

Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons



Guillermo Gallego; Garrett van Ryzin

Management Science, Vol. 40, No. 8 (Aug., 1994), 999-1020.

Stable URL:

<http://links.jstor.org/sici?sici=0025-1909%28199408%2940%3A8%3C999%3AODPOIW%3E2.0.CO%3B2-N>

Management Science is currently published by INFORMS.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/informs.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons

Guillermo Gallego • Garrett van Ryzin

Department of Industrial Engineering and Operations Research, Columbia University, New York, New York 10027
Graduate School of Business, Columbia University, New York, New York 10027

In many industries, managers face the problem of selling a given stock of items by a deadline. We investigate the problem of dynamically pricing such inventories when demand is price sensitive and stochastic and the firm's objective is to maximize expected revenues. Examples that fit this framework include retailers selling fashion and seasonal goods and the travel and leisure industry, which markets space such as seats on airline flights, cabins on vacation cruises, and rooms in hotels that become worthless if not sold by a specific time.

We formulate this problem using intensity control and obtain structural monotonicity results for the optimal intensity (resp., price) as a function of the stock level and the length of the horizon. For a particular exponential family of demand functions, we find the optimal pricing policy in closed form. For general demand functions, we find an upper bound on the expected revenue based on analyzing the deterministic version of the problem and use this bound to prove that simple, fixed price policies are asymptotically optimal as the volume of expected sales tends to infinity. Finally, we extend our results to the case where demand is compound Poisson; only a finite number of prices is allowed; the demand rate is time varying; holding costs are incurred and cash flows are discounted; the initial stock is a decision variable; and reordering, overbooking, and random cancellations are allowed.

(Dynamic Pricing; Inventory; Yield Management; Intensity Control; Stochastic Demand; Optimal Policies; Heuristics; Finite Horizon; Stopping Times)

1. Introduction and Motivation

Given an initial inventory of items and a finite horizon over which sales are allowed, we are concerned with the tactical problem of dynamically pricing the items to maximize the total expected revenue. Two key properties of this problem are the lack of short-term control over the stock and the presence of a deadline after which selling must stop. Demand is modeled as a price-sensitive stochastic point process with an intensity that is a known decreasing function of the price; revenues are collected as the stock is sold; no backlogging of demand is allowed; unsold items have a given salvage value; and all costs related to the purchase or production of items are considered sunk costs.

This generic problem arises in a variety of industries. Retailers that sell seasonal and style goods are an ex-

ample (cf. Pashigan 1988 and Pashigan and Bowen 1991). For instance, the authors recently have worked with a major New York fashion producer-retailer that designs, produces (via subcontractors), and sells fashion apparel through its own line of retail outlets. (Other similar retailer-producers include The GAP and The Limited.) The firm is known for the subtle and unique colors of its garments, which are achieved using custom-made fabrics. To produce its garments, the firm must special order fabric directly from mills. The raw bolts of fabric are then shipped off-shore (usually by surface freight) to a subcontractor that cuts and assembles the various styles. The finished garments are then shipped back to the U.S. (again, often by surface freight) where they are sorted, boxed, and delivered to individual stores. This entire production process takes from six to

eight months to complete, yet the firm plans to "sell-through" garments in as little as *nine weeks!*

The basic assumptions of the model fit this situation quite well. There is a deadline for the sales period (nine weeks), and for all practical purposes the company has *no resupply option* during the sales season. Further, it is clear that the once the items are on the rack, the entire production decision is sunk. Leftover garments are sold through an affiliated outlet store yielding a given salvage value. (The *salvage value* does impact the pricing decision, as we discuss below.) Demand for garments is uncertain but is influenced by price. The merchandise manager's job is to adjust the price (via markdowns or periodic sales) throughout the selling season in response to the realized demand to maximize revenues. Similar, though perhaps less extreme, instances of the problem occur when selling seasonal appliances such as snow blowers, air conditioners, etc.

The problem is also a fundamental one in the travel and leisure industry. Managers in that industry face hard time constraints and have almost no control in the short run over available space. For example, airlines have a specified number of seats available on each flight, and empty seats are worthless after the plane departs. To increase their revenues, airlines give customers incentives to book in advance. These incentives typically are adjusted in response to the realized demand by opening and closing the various fare classes available at any given point in time. This practice, known as *yield management*, is now used by all major airlines and is increasingly adopted by major hotel chains and even by some car rental companies and cruise ship lines (Kimes 1989). The benefits of yield management are often staggering; American Airlines reports a five-percent increase in revenue, worth approximately \$1.4 billion dollars over a three-year period, attributable to effective yield management (Smith et al. 1992).

Yet, despite its growing importance, there appears to be a certain confusion about precisely what phenomenon yield management actually is trying to exploit. Indeed, in a recent survey, Weatherford and Bodily (1992) conclude that "Several definitions of yield management have been put forward, but to date no agreement exists on its meaning." They point to market segmentation through time-of-purchase mechanisms (e.g., advance purchase requirements, cancellation penalties,

Saturday-night stays, etc.) as one possibility. Though it is certainly an important factor, market segmentation provides only a partial—and perhaps not the most central—explanation for the benefits of yield management. This explanation appears somewhat biased by the business-traveler/vacation-traveler division of the customer population particular to the airline industry. In fact, for resort hotels, cruiseship lines, and theaters, yield management mechanisms seem to be beneficial even though the customer population is arguably much more homogeneous.

Our results provide some important insights on this issue. In particular, they suggest two alternative explanations for the benefits of yield management: (1) *Yield management is an attempt to adjust prices to compensate for "normal" (to be made precise below) statistical fluctuations in demand.* For this first explanation, we have a negative result. Namely, under some rather mild assumptions, we prove that if demand as a function of price is known and prices are unconstrained, then a single fixed-price policy is very nearly optimal. Thus, offering multiple prices can at best capture only second-order increases in revenue due to the statistical variability in demand. (Of course, even second-order increases in revenue may be significant in practice, so this explanation cannot be totally discounted.) Also, the relative fluctuations of an optimal pricing policy appear to be small (on the order of 10% or less), while those found in the airline industry in particular can differ by 100% or more.

The second explanation revealed by our analysis is more compelling: (2) *Yield management is an attempt to "synthesize" a range of optimal prices from a small, static set of prices in response to a shifting demand function.* The above fixed-price results hold only when the firm knows the demand function in advance and can price each instance of the problem (e.g., day/flight/voyage) individually. In most applications, these conditions do not hold. Airlines and hotels must, for a variety of operational and customer-relations reasons, offer a limited number of fares that remain relatively static, at least in the sense of spanning several problem instances. Further, demand may shift significantly during the week or over holidays, and also may not be easy to predict in advance. In such a setting, we prove that a near-optimal policy is to allocate an appropriate fraction of

time and capacity to each fare class, much as is done in conventional yield management practice. In this way, a static set of fare classes together with a dynamic allocation scheme can be used to synthesize different prices for each instance. This interpretation better explains both the magnitude of revenue increases and the disparity in fare prices found in yield management practice.

Finally, we note that one important consideration which is ignored in our formulation is the cost of price changes. Often, these costs are small. For example, travel agents provide customers with current price quotes based on information obtained from computer databases which can easily be updated. In retailing, items may be bar-coded, and thus the cost of a change involves only a computer entry and a change in the displayed price. In such a case, assuming no cost for price changes is a reasonable approximation. In many businesses, however, substantial advertising or ticketing costs are associated with a price change. In these cases, more stable pricing strategies are needed. We show, however, that policies that have *no price changes* are asymptotically (as the expected volume of sales increases) optimal over the class of policies that allow an *unlimited number* of price changes at no cost. This, of course, implies asymptotic optimality for the problem with price change costs as well. Further, we bound the additional expected revenue one can obtain from a dynamic pricing policy over a fixed-price policy. This bound can then be used in conjunction with cost information on price changes to help determine if dynamic pricing is cost effective.

1.1. Literature Review

Research on pricing policies has been pursued by economists, marketing scientists, and operations researchers from a range of perspectives. A considerable body of work has evolved on joint ordering/production and pricing models. A recent and comprehensive survey of this area is given by Eliashberg and Steinberg (1991). In contrast, the main applications and models we study fundamentally have few or no options for reordering. However, in §5 we do analyze extensions to our model that consider initial inventory decisions, reordering, holding costs, and discounting under specialized (unit) cost structures. These extensions relate more closely to the production-pricing literature.

Production-pricing problems are broadly categorized in Eliashberg and Steinberg (1991) into convex and concave ordering cost cases. We shall adopt this classification as well. In the convex case, several discrete-time stochastic models have been investigated in which ordering and pricing decisions are allowed in each period. Single-period models are analyzed by Hempenius (1970), Karlin and Carr (1962), Mills (1959), and Whitin (1955) (his style goods model). These single-period models are essentially price-sensitive versions of the classic "news-boy" problem and are similar to our initial-order-quantity extension discussed in §5.4. The difference is that these models assume static prices and demand, while our model involves a continuous, dynamic demand process and allows dynamic pricing decisions throughout the period. Lazear (1986) considers a model of retail pricing with a single ordering decision and one recourse option to change the price. He formulates a simple, two-stage dynamic program to solve the problem. Pashigan (1988) and Pashigan and Bowen (1991) investigate this model empirically.

Multi-period, finite-horizon models with convex costs are considered by Hempenius (1970), Thowsen (1975) and Zabel (1972). Veinott (1980) uses the theory of lattice programming to investigate monotonicity properties of a class of deterministic, multi-period problems. Our reorder option extension in §5.5 fits broadly in this class, though it is a continuous time model and only unit ordering costs are considered. Karlin and Carr also analyze a stationary, infinite-horizon discounted cost problem, a problem which is also briefly discussed by Mills (1959).

To our knowledge, Li (1988) is the only other paper that considers a continuous time model where demand is a controlled Poisson processes. (Our reorder process in §5.5 is deterministic while Li's is Poisson.) The objective in his paper is to maximize expected discounted profit over an infinite horizon. There is a cost for production capacity, production and holding costs are linear, and both production and pricing decisions are considered. Li's main result is that a barrier policy is optimal for the production decision. He also gives an implicit characterization of the optimal pricing policy when dynamic pricing is allowed.

Concave order costs, usually due to the presence of fixed order costs, are more difficult to analyze and most

work has been confined to deterministic models. EOQ models with price sensitive demand are investigated in Keunreuther and Richard (1971) and Whitin (1955). Cohen (1977) and Rajan et al. (1992) consider problems with decaying inventories. Thomas (1970), Wagner (1960), and Wagner and Whitin (1958) analyze discrete-time, multi-period models with concave costs. To our knowledge, Thomas (1974) is the only paper that studies a stochastic, multi-period model with fixed order costs.

In marketing science, dynamic models of pricing date back to Robinson and Lakhani (1975) and the subsequent work of Bass (1980), Dolan and Jeuland (1981), Jeuland and Dolan (1982), and Kalish (1983). (See Rao 1984 for an overview.) This research, however, focuses on strategic issues of life cycle pricing based on *deterministic* models of how firm economics and consumer behavior change with time. Several marketing scientists have looked at tactical, dynamic pricing problems. Chakravarty and Martin (1989) examine setting optimal quantity discounts in the face of deterministic, dynamically changing demand. Kinberg and Rao (1975) model consumer purchase behavior as a Markov chain and examine the problem of selecting the optimal duration for a price promotion. (See also Nagle 1987 and Oren 1984.)

We have already mentioned that the area of yield management is quite related to our problem. The study of yield management problems in the airlines dates back to the work of Littlewood (1972) for a stochastic two-fare, single-leg problem and to Glover et al. (1982) for a deterministic network model. Belobaba (1987, 1989) proposed and tested a multiple-fare-class extension of Littlewood's rule, which he termed the *expected marginal seat revenue* (EMSR) heuristic. Extensions and refinements of the multiple-fare-class problem include recent papers by Brumelle et al. (1990), Curry (1989), Robinson (1991), and Wollmer (1992). Kimes (1989) gives a general overview of yield management practice in the hotel industry. (See Bitran and Gilbert 1992, Liberman and Yechiali 1978, and Rothstein 1974 for analytical models of hotel problems.) A recent review of research on yield management is given by Weatherford and Bodily (1992), where they adopt the term *perishable asset revenue management* (PARM) to describe this class of problems. Our problem can certainly be considered a continuous-time PARM problem.

Lastly, we mention three papers that address sufficiency conditions for problems similar to our basic model.¹ Miller (1968) studies a finite horizon, continuous-time Markov decision process where only finitely many actions (prices) are allowed. He obtains sufficient conditions for optimality and shows that optimal policies are piecewise constant. Kincaid and Darling (1963) analyze a problem that is functionally equivalent to the basic single-commodity version of our problem. By studying the problem from first principles, they again obtain sufficient conditions; recently Stadje (1990) independently re-derived a similar set of results. Unfortunately, the sufficient conditions derived in these papers *rarely* lead to a solution; indeed, even for the basic version of the problem few practical results have been obtained using these exact approaches.

