there exists a state space Ω such at $N_s(\omega) \geq N_{s+1}(\omega)$ almost surely. We are implicitly adopting this state space to conclude the stopping time is ranked almost surely.

Combining both parts of the proof, we thus have $D\Pi^k_{t+1}(y^1) \geq \hat{D}\Pi^k_{t+1}(y^1) \geq D\Pi^k_t(y^1)$ as required.

A.3 Proof of Theorem 3

Claim A: For $t \leq T - 2h$, there exists positive constants α, h_0 such that $x^k_t - x^k_{t+h} \leq \alpha h$ for $h < h_0$.

Proof: Let $\Lambda_{t+h} = \int_t^{t+h} \lambda_s ds$ be arrival rate over $(t, t+h]$, and let N_{t+h} be the realized number of arrivals in period $t+h$. We then have

$$\Pi^k_t(\text{sell 1 today}) = m(y^1) + e^{-rh}e^{-\Lambda_{t+h}}\Pi^{k-1}_{t+h}(\varnothing) + e^{-rh}(1-e^{-\Lambda_{t+h}})E_{t+h|N_{t+h}\geq 1}\left[\Pi^{k-1}_{t+h}(\mathbf{v_{t+h}})\right],$$
$$\Pi^k_t(\text{sell} \geq 1 \text{ tomorrow}) = e^{-rh}e^{-\Lambda_{t+h}}\left[m(y^1) + \Pi^{k-1}_{t+h}(\varnothing)\right] + e^{-rh}(1-e^{-\Lambda_{t+h}})E_{t+h|N_{t+h}\geq 1}\left[\max\{m(y^1), m(v^1_{t+h})\} + \Pi^{k-1}_{t+h}(\{y^1, \mathbf{v_{t+h}}\}^2_k)\right].$$

Subtracting the second line from the first,

$$D\Pi^k_t(y^1) = (1-e^{-rh})m(y^1) + e^{-rh}(1-e^{-\Lambda_{t+h}})E_{t+h|N_{t+h}\geq 1}\left[m(y^1) + \Pi^{k-1}_{t+h}(\mathbf{v_{t+h}}) - \max\{m(y^1), m(v^1_{t+h})\} - \Pi^{k-1}_{t+h}(\{y^1, \mathbf{v_{t+h}}\}^2_k)\right], \tag{A.4}$$

where the term in the square brackets is between 0 and $-m(\overline{v})$. We would like (1) a lower bound on how $D\Pi^k_t(y^1)$ changes in y^1, and (2) an upper bound on how $D\Pi^k_t(y^1)$ changes over time.

For (1), let $\underline{m}' := \inf_{v\in[\underline{v},\overline{v}]} m'(v)$; this is strictly positive because $m(v)$ is strictly increasing and continuously differentiable. Differentiating (A.4),

$$\frac{d}{dy^1}D\Pi^k_t(y^1) \geq (1-e^{-rh})m'(y^1) \geq \frac{1}{2}rh\underline{m}' \tag{A.5}$$

for $h \leq h_0 := (\ln 2)/r$.

For (2), note that $\Pr(N_t = 1) = \Lambda_t e^{-\Lambda_t}$ and $\Pr(N_t \geq 2) = 1 - e^{-\Lambda_t}(1+\Lambda_t) \leq \Lambda_t^2 \leq \overline{\lambda}^2 h^2$, using the fact that $1 - e^{-x} \leq x$ for $x \geq 0$ and $\overline{\lambda} := \max_{t\in[0,T]} \lambda_t$. Splitting (A.4) into the case when there is one entrant and that where there are multiple entrants,

$$D\Pi^k_t(y^1) \geq (1-e^{-rh})m(y^1) + e^{-rh}\Lambda_{t+h}e^{-\Lambda_{t+h}}E_v\left[m(y^1) + \Pi^{k-1}_{t+h}(v) - \max\{m(y^1), m(v)\} - \Pi^{k-1}_{t+h}(\min\{y^1, v\})\right] - \overline{\lambda}^2 m(\overline{v})h^2, \tag{A.6}$$

where E_v is the expectation over the value of a single entrant. Advancing one period,

$$D\Pi^k_{t+h}(y^1) \leq (1-e^{-rh})m(y^1) + e^{-rh}\Lambda_{t+2h}e^{-\Lambda_{t+2h}}E_v\left[m(y^1) + \Pi^{k-1}_{t+2h}(v) - \max\{m(y^1), m(v)\} - \Pi^{k-1}_{t+2h}(\min\{y^1, v\})\right].$$

