OPERATIONS RESEARCH
Vol. 56, No. 2, March–April 2008, pp. 326–343
ISSN 0030-364X | EISSN 1526-5463 | 08 | 5602 | 0326

DOI 10.1287/opre.1070.0438

Risk in Revenue Management and Dynamic Pricing

Yuri Levin, Jeff McGill, Mikhail Nediak
School of Business, Queen's University, Kingston, Ontario, Canada K7L 3N6
{ylevin@business.queensu.ca, jmcgill@business.queensu.ca, mnediak@business.queensu.ca}

We present a new model for optimal dynamic pricing of perishable services or products that incorporates a simple risk measure permitting control of the probability that total revenues fall below a minimum acceptable level. The formulation assumes that sales must occur within a finite time period, that there is a finite—possibly large—set of available prices, and that demand follows a price-dependent, nonhomogeneous Poisson process. This model is particularly appropriate for applications in which attainment of a revenue target is an important consideration for managers; for example, in event management, in seasonal clearance of high-value items, or for business subunits operating under performance targets. We formulate the model as a continuous-time optimal control problem, obtain optimality conditions, explore structural properties of the solution, and report numerical results on problems of realistic size.

Subject classifications: inventory/production: perishable/aging items, marketing/pricing, uncertainty; dynamic programming/optimal control: applications; probability: stochastic model applications.
Area of review: Manufacturing, Service, and Supply Chain Operations.
History: Received December 2005; revisions received August 2006, February 2007; accepted March 2007. Published online in *Articles in Advance* December 3, 2007.

1. Introduction

In recent years, the widely reported successes of revenue management in the airline and hotel industries and the expansion of online booking and retail sales systems have stimulated interest in revenue management and dynamic pricing in many new areas. Revenue management applications have been implemented or proposed in such diverse areas as freight transportation, automobile rental services, broadcast advertising, sports and entertainment event management, style goods inventory clearance, medical services, and real estate. Details of implementation across these areas differ, but most share a basic set of attributes: a finite supply of some inventory item, a fixed time span over which sales can occur, and stochastically varying demand. A typical objective is to maximize expected revenues by one or a combination of two approaches: (1) offering multiple product "classes" at different prices and varying the allocation of fixed inventory to those classes over time, or (2) offering a single product class and dynamically varying the price over time.

There is an extensive literature on revenue management and related practices. For surveys, see McGill and van Ryzin (1999), Belobaba (1987), Weatherford and Bodily (1992), Bitran and Caldentey (2003), Chan et al. (2004), Elmaghraby and Keskinocak (2003), and Yano and Gilbert (2004). A broad discussion of various aspects of revenue management can by found in books by Talluri and van Ryzin (2004) and Phillips (2005).

Traditional revenue management optimization models in transportation and accommodation services are *risk neutral*—the objective is to maximize expected revenues at the end of the disposal period without consideration for the variability of final revenues across realized problem instances. This is often appropriate in these applications because inventory control and pricing strategies are implemented over hundreds or thousands of problem instances (e.g., successive flight departures), and no single realization has a potential for severe impact on the company revenues. In these situations, the law of large numbers ensures that long-term average revenues will be maximized as long as good risk-neutral strategies are employed.

This is not the case for other potential applications. For example, consider the case of an event promoter who may organize a few large events per year. In this case, the promoter faces a very high fixed cost for rental of a stadium or concert hall, and her first priority will be to recover this cost. Also, problem instances are infrequent, and a single poor realization can have a major impact on the financial condition of the business. In other businesses, a manager's primary concern may be maintenance or expansion of market share because a slide in share can lead to negative stock-market assessments that can far outweigh the marginal revenues involved in revenue management decisions. Both of these circumstances are encountered by consultants in pricing and revenue management (Phillips 2006). In their description of a potential implementation of revenue management in the health care industry, Secomandi et al. (2002, p. 7) comment: "...the hospital's most important control of its revenue stream may be its ability to negotiate and implement contracts with payors to meet its revenue goals while minimizing risk." Furthermore, businesses with multiple divisions or profit centers typically set performance targets for their subunits as a part of overall strategy. Managers of subunits cannot afford to ignore minimum targets in their search for higher revenues. In these situations and others, managers may be willing to pursue pricing policies that sacrifice some expected revenues in return for a higher probability of attaining a minimum revenue goal.

