Revenue Management with Forward-Looking Buyers

Simon Board

University of California, Los Angeles

Andrzej Skrzypacz

Stanford University

A seller wishes to sell multiple goods by a deadline, for example, the end of a season. Potential buyers enter over time and can strategically time their purchases. Each period, the profit-maximizing mechanism awards units to the buyers with the highest valuations exceeding a sequence of cutoffs. We show that these cutoffs are deterministic, depending only on the inventory and time remaining; in the continuous-time limit, the optimal mechanism can be implemented by posting anonymous prices. When incoming demand decreases over time, the optimal cutoffs satisfy a one-period-look-ahead property and prices are defined by an intuitive differential equation.

I. Introduction

Each autumn, retailers stock up on coats that they seek to sell over the subsequent season. The unsold units are then put on a sequence of sales in January, as retailers make room for spring clothing, with any remain-

We are grateful to the editor, Jesse Shapiro, and the referees for many excellent comments. 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 suggestions. We also thank seminar audiences at Bologna, Center for Studies in Economics and Finance (Naples), Einaudi, Econometric Society World Congress, European University Institute,

Electronically published June 30, 2016
[*Journal of Political Economy*, 2016, vol. 124, no. 4]

ing inventory scrapped (i.e., given to charity, recycled, or sold at discount retailers). If a customer discovers a coat he likes in December, he must therefore choose not only whether to buy, but also when to buy. Delaying means that the customer pays a lower price, but he also has fewer opportunities to wear the coat and risks its selling out. This implies that high-value customers buy immediately while low-value customers postpone their purchase, consistent with surveys that report around 60 percent of consumers "wait for a sale to buy what they want."[1] The possibility of delay is important for retailers since price reductions lead to sales from both new customers and the reservoir of old customers. For example, figure 1 shows the sales pattern for a typical women's coat; price decreases lead to large spikes in demand that quickly fade, indicating that a stock of buyers wait for the price to fall. Moreover, such consumer delay means that the retailer must take into account how its sales strategy at the end of the season affects consumers' decisions earlier in the season. For example, JC Penney's customers became accustomed to "endless sales promotions," meaning that revenue dropped by 25 percent when it experimented with a flatter pricing policy (Robinson 2014).

In this paper, we derive the optimal sales strategy for a seller facing long-lived buyers who have rational expectations (i.e., are "forward-looking"). The seller can choose any feasible mechanism, allowing her to run a series of auctions, issue coupons to buyers who arrive early, or let the price paid by one buyer depend on reports of others who are waiting to buy. Despite all these options, we show that it is optimal for the seller to choose a sequence of anonymous posted prices and let buyers reveal their existence only when they purchase. When incoming demand decreases over time, the optimal prices can then be characterized via an intuitive differential equation.

This paper contributes to the field of revenue management, which studies how to sell inventory to customers entering a market over time. Typical revenue management models assume that buyers are short-lived, exiting the market if they do not immediately buy (see the book by Talluri and van Ryzin [2004]), and it is a well-known open problem to characterize optimal pricing with forward-looking consumers. This paper finds a natural setting in which we can use the tools of mechanism design to trac-

Microsoft, Midwest Meetings, New York University, Stanford, Southwest Workshop in Economic Theory, University of California, Los Angeles Anderson, Yonsei, Yahoo!, and Yale. A previous version of this paper went by "Optimal Dynamic Auctions for Durable Goods: Posted Prices and Fire-Sales."

[1] "America's Bargain-Hunting Habits," *Consumer Reports* (April 30, 2014). See also the survey of American Research Group on "2013 Christmas Gift Spending Plans Stall" (November 15, 2013) or the Acosta Mosaic Group report "Hot Topic Report—a Shift in the Lift" (November 2012).

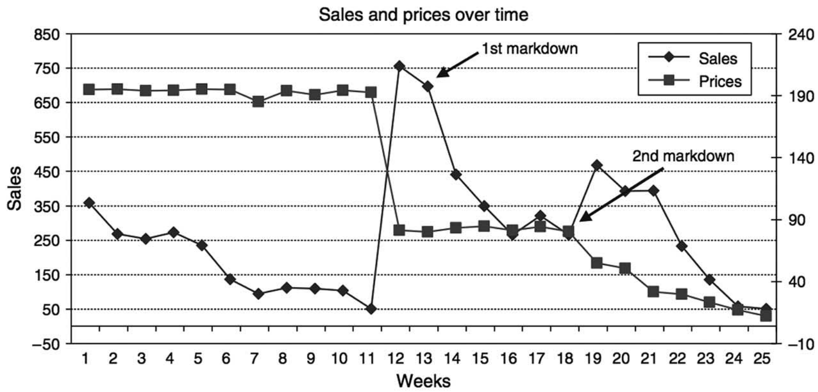

Fig. 1.—Prices and sales for a typical women's coat. Source: Soysal and Krishnamurthi (2012).

tably characterize the optimal prices. From a normative perspective, our model can therefore help design sales policies in a wide variety of markets in which revenue management is increasingly prevalent, such as online advertising, package holidays, and concert tickets. From a positive perspective, the paper provides a fully solved benchmark that complements the growing interest in estimating models of dynamic pricing (e.g., Gowrisankaran and Rysman 2012; Sweeting 2012; Lazarev 2013). The predictions concerning prices and sales can help make inferences in particular markets (e.g., whether customers are long- or short-lived); the model can also provide a basis for counterfactuals (e.g., the impact of resale). In addition, the paper sheds light on common business practices. We show that profits are higher if buyers are long-lived, which explains why firms such as Nordstrom benefit from having a predictable sales cycle and suggests that retailers should embrace price alerts (e.g., Camelcamelcamel) and price predictors (e.g., Kayak.com). We also show that firms have no incentive to hide their inventory, consistent with retailers' willingness to inform customers of their remaining stock.

The paper also elucidates the puzzle of why most goods are sold via posted prices rather than auctions (e.g., Einav et al. 2013). While posted prices can be optimal in large markets (Segal 2003), one would expect auctions to perform much better when there are a few items to sell and a few heterogeneous buyers in the market at the same time. Our paper shows that this need not be the case in a dynamic market: In our model, buyers accumulate over time; nevertheless, posted prices implement the "optimal auction." This result is particularly important for revenue management since the literature typically assumes that the seller uses posted prices (e.g., Su 2007; Aviv and Pazgal 2008). Our analysis indicates when this assumption is without loss and when the seller can do better.

In the model, the seller wishes 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. Buyers' values are drawn from a common distribution, but the number of entering buyers may vary over time. Once he arrives, a buyer can delay his purchase, incurring a costly delay and risking a stock-out in the hope of lower prices. We model the social loss of delay by assuming proportional discounting; in an extension, we show that analogous results obtain if instead the seller incurs 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. Tackling the problem in two stages simplifies the analysis: 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 single-agent dynamic programming problem.

In Section IV, we characterize optimal allocations. We first show that the seller awards a good to the buyer with the highest valuation if his value exceeds a cutoff x_t^k, where k is the inventory. We apply this principle repeatedly within a period, so multiple units may be allocated at one time if the highest value exceeds x_t^k, the second-highest value exceeds x_t^{k-1}, 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 depends only 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. This means that the seller will bring forward the sales period if a good has unusually high inventory and postpone the sales period if the inventory is low. Intuitively, if the seller delays awarding the kth 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 V, 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; that is, the seller chooses a single price at each point in time and buyers reveal their existence only when they purchase a unit. Prices are limiting in that they (i) do not discriminate on the basis of arrival times as in a coupon system, (ii) do not allow the price paid by one buyer to depend on the reports of waiting buyers, and (iii) do not adjust within a period as in an auction. In our model, the seller does not benefit from such mechanisms because the optimal cutoffs (i) are the same for each cohort, (ii) are deterministic and so independent of others' values, and (iii) are continuous, so simultaneous sales do not occur when periods are short.

When the arrival rate of buyers 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. When compared to a model with short-lived buyers, cutoffs are relatively constant and then drop rapidly, and sales are back-loaded. This helps us understand the importance of end-of-season sales for retailers and the role of discount websites for package holidays and concert tickets (e.g., Lastminute.com, Goldstar).

In Section VI, we consider a number of extensions. We first show that the spirit of the main results continues to apply if impatience comes from inventory costs or if units arrive and expire over time (e.g., for a retailer for which shelf space is costly and fashion trends change). Second, if there are different classes of buyers (e.g., rich media vs. static ads in online display advertising) 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, and the seller charges different prices for different types of buyers. Moreover, in the airline example, we propose a novel implementation of the optimal mechanism whereby a customer is issued a coupon when he registers his interest in a flight that can be redeemed when the customer makes a purchase. 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 indicative auctions in which all buyers bid and the seller makes a counteroffer to the highest rather than using posted prices.

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 that interarrival 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; Das Varma and Vettas 2001) 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 in which a seller has multiple units and wishes to sell them over finite time. This model is very rich, allowing buyers to arrive and leave the market over time and the distribution of entering buyers to vary. 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 allow the seller to choose any mechanism, show when posted prices are optimal, and then characterize the optimal prices; interestingly, such prices will tend to rise after sales, so the seller can do better than a series of markdowns.

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 independently and identically distributed (IID), the cutoff value is constant, except for the last period (see Bertsekas 2005, 185). Ross (1971) studies this problem directly in continuous time. 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 to future prices or mechanisms (e.g., Hörner and Samuelson 2011). There may be heterogeneous 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) and a smaller one on selling durable goods in competitive marketplaces (e.g., Deneckere and Peck 2012).

II. Model

Basics.—A seller (she) has K goods to sell to buyers (he) arriving over time. Time is discrete and finite, $t \in \{1, \ldots, T\}$.

Entrants.—At the start of period t, N_t buyers arrive. The number of arrivals 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 that 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.[2]

Preferences.—After a buyer enters 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 is 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. Motivated by the retailing application, a buyer's value declines throughout the season: If the buyer purchases at time s for price p_s, his utility is $v\delta^s - p_s$, where $\delta \in (0, 1)$. Let v_t^k denote the kth-highest order statistic of the buyers entering at time t.