Subtracting these and completing the square gives us,

$$D\Pi^k_{t+h}(y^1) - D\Pi^k_t(y^1) \le e^{-rh}(\Lambda_{t+2h}e^{-\Lambda_{t+2h}} - \Lambda_{t+h}e^{-\Lambda_{t+h}})E_v\left[m(y^1) + \Pi^{k-1}_{t+h}(v) - \max\{m(y^1), m(v)\} - \Pi^{k-1}_{t+h}(\min\{y^1, v\})\right]$$
$$+ e^{-rh}\Lambda_{t+2h}e^{-\Lambda_{t+2h}}E_v\left[\left(\Pi^{k-1}_{t+2h}(v) - \Pi^{k-1}_{t+2h}(\min\{y^1, v\})\right) - \left(\Pi^{k-1}_{t+h}(v) - \Pi^{k-1}_{t+h}(\min\{y^1, v\})\right)\right] + \overline{\lambda}^2 m(\overline{v})h^2.$$

Consider the first term on the right-hand-side. If $\Lambda_{t+2h} \ge \Lambda_{t+h}$ the entire term is negative, and so is bounded above by zero. Conversely, assume $\Lambda_{t+2h} < \Lambda_{t+h}$. Using the mean value theorem, let $\Lambda_{t+h} = \tilde{\lambda}_{t+h}h$, for some $\tilde{\lambda}_{t+h}$ in the range of $\{\lambda_t : t \in [t, t+h]\}$, and similarly for Λ_{t+2h}. And since λ_t is Lipschitz continuous, let the bound on its derivative be denoted β. The first right-hand-side term is bounded above by

$$(\Lambda_{t+h}e^{-\Lambda_{t+h}} - \Lambda_{t+2h}e^{-\Lambda_{t+2h}})m(\overline{v}) \le (\Lambda_{t+h} - \Lambda_{t+2h})(1 - \Lambda_{t+2h})e^{-\Lambda_{t+2h}}m(\overline{v}) \le (\tilde{\lambda}_{t+h} - \tilde{\lambda}_{t+2h})m(\overline{v})h^2 \le 2\beta m(\overline{v})h^2,$$

where the first inequality uses the fact that ze^{-z} is increasing and concave on $z \in [0, 1]$ and so can be bounded by its tangent through $z = \Lambda_{t+2h}$. With the second right-hand-side term, we claim that

$$\Pi^{k-1}_{t+h}(v) - \Pi^{k-1}_{t+h}(\min\{y^1, v\}) = E_{t+2h}\left[\int_{\min\{y^1,v\}}^{v} m'(z)e^{-r(\tau^{k-1}_{t+h}(z)-t-h)}dz\right].$$

using the the envelope theorem as in (A.1), where $\tau^{k-1}_{t+h}(z)$ is the purchasing time of the single buyer present at time $t + h$. Subtracting these two integrals, we claim the second term is

$$e^{-rh}\Lambda_{t+2h}e^{-\Lambda_{t+2h}}\left[E_{t+3h}\left[\int_{\min\{y^1,v\}}^{v} m'(z)e^{-r(\tau^{k-1}_{t+2h}(z)-t-2h)}dz\right] - E_{t+2h}\left[\int_{\min\{y^1,v\}}^{v} m'(z)e^{-r(\tau^{k-1}_{t+h}(z)-t-h)}dz\right]\right]$$
$$\le \overline{\lambda}h\left[(1 - e^{-\Lambda_{t+2h}})m(\overline{v}) + e^{-\Lambda_{t+2h}}(1 - e^{-rh})m(\overline{v})\right]$$
$$\le \overline{\lambda}(\overline{\lambda} + r)m(\overline{v})h^2.$$

The first inequality comes from considering two cases. If there is entry over $(t + h, t + 2h]$ then this might lead to $\tau^{k-1}_{t+h}(z) = \infty$, yielding an upper bound of $m(\overline{v})$. If there is no entry, then $\tau^{k-1}_{t+h}(z) \le \tau^{k-1}_{t+2h}(z)$ implying an upper bound of $(1 - e^{-rh})m(\overline{v})$. The second inequality uses the fact that $1 - e^{-x} \le x$ for $x \ge 0$. Putting all this together, we have

$$D\Pi^k_{t+h}(y^1) - D\Pi^k_t(y^1) \le (2\beta + 2\overline{\lambda}^2 + \overline{\lambda}r)m(\overline{v})h^2. \tag{A.7}$$

To finish the proof, suppose that $x^k_{t+h} \leq x^k_t$, else there is nothing to prove. We now claim:

$$\begin{aligned}(2\beta + 2\overline{\lambda}^2 + \overline{\lambda}r)m(\overline{v})h^2 &\geq D\Pi^k_{t+h}(x^k_t) - D\Pi^k_t(x^k_t) \\ &= \int_{x^k_{t+h}}^{x^k_t} \frac{d}{dy^1} D\Pi^k_{t+h}(y^1)dy^1 + D\Pi^k_{t+h}(x^k_{t+h}) \\ &\geq (x^k_t - x^k_{t+h})\frac{1}{2}rh\underline{m}'\end{aligned}$$

for $h \leq h_0$. The first inequality comes from (A.7), the second line uses $D\Pi^k_t(x^k_t) = 0$, the third lines uses $D\Pi^k_{t+h}(x^k_{t+h}) \geq \Delta\Pi^k_{t+h}(x^k_{t+h}) = 0$ and (A.5). Rearranging then implies that there exists $\alpha > 0$ such that $(x^k_t - x^k_{t+h}) \leq \alpha h$ for $h \leq h_0$.