In this paper, we consider the problem of dynamically pricing a stock of items over a finite time horizon so that both the expected revenues and the risk of poor performance are taken into account. We assume that demand for the product follows a nonhomogeneous Poisson process in which the demand rate is price sensitive, where prices are drawn from an arbitrary but finite predetermined set. We formulate a dynamic pricing model incorporating risk, which we analyze as a problem of optimal control.

Risk is introduced to the model by augmenting the expected revenue objective with a penalty term for the probability that total revenues fall below a desired level of revenue—a *loss-probability risk measure*. This risk measure has significant advantages: First, it is much more readily interpreted by managers than the conventional variance measure; second, it is appropriate for both symmetric and skewed total revenue distributions; and third, it leads in the present case to a tractable formulation that can be implemented on problems of realistic size.

The introduction of risk in this form does, however, have a significant computational drawback. In conventional dynamic inventory control problems of this type, the state of the system at any point in time can be summarized by a single state variable—remaining inventory—whereas in the present case, it is necessary to augment the state description with a second variable—current revenues. Nonetheless, we show in this paper that a computationally tractable formulation is possible even with a two-dimensional state space.

Prior studies most relevant to this paper are those on dynamic pricing by Gallego and van Ryzin (1994), Feng and Xiao (2000a, b), Zhao and Zheng (2000), and Chatwin (2000). In particular, Feng and Xiao (1999) introduced a risk factor in a dynamic pricing framework. They considered a model with two predetermined prices, and instead of looking at the expected revenue alone, used an objective function that reflects changes in the revenue variance as a result of price changes. Chen et al. (2006) used utility theory to develop a general framework for incorporating risk aversion in multiperiod inventory models as well as multiperiod models that coordinate inventory and pricing strategies. Other references on risk-averse inventory models include Agrawal and Seshadri (2000), Chen and Federgruen (2000), Eeckhoudt et al. (1995), and Martínez-de Albéniz and Simchi-Levi (2006). Integration of risk attitudes into dynamic pricing has also been recently addressed with different approaches by Lim and Shanthikumar (2007) and Feng and Xiao (2004).

The main contributions of this paper are:

- introduction of a practical risk factor into a general problem of optimal dynamic pricing of perishable items with stochastic demand over finite horizons, and
- explicit inclusion of the revenue process in the description of the system's state to formulate a class of simple Markovian models incorporating risk, which are implementable in problems of realistic size.

This paper is organized as follows. We present the notation and general model in §2, and derive the optimality conditions in §3. We explore some of the structural properties of the optimal policies in §4, present a few possible generalizations of the model in §5, and discuss approximate solution methods in §6. We provide the results of extensive numerical experiments in §7. Finally, §8 contains concluding remarks.

2. The Model

2.1. Statement

We present the model as a stochastic optimal control problem in continuous time over the fixed planning horizon $[T, 0]$ (i.e., the negative value T is the beginning, and 0 is the end of the planning interval). The notation $[T, 0]$ for the planning horizon is consistent with Zhao and Zheng (2000). Suppose that the demand process $N'(t)$ is a nonhomogeneous Poisson process. The rate $\lambda(t, p)$ of this process is nonincreasing in the current price p, and is a bounded, continuously differentiable function of time t for each p. Such a model for the demand process is standard in the dynamic pricing literature (see, for example, Gallego and van Ryzin 1994, Zhao and Zheng 2000).

The sales process $N(t)$ has its value limited by the initial inventory Y_T; thus, $N(t) = \min\{N'(t), Y_T\}$. If the price process is $p(t)$, then the revenue process $R(t)$ can be defined as a stochastic integral of the form

$$R(t) = \int_T^t p(\tau)\, dN(\tau).$$

The risk-neutral objective is to simply

$$\max E[R(0)] \tag{1}$$

without constraints, except possibly ones on the set of allowed prices. To incorporate risk, we add a constraint of the form

$$P[R(0) \geqslant z] \geqslant \pi_0, \tag{2}$$

where z is a desired minimum level of revenues and π_0 is the minimum acceptable probability with which we want this level to be reached. Maximizing $E[R(0)]$ (a primary objective) is appropriate for a decision maker who can sell the product repeatedly in consecutive planning horizons and is fully rational in the long term. On the other hand, reaching the level z represents a short-term (secondary) objective of the "satisficing" type. The term "satisfice" as an alternative to "optimize" was used by Simon (1956, p. 131) in a model of a decision maker with "a definite, fixed aspiration level." Formulation (1)–(2) combines both of these objectives in a single model. If π_0 is varied, the problem will have different optimal solutions resulting in an *efficient frontier* in the plane of optimal $P[R(0) \geqslant z]$ and $E[R(0)]$; that is, the efficient frontier is the set of attainable expected revenue/probability pairs that are not dominated by any other attainable pair. An alternative way to obtain this frontier is by solving the problem