1.2. Overview and Outline

In §2 we discuss our assumptions, formulate our basic model, provide structural results, and find an exact solution for an exponential demand function. We show that the stochastic optimal policy changes prices continuously and thus may be undesirable in practice. This leads us to try approximate methods. In §3 we find upper bounds on the optimal revenue by considering a deterministic version of the problem. We solve the deterministic problem and show that the optimal policy is to set a *fixed* price throughout the horizon. Further, this deterministic fixed-price policy is asymptotically optimal for the stochastic problem as the volume of expected sales increases or as the time horizon tends to zero. Numerical examples are given that indicate the performance of fixed-price policies is quite good even when the expected volume of sales is moderate. In §4 we analyze the case where only a finite set of prices is allowed, a variant of the problem that is most closely related to the yield management problem.

Finally, in §5 we examine several extensions to the basic problem. First, we generalize our results to the case where demand is a compound Poisson process. We then consider the case where the demand function varies in time through a multiplicative seasonality factor; holding costs are incurred and cash flows are discounted; the initial inventory is a decision variable; additional items can be obtained at a unit cost after the

¹ We are indebted to Sid Browne, Cyrus Derman and Arthur F. Veinott, Jr., for pointing out these references to us.

initial inventory is depleted. The last extension allows for overbooking and cancellations. For all these cases, we find asymptotically optimal heuristics. Our conclusions and thoughts for future research in this area are given in §6.

2. Assumptions, Formulation and Preliminary Results

2.1. Economic and Modeling Assumptions

We assume our firm operates in a market with imperfect competition. For example, the firm may be a monopolist, the product may be new and innovative, in which case the firm holds a temporary monopoly, or the market may allow for product differentiation. Under imperfect competition, a firm can influence demand by varying its price, p . We express the demand as a rate (# items / time) that depends only on the current price p through a function $\lambda(p)$. In the monopoly or new product case, $\lambda(p)$ is the market demand and is assumed to be non-increasing in p due to substitution effects. For example, if the arrival rate of customers is a , and each customer has an i.i.d. reservation price with tail probability $\bar{F}(p)$, then the expected demand rate at price p is $a\bar{F}(p)$. In the case of product differentiation, the demand function is unique to the firm and is assumed to be non-increasing in p due to both lost sales to competitors and substitution effects. In this case, for example, the demand rate seen by the firm as a function of its price and those of its competitors may be modeled using a multinomial logit (cf. Anderson et al. for a fairly extensive treatment of discrete choice theory of product differentiation). Here, we assume that $\lambda(p)$ is given and do not explicitly model the competitive forces that give rise to this demand function. (See Eliashberg and Steinberg 1991 and Dockner and Jorgensen 1988 for examples of dynamic pricing models that represent competition explicitly.)

The assumption that consumers respond only to the current price is, of course, somewhat restrictive. In particular, it does not account for the fact that consumers may act strategically, adjusting their buying behavior in response to the firm's pricing strategy. To do so would require a game theoretic formulation, which is beyond the scope of our analysis. The current-price assumption is approximately true when "impulse purchases" are common (e.g., fashion items). Further, the fact that near-optimal strategies use very stable prices makes this

assumption reasonable in other applications as well. (See Lazear 1986, p. 28 for further discussion of the importance of strategic behavior.)

Realized demand is stochastic and modeled as a Poisson process with intensity $\lambda(p)$. Thus, if the firm prices at p over an interval δ , it sells one item with probability $\lambda(p)\delta + o(\delta)$, no items with probability $1 - \lambda(p)\delta - o(\delta)$ and more than one item with probability $o(\delta)$. In §5.2 we study the case where the demand rate can also depend on the time. We initially consider the case where no backlogging of demand is allowed, so once the firm runs out of stock it collects no further revenues.

Several mild assumptions concerning the demand function are imposed: First, we assume there is a one-to-one correspondence between prices and demand rates so that $\lambda(p)$ has an inverse, denoted $p(\lambda)$. One can therefore alternatively view the intensity λ as the decision variable; the firm determines a target sales intensity λ (i.e., an output quantity) and the market determines the price $p(\lambda)$ based on this quantity. From an analytical perspective, the intensity is more convenient to work with.

We assume the revenue rate,

$$r(\lambda) \doteq \lambda p(\lambda), \tag{1}$$

satisfies $\lim_{\lambda \rightarrow 0} r(\lambda) = 0$, is continuous, bounded and concave, and has a bounded least maximizer defined by $\lambda^* = \min \{ \lambda : r(\lambda) = \max_{\lambda \geq 0} r(\lambda) \}$. Continuity, boundedness of the revenue rate and the maximizer λ^* , and the condition $\lim_{\lambda \rightarrow 0} r(\lambda) = 0$ are all reasonable requirements. Concavity of $r(\lambda)$ stems from the standard economic assumption that marginal revenue is decreasing in output.

Cohen and Karlin and Carr consider demand functions with similar conditions. Specifically, the condition $\lim_{\lambda \rightarrow 0} r(\lambda) = 0$ implies the existence of what Karlin and Carr [22] term a *null price* p_∞ (possibly $+\infty$) for which $\lim_{p \rightarrow p_\infty} \lambda(p) = 0$ and $\lim_{p \rightarrow p_\infty} p\lambda(p) = 0$. (Cohen requires the existence of a null price as well, though he does not give it this term.) In our case, the null price allows us to model the out-of-stock condition as an implicit constraint that forces the firm to price at $p = p_\infty$ when inventory is zero. Note that this modeling artifact partially blurs the distinction between demand and sales, since in reality we can certainly have demand for items without a corresponding sale when the firm is out

of stock. However, in the context of the model, no generality is lost by making this assumption.

We call a function $\lambda(p)$ that satisfies all of the assumptions above a *regular* demand function. An example of a regular demand function is the exponential class $\lambda(p) = ae^{-p}$. One can verify that this function is decreasing in p , has a unique inverse $p(\lambda) = \log(a/\lambda)$ and results in a concave revenue rate $r(\lambda) = \lambda \log(a/\lambda)$ with unique maximizer $\lambda^* = ae^{-1}$. The null price in this case is $p_\infty = +\infty$. Linear demand functions are also regular.

2.2. Formulation

The pricing problem is formulated as follows: At time zero, the firm has a stock n (a nonnegative integer) of items and a finite time $t > 0$ to sell them. The firm controls the intensity of the Poisson demand $\lambda_s = \lambda(p_s)$ at time s using a non-anticipating pricing policy p_s . The intensity $\lambda(\cdot)$ is assumed to be a regular demand function. Let N_s denote the number of items sold up to time s . A demand is realized at time s if $dN_s = 1$, in which case the firm sells one item and receives revenue of p_s .

The price p_s must be chosen from the set of allowable price $\mathcal{P} = \mathbb{R}^+ \cup \{p_\infty\}$. The set of allowable rates is denoted $\Lambda = \{\lambda(p) : p \in \mathcal{P}\}$. Note that since $p_\infty \in \mathcal{P}$, we always have $0 \in \Lambda$. We consider other sets of allowable prices \mathcal{P} in §5. We denote by \mathcal{U} the class of all non-anticipating pricing policies which satisfy

$$\int_0^t dN_s \leq n \quad (\text{a.s.}) \quad (2)$$

and

$$p_s \in \mathcal{P} \Leftrightarrow \lambda_s \in \Lambda \quad \forall s. \quad (3)$$

Constraint (2) is the modeling artifact mentioned above. It acts to "turn off" the demand process when the firm runs out of items to sell. The existence of the null price p_∞ in the set \mathcal{P} guarantees that it can always be satisfied.

Without loss of generality, we assume the salvage value of any unsold items at time t is zero, since for any positive salvage value q we can always define a new regular demand function $\lambda(p) \leftarrow \lambda(p - q)$ and a new price $p \leftarrow p - q$ (the excess over salvage value) that transforms the problem into the zero-salvage-value case. We also assume all costs related to the purchase and production of the product are sunk.

Given a pricing policy $u \in \mathcal{U}$, an initial stock $n > 0$, and a sales horizon $t > 0$, we denote the expected revenue by

$$J_u(n, t) \doteq E_u \left[\int_0^t p_s dN_s \right], \quad (4)$$

where

$$J_u(n, 0) \doteq 0 \quad \forall n \quad (5)$$

and

$$J_u(0, t) \doteq 0 \quad \forall t. \quad (6)$$

The firm's problem is to find a pricing policy u^* (if one exists) that maximizes the total expected revenue generated over $[0, t]$, denoted $J^*(n, t)$. Equivalently,

$$J^*(n, t) \doteq \sup_{u \in \mathcal{U}} J_u(n, t). \quad (7)$$

2.2.1. Optimality Conditions and Structural Results. One can informally derive the Hamilton-Jacobi sufficient conditions for J^* by considering what happens over a small interval of time δt . Since by selecting the intensity λ (i.e., pricing at $p(\lambda)$) we sell one item over the next δt with probability approximately $\lambda \delta t$ and no items with probability approximately $1 - \lambda \delta t$, by the Principle of Optimality,

$$J^*(n, t) = \sup_{\lambda} [\lambda \delta t (p(\lambda) + J^*(n-1, t-\delta t)) + (1 - \lambda \delta t) J^*(n, t-\delta t) + o(\delta t)].$$

Using $r(\lambda) \doteq \lambda p(\lambda)$, rearranging and taking the limit as $\delta t \rightarrow 0$, we obtain

$$\frac{\partial J^*(n, t)}{\partial t} = \sup_{\lambda} [r(\lambda) - \lambda (J^*(n, t) - J^*(n-1, t))] \quad \forall n \geq 1, \quad \forall t > 0. \quad (8)$$

with boundary conditions $J^*(n, 0) = 0, \forall n$ and $J^*(0, t) = 0, \forall t$. The above argument is not rigorous because we have not justified interchanging \sup_{λ} and $\lim_{\delta t \rightarrow 0}$; however, these conditions can be justified formally using Theorem II.1 in Bremaud, where general intensity control problems are studied. Thus, a solution to equation (8) is indeed the optimal revenue $J^*(n, t)$ and the intensities $\lambda^*(n, t)$ that achieve the supremum form an optimal intensity control. Equivalent conditions were derived in Kinkaid and Darling (1963), Miller (1986), and Stadje (1990) without using the theory of intensity control.

The existence of a unique solution to equation (8) is resolved by the following proposition, which is proved in the appendix:

PROPOSITION 1. *If $\lambda(p)$ is a regular demand function, then there exists a unique solution to equation (8). Further, the optimal intensities satisfies $\lambda^*(n, s) \leq \lambda^*$ for all n and for all $0 \leq s \leq t$.*

Although Proposition 1 guarantees the existence of a unique solution to equation (8), obtaining it in closed form is quite difficult—if not impossible—for arbitrary regular demand functions. However, we can make a number of qualitative statements about the optimal expected revenue, intensities and prices. We summarize these in the following theorem.

THEOREM 1. *$J^*(n, t)$ is strictly increasing and strictly concave in both n and t . Furthermore, there exists an optimal intensity $\lambda^*(n, t)$ (resp., price $p^*(n, t)$) that is strictly increasing (resp., decreasing) in n and strictly decreasing (resp., increasing) in t .*

This theorem shows that more stock and/or time leads to higher expected revenues. Further, at a given point in time, the optimal price drops as the inventory increases; conversely, for a given level of inventory, the optimal price rises if we have more time to sell. These properties are not only intuitively satisfying, but they are also useful if one wants to compute the optimal policy numerically because they significantly reduce the set of policies over which one needs to optimize. A proof of a slightly weaker version of Theorem 1 is implied by a sequence of results in Kincaid and Darling (1963). A compact proof of Theorem 1 is presented in the appendix.

2.3. An Optimal Solution for $\lambda(p) = ae^{-\alpha p}$

We can find an exact solution for the demand function $\lambda(p) = ae^{-\alpha p}$, where $a > 0$, $\alpha > 0$ are arbitrary parameters. The solution is useful if one can adequately fit demand to this particular function. More importantly, however, the solution provides interesting insights into the behavior of the optimal policy.