Mechanisms.—At time 0 the seller chooses a mechanism. Applying the revelation principle, it is without loss of generality to consider communication mechanisms in which each buyer makes report $\tilde{v}_i$ when he enters the market, and the seller tells him only when he is awarded a good (Myerson 1986). Intuitively, giving any information about the history of the game as it unfolds (e.g., the number of objects available, the reports of other agents) will add more incentive constraints. A mechanism consists of an allocation and payment rule $\langle \tau_i, TR_i \rangle$ that maps buyers' reports and birth dates into a purchasing time τ_i for buyer i and expected transfer TR_i. 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).[3]

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

$$u_i(\tilde{v}_i, v_i, t_i) = E_0[v_i \delta^{\tau_i(\tilde{v}_i, \mathbf{v}_{-i}, \mathbf{t})} - TR_i(\tilde{v}_i, \mathbf{v}_{-\mathbf{i}}, \mathbf{t}) | v_i, t_i], \tag{1}$$

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,

[2] When N_t is unobserved by the seller, one can think of our solving the "relaxed" problem, ignoring the incentive compatibility (IC) constraints on birth dates. In Sec. IV.A, 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.

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

and $\mathbf{t}$ is the vector of their birth dates. A mechanism is incentive compatible if the buyer wishes to tell the truth given that 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

$$\text{Profit} = E_0\left[\sum_i TR_i(\mathbf{v}, \mathbf{t})\right] \tag{2}$$

subject to incentive compatibility and individual rationality.

Remarks.—The interpretation of the model depends on the application at hand. For retailing, time T can be interpreted as the date the seller ships unsold goods to a discount retailer (e.g., TJ Maxx) or to a charity (e.g., the New York Clothing Bank) or recycles them (e.g., into cushion filler). We normalize the value of these unsold goods to zero. For an online ad, package holiday, or concert, time T is the date of the event. Time T can also be interpreted as the last time buyers enter the market since, in the optimal mechanism, no sales will occur after this point.

The discount factor δ represents the loss of value that results from delaying allocation.[4] With retailing, impatience comes from having fewer opportunities to wear the good. With an online ad, package holiday, or concert, it comes from having less time to make complementary decisions (e.g., organizing an advertising campaign, taking vacation days from work, inviting friends to the concert). As discussed in Section VI.A, one obtains analogous results if impatience comes from the seller's inventory costs rather than the buyer's discounting.

The model makes a couple of notable assumptions. We assume that buyers do not know the number of units remaining in the mechanism (indeed, they know only their value and birth date). However, when implementing the optimal allocation, the seller will tell buyers her remaining inventory and the history of past prices, so the seller does not benefit from hiding this information.

We also assume that the seller can commit to a mechanism. We think this is reasonable with applications such as retailing, online ads, and concerts in which the seller automates the pricing scheme and uses it repeatedly. It is also appropriate when using the model from a normative perspective to design the dynamic pricing strategy of, say, an airline. Such sellers have some degree of commitment since they often do not sell all their capacity: clothing stores shred or donate capacity rather than lower

[4] One could allow the decay rate δ to depend on time; this corresponds to rescaling time in the current model. One can also reinterpret our results using the durable-goods utility specification $(v - \tilde{p}_t)\delta^t$, where time preference comes from a standard discount factor. Whether or not we discount money does not affect allocations, although prices must be rescaled under this specification, with the new price given by $\tilde{p}_t = \delta^{-t}p_t$.

the price, while airlines often fly with empty seats;[5] we take the extreme case and suppose that they can fully commit. One can thus view this as an upper bound on the profit a seller can obtain. If the seller cannot commit, the problem is much harder to study (e.g., Fuchs and Skrzypacz 2010; Hörner and Samuelson 2011; Dilme and Li 2012).

Finally, we solve for the profit-maximizing allocation. This is mainly for practical relevance; however, if one replaces marginal revenues (defined below) with values, all our results apply to the welfare-maximizing allocation. Indeed, our results apply to any single-agent decision problem in which the agent has K units to allocate or acquire; it can therefore be applied to a consumer who searches for K goods or a firm that wishes to hire K employees.

Preliminaries.—When a buyer enters the market, he chooses his report $\tilde{v}_i$ to maximize his utility (1). As shown by Mas-Colell, Whinston, and Green (1995, proposition 23.D.2), an allocation rule is incentive compatible if and only if (a) the discounted allocation probability

$$E_0[\delta^{\tau_i(\mathbf{v},\mathbf{t})}|v_i, t_i] \tag{3}$$

is increasing in v_i and (b) applying the envelope theorem to (1), equilibrium utility is

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

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

$$E_0[u_i(v_i, v_i, t_i)] = E_0\left[\delta^{\tau_i(\mathbf{v},\mathbf{t})}\frac{1-F(v_i)}{f(v_i)}\right]. \tag{4}$$

Profit (2) equals welfare minus buyers' utilities. Summing utility (4) over all buyers, we obtain

$$\text{Profit} = E_0\left[\sum_i \delta^{\tau_i(\mathbf{v},\mathbf{t})} m(v_i)\right], \tag{5}$$

where the *marginal revenue* of buyer i is given by $m(v_i) \coloneqq v_i - [1 - F(v_i)]/f(v_i)$. Throughout we assume that $m(v)$ 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 (3). We also assume that $m(\underline{v}) < 0$, so the optimal cutoff is interior.

[5] For clothing stores, see "Where Unsold Clothes Meet People in Need," *New York Times*, January 8, 2010. For airlines, Ryanair has unsold capacity on 80–90 percent of its flights (Malighetti, Paleari, and Redondi 2009).

One should note that the profit equation (5) allows buyers' arrival times to be correlated within and across periods. That is, a buyer's arrival time may give him information about when other buyers arrive, as in McAfee and McMillan (1987). Intuitively, from each buyer's perspective, his rents are determined by the expected discounted purchasing time; it does not matter what is the underlying stochasticity. To characterize optimal allocations, we require only that N_t is independently distributed across periods, allowing us to use backward induction with the state variable equal to the vector of buyers' values. When we implement these allocations through prices, we assume that entry is Poisson, so buyers have symmetric expectations about the competition they face from other buyers.

III. Single-Unit Example

Before launching into the main analysis we develop some intuition by heuristically deriving the optimal allocation and prices for the case in which $K = 1$ and N_t is IID. As an example, suppose that a high-end retailer is selling a limited-edition jacket or 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, 185), the optimal cutoffs are constant up to the penultimate period, $x_t = x^*$ for $t < T$, and jump down to the usual monopoly cutoff in the final period, $x_T = m^{-1}(0)$.[6] Period T is identical to a standard auction, so the seller wants to sell to the buyer with the highest value as long as his marginal revenue is positive. In earlier periods, the cutoffs are uniquely given by

$$m(x^*) = \delta E_{t+1}[\max\{m(v^1_{t+1}), m(x^*)\}]. \tag{6}$$

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). In the penultimate period, the cutoff is clearly given by (6). In period $T - 2$, the seller is indifferent between awarding the good today and waiting until $T - 1$; if she waits then she has exactly the same trade-off tomorrow and is indifferent again, so we can assume that she sells at period $T - 1$, yielding (6). Working backward, the cutoffs are thus constant for all $t < T$.[7]

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 ar-

[6] This result is also a special case of theorem 2.

[7] This logic depends on entering demand N_t being IID. We consider more general demand processes in Sec. IV.

rivals. 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 multiunit case, as shown in Sections IV.A and IV.B.

To consider the continuous-time limit, suppose that buyers enter with Poisson rate λ and the instantaneous discount rate is r. In the discretized version with periods of length h, the arrival rate is λh and the discount factor is $\delta = e^{-rh}$. As the period length h becomes short, then (6) becomes

$$rm(x^*) = \lambda E_v[\max\{m(v) - m(x^*), 0\}], \tag{7}$$

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 $\delta^T m^{-1}(0)$. 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 that type x^* is indifferent between buying at price p_T and entering the auction. That is,

$$p_T = \delta^T E_0[\max\{y^2, m^{-1}(0)\}|y^1 = x^*], \tag{8}$$

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

$$\frac{dp_t}{dt} = -rx^* - (x^* - p_t)\lambda[1 - F(x^*)]. \tag{9}$$

When a buyer waits a little, he gains from the falling prices (the left-hand side) but loses the rental value of the good and risks 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 stock-out. 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* whereby a buyer bids at time t and is allocated a unit at a later date if no one offers a better price beforehand. Formally, suppose that the price path p_t is given by (9) with boundary condition (8). In the game, a buyer bids b at any time after he enters. If the bid is entered at time t and $b \geq p_t$, then he immediately purchases the good at price p_t. If $p_t \geq b \geq p_T$, 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.[8] 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 buyer to reserve the ad slot at a discount if no buyer is willing to pay the full price.

IV. Optimal Allocations

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

A. *General Case*

The seller's problem is to choose a feasible allocation rule $\langle \tau_i \rangle$ to maximize profits (5). Since this is a single-agent optimization problem, the principle of optimality means we can solve it by backward induction.

Suppose that 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 his allocation (and the allocation of all other buyers) is independent of t_i. This follows because the birth date enters only 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 his 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 $t'' > t'$. When we

[8] 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.

swap these two buyers but leave 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, when solving the seller's problem, we need keep track of the values of only 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[9]

$$\Pi_t^k(\mathbf{y}) := \max_{\tau_i \geq t} E_{t+1}\left[\sum_i \delta^{\tau_i - t} m(v_i)\right], \tag{10}$$

where 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 by

$$\tilde{\Pi}_{t+1}^{k-j}(\mathbf{y}^{-\mathbf{j}}) := \max_{\tau_i \geq t+1} E_{t+1}\left[\sum_i \delta^{\tau_i-(t+1)} m(v_i)\right] \tag{11}$$

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

$$\Pi_t^k(\mathbf{y}) = \max_{j \in \{0, \ldots, k\}}\left[\sum_{i=1}^{j} m(y^i) + \delta \tilde{\Pi}_{t+1}^{k-j}(\mathbf{y}^{-\mathbf{j}})\right]. \tag{12}$$

The proof of the following lemma shows that allocation is monotone in buyers' values $\mathbf{y}$. As a result, a mechanism can be characterized by *cutoffs* $x_t^j(\mathbf{y}^{-(\mathbf{k}-\mathbf{j}+\mathbf{1})})$, $j \leq k$, which describe the lowest type that is awarded the jth unit, assuming that the previous $k-j$ units have been sold.[10] Within a period, several units may be allocated. We allocate unit k to the highest

[9] In this equation, Π_t^k depends on $\mathbf{y}$ because the seller has the choice of allocating any good to a current buyer or to a future entrant. One should also note that 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 receives only $m(v)$ since all higher types gain rents, even those that purchase prior to period t.