$$\max E[R(0)] - CP[R(0) < z] \tag{3}$$

for a range of values of the coefficient $C \in [0, +\infty)$. The penalty parameter C can be interpreted as the maximum cost associated with not meeting the desired level of revenue z. In some cases, z may correspond to sunk costs that must be recovered to avoid negative consequences of a severity represented by the value C. In this case, z represents a reference point from which losses or gains are computed. Furthermore, the hybrid objective in (3), adjusted by subtracting the reference point z, is equal to $E[u(R(0)-z)]$ for a value function of the form $u(x) = x - CI(x < 0)$ applied to the gain/loss $x = R(0) - z$. (Here, $I(A)$ is an indicator function of event A.) Note that $u(x)$ is a somewhat extreme case of the value function proposed by the prospect theory of Kahneman and Tversky (1979): concave (linear in our case) for gains, convex for losses, and steeper for losses than for gains. We note that for many decision makers the values of the parameters z, π_0, and C may be unknown or available only as rough estimates; however, by optimizing the system over ranges of values of these parameters, managers can gain insight into their personal assessments of the risk/return trade-offs. For example, selecting an appropriate value for C reduces to selecting a point on the efficient frontier, which is quite similar to choosing an appropriate trade-off between the return and "risk" for a portfolio of securities. Furthermore, different values of z will result in different ranges of probabilities $P(R(0) \geqslant z)$ that can be attained on the same curve. Typically, a decision maker will have some idea of the range of probabilities he/she is interested in. Due to the equivalence of the families of problems represented by (1)–(2) and (3), and our particular interest in the efficient frontiers, we will focus on (3). For future reference, we will denote the optimal value of (3) as $J(Y_T, T, C, z)$, and we will refer to this problem as the *risk-adjusted maximization problem* with initial inventory Y_T, time horizon T, cost of risk C, and desired level of revenue z.

In the terminology of Gīhman and Skorohod (1979), the system can be described as a jump Markovian controlled process. Indeed, for any fixed-price trajectory $p(t)$, $t \in [T, 0]$, a pair $(Y(t), R(t))$ consisting of the remaining inventory $Y(t) = Y_T - N(t)$ and the revenues $R(t)$ obtained to date, is, by construction, a Markov jump process. In standard continuous-time dynamic pricing models, it is usually enough to only include remaining inventory $Y(t)$ in the state description. In our case, due to the presence of revenue targets, the state needs to be described as $(Y(t), R(t))$; that is, we need to keep track of current inventory and the value of revenues to date.

In this paper, we will consider admissible price controls in a class of state-feedback (Markovian) controls, which are functions of time t, remaining inventory $Y(t)$, and current value of revenue $R(t)$; that is, price control: $p = p(t, Y(t), R(t))$. Given the time point t and the value of $Y(t)$, it is impractical to compute the optimal price for each value of $R(t)$ if it is allowed to take a continuum of values. Therefore, it is desirable to have a discrete-valued revenue process $R(t)$. In fact, the price of items is always discrete in practice. Thus, we assume that there is a smallest price unit (e.g., one cent), and prices belong to a finite discrete set with the maximum price $p_{\max}$. This results in admissible controls taking values in the set $\mathcal{P} = \{0, \ldots, p_{\max}\}$. Discrete sets of allowed prices were considered, for example, in Feng and Xiao (2000a, b), and one can further restrict the set of allowed prices to some proper subset of $\mathcal{P}$. Because the price at each time t belongs to a fixed discrete set, the revenue process $R(t)$ is also discrete, taking values only in the set $\{0, \ldots, Y_T p_{\max}\}$. Thus, the state of the system at time t can be specified by a pair in the finite set $\bar{\mathcal{S}} = \{(n, r)\colon n \in \{0, 1, \ldots, Y_T\}, r \in \{0, \ldots, (Y_T - n)p_{\max}\}\}$. The states with $n = 0$ are absorbing. We will denote the set of all nonabsorbing states by $\mathcal{S}$.