First note that without loss of generality we can take $\alpha = 1$ by simply changing units of price to $p' \leftarrow \alpha p$. The maximizer of $r(\lambda)$ in the case $\alpha = 1$ is $\lambda^* = a/e$ and $p^* = p(\lambda^*) = 1$. It is not hard to verify (see also [25] and [44]) that the solution to equation (8) in this case is

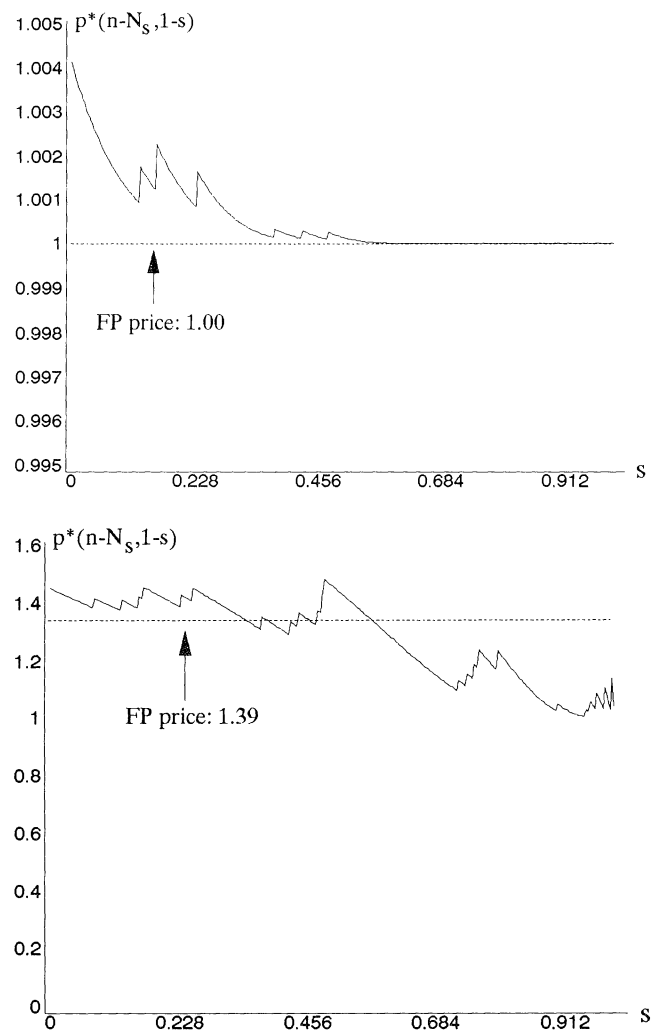
$$J^*(n, t) = \log \left(\sum_{i=0}^n (\lambda^* t)^i \frac{1}{i!} \right), \quad (9)$$

and the optimal price $p^*(n, t)$ is given by

$$p^*(n, t) = J^*(n, t) - J^*(n-1, t) + 1. \quad (10)$$

Some sample paths of this optimal price $p^*(n, t)$ are shown in Figure 1 for a problem with 25 items and a unit time horizon. (The line marked FP is explained below in §3.) The top graph shows a sample price path when demand is low relative to the initial stock ($a = 40$, $\lambda^* t = ae^{-1} \approx 9.5$), while the bottom graph shows a price path for the same 25 items when demand

Figure 1 Sample Paths of $p^*(n - N_s, 1 - s)$ over $s \in [0, 1]$: Top: $n = 25$, $a = 20$. Bottom: $n = 25$, $a = 100$



is high relative to the initial stock ($a = 100$, $\lambda^*t = ae^{-1} \approx 36.8$). There are several interesting things to note about these graphs. First, the upward jumps in price correspond to sales ($dN_s = 1$). After a sale, the price decays until another sale is made, at which point the price takes another jump. This behavior follows from Theorem 1. The upward jumps are due to the fact that the firm prices higher if it has fewer items to sell over a given interval t . The decaying price between sales can be thought of as a price promotion and follows from the fact $p^*(n, t)$ is decreasing in t for fixed n . The firm gradually reduces the price as time runs out in order to induce buying activity.

3. Bounds and Heuristics

For regular demand functions other than $\lambda(p) = ae^{-\alpha p}$ it is quite difficult, if not impossible, to find closed form solutions to equation (8). Further, by Theorem 1, the optimal price varies continuously over time. Yet in many applications this degree of price flexibility is either impossible or prohibitively expensive to implement. Therefore, one might often prefer more stable policies that are close to optimal over a "jittery" optimal policy. In this section, we propose heuristics that meet these criteria. They are easy to implement and also provably close to optimal in many cases. Our approach is to first construct an upper bound based on a deterministic version of the problem. The solution to this deterministic problem then suggests a simple fixed-price heuristic that we show is provably good when the volume of expected sales is large.

3.1. An Upper Bound Based on a Deterministic Problem

3.1.1. Formulation of Deterministic Problem.

Consider the following deterministic version of the problem: At time zero, the firm has a stock x , a continuous quantity, of product and a finite time $t > 0$ to sell it. The instantaneous demand rate is deterministic and a function of the price at time s , $p(s)$, again denoted $\lambda(p(s))$. (Our notation distinguishing a deterministic policy follows that in Bremaud (1980).) We assume $\lambda(\cdot)$ is a regular demand function. As before, without loss of generality, we assume the salvage value of the product at time t is zero and that all other costs are sunk. The price $p(s)$ must again be chosen from a set \mathcal{P} of allowable prices. As before, we can equivalently

view the firm as setting the rate $\lambda(s) \in \Lambda$, which implies charging a price $p(s) = p(\lambda(s)) \in \mathcal{P}$.

The firm's problem is to maximize the total revenue generated over $[0, t]$ given x , denoted $J^D(x, t)$.

$$J^D(x, t) = \max_{\{\lambda(s)\}} \int_0^t r(\lambda(s)) ds \quad (11)$$

subject to

$$\int_0^t \lambda(s) ds \leq x$$

$$\lambda(s) \in \Lambda.$$

3.2. Optimal Solution of the Deterministic Problem

We begin with some definitions. Define the *run-out rate*, denoted λ^0 , by $\lambda^0 \doteq x/t$, the *run-out price*, denoted p^0 , by $p^0 \doteq p(\lambda_0)$, and the *run-out-revenue rate* $r^0 \doteq p^0\lambda^0$. Notice that λ^0 (resp., p^0) is the fixed intensity (resp., price) at which the firm sells exactly its initial stock x over the interval $[0, t]$. Recall that λ^* is the least maximizer of the revenue function $r(\lambda) = \lambda p(\lambda)$. We find it convenient to define $p^* \doteq p(\lambda^*)$, and $r^* \doteq p^*\lambda^*$. The quantity r^* is the maximum instantaneous revenue rate. These definitions allow us to state the following proposition, which is proved in the appendix:

PROPOSITION 2. *The optimal solution to the deterministic problem (11) is $\lambda(s) = \lambda^D \doteq \min\{\lambda^*, \lambda^0\}$, $0 \leq s \leq t$. In terms of price the optimal policy is $p(s) = p^D \doteq \max\{p^*, p^0\}$, $0 \leq s \leq t$. Finally, the optimal revenue is*

$$J^D(x, t) = t \min\{r^*, r^0\} \quad (12)$$

The intuition for this solution is the following: If the firm has a large number of items to sell ($x \geq \lambda^*t$), it ignores the problem of running out of stock and prices at the level that maximizes the revenue rate. In this case the firm ends with $x - \lambda^*t$ unsold units. If the items are scarce ($x < \lambda^*t$), the firm can afford to price higher, and it indeed prices at the highest level that still enables it to sell all the items. Note that in both cases the solution is to set a fixed price for the entire interval.

3.2.1. The Deterministic Revenue as an Upper Bound. Intuitively, one would expect that the uncertainty in sales in the stochastic problem results in lower expected revenues. The following theorem formalizes this idea:

THEOREM 2. *If $\lambda(p)$ is a regular demand function, then for all $0 \leq n < +\infty$ and $0 \leq t < +\infty$,*

$$J^*(n, t) \leq J^D(n, t).$$

PROOF OF THEOREM 2. As shown in Proposition 1, $\lambda_s \leq \lambda^* < \infty$, which implies $\int_0^t \lambda_s ds < \infty$ almost surely for all $t \geq 0$. Recall \mathcal{U} denotes the class of policies that satisfy $\int_0^t dN_s \leq n$ (a.s.); therefore by Bremaud 1980, Theorem II,

$$E_u \left[\int_0^t dN_s \right] = E_u \left[\int_0^t \lambda_s ds \right] \leq n. \quad (13)$$

Since the demand intensity in the control problem (7) is Markovian, it is sufficient to consider only Markovian policies u (Bremaud 1980, Corollary II.2). That is, policies for which the price at time s is a function $p_s = p_u(n - N_s, s)$ only. (Equivalently, the intensity at time s is a function $\lambda_s = \lambda(n - N_s, s)$.) By Bremaud 1980, Theorem II, we can write,

$$\begin{aligned} J_u(n, t) &= E_u \left[\int_0^t p_u(n - N_s, s) dN_s \right] \\ &= E_u \left[\int_0^t r(\lambda_s) ds \right] \end{aligned} \quad (14)$$

and

$$J^*(n, t) = \sup_{u \in \mathcal{U}} J_u(n, t).$$

Now for $\mu \geq 0$ we define the augmented cost functional

$$\begin{aligned} J_u(n, t, \mu) &= E_u \left[\int_0^t (r(\lambda_s) - \mu \lambda_s) ds \right] \\ &+ n\mu \geq J_u(n, t), \end{aligned} \quad (15)$$

and the augmented deterministic cost function

$$J^D(n, t, \mu) = \max_{\lambda(s) \in \Lambda} \int_0^t (r(\lambda(s)) - \mu \lambda(s)) ds + n\mu. \quad (16)$$

We claim the following:

LEMMA 1.

$$J_u(n, t, \mu) \leq J^D(n, t, \mu) \quad \forall u \in \mathcal{U}, \mu \geq 0.$$

PROOF. This follows by viewing the integrand inside the expectation in equation (15) as purely a function of λ and maximizing pointwise:

$$\begin{aligned} J_u(n, t, \mu) &\leq \int_0^t \max_{\lambda(s) \in \Lambda} \{r(\lambda(s)) - \mu \lambda(s)\} ds + n\mu \\ &= \max_{\{\lambda(s) \in \Lambda\}} \int_0^t (r(\lambda(s)) - \mu \lambda(s)) ds + n\mu \\ &\doteq J^D(n, t, \mu). \quad \square \end{aligned}$$

Since Lemma 1 holds for all $u \in \mathcal{U}$ and $\mu \geq 0$, we have by equation (15)

$$J^*(n, t) \leq \inf_{\mu \geq 0} J^D(n, t, \mu).$$

Theorem 2 then follows by noting that the quantity on the right above is the optimal dual value of the infinite dimensional program

$$\begin{aligned} J^D(n, t) &= \max_{\lambda(s) \in \Lambda} \int_0^t r(\lambda(s)) ds \\ &\text{subject to} \\ &\int_0^t \lambda(s) ds \leq n. \end{aligned}$$

Since this is a convex program, and the null price together with the fact that $n > 0$ implies that $\lambda(s) = 0, 0 \leq s \leq t$ is a strictly interior solution, there exists a multiplier μ^* for which the duality gap is zero and $J^D(n, t) = \inf_{\mu \geq 0} J^D(n, t, \mu) = J^D(n, t, \mu^*)$. (See Luenberger 1969). \square

Theorem 2 is useful for several reasons. It suggests that the solution of the deterministic problem may provide insight into optimal or near-optimal pricing strategies for the stochastic problem. It also provides performance guarantees on the cost of such pricing strategies. Together, these results can be used to establish a strong relationship between the stochastic and deterministic problems, as we show next.

3.3. Asymptotically Optimal Fixed-price Heuristics

The deterministic optimal solution suggests a simple *fixed price* (FP) heuristic, namely, for the entire horizon set the price to $p^D = \max \{p^0, p^*\}$. Of course, one could improve on this heuristic by choosing the best *fixed price*; that is, the one maximizing $pE[\min \{n, N_{\lambda(p), t}\}]$, where N_α denotes a Poisson random variable with mean α . In general the best fixed price cannot be found analytically; however, it is easy to find numerically. We let OFP denote this optimal fixed-price heuristic and $J^{\text{OFP}}(n, t)$ (resp., $J^{\text{FP}}(n, t)$) denote the revenue of the OFP (resp., FP) heuristic. We will use the fact that $J^{\text{OFP}}(n, t) \geq J^{\text{FP}}(n, t)$.

Note that using a fixed price for the entire horizon is quite convenient since there is no effort involved in monitoring time and inventory levels and no cost incurred for changing prices. Further, the performance of

the heuristics turn out to be quite good in several cases, one of which is shown by the following theorem:

THEOREM 3.

$$\frac{J^{\text{OFP}}(n, t)}{J^*(n, t)} \geq \frac{J^{\text{FP}}(n, t)}{J^*(n, t)} \geq 1 - \frac{1}{2\sqrt{\min\{n, \lambda^*t\}}}.$$

PROOF. The first inequality follows from the definition of the heuristics. To show the second, note the expected revenue obtained when the price is fixed at p is

$$pE[N_{\lambda(p)t} - (N_{\lambda(p)t} - n)^+]. \quad (17)$$

Gallego (1992) shows that for any random variable N with finite mean μ and finite standard deviation σ , and for any real number n ,

$$E[(N - n)^+] \leq \frac{\sqrt{\sigma^2 + (n - \mu)^2} - (n - \mu)}{2}, \quad (18)$$

where $x^+ \doteq \max(x, 0)$. Consider first the case $\lambda^*t > n$. That is, the case where items are scarce and the FP heuristic uses the run-out price p^0 . Using the above inequality in equation (17) and noting that when pricing at p^0 , $\mu = \sigma^2 = n$, we obtain

$$J^{\text{FP}}(n, t) \geq np^0 \left(1 - \frac{1}{2\sqrt{n}}\right) = r^0 t \left(1 - \frac{1}{2\sqrt{n}}\right).$$

In the case where $\lambda^*t \leq n$ we price at p^* , by the same reasoning we obtain

$$\begin{aligned} J^{\text{FP}}(n, t) &\geq p^* \left(\lambda^*t - \frac{\sqrt{\lambda^*t + (n - \lambda^*t)^2} - (n - \lambda^*t)}{2} \right) \\ &\geq p^* \lambda^*t \left(1 - \frac{1}{2\sqrt{\lambda^*t}}\right) = r^* t \left(1 - \frac{1}{2\sqrt{\lambda^*t}}\right). \end{aligned} \quad (19)$$

Comparing these two cases to the deterministic revenue (12) and using Theorem 2 completes the proof. \square

REMARK. When $\lambda^*t > n$, one can determine the exact cost of the FP heuristic by noting that $E(N_n - n)^+ = n(1 - P\{N_n = n\})$, which implies

$$J^{\text{FP}}(n, t) = np^0 \left(1 - \frac{n^n}{n!} e^{-n}\right).$$

This provides a slightly better guarantee for small n , though it has an identical rate of convergence since by Stirling's formula $(n^n/n!)e^{-n} \sim 1/\sqrt{2\pi n}$.