[10] Fixing $\mathbf{y}^{-(\mathbf{k}-\mathbf{j}+\mathbf{1})}$, the cutoff is well defined if there are some types $y^{(k-j+1)} \geq y^{(k-j+2)}$ who are not allocated a unit. If all $y^{(k-j+1)} \geq y^{(k-j+2)}$ are allocated a unit, then we define $x_t^j(\mathbf{y}^{-(\mathbf{k}-\mathbf{j}+\mathbf{1})}) = \underline{v}$. Below, we show that the optimal allocation can be implemented by deterministic cutoffs that are independent of $\mathbf{y}^{-(\mathbf{k}-\mathbf{j}+\mathbf{1})}$. At that point, we no longer need to condition on whether or not $y^{k-j+1} > y^{k-j+2}$.

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

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

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

Subtracting the first equation from the second,

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

where the second inequality uses the fact that 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 that the allocation of unit j is monotone in y^{k-j+1}. Fixing $\mathbf{y}^{-(\mathbf{k}-\mathbf{j}+\mathbf{1})}$, we can thus define $x_t^j(\mathbf{y}^{-(\mathbf{k}-\mathbf{j}+\mathbf{1})})$ as the lowest $y^{k-j+1} > y^{k-j+2}$ that is allocated a unit; if all such $y^{k-j+1} > y^{k-j+2}$ are allocated units, then set $x_t^j(\mathbf{y}^{-(\mathbf{k}-\mathbf{j}+\mathbf{1})}) = \underline{v}$. QED

If the seller starts period t with k units, she sells the jth 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 jth 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

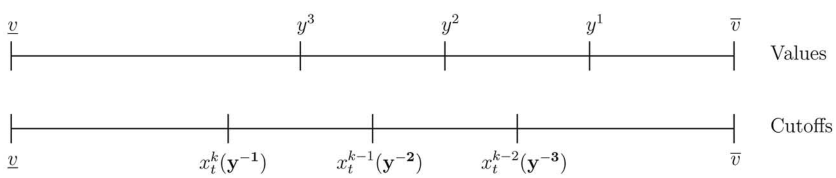

FIG. 2.—Allocation within a period. This figure shows the top three value buyers and the cutoffs for three units. Unit k is allocated to y^1 since $y^1 \geq x_t^k(\mathbf{y^{-1}})$. Similarly, unit $k-1$ is allocated to y^2 since unit k was sold and $y^2 \geq x_t^{k-1}(\mathbf{y^{-2}})$. However, unit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

and we can treat each unit separately, simply comparing the jth cutoff to the corresponding buyer's valuation (as in fig. 2).

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

Proof. If unit j is allocated, then the corresponding buyer's value must exceed the cutoff. Conversely, if $y^{k-j+1} \geq x_t^j(\mathbf{y}^{-(\mathbf{k}-\mathbf{j}+1)})$, then

$$y^{k-\ell+1} \geq y^{k-j+1} \geq x_t^j(\mathbf{y}^{-(\mathbf{k}-\mathbf{j}+1)}) \geq x_t^\ell(\mathbf{y}^{-(\mathbf{k}-\ell+1)})$$

for all $\ell \geq j$, since the cutoffs are decreasing. Hence all units $\ell \geq j$ are allocated to their corresponding buyer. QED

The next step is to observe that we can simplify notation. 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, 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.

Define the profit if the seller sells zero units today or if she sells only one:

$$\Pi_t^k(\text{sell 0 today}) = \delta\tilde{\Pi}_{t+1}^k(y^1, \mathbf{y}^{-1}), \tag{13}$$

$$\Pi_t^k(\text{sell 1 today}) = m(y^1) + \delta\tilde{\Pi}_{t+1}^{k-1}(\mathbf{y}^{-1}). \tag{14}$$

Denote the difference function by

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

which reflects the incentives to sell today rather than wait. These definitions are useful because if cutoffs are decreasing in k, then a seller who is indifferent between selling to buyer y^1 today and waiting weakly prefers not to sell a second unit today.

The next result, lemma 3, establishes some basic properties of $\Delta\Pi_t^k(y^1, \mathbf{y}^{-1})$. We say that the cutoffs $x_t^k(\mathbf{y}^{-1})$ are *deterministic* if they are independent of $\mathbf{y}^{-1}$.

LEMMA 3. Suppose that future cutoffs $\{x_s^j\}_{s\geq t+1}$ are deterministic and decreasing in $j \leq k$. Then

a. $\Delta\Pi_t^k(y^1, \mathbf{y}^{-1})$ is independent of $\mathbf{y}^{-1}$,
b. $\Delta\Pi_t^k(y^1)$ is continuous and strictly increasing in y^1, and
c. $\Delta\Pi_t^k(y^1)$ is increasing in k.

Proof. See the Appendix.

Part *a* says that $\Delta\Pi_t^k(y^1, \mathbf{y}^{-1})$ is independent of $\mathbf{y}^{-1}$. 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 that the seller has *k* goods in period *t*. The optimal allocation rule awards a unit to the highest remaining buyer if his 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 that 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.[11]

Continuing by induction, fix $k > 1$ and suppose that $x_t^j \leq x_t^{j-1}$ 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 $\mathbf{y}^{-1}$.[12] 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 $x_t^k(\mathbf{y}^{-1}) \leq x_t^{k-1}$. Fix $\mathbf{y}^{-1}$ and consider two cases. If the seller is indifferent at $y^1 = x_t^k(\mathbf{y}^{-1}) \geq y^2$, then she weakly prefers not to allocate a second unit $y^2 \leq x_t^{k-1}$, and the cutoff is determined by $\Delta\Pi_t^k(x_t^k(\mathbf{y}^{-1})) = 0$. Lemma 3($a$) then implies that the cutoff x_t^k is independent of $\mathbf{y}^{-1}$. If the seller prefers to allocate to all $y^1 \geq y^2$, then $\Delta\Pi_t^k(y^1) \geq 0$.[13] We can therefore

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

[12] By definition of the cutoff, $x_t^k(\mathbf{y}^{-1}) > x_t^{k-1}$ implies that $x_t^k(\mathbf{y}^{-1}) > y^2$ (see fn. 10). Since type $\overline{v}$ is immediately awarded a good, we can assume that $x_t^k(\mathbf{y}^{-1}) \in (y^2, \overline{v})$.

[13] To see this, consider two cases. If the seller wishes to sell only one good, then we have $\Delta\Pi_t^k(y^1) \geq 0$. If she wishes to sell two or more, then $y^2 > x_t^{k-1}$, so $\Delta\Pi_t^{k-1}(y^2) \geq 0$ and parts *b* and *c* of lemma 3 imply $\Delta\Pi_t^k(y^1) \geq 0$.

let the cutoff be the solution of $\Delta\Pi_t^k(x_t^k(\mathbf{y}^{-1})) = 0$, which, by lemma 3(*a*) is independent of $\mathbf{y}^{-1}$. QED

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 that 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. 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 decrease when there are more units available. Intuitively, if the seller delays awarding a unit by one period, then she can allocate it to an entrant rather than to 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.

B. *Weakly Decreasing Demand*

In this section we suppose that the arrival rate of incoming demand is weakly decreasing over time. This captures the idea that the pool of potential new customers falls over time (e.g., if Zara launches a line of coats). It also includes the canonical case of IID arrivals. We show that this property means that we can characterize the optimal cutoffs by intuitive difference equations (which will become differential equations in the continuous-time limit).

In theorem 1, a seller is indifferent between selling to buyer x_t^k today and waiting for future entry. We say that 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 $\mathbf{v_t} := \{v_t^1, \ldots, v_t^k\}$, and let $\{y^1, \mathbf{v_{t+1}}\}_k^2$ represent the ordered vector of the second- to kth-highest choices from $\{y^1, \mathbf{v_{t+1}}\}$.[14] Define

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

and

[14] If there are fewer than k buyers who enter or are present in the market, then the corresponding entries equal zero.

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

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 $y^1 \geq x_t^k \geq x_{t+1}^k$.

Lemma 4. Suppose that 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 $s \in \{t+1, \ldots, T-1\}$ and $j \leq k$. Then $D\Pi_{t+1}^k(y^1) \geq D\Pi_t^k(y^1)$.

Proof. See the Appendix.

Theorem 2. Suppose that 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

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

In this equation, the second inequality follows from the envelope theorem analogous to the proof of lemma 3(b); intuitively, an increase in y^1 will raise the profit from selling today by $m'(y^1)$ and raise the profit by waiting by at most $\delta m'(y^1)$, so the difference $D\Pi_t^k(y^1)$ increases. 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 $D\Pi_t^k(x_t^k) = 0$. QED

Intuitively, if the seller delays awarding the kth unit by one period, then she can allocate it to an entrant rather than to 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,

$$m(x_T^k) = 0. \tag{15}$$

In period $t = T-1$, the seller balances the revenue from allocating the kth good against the opportunity cost derived from the possibility of denying the good to the kth-highest new entrant. Hence,

$$m(x_{T-1}^k) = \delta E_T[\max\{m(x_{T-1}^k), m(v_T^k)\}]. \tag{16}$$

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 sells only 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,

$$m(x_t^k) + \delta E_{t+1}[\Pi_{t+1}^{k-1}(\mathbf{v_{t+1}})] = \delta E_{t+1}[\max\{m(x_t^k), m(v_{t+1}^1)\}] + \delta E_{t+1}[\Pi_{t+1}^{k-1}(\{x_t^k, \mathbf{v_{t+1}}\}_k^2)]. \tag{17}$$

In equation (17), we have set $\mathbf{y}^{-1} = 0$ because cutoffs are deterministic.[15]

V. Implementation

In this section we show that the optimal cutoffs can be implemented with posted prices as periods become short. The seller uses *posted prices* if she announces how many goods are remaining and charges a single price in each period. The buyers reveal their existence only 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.