Claim B: If $k \geq 2$, there exist positive constants α, h_0 such that $x^k_{T-h} - x^k_T \leq \alpha h$ for $h \leq h_0$.

Proof: In period $t = T - h$, if $m(y^1) \geq 0$ then (A.6) becomes

$$D\Pi^k_{T-h}(y^1) \geq (1 - e^{-rh})m(y^1) - \overline{\lambda}^2 m(\overline{v})h^2, \tag{A.8}$$

since the term in square brackets in (A.6) is zero.

In period T, we have $x^k_T = m^{-1}(0)$. Since the seller will never sell to a buyer with negative marginal revenue, we have $x^k_{T-h} \geq x^k_T$. We now claim:

$$\begin{aligned}\overline{\lambda}^2 m(\overline{v})h^2 &\geq D\Pi^k_{T-h}(x^k_{T-h}) - D\Pi^k_{T-h}(x^k_T) \\ &= \int_{x^k_T}^{x^k_{T-h}} \frac{d}{dy^1} D\Pi^k_{T-h}(y^1)dy^1 \\ &\geq (x^k_{T-h} - x^k_T)\frac{1}{2}rh\underline{m}'\end{aligned}$$

for $h \leq h_0$. The first line uses (A.8), $m(x^k_T) = 0$ and $D\Pi^k_{T-h}(x^k_{T-h}) = 0$. The second line follows from the fundamental theorem of calculus. The third line uses (A.5). Rearranging yields the result.

Claim C: The lost profits from using prices with a second-price auction for the last unit at time T is of order $O(h)$.

Proof: We use the following mechanism: In each period the seller chooses a price p^k_t and allocates the good to anyone willing to pay; the only exception is in period T if there is a single unit, when she runs a second-price auction with reserve $m^{-1}(0)$. If there is more demand than supply in a given period, allocations are randomized.

First, we claim that these prices induce a series of cutoffs x_t^k, such that buyers wish to buy if their value exceeds the cutoff, where k is the number of units at the start of the period. To see this observe that, since buyers enter according to a Poisson process, each type (v,t) has the same expectation over prices. A buyer with type (v,t) thus chooses a (random) purchasing time τ after his entry date t to maximize

$$u(v,t,\tau) = E_0\left[v\mathbf{1}_{\tau\geq t}e^{-r\tau} - p_\tau\right]. \tag{A.9}$$

Here, the price p_t is a random variable, depending on the sales to other buyers. If other buyers' demand as many or more units than the seller has to offer, the price may also rise to ∞ depending on the priority of the buyer at the rationing stage; a choice of $\tau = \infty$ then indicates that the buyer does not buy. The function $u(v,t,\tau)$ has strictly decreasing differences in (v,τ) since $r > 0$, and is (weakly) supermodular in τ. Hence every optimal selection $\tau^*(v,t)$ is decreasing in v by Topkis (1998, Theorem 2.8.4) and we can let $x_t^k = \inf\{v : \tau^*(v,t) = t\}$ be the lowest type who buys in period t.

Conversely, we claim that any sequence of cutoffs can be implemented by prices. These prices can be constructed by backward induction. Alternatively, one can consider the utility of a buyer with type x_t^k who enters at time t,

$$e^{-rt}(x_t^k - p_t^k) = E_0\left[\int_{\underline{v}}^{x_t^k} e^{-r\tau(z,t)}dz\right],$$

where the left-hand side is his direct utility, and the right-hand side comes from applying the envelope theorem to utility (A.9). In this equation, $\tau(z,t)$ is the (random) purchasing time of a buyer with value z born at time t induced by the cutoffs $\{x_t^k\}$ and the rationing rule.