Theorem 3 shows that the FP heuristic, and consequently the OFP heuristic, are asymptotically optimal in two limiting cases: (1) the number of items is large ($n \gg 1$) and there is plenty of time to sell them ($n < \lambda^*t$); or (2) there is the potential for a large number of sales at the revenue maximizing price ($\lambda^*t \gg 1$), and there are enough items in stock to satisfy this potential demand ($n \geq \lambda^*t$). Thus, we see that if the volume of expected sales is large, the heuristics perform quite well.

One can gain an intuitive understanding of this result by examining Figure 1, which shows the FP price and the optimal price for two sample paths of the example in §2.3. Note that the optimal price paths in this figure appear roughly centered about the FP price shown by the horizontal lines in Figure 1. Also, on a relative basis the variations about the FP price appear small. Thus, it seems the FP price is a reasonable approximation to the optimal policy.

An example serves to illustrate the utility of the bounds in Theorem 3: Consider a firm that has 400 items and enough time to price at the run-out price p^0 ($\lambda^*t > n$). Theorem 3 then guarantees that the expected revenue collected by simply offering a fixed price of p^0 is at least 97.5% of what could be obtained by using an optimal state-dependent strategy. For 100 items, the guarantee drops to 95%, while for 25 items, it is only 90%. However, as we illustrate in the next subsection, these guarantees are in fact quite pessimistic, and the actual performance of fixed-price policies is good even for small (≈ 10 items) problems.

As a last example where fixed-price heuristics are asymptotically optimal, we state without proof

THEOREM 4.

$$\lim_{t \rightarrow 0} \frac{J^{\text{OFP}}(n, t)}{J^*(n, t)} \geq \lim_{t \rightarrow 0} \frac{J^{\text{FP}}(n, t)}{J^*(n, t)} = 1 \quad \forall n > 0.$$

3.4. Numerical Example of the Performance of Fixed-Price Heuristics

For the case where the demand function is $ae^{-\alpha p}$ we have a closed form expression for the optimal cost, which allows us to examine the performance of the fixed-price heuristics for problems of moderate size. Table 1 shows the prices and resulting revenues for a series of problems with a unit horizon, $\lambda^*t = 10$ and starting inventories n ranging from 1 to 20. Note that the optimal fixed price (p^{OFP}) is initially lower than the deterministic

Table 1 Prices and Revenues for the Case $\lambda^*t = 10$

n	p^{OFP}	p^{FP}	J^*	J^{OFP}/J^*	J_{FP}/J^*
1	2.74	3.30	2.40	0.945	0.871
2	2.36	2.61	4.11	0.947	0.926
3	2.10	2.20	5.43	0.950	0.945
4	1.90	1.92	6.47	0.954	0.954
5	1.74	1.69	7.30	0.958	0.956
6	1.61	1.51	7.96	0.962	0.956
7	1.50	1.35	8.49	0.967	0.952
8	1.41	1.22	8.89	0.971	0.946
9	1.33	1.11	9.22	0.976	0.937
10	1.26	1.00	9.46	0.980	0.925
11	1.21	1.00	9.64	0.985	0.951
12	1.16	1.00	9.77	0.989	0.970
13	1.12	1.00	9.85	0.992	0.982
14	1.08	1.00	9.91	0.995	0.990
15	1.05	1.00	9.95	0.997	0.995
16	1.04	1.00	9.97	0.998	0.997
17	1.02	1.00	9.99	0.999	0.999
18	1.01	1.00	9.99	0.999	0.999
19	1.01	1.00	10.00	1.000	1.000
20	1.00	1.00	10.00	1.000	1.000

price (p^{FP}) when there are few items to sell, but for $n > 5$ it is higher. Thus, the OFP price seems to smooth the transition between the low and high demand price extremes, p^* and p^0 . Note also that the *worst* relative performance of the OFP heuristic is only 5.5%, and when $n > 12$ it is within 1% of the optimal revenue. Indeed, in numerical experiments on many different examples we never once observed a value of J^{OFP} that was more than seven percent less than the optimal revenue. The relative performance of the FP heuristic, on the other hand, is poorest at $n = 1$ (12.9% below the optimal revenue), though for $n > 15$ its revenue is comparable to that of the OFP heuristic.

These results suggest that even for moderate sized problems the FP heuristic, and especially the OFP heuristic, perform quite well. They also suggest that dynamic pricing in response to the sort of statistical variations in demand modeled here can at best provide only minimal increases in revenue—on the order of one percent or less for moderate to large problems. For this reason, we conclude that if demand functions are well known and prices can be set freely, then one should not see great benefits from the highly dynamic pricing practices, such as those found in fashion retailing and yield management practice. Other explanations of the

benefits of these practices, one of which we propose in §4, are needed. (An explanation based on the producer's imperfect knowledge of customers' reservation prices is proposed by Lazear.)

3.5. Some Structural Observations for the Case

$$\lambda(p) = ae^{-p}$$

We next show that for $\lambda(p) = ae^{-p}$ the optimal intensity (resp., price) for the stochastic problem is always smaller (resp., larger) than the corresponding optimal intensity (resp., price) for the deterministic problem.

PROPOSITION 3. If $\lambda(p) = ae^{-p}$, then $\forall n \geq 0, \forall t \geq 0$,

$$\lambda^*(n, t) \leq \lambda^D(n, t)$$

and

$$p^*(n, t) \geq p^D(n, t).$$

PROOF. We can write the optimal intensity as

$$\lambda^*(n, t) = \lambda^* \frac{P\{N_{\lambda^*t} \leq n-1\}}{P\{N_{\lambda^*t} \leq n\}} \leq \lambda^*.$$

Since $\lambda^D(n, t) = n/t$ for $t \geq n/\lambda^*$ and $\lambda^D(n, t) = \lambda^*$ otherwise, and by Proposition 1, $\lambda^*(n, t) \leq \lambda^*$ always, we only need to show that $\lambda^*(n, t) \leq n/t$, for $t \geq n/\lambda^*$. Equivalently,

$$\lambda^*tP\{N_{\lambda^*t} \leq n-1\} \leq nP\{N_{\lambda^*t} \leq n\}.$$

But this holds since the left-hand side can be written $\sum_{i=0}^n ie^{-\lambda^*t}(\lambda^*t)^i/i!$, which is clearly less than $nP\{N_{\lambda^*t} \leq n\}$. The corresponding properties for $p^*(n, t)$ follow in a similar way. \square

This proposition helps address a question raised by Mills (1959) about the relation between the optimal price for a stochastic model and its deterministic counterpart. Karlin and Carr (1962) and Thowsen (1975) analyzed this question for fixed price models under additive or multiplicative uncertainty. They showed the optimal stochastic price is always higher (resp., lower) under multiplicative (resp., additive) uncertainty than the optimal deterministic price. Note that the demand uncertainty in the exponential model is neither additive nor multiplicative. Though Proposition 3 suggests that the optimal stochastic price is always higher, this is not true for the revenue function $r(\lambda) = 1 - (\lambda - 1)^2$. It is true, however, that for all regular demand functions over short time horizons, i.e., $t \leq n/\lambda^*$, the optimal stochastic price (resp., intensity) is always higher (resp.,

lower) than the optimal deterministic price (resp., intensity).

4. Discrete Price Case

Consider the case where the set of allowable prices is $\mathcal{P} = \{p_1, \dots, p_K, p_\infty\}$, a discrete set. The restriction to a discrete set of prices may arise if a firm decides, at a strategic level, to restrict itself to a given set of prices in order to achieve market segmentation. Alternatively, the discrete set of prices may be the result of an explicit or implicit consensus at the industry level (e.g., "price points"). We also suggest below that a discrete price scheme together with dynamic allocation of the units allows firms to synthesize a wide range of effective prices to accommodate shifting demand functions while retaining the appeal and practicality of having only a small, stable set of prices for their product or service. We propose this as one plausible explanation for the practice of yield management.

As before, we have n items and t units of time to sell them. Corresponding to price p_k , we have a known demand rate λ_k , $k = 1, \dots, K$. Without loss of generality, we assume that $p_1 < p_2 < \dots < p_K$, and $\lambda_1 > \lambda_2 > \dots > \lambda_K$. Equality in the demand rates is ruled out since equation (8) selects the largest price corresponding to any given demand rate.

Let $r_k \doteq p_k \lambda_k$, denote the *revenue rate* associated with price p_k , $k = 1, \dots, K$. We assume that the revenue rates are monotonically decreasing:

$$r_1 > r_2 > \dots > r_K.$$

This assumption is again without loss of generality since $0 \leq \lambda \leq \lambda^*$ by Proposition 1, and because $r(\cdot)$ is increasing over this region.

4.0.1. Optimal Solution of the Deterministic Problem. The proof of Theorem 2 goes through unchanged for the discrete price case; thus, we can again use the deterministic revenue as an upper bound. The deterministic solution also gives an asymptotically optimal heuristic, though the resulting heuristic is no longer a *fixed-price* heuristic, but consists of pricing at some price, p_{k^*} , for a specified period of time and at a neighboring price, p_{k^*+1} , for the balance of the horizon.

To solve the deterministic pricing problem let $t_k = \int_0^t 1(p_s = p_k) ds$ denote the amount of time we price the items at p_k , $k = 1, \dots, K$, over the horizon $[0, t]$. Here $1(p_s = p_k) = 1$ if the price at time s is p_k and zero

otherwise. Then equation (11) reduces to a linear program. For convenience, set $\lambda_0 \doteq \infty$, $\lambda_{K+1} \doteq 0$, and $r_0 = r_{K+1} \doteq 0$. The next proposition states, without proof, that this linear program can be solved in closed form:

PROPOSITION 4. For any (n, t) , let k^* be such that $\lambda_{k^*} t \geq n > \lambda_{k^*+1} t$, then the solution to the linear program is given by $t_j = 0$ for $j \notin \{k^*, k^* + 1\}$, and

$$t_{k^*} = \frac{n - \lambda_{k^*+1} t}{\lambda_{k^*} - \lambda_{k^*+1}}$$

$$t_{k^*+1} = \frac{\lambda_{k^*} t - n}{\lambda_{k^*} - \lambda_{k^*+1}},$$

where $t_k \doteq 0$, $t_{k+1} \doteq t$ when $k^* = 0$ and $t_k \doteq n/\lambda_K$, $t_{k+1} \doteq 0$ when $k^* = K$.

REMARK. If there exists a salvage value $q > 0$ then the above results continue to hold provided (1) we eliminate all prices $p < q$, and all prices p_i such that $q(\lambda_i - \lambda_j) > r_i - r_j$ for some $p_j > p_i > q$, and (2) we set $r_i = \lambda_i(p_i - q)$ in the linear program.

Notice that the solution prices at p_{k^*} for αt units of time and at p_{k^*+1} for $(1 - \alpha)t$ units of time where $\alpha \in [0, 1)$ satisfies $\alpha \lambda_{k^*} t + (1 - \alpha) \lambda_{k^*+1} t = n$. Thus, $\alpha \lambda_{k^*} t$ and $(1 - \alpha) \lambda_{k^*+1} t$ are approximately the number of items allocated to prices p_{k^*} and p_{k^*+1} respectively, and $(\alpha r_{k^*} + (1 - \alpha) r_{k^*+1}) t / n$ is the *effective price* paid, averaged across the n items. By adjusting the allocations in this way, one can synthesize effective prices for many different demand functions. There are many practical advantages to such a scheme. It allows a firm to offer only a small set of stable prices, which are easy for consumers to interpret and for the firm to advertise and manage. Yet it also enables the firm to respond to short-term variations in demand, such as those caused by day-of-the-week cycles, holidays, seasonalities, etc. This may be one reason why industries with highly variable demand patterns, such as airlines, hotels and cruise-ships, have adopted the fixed-fare-classes, dynamic allocation policies of yield management.