The fact that prices are optimal is striking since there are many reasonable pricing tactics that might raise revenue. These include a series of auctions (e.g., Priceline), pricing as a function of the number of interested buyers (e.g., using flight query data), the issue of coupons when buyers register (e.g., Restaurant.com), or pricing as a function of buyers' indicative bids (e.g., house sales). Remarkably, theorem 3 proves that none of these tactics is useful in the benchmark model. This is far from obvious since all of these tactics are beneficial in variations of the model: Auctions are useful if the entry rate is discontinuous (see Sec. V.A), query-based pricing is useful if the number of entrants is public (see Sec. V.A), coupons are useful if different cohorts have different demand functions (see Sec. VI.D), and indicative bids are useful if buyers disappear over time (see Sec. VI.E). The traditional revenue management literature allows sellers only to charge posted prices; our analysis shows when this is without loss and when the seller can do better.[16]

[15] Since we know future cutoffs, the value functions in (17) can be calculated via the sequence problem (10) or the Bellman equation (12).

[16] While we show that implementation in continuous time is relatively easy, the problem is much harder in discrete time. With a single good, $K = 1$, the optimal cutoffs can be implemented via a sequence of second-price auctions (Board and Skrzypacz 2010). With more goods, Li (2011) shows that the seller can use a sequence of ascending auctions in which buyers compete against a robot that acts like the cutoff type. The basic problem in the discrete-time game is that more is known about older buyers' values, implying that a new and an 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.

The ordering of this section is parallel to that in Section IV. In Section V.A, we first consider general sequences of the demand process, showing that the optimal allocations can be implemented by prices. In Section V.B, we assume that the entry rate is weakly decreasing and show that the prices are given by an intuitive differential equation.

A. *General Case*

Suppose that time is continuous and, motivated by the law of rare events, buyers enter the market continuously according to a Poisson process with nonhomogeneous arrival rate λ_t. Let r be the instantaneous discount rate. Consider the discretized problem in which sales occur at discrete intervals of length h, agents arrive at rate $\int_t^{t+h}\lambda_s ds$, and the discount factor is $\delta = e^{-rh}$. Define $\Pi^*(h)$ as the optimal profits derived in Section IV.A and $\Pi^* = \lim_{h\to 0}\Pi^*(h)$ as the continuous-time profits.

THEOREM 3. Suppose that λ_t is Lipschitz continuous in t. If the seller uses posted prices with a second-price auction for the last unit at time T, then she can obtain $\Pi^* - O(h)$.

Proof. See the Appendix.

Theorem 3 is based on the fact that the cutoffs are deterministic, which means that 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 three 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, these issues arise only if there are two sales in a single period. Since cutoffs do not jump down much, the probability of two sales within any given period is $O(h^2)$, and the seller can obtain $\Pi^*(h) - O(h)$ for sufficiently small h. Third, a discrete-time seller can always replicate the strategy of the continuous-time seller delayed by at most one period, implying that $\Pi^* - \Pi^*(h) = O(h)$.

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 number of arrivals to the market or the reports of the buyers (else it would not be a posted price mechanism). It is also notable that the seller publicly announces her inventory, so she does not gain from keeping k private.[17] The general idea is that the seller wishes to implement cutoffs that depend only on (k, t); if buyers have any information that helps them predict when future units are sold (e.g., the timing of past sales, which are indicative of the number of waiting buyers), then the seller must condition prices on this information in order to "cancel it out."[18]

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 such auction by quickly declining prices, analogous to a Dutch auction. Similarly, one can replace the "final auction" for the last unit with a rapidly declining series of prices. Hence the firm's maximal profit can be approximated arbitrarily closely by a sales mechanism that uses only prices.[19]

The assumption of Poisson entry is more important since it implies that a buyer's entry time tells him 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 his entry time would give him information about other entrants' existence and even their values. For example, if buyers enter in pairs, then knowing that he entered earlier and no sale had occurred implies that a buyer's "partner" has a lower valuation.

Finally, we assume that the seller uses a second-price auction, but a first-price auction with reserve $e^{-rT}m^{-1}(0)$ 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 with increasing bidding strategies.

B. *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.

[17] 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.

[18] As an example, if buyers observed other buyers' entry into the market, this would provide information about the competition faced by a buyer, and so prices would also have to depend on N_t.

[19] Intuitively, consider the continuous-time limit and take any sequence of cutoffs x_t^k, which may include jumps. By Luzin's theorem, we can approximate this by continuous cutoffs that are almost identical almost everywhere; since the approximation is close in L^1, the lost profit (5) is small. The continuous cutoffs and Poisson entry imply that the probability of two simultaneous sales is zero; analogous to claim C of Sec. C in the Appendix, we can now implement these continuous cutoffs via prices.

First, consider optimal allocations. In period T, the optimal cutoffs are given by $m(x_T^k) = 0$. In period $t < T$, equation (17) becomes

$$rm(x_t^k) = \lambda_t E_v[\max\{m(v) - m(x_t^k), 0\} + \Pi_t^{k-1}(\min\{v, x_t^k\}) - \Pi_t^{k-1}(v)], \tag{18}$$

where E_v is the expectation over $v \sim F(\cdot)$. Equation (18) states that the seller is indifferent between selling today and delaying a little. The cost of delay is the forgone rental value (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.[20]

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 Sec. III). 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 over time.

The optimal cutoffs can be implemented by a sequence of decreasing prices p_t^k with an auction for the last unit in period T. These prices can be derived backward, starting at time T. When $k = 1$, the seller can use a second-price auction with reserve $e^{-rT}m^{-1}(0)$ 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

$$p_T^1 = e^{-rT}E_0\left[\max\{y^2, m^{-1}(0)\}|y^1 = \lim_{h\to 0} x_{T-h}^1, \{s_T(x)\}_{x\leq y^1}\right], \tag{19}$$

where $s_T(x)$ denotes the last time the cutoff went below x when looking back from time T.[21] 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 cares only 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,

[20] The proof of theorem 3 shows that the properties of the cutoffs as $t \to T$ do not depend on having weakly decreasing demand.

[21] That is, if $k(t)$ is the realized number of units left at time t, then $s_T(x) = \max\{t \leq T : x_t^{k(t)} \leq x\}$.

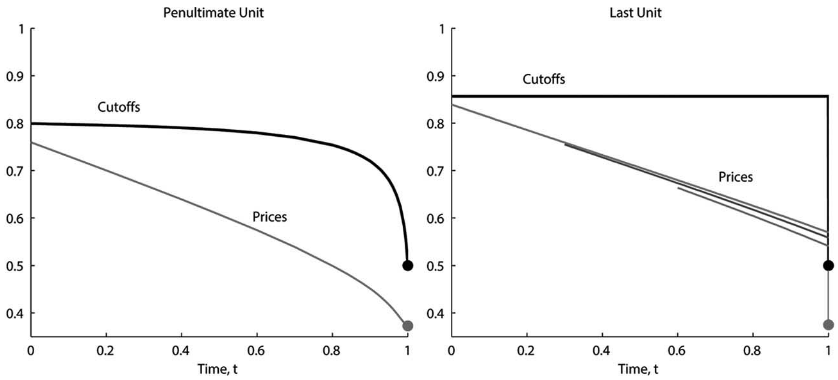

FIG. 3.—Optimal cutoffs and prices with two units. The left panel shows the optimal cutoffs/prices when the seller has two units remaining. The right panel shows the optimal cutoffs/prices when the the seller has one unit remaining. The three price lines illustrate the seller's strategy when she sells the penultimate unit at times $t = 0$, $t = 0.3$, and $t = 0.6$. In this figure, buyers enter continuously with Poisson parameter $\lambda = 10$, meaning that 10 interested buyers enter during an average season. They have values $v \sim U[0, 1]$, so the static monopoly cutoff is 0.5. Total time is $T = 1$ and the interest rate is $r = \ln(4/3)$, so a good loses one-fourth of its value over the season.

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 outcome, $x_t^k \to m^{-1}(0)$, as does the price, $p_t^k \to e^{-rT} m^{-1}(0)$.

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

$$\frac{dp_t^k}{dt} = -rx_t^k + \left[\frac{dx_t^k}{dt} f(x_t^k) \int_{s_t(x_t^k)}^{t} \lambda_s ds - \lambda_t [1 - F(x_t^k)]\right] \times [x_t^k - p_t^k - U_t^{k-1}(x_t^k, \{s_t(x)\}_{x \leq x_t^k})], \tag{20}$$

where $U_t^{k-1}(x_t^k, \{s_t(x)\}_{x \leq x_t^k})$ is the utility of buyer type x_t^k at time t when there are $k - 1$ goods left, conditional on x_t^k being the highest-value buyer at time t.[22] Intuitively, when a buyer waits a little, he gains from the falling prices (the left-hand side) but loses the rental value of the good and risks 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 stock-out 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.

[22] Buyer type x_t^k's utility depends on how much competition he believes he faces from existing buyers and hence depends on the history of cutoffs.

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 to make up for the rental value. But since he is also risking the arrival of new competition, the price has to fall faster. While cutoffs are deterministic, depending only 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 that there may be many other buyers in the market waiting for the price to drop, meaning that the seller can charge a higher price to implement the same cutoff. When there are two units (the left panel), the prices fall over time reflecting the declining cutoffs, buyers' impatience to buy 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 as the cutoff rapidly declines.