Next, we claim that the profits lost from using prices are of order $O(h)$. Fix a realization of cutoffs up to time $t-h$. Let $s_t(x)$ denote the last time the cutoff went below x when looking back from time t. Over the time $(t-h,t]$, the next sale arrives according to a non-homogeneous Poisson process in which the integral of the arrival rate is

$$\Phi_t = \int_{x_t^k}^{x_{t-h}^k}\int_{s_t(z)}^{t-h}\lambda_s\,ds\,dF(z) + \left(\int_{t-h}^{t}\lambda_s ds\right)(1 - F(x_t^k)) \leq \overline{\lambda}T\overline{f}\alpha h + \overline{\lambda}h =: \gamma h$$

for $h \leq h_0$, where the inequality uses Claims A and B, and $\overline{f}$ is the upper bound on the continuous density. The probability of two or more sales over $(t-h,t]$ is $1 - e^{\Phi_t}(1+\Phi_t) \leq \Phi_t^2 \leq \gamma^2h^2$. The probability of two or more entrants in any periods is bounded above

$$1 - (1-\gamma^2h^2)^{T/h} \leq \frac{T}{h}\gamma^2h^2 = T\gamma h$$

for $h \leq h_0$, as required.

Finally, if $t = T$ and $k = 1$ then it is a weakly dominant strategy for the buyers to bid their true value in the second-price auction. Hence the unit is allocated to the buyer with the highest value.

References

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. (2005): *Dynamic Programming and Optimal Control*, vol. I. Athena Scientific, third 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.

Board, S., and A. Skrzypacz (2010): "Revenue-Management with Forward-Looking Buyers," Working paper, UCLA.

Bruss, F. T., and T. S. Ferguson (1997): "Multiple buying or selling with vector offers," *Journal of Applied Probability*, pp. 959–973.

Bulow, J. (1982): "Durable-Goods Monopolists," *Journal of Political Economy*, 90(2), 314–332.

Dilme, F., and F. Li (2012): "Revenue Management Without Commitment: Dynamic Pricing and Periodic Fire Sales," Working paper, Penn.

Elmaghraby, W., A. Gülcü, and P. Keskinocak (2008): "Designing optimal preannounced markdowns in the presence of rational customers with multiunit demands," *Manufacturing & Service Operations Management*, 10(1), 126–148.

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*, 1(2), 168–198.

——— (2009b): "Learning about the Future and Dynamic Efficience," *American Economic Review*, 99(4), 1576–1588.

Gershkov, A., B. Moldovanu, and P. Strack (2013): "Dynamic Allocation and Learning with Strategic Arrivals," Working paper, Bonn.

Graddy, K. (1995): "Testing for imperfect competition at the Fulton fish market," *The RAND Journal of Economics*, pp. 75–92.

——— (2006): "Markets: the Fulton fish market," *The Journal of Economic Perspectives*, 20(2), 207–220.

Graddy, K., and G. Hall (2011): "A dynamic model of price discrimination and inventory management at the Fulton Fish Market," *Journal of Economic Behavior & Organization*, 80(1), 6–19.

Hörner, J., and L. Samuelson (2011): "Managing Strategic Buyers," *Journal of Political Economy*, 119(3), 379–425.

Lazear, E. P. (1986): "Retail Pricing and Clearance Sales," *The American Economic Review*, 76(1), 14–32.

Li, A. (2011): "Selling Storable Goods to a Dynamic Population of Buyers: A Mechanism Design Approach," Working paper, Stanford University.

Li, J., N. Granados, and S. Netessine (2011): "Are consumers strategic? Structural estimation from the air-travel industry," Discussion paper, Wharton.

MacQueen, J., and R. Miller (1960): "Optimal persistence policies," *Operations Research*, 8(3), 362–380.

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.

McAfee, R. P., and V. Te Velde (2006): "Dynamic pricing in the airline industry," *in Handbook on Economics and Information Systems*.

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

Myerson, R. B. (1981): "Optimal Auction Design," *Mathematics of operations research*, 6(1), 58–73.

Pai, M., and R. Vohra (2013): "Optimal Dynamic Auctions and Simple Index Rules," *Mathematics of Operations Research*, forthcoming.

Ross, S. M. (1971): "Infinitesimal look-ahead stopping rules," *The Annals of Mathematical Statistics*, 42(1), 297–303.

Said, M. (2012): "Auctions with dynamic populations: Efficiency and revenue maximization," *Journal of Economic Theory*, 147(6), 24192438.

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

Soysal, G. P., and L. Krishnamurthi (2012): "Demand Dynamics in the Seasonal Goods Industry: An Empirical Analysis," *Marketing Science*, 31(2), 293–316.

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

Su, X. (2007): "Intertemporal pricing with strategic customer behavior," *Management Science*, 53(5), 726–741.

Sweeting, A. (2012): "Dynamic Pricing Behavior in Perishable Goods Markets: Evidence from Secondary Markets for Major League Baseball Tickets," *Journal of Political Economy*, 120(6), 1133–1172.

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

Topkis, D. M. (1998): *Supermodularity and Complimentarity*. Princeton University Press, New Jersey.