4.1. An Asymptotically Optimal Heuristic

The deterministic solution suggests a stopping-time (ST) heuristic for the stochastic problem. Let $m \doteq \lceil \lambda_{k^*} t_{k^*} \rceil$, T_m be the (random) time the m th item is demanded when the price is fixed at p_{k^*} , and let $t_m = m / \lambda_{k^*}$, be the time it takes to sell m items at price p_{k^*} when demand is deterministic. The heuristic is defined as follows:

ST Heuristic: Start pricing at p_{k^*} and switch to p_{k^*+1} at (random) time

$$\tau = \min(T_m, t_m).$$

Let $J^{\text{ST}}(n, t)$ denote the expected revenue for the ST heuristic. The following theorem is proved in the appendix:

THEOREM 5. *Suppose $n \rightarrow \infty$ and $t \rightarrow \infty$ such that $\lambda_k t \geq n > \lambda_{k+1} t$. Then,*

$$\lim_{t \rightarrow \infty} \frac{J^{\text{ST}}(n, t)}{J^{\text{D}}(n, t)} = 1.$$

As an example illustrating the rate of convergence, consider a flight with $n = 300$ seats that is open for sale $t = 360$ days before the departure of the flight. Assume that at the promotional fare $p_1 = \$198$, the average demand rate is $\lambda_1 = 1$ seats per day, and that at the regular fare $p_2 = \$358$, the average demand is $\lambda_2 = 0.5$ seats per day. Then $m = t_1 = 240$, and the promotional fare is stopped when 240 seats are sold or when 240 days elapse, whichever occurs first. Using the bounds in the appendix, we obtain

$$\$66,080 \leq J^{\text{ST}}(n, t) \leq J^*(n, t) \leq \$69,000.$$

To assess the performance of the ST heuristic we simulated 300 flights with the above data. The expected revenue of the ST heuristic was estimated to be \$67,546, or about 98% of the deterministic upper bound.

The analysis of the ST heuristic can be sharpened when $n \geq \lambda_1 t$, and when $n \leq \lambda_k t$ in the sense that the absolute, rather than the relative, error goes to zero as n and/or t goes to infinity. The first case occurs when demand is so low that we cannot expect to sell all the items even at the lowest price (p_1); the second occurs when demand is so high that we can expect to sell all the items at the highest price (p_k). In both cases the heuristic reserves the entire stock to the lowest (resp., the highest) price.

Finally, we point out that a heuristic with the same asymptotic properties can be constructed whereby initially the items are priced at p_{k^*+1} and subsequently reduced to p_{k^*} . Thus, both the low-to-high and the high-to-low stopping-time heuristics are asymptotically optimal. For example, in air travel the desirability of a seat usually increases as the date of flight is approached,

while in fashion retailing the desirability of garments decreases as the season draws to a close; thus airlines price low-high, while fashion retailers price high-low. See Feng and Gallego (1992) for structural results and algorithms to compute *optimal* stopping-time rules in situations that allow, at most, one price change.

5. Extensions to the Basic Problem

We next examine several extensions of the basic problem. The first extension allows demand to be compound Poisson. Next we consider a demand function that varies with time according to a multiplicative seasonality factor. Then, we extend our results to the case where there are holding costs and cash flows are discounted. We then allow the initial stock n to be a decision variable along with price. Finally, we allow a resupply option in the presence of overbooking and random cancellations. For all these cases, we find asymptotically optimal heuristics and, in some instances, a closed-form optimal policy for the exponential demand case.

5.1. Demand is a Compound Poisson Processes

Let N_s be a Poisson Process with random intensity $\{\lambda_u: 0 \leq u \leq s\}$ and let T_k be the epoch of the k th arrival of N_s . That is, $N_s = k$ for $T_k \leq s < T_{k+1}$. At time T_k we see a demand of size X_k where the X_k 's are i.i.d. random variables with $EX > 0$, and $EX^2 < \infty$. Let \mathcal{U} be the set of nonanticipatory policies such that $\int_0^t X_{N_s} dN_s \leq n$ almost surely. The expected revenue can be written as

$$J_u(n, t) = E_u \int_0^t p(\lambda_s) X_{N_s} dN_s.$$

Let $J^{\text{D}}(n, t) = EX \max_{\lambda(s) \in \Lambda} \int_0^t (r(\lambda(s))) ds$ subject to $\int_0^t X_{N_s} dN_s \leq n$ be the optimal revenue for the deterministic problem. We next show

THEOREM 6.

$$J_u(n, t) \leq J^{\text{D}}(n, t).$$

PROOF. For $\mu \geq 0$, we define

$$J_u(n, t, \mu) \doteq J_u(n, t) + \mu E_u \left(n - \int_0^t X_{N_s} dN_s \right) \geq J_u(n, t).$$

Because $X_k = X_{N_{T_k}}$ is independent of $N_s, s \leq T_k$, we can write

$$\begin{aligned}
 J_u(n, t, \mu) &= E_u \sum_{k=1}^{\infty} (p(\lambda_{T_k}) - \mu) X_k 1\{T_k \leq t\} + n\mu \\
 &= \sum_{k=1}^{\infty} EX_k E_u(p(\lambda_{T_k}) - \mu) 1\{T_k \leq t\} + n\mu \\
 &= EX \sum_{k=1}^{\infty} E_u(p(\lambda_{T_k}) - \mu) 1\{T_k \leq t\} + n\mu \\
 &= EX E_u \int_0^t (p(\lambda_s) - \mu) dN_s + n\mu \\
 &= EX E_u \int_0^t (r(\lambda_s) - \mu \lambda_s) ds + n\mu \\
 &\leq EX \int_0^t \max_{\lambda(s) \in \Lambda} (r(\lambda(s)) - \mu \lambda(s)) ds + n\mu \\
 &= EX \max_{\lambda(s) \in \Lambda} \int_0^t (r(\lambda(s)) - \mu \lambda(s)) ds + n\mu \\
 &\doteq J^D(n, t, \mu).
 \end{aligned}$$

Consequently,

$$J_u(n, t) \leq J_u(n, t, \mu) \leq \inf_{\mu \geq 0} J^D(n, u, \mu) = J^D(n, t). \quad \square$$

Thus, again the deterministic problem provides an upper bound. The solution to the deterministic problem is easily seen to be $\lambda_s = \lambda^D$, $0 \leq s \leq t$, where $\lambda^D \doteq \min\{\lambda^*, \lambda^0\}$ and $\lambda^0 \doteq n/(tEX)$. Consequently, $J^D(n, t) = \min\{r^*, r^0\}tEX$, where $r^0 \doteq \lambda^0 p(\lambda^0)$. As before, we can use the deterministic solution as a heuristic for the stochastic problem. Let

$$J^{\text{FP}}(n, t) = p^D E \min \left\{ n, \sum_{k=1}^{N_{\lambda^D}} X_k \right\}.$$

Following the arguments used in Theorem 3 to establish the asymptotic optimality of the fixed price heuristic, we obtain

THEOREM 7.

$$\frac{J^{\text{OFP}}(n, t)}{J^*(n, t)} \geq \frac{J^{\text{FP}}(n, t)}{J^*(n, t)} \geq 1 - \frac{\sqrt{\rho}}{2\sqrt{\min\{n, \lambda^* t\}}}$$

where $\rho \doteq EX^2/EX$.

5.2. Time Varying Demand

Assume now that the demand rate $\lambda(p, s)$ depends both on the price p and the time elapsed s since the start of

the selling season. Assume further that dependence in time is through a positive multiplicative factor $g(s)$, so

$$\lambda(p, s) = \lambda(p)g(s) \quad 0 \leq s \leq t.$$

For example, $g(s)$ may be a concave function peaking near the middle of the selling season. A simple method allows us to transform this problem into one in which demand is time homogeneous. Let

$$u = G(s) = \int_0^s g(z) dz, \quad 0 \leq s \leq t,$$

and define

$$\tilde{\lambda}(p, u) \doteq \lambda(p), \quad 0 \leq u \leq G(t).$$

Then, for all $s < s'$, let $u = G(s)$, and $u' = G(s')$, and note that

$$\begin{aligned}
 \int_s^{s'} \lambda(p, z) dz &= \lambda(p) \int_s^{s'} g(z) dz = \lambda(p)[G(s') - G(s)] \\
 &= \lambda(p)[u' - u] = \int_u^{u'} \tilde{\lambda}(p, v) dv.
 \end{aligned}$$

Thus by using the clock $u = G(s)$, $0 \leq u \leq G(t)$, instead of the clock $0 \leq s \leq t$, we have transformed the problem into one where demand is time homogeneous. Consequently, all of our results apply to the transformed problem. In particular, the FP heuristic becomes:

$$p^{\text{FP}} = \max\{p^*, p(n/G(t))\}.$$

By Theorem 3, the performance guarantee of the FP heuristic is

THEOREM 8.

$$\frac{J^{\text{FP}}(n, t)}{J^D(n, t)} \geq 1 - \frac{1}{\sqrt{\min(n, \lambda^* G(t))}}.$$

The above procedure can also be used in the discrete price case as well. Indeed as in §4, let k^* be such that $\lambda_{k^*} G(t) \geq n > \lambda_{k^*+1} G(t)$. Then the optimal solution to the transformed deterministic problem is to price at p_{k^*} for

$$u_{k^*} = \frac{n - \lambda_{k^*}^* G(t)}{\lambda_{k^*} - \lambda_{k^*+1}}$$

units of time, and to price at p_{k^*+1} for

$$u_{k^*+1} = \frac{\lambda_{k^*} G(t) - n}{\lambda_{k^*} - \lambda_{k^*+1}}$$

units of time. The stopping-time heuristic for the original problem can be constructed by pricing at p_k^* for $s_{k^*} = G^{-1}(u_{k^*})$ units of time in the original clock, and by pricing at p_{k^*+1} for $t - s_{k^*}$ units of time, again in the original clock.

Now, let $\tilde{\lambda}^*(n, u)$ denote the optimal intensity for the transformed problem when u units of time have elapsed (with respect to the new clock) and there are n units in inventory. We know from Theorem 1 that $\tilde{\lambda}^*(n, u)$ is increasing in u . The optimal intensity $\lambda^*(n, s)$ when s units of time have elapsed (with respect to the original clock) and there are n units in inventory is related to $\tilde{\lambda}^*(n, u)$ by

$$\lambda^*(n, s) = \tilde{\lambda}^*(n, G(s))g(s).$$

Now let $p^*(n, s)$ denote the optimal price when s units of time have elapsed with respect to the original clock and there are n units in inventory. Then, by definition,

$$\lambda^*(n, s) = \lambda(p^*(n, s), s) = \lambda(p^*(n, s))g(s).$$

Consequently, we have

$$\lambda(p^*(n, s)) = \tilde{\lambda}^*(n, G(s)).$$

Now since $\lambda(p)$ is a decreasing function of p , it follows that $p^*(n, s)$ is decreasing in s . This is consistent with Theorem 1 viewing s as elapsed time. We note, however, that the analogous result does not hold for $\lambda^*(n, s)$ since its behavior also depends on $g(s)$.

5.3. Holding Cost and Discount Rate

Now suppose cash flows are discounted at rate β , and a linear holding cost h is charged on existing inventories. Let $Z(s)$ be the inventory level at time s . Then

$$Z(s) = n - N_s$$

where N_s is a Poisson process with random intensity $\{\lambda_u, 0 \leq u \leq s\}$. The intensity λ_s is set to zero whenever $Z(s) = 0$.

Recall \mathcal{U} denotes the class of nonanticipatory policies that satisfy $\int_0^t dN_s \leq n$ almost surely. For any $u \in \mathcal{U}$, the expected discounted revenue is given by

$$E_u \int_0^t e^{-\beta s} p(\lambda_s) dN_s = E_u \int_0^t e^{-\beta s} r(\lambda_s) ds.$$

The expected discounted holding cost is given by

$$hE_u \int_0^t e^{-\beta s} Z(s) ds = hE_u \int_0^t e^{-\beta s} \left(n - \int_0^s \lambda_u du \right) ds.$$

Integrating the last expression by parts, we obtain

$$hE_u \int_0^t e^{-\beta s} \left(n - \frac{1}{\beta} (1 - e^{-\beta(t-s)}) \lambda_s \right) ds.$$

Consequently, the net expected discounted revenue $J_u(n, t)$ is given by

$$J_u(n, t) = E_u \int_0^t e^{-\beta s} \left[r(\lambda_s) + \frac{h}{\beta} (1 - e^{-\beta(t-s)}) \lambda_s - hn \right] ds.$$

Let $J^*(n, t) = \max_{u \in \mathcal{U}} J_u(n, t)$ denote the maximal expected net revenue among policies in \mathcal{U} . Let $\hat{r}(\lambda_s) = e^{-\beta s} [r(\lambda_s) + h/\beta(1 - e^{-\beta(t-s)})\lambda_s - hn]$, and let $J^D(n, t) = \max_{\lambda_s} \int_0^t \hat{r}(\lambda_s) ds$ subject to $\int_0^t \lambda_s ds \leq n$ denote the maximal net discounted revenue when demand is deterministic. Note that $\hat{r}(\lambda_s)$ inherits the concavity of $r(\lambda_s)$, so by Theorem 2 we have

THEOREM 9.

$$J^*(n, t) \leq J^D(n, t).$$

Again, one can show that the deterministic solution is an asymptotically optimal heuristic for the stochastic problem, though in the presence of holding cost and/or discount rates, it is no longer time invariant. Indeed, let $J^D(n, t, \mu) = \max_{\lambda_s} \int_0^t [\hat{r}(\lambda_s) - \mu \lambda_s] ds + n\mu$, then $J^D(n, t) = \inf_{\mu \geq 0} J^D(n, t, \mu)$. Let μ^* denote the optimal dual variable. Then for each $s \in (0, t)$ we have $\hat{r}'(\lambda_s) = \mu^*$. Or equivalently,

$$\lambda_s = g \left(e^{\beta s} \left(\mu^* + \frac{h}{\beta} e^{-\beta t} \right) - \frac{h}{\beta} \right)$$

where $g(\cdot) \doteq \hat{r}'^{-1}(\cdot)$. Now, since $g(\cdot)$ is a decreasing function, and the argument $e^{\beta s}(\mu^* + (h/\beta)e^{-\beta t}) - h/\beta$ is increasing in s , it follows that the optimal intensity λ_s (resp., price p_s) is monotonically decreasing (resp., increasing) in $s \in (0, t)$.