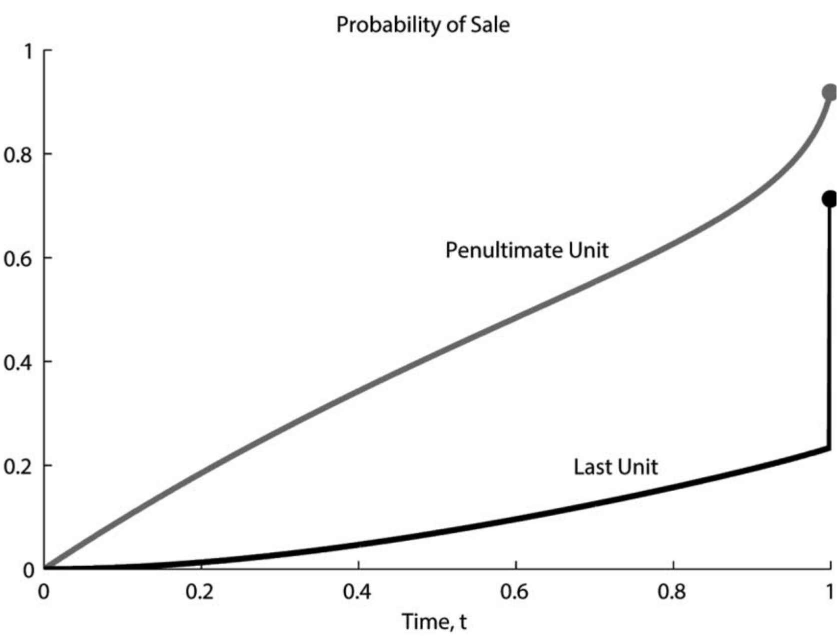

FIG. 4.—Probability of sale with two units. This figure shows the unconditional probability the the last/penultimate unit is sold by time t. The parameters are the same as in figure 3.

The existence of the last-minute sale comes from the concavity of the path of cutoffs (see fig. 3) and the stock of buyers building up, waiting to buy. This pattern of posted prices and a "last-minute" sale is qualitatively consistent with the sale of online ads or package holidays. Similarly, in the secondary market for baseball tickets, Sweeting (2012) shows that prices decline by 60 percent 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.

C. *Short-Lived versus Long-Lived Buyers*

Typical revenue management models assume that buyers are short-lived, 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 seller's 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.[23]

Figure 5 illustrates the optimal cutoffs/prices and the probability of sale when buyers are short-lived, under the same parameters as in figures 3 and 4. A first observation is that profits are higher when buyers are forward-looking.[24] This initially might seem surprising since forward-looking buyers can time their purchase to lower their payments. For example, fixing the retail prices, Soysal and Krishnamurthi (2012) found that profits for women's coats would be 9 percent higher if customers were short-lived. However, when the seller is choosing the optimal mechanism, the ability to delay means that 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 short-lived 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.

[23] One may also consider a third case: long-lived but myopic buyers, analogous to Lazear (1986). With a fixed population of buyers, the seller can fully extract by having prices start high and quickly fall; buyers then purchase as soon as their valuation equals the price. With buyers entering over time, the seller can maintain high prices most of the time and hold regular fast, deep sales to allocate units. This would extract the buyers' information rents while allowing the seller to approximately implement the welfare-optimal cutoffs.

[24] Proof: Since the arrival time is observable, the seller could replicate the short-lived allocation. Yet lemma 1 shows that it is optimal to treat all generations of buyers equally.

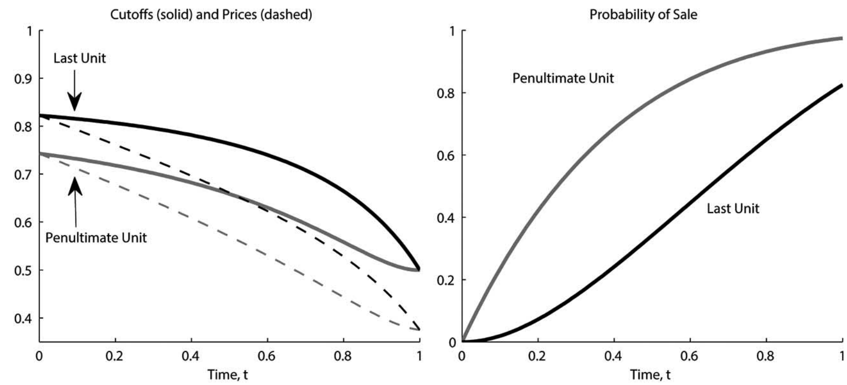

FIG. 5.—Short-lived buyers with two units. The left panel shows the optimal cutoffs and prices when the seller starts with two units and buyers are short-lived. The right panel shows the unconditional probability of sale of one/both units. The parameters are the same as in figure 3.

Third, sales occur later when buyers are forward-looking. When buyers are short-lived, sales occur fairly evenly throughout $[0, T]$, as seen in figure 5. When buyers are forward-looking, the combination of concave cutoffs (fig. 3 vs. fig. 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 short-lived buyers.[25] 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.

Overall, these results suggest that, when retailers are faced with forward-looking buyers, cutoffs are relatively constant and then drop rapidly, and sales are back-loaded. They also suggest that firms should encourage buyers to be forward-looking. This could mean having regular, predictable sales (e.g., Nordstrom's half-yearly sale) and notifying buyers when a sale is about to take place. Sellers could also embrace tools that help customers time their purchases including price prediction tools and price alerts (e.g., Kayak) and price-freezing options (e.g., United).

VI. Extensions

In this section we consider a number of extensions. This has the dual purpose of allowing us to explore the robustness of the main results as well as illustrating the applicability of the model.

[25] Proof: The forward-looking cutoffs exceed the short-lived 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 a sale is higher under myopia.

A. Inventory Costs

In the benchmark model we assume that impatience comes from proportional discounting. However, in some applications, a primary cost of delay comes from the cost of maintaining inventories. For example, a retailer prefers 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 that buyers do not discount over the shopping season, $\delta = 1$, but the seller pays a per-unit inventory cost, c_t; so a firm selling in period t gets profit $p_t - c_t$, where c_t increases in t.[26] We can adapt (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 $\Delta c_t \coloneqq c_{t+1} - c_t$ are Lipschitz continuous, then the optimal cutoffs can be implemented with prices (theorem 3).[27]

To see the effect of inventory costs, suppose that entry is decreasing and costs c_t are convex, so the marginal cost Δc_t weakly increases in t. Adapting (17), the one-period-look-ahead property implies that cutoffs are determined by

$$\begin{aligned} m(x_t^k) + E_{t+1}[\Pi_{t+1}^{k-1}(\mathbf{v}_{t+1})] = {} & E_{t+1}[\max\{m(x_t^k), m(v_{t+1}^1)\}] \\ & + E_{t+1}[\Pi_{t+1}^{k-1}(\{x_t^k, \mathbf{v}_{t+1}\}_k^2)] - \Delta c_t \end{aligned}$$

for $t < T$, with $m(x_T^k) = -\Delta c_T$. The resulting cutoffs are decreasing over time because the 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 (18) to obtain

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

[26] We eliminate discounting for simplicity; one could include both inventory costs and discounting.

[27] 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 parts b and c of lemma 3 as $\mathbf{1}_{\tau \le T}$; this is still strictly less than one in expectation 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.

Adapting (20), prices are then determined by the differential equation

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

with boundary condition analogous to (19). 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.

B. *Arriving and Expiring Supply*

In the benchmark model, there are K units of a good that can be sold over time $\{1, \ldots, T\}$. However, in some applications, supply arrives and departs over time. For example, consider the Fulton Fish Market, where dealers must sell their fish to customers arriving over time prior to the end of the week (Graddy 2006). During the week, new stock arrives, largely determined by weather conditions, while fish expire after a few days, resulting in an exogenous death process. A similar issue arises with airlines or display ads, where the seller sells a sequence of goods with different flight/broadcast times.

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 (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 in which 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).[28] 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 weakly decreasing

[28] 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 on this result.

over time and a_t, b_t are deterministic, then cutoffs fall over time and the one-period-look-ahead property holds (theorem 2).[29] 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, one can use prices to implement the optimal allocations (theorem 3) with an auction for the last unit to expire in a period, since this is when the cutoffs jump down. However, as discussed in Section V.A, one can approximate these auctions with posted prices analogous to a Dutch auction.

C. *Third-Degree Price Discrimination*

The benchmark model assumes that all customers look alike to the seller. However, in some applications, firms can divide their customers into multiple groups. For example, Yahoo! sells ad space to movie studios that demand rich media ads (e.g., video, flash) and to insurance companies that buy static display ads. Similarly, Graddy (2006) discusses how dealers in the Fulton Fish Market discriminate between white and Asian buyers, who often resell to very different markets and therefore have different demand functions.

To model group pricing, suppose that buyers of rich and static ads are drawn from distributions f_R and f_S, inducing marginal revenues m_R and m_S, and the seller knows from which distribution a given buyer is drawn. The seller has K units that she 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 $\sum_i \mathbf{1}_{\tau_i \le T} \le K$, 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 the 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 entrants is 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).

[29] For theorem 2, the key step is to observe that if future cutoffs and N_t are decreasing over time, then $\tau^{k-1}_{t+1}(z) - (t+1) \ge \tau^{k-1}_{t+2}(z) - (t+2)$ in lemma 4, even if units expire at different times. This means that the option value of delay falls over time, along with the cutoff.

The seller can then implement the optimal cutoffs with two different price paths for the two types of buyers (theorem 3). These are related through the inventory, so a sale of a static ad raises the prices for both types of buyers.[30]

D. *Changing Distributions of Incoming Values*

The benchmark model assumes that demand is IID, no matter when a buyer enters the market. However, in some applications, the distribution of entrants' values may change over time. For example, in the airline market, McAfee and Te Velde (2006) observe that prices rise in the last few days before the flight date, suggesting that late-arriving buyers have higher values.[31] If we suppose that buyers who enter in period t have distribution F_t, each generation is associated with a different marginal revenue and equation (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 VI.C, the seller would thus like to bias allocation toward 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 $e^{-rt}\mu_{t_i}$. Even if the seller cannot discriminate between different cohorts, this mechanism is still incentive compatible if the distribution F_t gets stronger over time, as with airlines. In this case, the seller would like to bias allocation toward earlier generations, and the intergenerational 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 implement

[30] Theorem 1 relies on the fact that a buyer's allocation depends only on his marginal revenue and his rank; the distribution and arrival rate of marginal revenues can vary over time, as long as they are independent of past realizations (which ensures that the marginal revenues are a sufficient state variable). The proof of theorem 2 is identical. For theorem 3, observe that the cutoffs in marginal revenue space correspond to different cutoffs in value space for the different buyers. These are Lipschitz continuous if λ_i is Lipschitz continuous and can be implemented in prices, with a different price sequence for each type of buyer.