At this point it is useful to isolate the holding and discounting effects. If there were no discounting, then

$$\lambda_s = g(\mu^* - h(t - s)),$$

and the argument is strictly increasing in $s \in (0, t)$ regardless of the value of μ^* . If there were no holding costs, then

$$\lambda_s = g(e^{\beta s} \mu^*)$$

and the argument is strictly increasing in $s \in (0, t)$ only if $\mu^* > 0$. From the holding cost point of view, the intuition is that we want to sell initially at a faster rate in

order to reduce the cost of holding inventories. From the discounting point of view, we are interested in the rate at which revenue is flowing in. If n is large enough so that $\mu^* = 0$, we want to sell at $\lambda^* \doteq g(0)$ to maximize the revenue rate $r(\lambda^*)$. If on the other hand, n is small enough so that $\mu^* > 0$, then we want to sell at a lower rate $\lambda_s = g(e^{\beta s} \mu^*) < \lambda^*$ to avoid running out of stock before time t . However, since we are discounting we start with higher revenue rates.

5.4. Initial Stock as a Decision Variable

Suppose we are allowed to determine the initial stock n , the order quantity, and also decide the subsequent pricing policy. If the initial stock can be purchased at a unit cost $c > 0$, we want to find the order quantity n^* that maximizes the expected profit

$$\Pi(n, t) = J^*(n, t) - cn.$$

This problem reduces to the classical newsboy problem if we replace J^* above by the expected revenue for a given fixed price. If we control this fixed price, then we obtain the problem studied by Karlin and Carr (1962) and Whitin (1955).

By Theorem 2, $J^*(n, t) \leq J^D(n, t)$; consequently an upper bound on $\Pi(n, t)$ (cf. equation (12)) is given by

$$\Pi_B(n, t) = \begin{cases} tr(\lambda^*) - cn & \text{if } n > \lambda^*t \\ tr(n/t) - cn & \text{otherwise.} \end{cases}$$

Treating n as a continuous variable, let n^c denote the maximizer of $\Pi^D(n, t)$. We see that for $n > \lambda^*t$, $\Pi^D(n, t)$ is strictly decreasing in n , so $n^c \leq \lambda^*t$. For $n \leq \lambda^*t$, $\Pi^D(n, t)$ is concave in n , so

$$n^c = \lambda^c t,$$

where $\lambda^c \doteq r'^{-1}(c) \leq \lambda^*$. Thus

$$\Pi(n^*, t) \leq \Pi^D(n^c, t) = t[r(\lambda^c) - c\lambda^c].$$

THEOREM 10. *The deterministic solution (n^c, λ^c) is asymptotically optimal as $t \rightarrow \infty$.*

PROOF OF THEOREM 10. By Theorem 3,

$$\begin{aligned} \Pi(n^c, t) &= J^*(n^c, t) - cn^c \\ &\geq \Pi^D(n^c, t) - \frac{1}{2} r(\lambda^c) \sqrt{t/\lambda^c}. \end{aligned}$$

Consequently,

$$\frac{\Pi(n^c, t)}{\Pi(n^*, t)} \geq \frac{\Pi(n^c, t)}{\Pi^D(n^c, t)} \geq 1 - \frac{r(\lambda^c)}{2(r(\lambda^c) - c\lambda^c)\sqrt{\lambda^c t}}.$$

Thus, we have

$$\lim_{t \rightarrow \infty} \frac{\Pi(n^c, t)}{\Pi(n^*, t)} = 1. \quad \square$$

REMARK. Using Theorem 1, one can show that n^* and λ^c are related in a rather interesting way, namely,

$$\lambda(n^*, t) \leq \lambda^c \leq \lambda(n^* + 1, t).$$

5.5. Resupply, Cancellations, and Overbooking

Suppose additional units can be secured at a unit cost $b > 0$, so the firm now has the option of selling beyond its initial inventory (overbooking). We view this option in one of two ways: (1) demand is satisfied by placing a special order every time a sale is made while out of stock, or (2) demand is backlogged and at time t the firm orders as many additional units as needed to satisfy the backlog. The first case is most common when items are hard goods (clothes, appliances, etc.), in which case b may represent unit transshipment costs or special handling charges. The second case applies to a model of *overbooking* in the airline and hotel industry, where b may correspond to the cost of a seat on an alternate flight or a room at an alternate hotel site (i.e., a secondary supply) or may also be a loss-of-goodwill penalty for not providing on time service.

Overbooking is often practiced to compensate for cancellations. Here we assume that each reservation is canceled independently, at time t , with probability $1 - \rho$. In addition, we assume that customers who cancel are refunded the purchase price less a penalty, which consists of a fixed plus variable component. Specifically, let c represent the fixed fee and β represent the fraction of the price paid that composes the variable fee. Thus, a customer who pays price p and cancels gets a refund of $p(1 - \beta) - c$. (See Bitran and Gilbert 1992 and Liberman and Yechiali 1978 for alternative models that consider cancellations and overbooking.)

Given a non-anticipating intensity control policy λ_s based on the initial inventory n and the current history of reservations, the number of reservations N_s is Poisson with random intensity $\int_0^s \lambda_v dv$. Let $\{T_k: k \geq 1\}$ denote the jump points of the counting process N_s , $0 \leq s \leq t$, and let $\{Z_k: k \geq 1\}$ be a sequence of independent Bernoulli random variables taking value 1 with probability ρ and taking value 0 with probability $1 - \rho$. By our above assumptions about the cancellation process, these

random variables are also independent of the counting process N_s .

We assume that revenues are collected as reservations are made and refunds for canceled reservations are paid at the end of the horizon. If we disregard the time value of money, the net expected revenue can be written as

$$E \sum_{k \geq 1} p(\lambda_{T_k}) 1(T_k \leq t) 1(Z_k = 1) = E \int_0^t \rho r(\lambda_s) ds.$$

If the firm imposes a $100\beta\%$ penalty of the price paid for each canceled reservation, then the expected net revenue is obtained by replacing ρ by $\rho + \beta(1 - \rho)$ above. In addition, if each canceled reservation is subject to a fixed penalty c , then we add to the expected net revenue the quantity

$$Ec \sum_{k \geq 1} 1(T_k \leq t) 1(Z_k = 0) = E \int_0^t \sum c(1 - \rho) \lambda_s ds.$$

The number of uncanceled reservations is

$$\sum_{k \geq 1} 1(T_k \leq t) 1(Z_k = 1) = \int_0^t d\bar{N}_s$$

where \bar{N}_s is Poisson with random intensity $\rho \int_0^s \lambda_v dv$. If $\int_0^t d\bar{N}_s > n$, we must purchase $(\int_0^t d\bar{N}_s > n)^+$ additional units at b dollars each. Therefore, the expected net revenue under a nonanticipating policy u is

$$V_u(n, t) = E_u \int_0^t (\rho + \beta(1 - \rho)) r(\lambda_s) ds + E_u \int_0^t c(1 - \rho) \lambda_s ds - bE_u \left(\int_0^t d\bar{N}_s - n \right)^+,$$

where

$$V_u(m, 0) \doteq \begin{cases} 0 & m \geq 0 \\ bE(X_m - n)^+ & m < 0, \end{cases} \quad (20)$$

n denotes the initial inventory (capacity), m denotes the possibly negative unsold capacity at time t before learning about cancellations, and X_m is a binomial random variable with parameters $n - m$ and ρ . Thus, $X_m - n$, if positive, is the number of uncanceled reservations in excess of the initial capacity.

As before, the firm's problem is to find a pricing policy u^* (if one exists) that achieves an expected revenue

$$V^*(n, t) = \sup_{u \in \mathcal{U}} V_u(n, t), \quad (21)$$

where we let \mathcal{U} denote the class of all Markovian policies satisfying $p_s \in \mathcal{P}$, $\forall s$.

In the next subsection we find a closed-form solution to the stochastic problems when demand is exponentially decaying and no cancellations occur ($\rho = 1$). We then solve the deterministic counterpart for the general case and present an asymptotically optimal heuristic.

5.5.1. An Optimal Policy for the Exponential Demand Function with no Cancellations. Let $\lambda(p) = ae^{-\alpha p}$ and $\rho = 1$, and $b > 0$. This case corresponds to having no cancellations and a unit reorder cost. It perhaps most appropriate for applications where items are hard goods and the cost b is a per-unit special-order cost or per-unit transshipment cost for obtaining additional units. One can verify that in this case $V^*(n, t)$ is the solution to equation (8) with boundary conditions $V^*(n, t) = 0$ if $n \geq 0$ and $V^*(n, t) = nb$ if $n < 0$. As before, without loss of generality we take $\alpha = 1$. Let $\lambda^b \doteq \operatorname{argmax} (r(\lambda) - \lambda b) = \lambda^* e^{-b}$. Then $V^*(n, t)$ is given by

$$V^*(n, t) = \begin{cases} \log \left(\sum_{i=0}^n \frac{(\lambda^* t)^i}{i!} + e^{nb} \sum_{i=n+1}^{\infty} \frac{(\lambda^* t)^i}{i!} \right), & n > 0 \\ \lambda^b t + nb & n \leq 0 \end{cases}$$

and the optimal price is given by $p^*(n, t) = V^*(n, t) - V^*(n - 1, t) + 1$.

Note, for $n > 0$ we can write,

$$\exp(V^*(n, t)) = \exp(J^*(n, t)) + e^{nb} \sum_{i=n+1}^{\infty} \frac{(\lambda^* t)^i}{i!},$$

where $J^*(n, t)$ is the optimal revenue with no reorder option (the basic problem), and the second term above is always nonnegative. Thus, the expected revenue is strictly greater than without the reorder option as expected. The price trajectory itself has characteristics similar to the basic problem ($b = \infty$), taking upward jumps as items are sold and decaying as time elapses without a sale. The exception is when the inventory drops to zero, at which point the policy switches to a fixed price of $b + 1$.

5.5.2. An Asymptotically Optimal Heuristic for the General Case. For the general case with both cancellations and reordering, the deterministic problem corresponding to equation (21) can be written as

$$V^D(x, t) = \max_{\lambda(s)} \int_0^t (\rho + \beta(1 - \rho))r(\lambda(s))ds + \int_0^t c(1 - \rho)\lambda(s)ds - b\left(\rho \int_0^t \lambda(s)ds - x\right)^+.$$

Now $V^*(x, t) \leq V^D(x, t)$ follows by applying Jensen's inequality to the third term of $V^u(x, t)$ and by viewing the integrand inside the expectation as purely a function of λ and maximizing pointwise.