[31] Lazarev (2013) estimates such a model with "business" and "leisure" customers and finds that, in order to justify realized price paths, business customers constitute a quarter of the market, have values up to six times higher than those of leisure customers, and tend to enter in the market in the last week.

the optimal mechanism by issuing a coupon worth $e^{-rt}(\mu_T - \mu_{t_i})$ to a buyer who registers for a flight in period t_i and buys at time t.[32]

In some 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 long-lived, forward-looking buyers may now lower the firm's profits. When buyers become long-lived, the seller gains from the option value to delay selling a unit but loses from intergenerational price discrimination becoming harder. If the distribution F_t gets stronger over time, profits are still higher with long-lived buyers because a delaying buyer has a higher marginal revenue than a younger buyer.[33] However, if the distribution F_t grows weaker over time, then profits may be lower with long-lived customers as stronger generations delay to merge with weaker generations.

E. *Disappearing Buyers*

The benchmark model assumes that buyers are long-lived, remaining in the market until they buy. However, in some applications, it would seem natural to allow buyers to exit probabilistically over time. For example, when an owner wishes to sell her house, the potential buyers will disappear if they purchase another property.

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 that 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

[32] If demand weakens over time, then the intergenerational IC constraints will bind. This is an interesting problem but is beyond the scope of this paper.

[33] Proof: Fix some optimal cutoffs with short-lived buyers x_t^k. Now, suppose that 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 have just increased the number of buyers in each period, this raises the seller's expected profit. Next, suppose that this buyer contributed marginal revenue $m_s(v) \geq m_t(v)$ if he buys 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.

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 factors, $\delta \in \{\delta_L, \delta_H\}$, 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 indicative bidding 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.

VII. Conclusion

We have considered a seller who wishes to sell multiple goods by a deadline 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, while prices are characterized by an intuitive differential equation.

The analysis provides a natural benchmark for a number of markets, such as retailing, ad auctions, and airlines. For retailing, our results make predictions about when retailers should initiate sales and how quickly prices should fall. If consumers enter uniformly over time, then, fixing the inventory, prices should be decreasing and concave. If consumers are very aware of when new products are released, entry will decrease over time and prices should fall even faster. And if the good is faddish, then agents are more impatient, sales are more front-loaded, and prices will decline faster. For ad auctions, our results make predictions about the price paths and popularity of ad exchanges used for last-minute excess capacity. Such ad exchanges are large and anonymous, selling slots on the basis of demographics and giving rise to the seller's reservation value, v_0. Our results then show that the price in the forward market should always converge to the same final price, independent at the remaining inventory, and that this last-minute price exceeds the clearing price in the exchange. Moreover, as prebooking becomes more important for buyers (e.g., movie releases), then sales shift forward, ad exchanges become less popular, and prices for a fixed inventory fall faster. For airlines, evidence suggests that the distribution of values increases over time (e.g., Lazarev 2013), meaning that anonymous prices are no longer optimal. In such a situation, our model indicates that the seller can raise profits by using a coupon-based mechanism, whereby buyers who register earlier get larger coupons giving them a discount on the market price.

Our analysis helps us understand broader features of pricing data. If buyers are long-lived, then price declines lead to temporary bursts of sales, price paths exhibit more last-minute discounts, sales are more back-loaded, and profits are higher than in a world with short-lived buyers. The results thus help researchers identify whether buyers are long- or short-lived in particular markets and assess the incentives for firms to encourage such behavior via “price alerts,” “contingent contracts,” or other mechanisms. The paper characterizes the optimal price path under commitment, and so also helps one understand how much commitment power firms possess. As a comparison, without commitment, one would expect the seller to dispose of much more of her inventory by date T and for prices to cycle as the seller discards excess capacity (Conlisk, Gerstner, and Sobel 1984; Dilme and Li 2012). Finally, by showing that the optimal mechanism can be implemented via posted prices, we help explain the limited appeal of more complicated auction-like mechanisms and also clarify when they are of use (e.g., when there is a mix of short- and long-lived buyers).

While our model provides a natural benchmark, 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 rate at which buyers enter, implying that 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, which are indicative of future entry. Indeed, Gershkov, Moldovanu, and Strack (2013) solve such a model and show that if buyers arrive according to a Poisson process, then it is incentive compatible for buyers to announce their arrival date truthfully. Second, for some markets, one should interpret our model of proportional discounting as a reduced form for “attention costs” or “coordination costs.” It would be interesting to model this 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, studying how buyers trade off intratemporal decisions (e.g., which type of ad to buy) and intertemporal decisions (when to buy).

Appendix

A. Proof of Lemma 3

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 whether or not to sell a unit at time t, y^j makes the same contribution to profits (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 (14),

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

Using equation (13) 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 is 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 that $\{x_s^k\}_{s\geq t+1}$ are deterministic and decreasing in k. We first prove a preliminary result. Let $\{y^1, \ldots, y^k\}$ and $\{\tilde{y}^1, \ldots, \tilde{y}^k\}$ be ordered vectors, where $y^j \geq \tilde{y}^j$ for each j. We claim that for time $\sigma \geq t+1$,

$$\begin{aligned}
&\Pi_\sigma^k(y^1, \ldots, y^k) - \Pi_\sigma^k(\tilde{y}^1, \ldots, \tilde{y}^k) \\
&\quad = E_{\sigma+1}\left[\delta^{-\sigma}\int_{\{\tilde{y}^1,\ldots,\tilde{y}^k\}}^{\{y^1,\ldots,y^k\}} [m'(z^1)\delta^{\tau_1^k(z^1)}, \ldots, m'(z^k)\delta^{\tau_k^k(z^k)}]\,d(z^1, \ldots, z^k)\right] \\
&\quad \geq E_{\sigma+1}\left[\delta^{-\sigma}\int_{\{\tilde{y}^1,\ldots,\tilde{y}^k\}}^{\{y^1,\ldots,y^k\}} [m'(z^1)\delta^{\tau_1^{k-1}(z^1)}, \ldots, m'(z^k)\delta^{\tau_k^{k-1}(z^k)}]\,d(z^1, \ldots, z^k)\right] \\
&\quad = \Pi_\sigma^{k-1}(y^1, \ldots, y^k) - \Pi_\sigma^{k-1}(\tilde{y}^1, \ldots, \tilde{y}^k).
\end{aligned} \tag{A1}$$

The first line applies the envelope theorem to equation (11), where τ_j^k is the purchasing time of the buyer in the jth 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 kth good. The final line again uses the envelope theorem.

Suppose that 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 jth-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 $\mathbf{v}_{(\mathbf{t},\sigma]} \coloneqq \{v_{(t,\sigma]}^1, \ldots, v_{(t,\sigma]}^k\}$, let $\{\mathbf{v}_{(\mathbf{t},\sigma]}\}_{k-1}^1$ be the ordered vector of the largest $k-1$ elements from $\mathbf{v}_{(\mathbf{t},\sigma]}$, and let $\{y^1, \mathbf{v}_{(\mathbf{t},\sigma]}\}_k^2$ be the second- to kth-highest choices from $\{y^1, \mathbf{v}_{(\mathbf{t},\sigma]}\}$. We claim that

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

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 (A1) and the fact that $\{y^1, \mathbf{v}_{(\mathbf{t},\sigma]}\}_{k-1}^1$ is pointwise larger than $\{y^1, \mathbf{v}_{(\mathbf{t},\sigma]}\}_k^2$. 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

$$\delta^{t+1}\tilde{\Pi}_{t+1}^{k-2}(\varnothing) = E_{t+1}[\delta^\sigma\Pi_\sigma^{k-2}(\{\mathbf{v}_{(\mathbf{t},\sigma]}\}_{k-1}^1)]$$

and

$$\delta^{t+1}\tilde{\Pi}_{t+1}^{k-1}(y^1) \geq E_{t+1}[\max\{m(y^1), m(v_{(t,\sigma]}^1)\} + \Pi_\sigma^{k-2}(\{y^1, \mathbf{v}_{(\mathbf{t},\sigma]}\}_k^2)].$$

QED

B. 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

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

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

$$\Pi_{t+2}^{k-1}(\{\mathbf{v}_{\mathbf{t+2}}\}_{k-1}^1) - \max\{m(y^1), m(v_{t+2}^1)\} - \Pi_{t+2}^{k-1}(\{y^1, \mathbf{v}_{\mathbf{t+2}}\}_k^2) \tag{A2}$$

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

$$m'(v_{t+2}^j)[\delta^{\hat{\tau}_j^1(v_{t+2}^j)} - \delta^{\hat{\tau}_j^0(v_{t+2}^j)}]\delta^{-(t+2)}, \tag{A3}$$

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 that future cutoffs are deterministic, $\hat{\tau}_j^0(v_{t+2}^j) \leq \hat{\tau}_j^1(v_{t+2}^j)$ and (A3) is negative. Hence (A2) is decreasing in $\mathbf{v}_{\mathbf{t+2}}$.

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 $\mathbf{v}_{\mathbf{t+2}}$ with $\hat{\mathbf{v}}_{\mathbf{t+2}}$ in $D\Pi^k_{t+1}(y^1)$, define

$$\begin{aligned}\hat{D}\Pi^k_{t+1}(y^1) = m(y^1) + \delta E_{t+2}[\Pi^{k-1}_{t+2}(\{\hat{\mathbf{v}}_{\mathbf{t+2}}\}^1_{k-1})] \\ - \delta E_{t+2}[\max\{m(y^1), m(\hat{v}^1_{t+2})\} + \Pi^{k-1}_{t+2}(\{y^1, \hat{\mathbf{v}}_{\mathbf{t+2}}\}^2_k)].\end{aligned}$$

Since N_t is decreasing in the usual stochastic order, $\hat{\mathbf{v}}_{\mathbf{t+2}}$ exceeds $\mathbf{v}_{\mathbf{t+2}}$ in the usual stochastic order, and since (A2) is decreasing in $\mathbf{v}_{\mathbf{t+2}}$, $D\Pi^k_{t+1}(y^1) \geq \hat{D}\Pi^k_{t+1}(y^1)$. Intuitively, the seller has more to gain from selling today if there are fewer entrants tomorrow.