To solve the deterministic problem we need to introduce notation that is pertinent only to this section. Let

$$\hat{r}(\lambda) \doteq (\rho + \beta(1 - \rho))r(\lambda) + c(1 - \rho)\lambda, \quad (22)$$

denote the modified revenue rate. Let $\lambda^0 \doteq x/(\rho t)$ be the expected-run-out rate. At rate λ^0 , we book $\lambda^0 t = x/\rho \geq x$ units over the horizon, of which $\rho \lambda^0 t = x$ show at time t . Let $p^0 \doteq p(\lambda_0)$ be the expected-run-out price, and $\hat{r}^0 \doteq \hat{r}(\lambda_0)$. Let λ^* denote the least maximizer of $\hat{r}(\lambda)$, $p^* \doteq p(\lambda^*)$ its corresponding price, and $\hat{r}^* \doteq \hat{r}(\lambda^*)$. Finally, let λ^b be the least maximizer of $\hat{r}(\lambda) - b\rho\lambda$, $p^b \doteq p(\lambda^b)$ its corresponding price, and $\hat{r}^b \doteq \hat{r}(\lambda^b)$. The following proposition is given without proof:

PROPOSITION 5. *The optimal solution to the deterministic problem (22) is*

$$p_D(s) = \begin{cases} p^* & \rho \lambda^* t \leq x \\ p^0 & \rho \lambda^b t \leq x < \rho \lambda^* t \\ p^b & x < \rho \lambda^b t \end{cases} \quad 0 \leq s \leq t, \quad (23)$$

$$\lambda_D(s) = \begin{cases} \lambda^* & \rho \lambda^* t \leq x \\ \lambda^0 & \rho \lambda^b t \leq x < \rho \lambda^* t \\ \lambda^b & x < \rho \lambda^b t \end{cases} \quad 0 \leq s \leq t, \quad (24)$$

and

$$V^D(x, t) = \begin{cases} \hat{r}^* t & \rho \lambda^* t \leq x \\ \hat{r}^0 t & \rho \lambda^b t \leq x < \rho \lambda^* t \\ (\hat{r}^b - b\rho \lambda^b)t + xb & x < \rho \lambda^b t. \end{cases} \quad (25)$$

Thus if capacity is high ($x \geq \rho \lambda^* t$), we price to maximize the modified revenue rate (equation (22)). If capacity is low ($x < \rho \lambda^b t$), we price at p^b , since in this case λ^b maximizes the modified profit rate $\hat{r}(\lambda) - b\rho\lambda$. For intermediate capacity ($\rho \lambda^b t \leq x < \rho \lambda^* t$) we price at the expected-run-out price.

REMARK. If $\rho b > \hat{r}'(0)$, then $\lambda^b = 0$, and the solution reduces to the case with no reorder option provided $x \geq 0$.

Notice that the deterministic solution consists of a fixed price over the entire horizon. Consider the fixed-price (FP) heuristic that prices according to the deterministic intensities $\lambda_D(s)$, $0 \leq s \leq t$. The expected value of the FP heuristic is given by

$$V^{FP}(x, t) = \int_0^t (\rho + \beta(1 - \rho))r(\lambda_D(s))ds + \int_0^t c(1 - \rho)\lambda_D(s)ds - Eb(\hat{N}_t - x)^+,$$

where \hat{N}_t is Poisson with intensity $\int_0^t \rho \lambda_D(v)dv$. Note that the first two terms of $V^{FP}(x, t)$ are equal to those of $V^D(x, t)$. To establish the asymptotic optimality of $V^{FP}(x, t)$, we need a slight variant of (18). Let N be a random variable with mean μ and variance σ^2 , writing $(N - x)^+ = \frac{1}{2}(|N - x| + (N - x))$, taking expectations and using the Cauchy-Schwartz inequality, we obtain

$$E(N - x)^+ \leq \frac{1}{2} \sqrt{\sigma^2 + (\mu - x)^2} + \frac{1}{2}(\mu - x) \leq \frac{1}{2}\sigma + \frac{1}{2}(|\mu - x| + (\mu - x)) = \frac{1}{2}\sigma + \frac{1}{2}(\mu - x)^+.$$

We can now state

THEOREM 11.

$$\frac{V^{FP}(n, t)}{V^*(n, t)} \geq \begin{cases} 1 - \frac{b\sqrt{\rho \lambda^* t}}{2\hat{r}^* t} & \rho \lambda^* t \leq n \\ 1 - \frac{b\sqrt{\rho \lambda^0 t}}{2\hat{r}^0 t} & \rho \lambda^b t \leq n < \rho \lambda^* t \\ 1 - \frac{b\sqrt{\rho \lambda^b t}}{2(\hat{r}^b - b\rho \lambda^b)t + 2nb} & n < \rho \lambda^b t. \end{cases}$$

PROOF. Applying the above bound to $E(\hat{N}_t - x)^+$ in $V^{FP}(n, t)$ we obtain

$$Eb(\hat{N}_t - n)^+ \leq \frac{1}{2} b\sqrt{\rho \lambda_D t} + \frac{1}{2} b(\rho \lambda_D t - n)^+.$$

Consequently, $V^{FP}(n, t) \geq V^D(n, t) - \frac{1}{2}b\sqrt{\rho\lambda_D}t$. The result follows after dividing by $V^D(n, t)$. \square

By observing that when $n \leq 0$, $E[(N_t - n)^+] = \lambda^b t - n$, we obtain the following corollary to Theorem 11:

COROLLARY 1. *If $n \leq 0$, then*

$$V^{FP}(n, t) = V^*(n, t).$$

That is, when there are no items in inventory, the optimal policy is to fix the price at p^b . The reason for this is that backlogged items represent a sunk cost that cannot be influenced by our pricing policy. Thus, we ignore n and simply try to maximize the net revenue rate $\hat{r}(\lambda) - \rho\lambda b$ over the remaining time, which implies pricing at p^b .

6. Conclusions

We have shown how a range of inventory pricing problems can be analyzed using intensity control theory, bounds, and heuristics. By analyzing the deterministic version of different versions of the basic problem, we were able to obtain both upper bounds on the expected revenue and insights into the form of near-optimal policies. Exact optimal policies were found in certain cases for a family of exponential demand functions. Perhaps the strongest conclusion from our results is that using simple fixed-price policies appears to work surprisingly well in many instances. This is encouraging since the optimal dynamic policies are quite jittery and require constant price adjustments, an undesirable characteristic in practical applications. In the discrete-price case, we showed that a policy that varies the allocation of units and time to two neighboring prices is nearly optimal. The policy provides a good explanation of the structure of current yield management practice.

We believe that this class of inventory pricing models represents a fertile area for future research. From a practical standpoint, revenue maximization holds the potential for dramatic improvements in profitability and thus is likely to be a topic of intense interest to managers in a wide range of industries. On a methodological level, we think that formulating problems in the framework of intensity control is a promising approach. Though exact solutions appear limited to special cases, one can easily obtain bounds similar to those in Theorem 2 that relate the stochastic and deterministic variants of the problem. No doubt other variants of the problem can

be attacked using precisely this approach. Similar bounds could potentially be useful for a wide range of intensity control problems in other application contexts as well.²

²We thank three anonymous referees and the associate editor for providing several references and many helpful comments. The research of G. Gallego was supported in part by the National Science Foundation grant DDM 9109636. The research of both authors was supported by the National Science Foundation grant SES93-09394.

Appendix

PROOF OF PROPOSITION 1. We first show that the supremum in equation (8) can be replaced by $\max_{\lambda_i \in [0, \lambda^*]}$. To do so, let λ_i be an arbitrary intensity satisfying $\lambda_i > \lambda^*$. By concavity of $r(\lambda)$ and the definition of λ^* , we have $r(\lambda^*) \geq r(\lambda_i)$, and since $J(n, t)$ is non-decreasing in n , we have

$$r(\lambda^*) - \lambda^*[J(n, t) - J(n-1, t)] \geq r(\lambda_i) - \lambda_i[J(n, t) - J(n-1, t)].$$

Hence the optimal choice of λ is always within the set $[0, \lambda^*]$, a compact set. Combining compactness with the fact that $r(\lambda)$ is continuous and bounded establishes the conditions required by Bremaud Theorem II.3 for the existence of a unique solution to equation (8).

PROOF OF THEOREM 1. The fact that $J^*(n, t)$ is strictly increasing in n and t is straightforward to show, so we omit the details. We next show by induction that $\lambda^*(n, t)$ is strictly decreasing in t and in so doing establish that $\lambda^*(n, t)$ is strictly increasing in n and that $J^*(n, t)$ is strictly concave in both n and t . The results for $p^*(n, t)$ follow from the fact that $\lambda(p)$ is a regular demand function.

We begin with the case $n = 1$. Note that from equation (8),

$$J^*(n, t) = J^*(n-1, t) + r'(\lambda^*(n, t)) \geq J^*(n-1, t) \quad (26)$$

with strict inequality holding when $t > 0$. For $n = 1$, observe that $J^*(1, t) = r'(\lambda^*(1, t))$. Thus, since $J^*(1, t)$ is strictly increasing in t , we have

$$0 < \frac{\partial J^*(1, t)}{\partial t} = r''(\lambda^*(1, t))\lambda^{*'}(1, t),$$

which along with the concavity of $r(\cdot)$ implies that $\lambda^*(1, t)$ is strictly decreasing in t . It also follows from equation (8) that

$$\frac{\partial J^*(1, t)}{\partial t} = r(\lambda^*(1, t)) - \lambda^*(1, t)J^*(1, t).$$

Again, combining this with the fact that $J^*(1, t) = r'(\lambda^*(1, t))$, we find

$$\frac{\partial^2 J^*(1, t)}{\partial t^2} = -\lambda^*(1, t) \frac{\partial J^*(1, t)}{\partial t} < 0,$$

which shows $J^*(1, t)$ is strictly concave in t . Thus, all of the claimed properties hold for $n = 1$.

Next assume that $\lambda^*(n-1, t)$ is strictly decreasing in t . From equation (26) we see that $r'(\lambda^*(n, t)) > 0$ for $t > 0$, implying $\lambda^*(n, t) < \lambda^*$. Note that as t approaches zero from the right in problem (26), $\lim_{t \rightarrow 0} r'(\lambda^*(n, t)) = 0$, so $\lambda^*(n, 0^+) = \lambda^*$. Hence, $\lambda^*(n, t)$ is initially strictly decreasing in t . Now assume for the sake of contradiction that $\lambda^*(n, t)$ is strictly decreasing over $[0, t_0]$ but is nondecreasing over a nonempty interval $[t_0, t_1]$. Taking derivatives with respect to t in equation (8), we find that over $[0, t_0]$

$$\frac{\partial J^*(n, t)}{\partial t} > \frac{\partial J^*(n-1, t)}{\partial t} \quad (27)$$

with the opposite inequality holding over $[t_0, t_1]$. From this and equation (8), it then follows that over $[0, t_0]$, $J^*(n, t) - J^*(n-1, t) < J^*(n-1, t) - J^*(n-2, t)$, and consequently $\lambda^*(n, t) > \lambda^*(n-1, t)$, again with the opposite inequalities holding over $[t_0, t_1]$. But this implies that $\lambda^*(n-1, t)$ must be nondecreasing in the neighborhood of t_0 , which contradicts the inductive hypothesis. Therefore, we conclude that $\lambda^*(n, t)$ must be strictly decreasing in t and that $\lambda^*(n, t) > \lambda^*(n-1, t)$.

We now use these facts to show concavity of $J^*(n, t)$. Indeed, the fact that $\lambda^*(n, t)$ is strictly increasing in n , the concavity of $r(\cdot)$ and (26) imply

$$J^*(n, t) - J^*(n-1, t) < J^*(n-1, t) - J^*(n-2, t),$$

so $J^*(n, t)$ is strictly concave in n . Also, equations (8) and (27) imply

$$\frac{\partial^2 J^*(n, t)}{\partial t^2} = \lambda^*(n, t) \left[\frac{\partial J^*(n, t)}{\partial t} - \frac{J^*(n-1, t)}{\partial t} \right] < 0,$$

so $J^*(n, t)$ is strictly concave in t . The claims for $p^*(n, t)$ follow directly from the results for $\lambda^*(n, t)$.

PROOF OF PROPOSITION 2. Consider the deterministic problem (11). Note that the integrand in problem (11) is simply the revenue function, $r(\lambda)$, which is concave by assumption. There are two cases. First, suppose the maximizer of $r(\lambda)$, λ^* , satisfies $\lambda^*t \leq x$, then clearly $\lambda_s = \lambda^*$, $0 \leq s \leq t$ is the optimal solution since this choice maximizes the integrand pointwise. In the second case, $\lambda^*t > x$, it follows from the fact that $r(\lambda)$ is concave that for a given value $y = \int_0^t \lambda_s ds$, $\lambda_s = y/t$, $0 \leq s \leq t$ maximizes the integral. The maximum revenue given y is therefore $t(y/t)p(y/t) = tp(y/t)$. Now since $y/t < \lambda^*$ and $r(\lambda)$ is increasing for $\lambda < \lambda^*$, it follows that $y = x$ in any optimal solution, and thus $\lambda_s = (x/t) = \lambda^0$, $0 \leq s \leq t$ maximizes the integral. Converting these rates to their corresponding prices and computing the corresponding total revenue associated with this solution establishes the proposition.

PROOF OF THEOREM 5. Notice that τ is a stopping time, since τ is finite with probability one and the event $\tau < s$, can be determined by the history of the arrivals up to time s . To obtain a lower bound on $J^{ST}(n, t)$ consider a wasteful heuristic that reserves m (resp., $n-m$) units to be priced at p_k (resp., p_{k+1}) over t_m (resp., $t-t_m$) units of time. Let $J^W(n, t)$ denote the expected revenue of the wasteful heuristic. Evidently, the wasteful heuristic is a lower bound on the stopping-time heuristic since if $\tau = T_m < t_m$ the wasteful heuristic delays the selling of the remaining $n-m$ units until time t_m . On the other hand, if $\tau = t_m < T_m$ more than $n-m$ units are left at time t_m ,

and the wasteful heuristic only makes $n-m$ of them available for sale at price p_{k+1} . In spite of these limitations, we will show that the wasteful heuristic, and consequently the ST heuristic, is asymptotically optimal.