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

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

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

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

Recall that 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 $\mathbf{v}_{\mathbf{t+1}}$ have the same distribution, we can suppose $\hat{\mathbf{v}}_{\mathbf{t+2}} = \mathbf{v}_{\mathbf{t+1}}$. If $\tau^{k-1}_{t+1}(z) = s$ for $s < T$, then $\tau^{k-1}_{t+2}(z) \leq s+1$ since future cutoffs decrease in t and N_t falls over time.[34] In addition, if $\tau^{k-1}_{t+1}(z) = T$, then $\tau^{k-1}_{t+2}(z) \leq T$ since more entrants enter over time. Putting this together, $\tau^{k-1}_{t+1}(z) - (t+1) \geq \tau^{k-1}_{t+2}(z) - (t+2)$ for all z. Taking expectations over the distribution of entrants, the integral equations then imply that $\hat{D}\Pi^k_{t+1}(y^1) \geq D\Pi^k_t(y^1)$. 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.

C. *Proof of Theorem 3*

This proof consists of several steps. For a small time interval h, lemma 5 shows that cutoffs do not jump down more than αh in any period $t < T$. Lemma 6 demonstrates that the cutoffs also do not jump down in the last period as long as $k \geq 2$. Finally, lemma 7 shows that the seller can use posted prices to obtain the profits from the optimal mechanism, Π^*.

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

Lemma 5. For $t \le T - 2h$, there exist positive constants α, h_0 such that $x_t^k - x_{t+h}^k \le \alpha h$ for $h < h_0$.

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

$$\begin{aligned}\Pi_t^k(\text{sell 1 today}) &= m(y^1) + e^{-rh}e^{-\Lambda_{t+h}}\Pi_{t+h}^{k-1}(\varnothing)\\ &\quad + e^{-rh}(1 - e^{-\Lambda_{t+h}})E_{t+h|N_{t+h}\ge 1}[\Pi_{t+h}^{k-1}(\mathbf{v_{t+h}})],\\ \Pi_t^k(\text{sell} \ge 1 \text{ tomorrow}) &= e^{-rh}e^{-\Lambda_{t+h}}[m(y^1) + \Pi_{t+h}^{k-1}(\varnothing)]\\ &\quad + e^{-rh}(1 - e^{-\Lambda_{t+h}})E_{t+h|N_{t+h}\ge 1}[\max\{m(y^1), m(v_{t+h}^1)\}\\ &\quad + \Pi_{t+h}^{k-1}(\{y^1, \mathbf{v_{t+h}}\}_h^2)].\end{aligned}$$

Subtracting the second line from the first,

$$\begin{aligned}D\Pi_t^k(y^1) = {}&(1 - e^{-rh})m(y^1) + e^{-rh}(1 - e^{-\Lambda_{t+h}})E_{t+h|N_{t+h}\ge 1}[m(y^1) + \Pi_{t+h}^{k-1}(\mathbf{v_{t+h}})\\ &- \max\{m(y^1), m(v_{t+h}^1)\} - \Pi_{t+h}^{k-1}(\{y^1, \mathbf{v_{t+h}}\}_k^2)],\end{aligned} \tag{A4}$$

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

For point 1, let $\underline{m}' \coloneqq \inf_{v\in[\underline{v},\bar{v}]} m'(v)$; this is strictly positive because $m(v)$ is strictly increasing and continuously differentiable. Differentiating (A4),

$$\frac{d}{dy^1}D\Pi_t^k(y^1) \ge (1 - e^{-rh})m'(y^1) \ge \frac{1}{2}rh\underline{m}' \tag{A5}$$

for $h \le h_0 \coloneqq (\ln 2)/r$.

For point 2, note that $\Pr(N_t = 1) = \Lambda_t e^{-\Lambda_t}$ and $\Pr(N_t \ge 2) = 1 - e^{-\Lambda_t}(1 + \Lambda_t) \le \Lambda_t^2 \le \bar{\lambda}^2 h^2$, using the fact that $1 - e^{-x} \le x$ for $x \ge 0$ and $\bar{\lambda} \coloneqq \max_{t\in[0,T]}\lambda_t$. Splitting (A4) into the case in which there is one entrant and that in which there are multiple entrants,

$$\begin{aligned}D\Pi_t^k(y^1) \ge {}&(1 - e^{-rh})m(y^1) + e^{-rh}\Lambda_{t+h}e^{-\Lambda_{t+h}}E_v[m(y^1) + \Pi_{t+h}^{k-1}(v)\\ &- \max\{m(y^1), m(v)\} - \Pi_{t+h}^{k-1}(\min\{y^1, v\})] - \bar{\lambda}^2 m(\bar{v})h^2,\end{aligned} \tag{A6}$$

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

$$\begin{aligned}D\Pi_{t+h}^k(y^1) \le {}&(1 - e^{-rh})m(y^1) + e^{-rh}\Lambda_{t+2h}e^{-\Lambda_{t+2h}}E_v[m(y^1) + \Pi_{t+2h}^{k-1}(v)\\ &- \max\{m(y^1), m(v)\} - \Pi_{t+2h}^{k-1}(\min\{y^1, v\})].\end{aligned}$$

Subtracting these and completing the square gives us

$$\begin{aligned}D\Pi_{t+h}^k(y^1) - D\Pi_t^k(y^1) \le {}&e^{-rh}(\Lambda_{t+2h}e^{-\Lambda_{t+2h}} - \Lambda_{t+h}e^{-\Lambda_{t+h}})\\ &\times E_v[m(y^1) + \Pi_{t+h}^{k-1}(v) - \max\{m(y^1), m(v)\} - \Pi_{t+h}^{k-1}(\min\{y^1, v\})]\\ &+ e^{-rh}\Lambda_{t+2h}e^{-\Lambda_{t+2h}}E_v[(\Pi_{t+2h}^{k-1}(v) - \Pi_{t+2h}^{k-1}(\min\{y^1, v\}))\\ &- (\Pi_{t+h}^{k-1}(v) - \Pi_{t+h}^{k-1}(\min\{y^1, v\}))] + \bar{\lambda}^2 m(\bar{v})h^2.\end{aligned}$$

Consider the first term on the right-hand side. If $\Lambda_{t+2h} \geq \Lambda_{t+h}$, 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

$$\begin{aligned}(\Lambda_{t+h} e^{-\Lambda_{t+h}} - \Lambda_{t+2h} e^{-\Lambda_{t+2h}}) m(\bar{v}) &\leq (\Lambda_{t+h} - \Lambda_{t+2h})(1 - \Lambda_{t+2h}) e^{-\Lambda_{t+2h}} m(\bar{v}) \\ &\leq (\tilde{\lambda}_{t+h} - \tilde{\lambda}_{t+2h}) m(\bar{v}) h^2 \leq 2\beta m(\bar{v}) h^2,\end{aligned}$$

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

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

using the envelope theorem as in (A1), where $\tau_{t+h}^{k-1}(z)$ is the purchasing time of the single buyer present at time $t+h$. Subtracting these two integrals, we claim that the second term is

$$\begin{aligned}&e^{-rh} \Lambda_{t+2h} e^{-\Lambda_{t+2h}} \left\{ E_{t+3h}\left[\int_{\min\{y^1,v\}}^{v} m'(z) e^{-r(\tau_{t+2h}^{k-1}(z) - t - 2h)} dz\right]\right. \\ &\qquad \left. - E_{t+2h}\left[\int_{\min\{y^1,v\}}^{v} m'(z) e^{-r(\tau_{t+h}^{k-1}(z) - t - h)} dz\right]\right\} \\ &\qquad \leq \bar{\lambda} h[(1 - e^{-\Lambda_{t+2h}}) m(\bar{v}) + e^{-\Lambda_{t+2h}}(1 - e^{-rh}) m(\bar{v})] \\ &\qquad \leq \bar{\lambda}(\bar{\lambda} + r) m(\bar{v}) h^2.\end{aligned}$$

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

$$D\Pi_{t+h}^{k}(y^1) - D\Pi_{t}^{k}(y^1) \leq (2\beta + 2\bar{\lambda}^2 + \bar{\lambda} r) m(\bar{v}) h^2. \tag{A7}$$

To finish the proof, suppose that $x_{t+h}^k \leq x_t^k$, else there is nothing to prove. We now claim that

$$\begin{aligned}(2\beta + 2\bar{\lambda}^2 + \bar{\lambda} r) m(\bar{v}) h^2 &\geq D\Pi_{t+h}^{k}(x_t^k) - D\Pi_{t}^{k}(x_t^k) \\ &= \int_{x_t^k+h}^{x_t^k} \frac{d}{dy^1} D\Pi_{t+h}^{k}(y^1) dy^1 + D\Pi_{t+h}^{k}(x_{t+h}^k) \\ &\geq (x_t^k - x_{t+h}^k) \frac{1}{2} rh\underline{m}'\end{aligned}$$

for $h \le h_0$. The first inequality comes from (A7), the second line uses $D\Pi_t^k(x_t^k) = 0$, and the third lines uses $D\Pi_{t+h}^k(x_{t+h}^k) \ge \Delta\Pi_{t+h}^k(x_{t+h}^k) = 0$ and (A5). Rearranging then implies that there exists $\alpha > 0$ such that $(x_t^k - x_{t+h}^k) \le \alpha h$ for $h \le h_0$. QED

Lemma 6. If $k \ge 2$, there exist positive constants α, h_0 such that $x_{T-h}^k - x_T^k \le \alpha h$ for $h \le h_0$.

Proof. In period $t = T - h$, if $m(y^1) \ge 0$, then (A6) becomes

$$D\Pi_{T-h}^k(y^1) \ge (1 - e^{-rh})m(y^1) - \bar{\lambda}^2 m(\bar{v})h^2, \tag{A8}$$

since the term in brackets in (A6) is zero.

In period T, we have $x_T^k = m^{-1}(0)$. Since the seller will never sell to a buyer with negative marginal revenue, we have $x_{T-h}^k \ge x_T^k$. We now claim that

$$\begin{aligned}\bar{\lambda}^2 m(\bar{v})h^2 &\ge D\Pi_{T-h}^k(x_{T-h}^k) - D\Pi_{T-h}^k(x_T^k)\\ &= \int_{x_T^k}^{x_{T-h}^k} \frac{d}{dy^1} D\Pi_{T-h}^k(y^1)\,dy^1\\ &\ge (x_{T-h}^k - x_T^k)\frac{1}{2} rh\underline{m}'\end{aligned}$$

for $h \le h_0$. The first line uses (A8), $m(x_T^k) = 0$, and $D\Pi_{T-h}^k(x_{T-h}^k) = 0$. The second line follows from the fundamental theorem of calculus. The third line uses (A5). Rearranging yields the result. QED

Lemma 7. The firm can obtain profits $\Pi^* - O(h)$ by using posted prices with a second-price auction for the last unit at time T.