To do this, let $t_{n-m} \doteq n-m/\lambda_{k+1}$. Note that t_{n-m} is the time it takes to sell $n-m$ items at price p_{k+1} when the demand is deterministic. Consequently, $t' = t_m + t_{n-m}$ is the total time it takes to dispose of the n items when the demand rates are deterministic. Observe that by our choice of m we have

$$t - \frac{\lambda_k - \lambda_{k+1}}{\lambda_k \lambda_{k+1}} < t' \leq t. \quad (28)$$

Consequently, a lower bound on the wasteful heuristic can be obtained by delaying the start of the sales by $t-t'$ so that effectively the horizon is shrunk to t' . Recall that the deterministic revenue is

$$J^D(n, t) = \frac{r_k - r_{k+1}}{\lambda_k - \lambda_{k+1}} n + \frac{\lambda_k r_{k+1} - \lambda_{k+1} r_k}{\lambda_k - \lambda_{k+1}} t,$$

so by equation (28) $J^D(n, t) < J^D(n, t') + (p_{k+1} - p_k)$. We thus have

$$\frac{J^{ST}(n, t)}{J^*(n, t)} \geq \frac{J^W(n, t)}{J^D(n, t)} \geq \frac{J^W(n, t')}{J^D(n, t') + (p_{k+1} - p_k)}. \quad (29)$$

Now if $t \rightarrow \infty$ with $\lambda_k t \geq n > \lambda_{k+1} t$. Then, by construction, $t' \rightarrow \infty$ with $\lambda_k t' \geq n > \lambda_{k+1} t'$. Evidently $J^W(n, t') \rightarrow \infty$ and $J^D(n, t') \rightarrow \infty$ as $t \rightarrow \infty$, hence, if we can show that

$$\lim_{t \rightarrow \infty} \frac{J^W(n, t')}{J^D(n, t')} = 1,$$

we can conclude that

$$\lim_{t \rightarrow \infty} \frac{J^W(n, t')}{J^D(n, t') + (p_{k+1} - p_k)} = 1$$

and by (29) that the ST heuristic is asymptotically optimal.

Notice that showing this first limit is equivalent to showing it holds for a subsequence $\{t'_m \doteq m/(\alpha\lambda_k), m = 1, \dots\}$ where

$$m = [\alpha\lambda_k t'_m] = \alpha\lambda_k t'_m.$$

Thus, we drop the prime notation and assume that $m = \alpha\lambda_k t$ and that $n-m = \bar{\alpha}\lambda_{k+1}t$ are integers.

Let N_{λ_s} be a Poisson random variable with rate λ_s . Clearly the expected revenue for the wasteful heuristic is

$$J^W(n, t) = p_k E \min \{N_{\lambda_k t}, \alpha\lambda_k t\} + p_{k+1} E \min \{N_{\bar{\alpha}\lambda_{k+1}t}, \bar{\alpha}\lambda_{k+1}t\}.$$

Noting that $E \min (N_{\lambda_s}, \lambda_s) = EN_{\lambda_s} - E(N_{\lambda_s} - \lambda_s)^+$, and that $EN_{\lambda_s} = \text{Var}(N_{\lambda_s}) = \lambda_s$, and using equation (18) we obtain the following lower bound on the performance of the wasteful heuristic.

$$J^W(n, t) \geq p_k[\alpha\lambda_k t - 1/2\sqrt{\alpha\lambda_k t}] + p_{k+1}[\bar{\alpha}\lambda_{k+1}t - 1/2\sqrt{\bar{\alpha}\lambda_{k+1}t}].$$

From the deterministic solution, we know that $J^D(n, t) = p_k\alpha\lambda_k t + p_{k+1}\bar{\alpha}\lambda_{k+1}t$. Taking ratios, we observe that

$$\frac{J^W(n, t)}{J^D(n, t)} \geq 1 - 1/2 \left[\frac{1}{\sqrt{\alpha\lambda_k t}} + \frac{1}{\sqrt{\bar{\alpha}\lambda_{k+1}t}} \right],$$

which establishes the result.

References

- Anderson, S. P., A. de Palma, and J. F. Thisse, *Discrete Choice Theory of Product Differentiation*, The MIT Press, Cambridge, MA, 1992.
- Bass, F. M., "The Relationship Between Diffusion Rates, Experience Curves, and Demand Elasticities for Consumer Durable Technological Innovations," *J. Business*, 53 (1980), S51-S67.
- Belobaba, P. P., "Airline Yield Management: An Overview of Seat Inventory Control," *Transportation Sci.*, 21 (1987), 63-73.
- , "Application of a Probabilistic Decision Model to Airline Seat Inventory Control," *Oper. Res.*, 37 (1989), 183-197.
- Bitran, G. R. and S. M. Gilbert, "Managing Hotel Reservations with Uncertain Arrivals," MIT Sloan School Working Paper (1992).
- Bremaud, P., *Point Processes and Queues, Martingale Dynamics*, Springer-Verlag, NY, 1980.
- Brumelle, S. L. et al., "Allocation of Airline Seats Between Stochastically Dependent Demand," *Transportation Sci.*, 24 (1990), 183-192.
- Chakravarty, A. K. and G. E. Martin, "Discount Pricing Policies for Inventories Subject to Declining Demand," *Naval Res. Logistics*, 36 (1989), 89-102.
- Cohen, M. A., "Joint Pricing and Ordering Policy for Exponentially Decaying Inventory with Known Demand," *Naval Res. Logistics*, 24 (1977), 257-268.
- Curry, R. E., "Optimal Airline Seat Allocation With Fare Classes Nested by Origins and Destinations," Aeronomics Incorporated Report, 1989.
- Dockner, E. and S. Jorgensen, "Optimal Pricing Strategies for New Products in Dynamic Oligopolies," *Marketing Sci.*, 7, 4 (1988), 315-333.
- Dolan, T. J. and A. P. Jeuland, "Experience Curves and Dynamic Demand Models: Implications for Optimal Pricing Strategies," *J. Marketing*, 45 (1981), 52-73.
- Eliashberg, J. and R. Steinberg, "Marketing-Production Joint Decision Making," in J. Eliashberg and J. D. Lilien (Eds.), *Management Science in Marketing*, Handbooks in Operations Research and Management Science, North Holland, 1991.
- and —, "Competitive Strategies for Two Firms with Asymmetric Production Cost Structures," *Management Sci.*, 37, 11 (1991), 1452-1473.
- Feng, Y. and G. Gallego, "Optimal Stopping (Starting) Times for Promotional Fares (Sales)," to appear in *Management Sci.*
- Gallego, G., "A Minmax Distribution Free Procedure for the (Q, R) Inventory Model," *Oper. Res. Letters*, 11 (1992), 55-60.
- Glover, F., R. Glover, J. Lorenzo, and C. McMillan, "The Passenger-Mix Problem in the Scheduled Airlines," *Interfaces*, 12 (1982), 73-79.
- Harris, M. and A. Raviv, "A Theory of Monopoly Pricing Schemes with Demand Uncertainty," *American Economic Review*, 71 (1981), 347-365.
- Hempenius, A. L., *Monopoly with Random Demand*, Rotterdam University Press, Rotterdam, The Netherlands, 1970.
- Jeuland, A. P. and R. J. Dolan, "An Aspect of New Product Planning: Dynamic Pricing," in A. A. Zoltners (Ed.), *Marketing Planning Models*, TIMS Studies in the Management Sciences, 18, North Holland, NY, 1982.
- Kalish, S., "Monopolist Pricing with Dynamic Demand and Production Costs," *Marketing Sci.*, 2 (1983), 135-160.
- Karlin, S. and C. R. Carr, "Prices and Optimal Inventory Policies," in K. J. Arrow, S. Karlin and H. Scarf (Eds.), *Studies in Applied Probability and Management Science*, Stanford University Press, Stanford, CA, 1962.
- Kimes, S. E., "The Basics of Yield Management," *Cornell H.R.A. Quarterly*, 30, 6 (1989), 14-19.
- Kinberg, Y. and A. G. Rao, "Stochastic Models of Price Promotion," *Management Sci.*, 21 (1975), 897-907.
- Kincaid, W. M. and D. Darling, "An Inventory Pricing Problem," *J. Mathematical Analysis and Applications*, 7 (1963), 183-208.
- Kunreuther, H. and J. F. Richard, "Optimal Pricing and Inventory Decisions for Non-Seasonal Items," *Econometrica*, 39, 1 (1971), 173-175.
- Lazear, D. P., "Retail Pricing and Clearance Sales," *American Economics Review*, 76, 1 (1986), 14-32.
- Li, L., "A Stochastic Theory of the Firm," *Mathematics of Operations Research*, 13, 3 (1988), 447-466.
- Lieberman, V. and U. Yechiali, "On the Hotel Overbooking Problem—An Inventory System with Stochastic Cancellations," *Management Sci.*, 24, 11 (1978), 1117-1126.
- Littlewood, K., "Forecasting and Control of Passenger Bookings," *AGIFORS Symposium Proceedings*, (1972), 95-117.
- Luenberger, D. G., *Optimization by Vector Space Methods*, John Wiley & Sons, NY, 1969.
- Miller, B. L., "Finite State Continuous Time Markov Decision Processes with a Finite Planning Horizon," *SIAM J. Control*, 6 (1968), 266-280.
- Mills, E. S., "Uncertainty and Price Theory," *Quarterly J. Economics*, 73 (1959), 117-130.
- Nagle, T. T., *The Strategy & Tactics of Pricing, A Guide to Profitable Decision Making*, Prentice Hall, Englewood Cliffs, NJ, 1987.
- Oren, S. S., "Comments on 'Pricing Research in Marketing: The State of the Art'," *J. Business*, 57 (1984), S61-S64.
- Pashigan, P. B., "Demand Uncertainty and Sales: A Study of Fashion and Markdown Pricing," *American Economic Review*, 78, 5 (1988), 936-953.
- and B. Bowen, "Why Are Products Sold on Sale?: Explanations of Pricing Regularities," *The Quarterly J. Economics*, Nov. (1991), 1014-1038.
- Rajan, A., Rakesh and R. Steinberg, "Dynamic Pricing and Ordering Decisions by a Monopolist," *Management Sci.*, 38, 2 (1992), 240-262.
- Rao, V. R., "Pricing Research in Marketing: The State of the Art," *J. Business*, 57 (1984), S39-S64.
- Robinson, B. and C. Lakhani, "Dynamic Pricing Models for New Product Pricing," *Management Sci.*, 21 (1975), 1113-1122.
- Robinson, L. W., "Optimal and Approximate Control Policies for Airline Booking with Sequential Fare Classes," Cornell University Working Paper, 1991.
- Rothstein, M., "Hotel Overbooking as a Markovian Sequential Decision Process," *Decision Sci.*, 5 (1974), 389-404.
- Smith, B., J. Leimkuhler, R. Darrow, and J. Samuels, "Yield Management at American Airlines," *Interfaces*, 1 (1992), 8-31.

GALLEGO AND VAN RYZIN
Optimal Dynamic Pricing of Inventories

- Stadje, W., "A Full Information Pricing Problem for the Sale of Several Identical Commodities," *Zeitschrift für Operations Research*, 34 (1990), 161-181.
- Thomas, L. J., "Price-Production Decisions with Deterministic Demand," *Management Sci.*, 16 (1970), 747-750.
- Thomas, L. J., "Price and Production Decisions with Stochastic Demand," *Oper. Res.*, 22 (1974), 513-518.
- Thowsen, G. T., "A Dynamic, Nonstationary Inventory Problem for a Price/Quantity Setting Firm," *Naval Res. Logistics*, 22 (1975), 461-476.
- Varian, H. R., *Microeconomic Analysis*, second ed., W. W. Norton & Company, NY, 1984.
- Veinott, A., Jr., Chapter 2, Unpublished Class Notes on Inventory Theory, 1980.
- Wagner, H. M., "A Postscript to 'Dynamic Problems in the Theory of the Firm,'" *Naval Res. Logistics*, 7 (1960), 7-13.
- and T. M. Whitin, "Dynamic Problems in The Theory of The Firm," *Naval Res. Logistics*, 5 (1958), 53-74.
- Weatherford, L. R. and S. E. Bodily, "A Taxonomy and Research Overview of Perishable-Asset Revenue Management: Yield Management, Overbooking and Pricing," *Oper. Res.*, 40, 5 (1992), 831-844.
- Whitin, T. M., "Inventory Control and Price Theory," *Management Sci.*, 2 (1955), 61-68.
- Wollmer, R. D., "An Airline Seat Management Model for a Single Leg Route When Lower Fare Classes Book First," *Oper. Res.*, 40, 1 (1992), 26-37.
- Zabel, E., "Multi-Period Monopoly Under Uncertainty," *J. Economic Theory*, 5 (1972), 524-536.

Accepted by Linda Green; received March 5, 1992. This paper has been with the authors 3 months for 2 revisions.