Proof. We use the following mechanism: In each period the seller chooses a price p_t^k 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 $e^{-rT}m^{-1}(0)$. 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^κ, such that buyers wish to buy if their value exceeds the cutoff, where κ 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

$$u(v, t, \tau) = E_0[v\mathbf{1}_{\tau \ge t}e^{-r\tau} - p_\tau]. \tag{A9}$$

Here, the price p_t is a random variable, depending on the sales to other buyers. If other buyers demand as many as or more units than the seller has to offer, the price may also rise to infinity 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 $x_t^\kappa = \inf\{v : \tau^*(v, t) = t\}$ be the lowest type who wishes to buy in period t.

Conversely, we claim that any sequence of cutoffs can be implemented by prices. These prices can be constructed by backward induction (e.g., Kruse and Strack

2014). Alternatively, one can consider the utility of a buyer with type x_t^κ who enters at time t,

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

where the left-hand side is his direct utility, and the right-hand side comes from applying the envelope theorem to utility (A9). 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 $\{x_t^\kappa\}$ and the rationing rule.

We next claim that a posted price mechanism attains profits $\Pi^*(h) - O(h)$, where $\Pi^*(h)$ is the seller's profits from the optimal mechanism in the discretized problem. To do this, consider the case in which $t < T$ or $t = T$ and $k \geq 2$ and assume that the seller chooses the prices so that the induced cutoff x_t^κ coincides with the optimal cutoffs when starting the period with κ units. This will implement the wrong allocation only if two buyers wish to buy in a single period, in which case the loss is bounded by $m(\overline{v})$. To show that two sales occur with small probability, 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 nonhomogeneous Poisson process in which the integral of the arrival rate is, for $h \leq h_0$,

$$\begin{aligned}\Phi_t &= \int_{x_t^\kappa}^{x_{t-h}^\kappa}\int_{s_t(z)}^{t-h}\lambda_s\,ds\,dF(z) + \left(\int_{t-h}^{t}\lambda_s\,ds\right)[1 - F(x_t^\kappa)]\\ &\leq \bar{\lambda}T\bar{f}\alpha h + \bar{\lambda}h =: \gamma h,\end{aligned}$$

where the first term captures sales from existing buyers, and the second term captures sales from new buyers. The inequality uses lemmas 5 and 6, and $\overline{f}$ is the upper bound on the continuous density. The probability of two or more sales over $(t - h, t]$ is $1 - e^{\Phi_t}(1 + \Phi_t) \leq \Phi_t^2 \leq \gamma^2 h^2$. The probability of two or more sales in any period is thus bounded above

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

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, as in the optimal mechanism.

The seller can thus attain profits $\Pi^*(h) - O(h)$ with posted prices. Since the discrete-time seller can mimic the continuous-time seller with a delay of at most one period, we have $\Pi^* - \Pi^*(h) \leq rh\Pi^*$. Putting these observations together implies that the price mechanism obtains $\Pi^* - O(h)$. QED

References

Aviv, Y., and A. Pazgal. 2008. "Optimal Pricing of Seasonal Products in the Presence of Forward-Looking Consumers." *Manufacturing and Service Operations Management* 10 (3): 339–59.

Bertsekas, D. P. 2005. *Dynamic Programming and Optimal Control*. Vol. 1. 3rd ed. Belmont, MA: Athena Scientific.
Board, S. 2007. "Selling Options." *J. Econ. Theory* 136 (1): 324–40.
———. 2008. "Durable-Goods Monopoly with Varying Demand." *Rev. Econ. Studies* 75 (2): 391–413.
Board, S., and A. Skrzypacz. 2010. "Revenue-Management with Forward-Looking Buyers." Working paper, Univ. California, Los Angeles.
Bruss, F. T., and T. S. Ferguson. 1997. "Multiple Buying or Selling with Vector Offers." *J. Appl. Probability* 34:959–73.
Conlisk, J., E. Gerstner, and J. Sobel. 1984. "Cyclic Pricing by a Durable Goods Monopolist." *Q.J.E.* 99 (3): 489–505.
Das Varma, G., and N. Vettas. 2001. "Optimal Dynamic Pricing with Inventories." *Econ. Letters* 72 (3): 335–40.
Deneckere, R., and J. Peck. 2012. "Dynamic Competition with Random Demand and Costless Search: A Theory of Price Posting." *Econometrica* 80 (3): 1185–1247.
Dilme, F., and F. Li. 2012. "Revenue Management without Commitment: Dynamic Pricing and Periodic Fire Sales." Working paper, Univ. Pennsylvania.
Einav, L., C. Farronato, J. Levin, and N. Sundaresan. 2013. "Sales Mechanisms in Online Markets: What Happened to Internet Auctions?" Working paper, Stanford Univ.
Elmaghraby, W., A. Gülcü, and P. Keskinocak. 2008. "Designing Optimal Preannounced Markdowns in the Presence of Rational Customers with Multiunit Demands." *Manufacturing and Service Operations Management* 10 (1): 126–48.
Fuchs, W., and A. Skrzypacz. 2010. "Bargaining with Arrival of New Traders." *A.E.R.* 100 (3): 802–36.
Gallego, G., and G. van Ryzin. 1994. "Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons." *Management Sci.* 40 (8): 999–1020.
Gallien, J. 2006. "Dynamic Mechanism Design for Online Commerce." *Operations Res.* 54 (2): 291–310.
Gershkov, A., and B. Moldovanu. 2009a. "Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach." *American Econ. J.: Microeconomics* 1 (2): 168–98.
———. 2009b. "Learning about the Future and Dynamic Efficiency." *A.E.R.* 99 (4): 1576–88.
Gershkov, A., B. Moldovanu, and P. Strack. 2013. "Dynamic Allocation and Learning with Strategic Arrivals." Working paper, Univ. Bonn.
Gowrisankaran, G., and M. Rysman. 2012. "Dynamics of Consumer Demand for New Durable Goods." *J.P.E.* 120 (6): 1173–1219.
Graddy, K. 2006. "Markets: The Fulton Fish Market." *J. Econ. Perspectives* 20 (2): 207–20.
Hörner, J., and L. Samuelson. 2011. "Managing Strategic Buyers." *J.P.E.* 119 (3): 379–425.
Kruse, T., and P. Strack. 2014. "Optimal Stopping with Private Information." Working paper, Univ. California, Berkeley.
Lazarev, J. 2013. "The Welfare Effects of Intertemporal Price Discrimination: An Empirical Analysis of Airline Pricing in U.S. Monopoly Markets." Working paper, New York Univ.
Lazear, E. P. 1986. "Retail Pricing and Clearance Sales." *A.E.R.* 76 (1): 14–32.
Li, A. 2011. "Selling Storable Goods to a Dynamic Population of Buyers: A Mechanism Design Approach." Working paper, Univ. Washington, St. Louis.

MacQueen, J., and R. Miller. 1960. "Optimal Persistence Policies." *Operations Res.* 8 (3): 362–80.
Malighetti, P., S. Paleari, and R. Redondi. 2009. "Pricing Strategies of Low-Cost Airlines: The Ryanair Case Study." *J. Air Transport Management* 15 (4): 195–203.
Mas-Colell, A., M. D. Whinston, and J. R. Green. 1995. *Microeconomic Theory*. Oxford: Oxford Univ. Press.
McAfee, R. P., and J. McMillan. 1987. "Auctions with a Stochastic Number of Bidders." *J. Econ. Theory* 43 (1): 1–19.
———. 1988. "Search Mechanisms." *J. Econ. Theory* 44 (1): 99–123.
McAfee, R. P., and V. Te Velde. 2006. "Dynamic Pricing in the Airline Industry." In *Economics and Information Systems*, edited by Terrence Hendershott. Amsterdam: Elsevier.
Mierendorff, K. 2009. "Optimal Dynamic Mechanism Design with Deadlines." Working paper, Univ. Bonn.
Myerson, R. B. 1981. "Optimal Auction Design." *Math. Operations Res.* 6 (1): 58–73.
———. 1986. "Multistage Games with Communication." *Econometrica* 54:323–58.
Pai, M., and R. Vohra. 2013. "Optimal Dynamic Auctions and Simple Index Rules." *Math. Operations Res.* 38 (4): 682–97.
Robinson, D. 2014. "J. C. Penney: Activist Investors and the Rise and Fall of Ron Johnson." Case series, Haas School, Univ. California, Berkeley.
Ross, S. M. 1971. "Infinitesimal Look-Ahead Stopping Rules." *Ann. Math. Statis.* 42 (1): 297–303.
Said, M. 2012. "Auctions with Dynamic Populations: Efficiency and Revenue Maximization." *J. Econ. Theory* 147 (6): 2419–38.
Segal, I. 2003. "Optimal Pricing Mechanisms with Unknown Demand." *A.E.R.* 93 (3): 509–29.
Soysal, G. P., and L. Krishnamurthi. 2012. "Demand Dynamics in the Seasonal Goods Industry: An Empirical Analysis." *Marketing Sci.* 31 (2): 293–316.
Stokey, N. L. 1979. "Intertemporal Price Discrimination." *Q.J.E.* 93 (3): 355–71.
Su, X. 2007. "Intertemporal Pricing with Strategic Customer Behavior." *Management Sci.* 53 (5): 726–41.
Sweeting, A. 2012. "Dynamic Pricing Behavior in Perishable Goods Markets: Evidence from Secondary Markets for Major League Baseball Tickets." *J.P.E.* 120 (6): 1133–72.
Talluri, K., and G. van Ryzin. 2004. *The Theory and Practice of Revenue Management*. Boston: Springer.
Topkis, D. M. 1998. *Supermodularity and Complementarity*. Princeton, NJ: Princeton Univ. Press.