The special case of our model with $C = 0$ or $z = 0$ is the discrete-price version of the model of Gallego and van Ryzin (1994), studied extensively by Feng and Xiao (2000a, b) and by Zhao and Zheng (2000) for the more general case of a nonstationary distribution of reservation price. In this paper, we will often refer to that model as the *risk-neutral* model.

We see that under the finite discrete-price assumption, the process $(Y(t), R(t))$ is an intensity controlled, nonhomogeneous, continuous-time, finite-state Markov chain. In §2.2, we will see that problem (3) over the class of state-feedback controls can be formulated as a deterministic control problem for the state probability distribution vector of the process $(Y(t), R(t))$ and analyzed using deterministic continuous-time optimal control techniques via the Pontryagin maximum principle (see, for example, Vinter 2000).

2.2. Solution by Deterministic Optimal Control Methods

Recall that the demand rate at time t and price level p is $\lambda(t, p)$. Suppose that a company sets the price $p(t, n, r)$ in the state (n, r) at time t. Then, our continuous-time Markov chain makes transitions from the state (n, r) to the state $(n-1, r+p)$ with a rate $\lambda(t, p(t, n, r))$. Let $P_{(n,r)}(t)$ be the probability that the system is in the state (n, r) at time t. Then, $P_{(n,r)}(t)$ is governed by the following system of ordinary differential equations, which follow from the Kolmogorov forward equations for nonhomogeneous processes (see, for example, Gnedenko 1967, p. 37):

$$\frac{dP_{(n,r)}}{dt} = -\lambda(t, p(t,n,r))P_{(n,r)}(t) + \sum_{r': r'=r-p(t,n+1,r')} \lambda(t, p(t,n+1,r'))P_{(n+1,r')}(t) \quad \text{for } 0 < n < Y_T,\ 0 \leqslant r \leqslant (Y_T - n)p_{\max}, \tag{4}$$

$$\frac{dP_{(0,r)}}{dt} = \sum_{r': r'=r-p(t,1,r')} \lambda(t, p(t,1,r'))P_{(1,r')}(t) \quad \text{for } 0 \leqslant r \leqslant Y_T p_{\max}, \tag{5}$$

$$\frac{dP_{(Y_T,0)}}{dt} = -\lambda(t, p(t, Y_T, 0))P_{(Y_T,0)}(t), \tag{6}$$

with the initial conditions

$$P_{(n,r)}(T) = 0 \quad \text{for } (n,r) \in \bar{\mathcal{S}} \backslash (Y_T, 0), \tag{7}$$

$$P_{(Y_T,0)}(T) = 1. \tag{8}$$

To be able to employ standard results from control theory (standard existence theorem and optimality conditions), we convexify the controls by embedding them into a larger policy space of randomized policies formed by time-dependent measures $\pi(p \mid t,n,r)$ on the set of allowable prices in all possible states (n,r). This extended policy space (a set of randomized state-feedback controls) can be interpreted as follows. Suppose that at time t, the company has inventory n and revenue r. When $n > 0$, the price p is offered to a customer with probability $\pi(p \mid t,n,r)$. When $n = 0$, no price is offered—the zero inventory state is absorbing. Both situations can be captured using normalizations

$$\sum_p \pi(p \mid t,n,r) = 1 \quad \text{for } (n,r) \in \mathcal{S}, \tag{9}$$

$$\sum_p \pi(p \mid t,0,r) = 0, \tag{10}$$

where, for brevity, we dropped the summation range condition $p \in \mathcal{P}$. (We also drop it in all subsequent summations over the index p for the same reason.) While this extended form of controls helps us establish existence, we emphasize that the extension is merely an analytical device, and the original space of discrete price controls remains our primary interest.

The corresponding rate of transition out of state (n,r) is $\sum_p \lambda(t,p)\pi(p \mid t,n,r)$; thus, for each p, $\lambda(t,p)\pi(p \mid t, n,r)$ transitions per unit of time are directed into the state $(n-1, r+p)$, and the system of equations governing $P_{(n,r)}(t)$ can be rewritten succinctly as

$$\frac{dP_{(n,r)}}{dt} = -\left(\sum_p \lambda(t,p)\pi(p \mid t,n,r)\right)P_{(n,r)}(t) + \sum_{p:(n+1,r-p)\in\mathcal{S}} \lambda(t,p)\pi(p \mid t,n+1,r-p) \cdot P_{(n+1,r-p)}(t) \quad \text{for } (n,r) \in \bar{\mathcal{S}}. \tag{11}$$

Note that $\pi(p \mid t,n,r)$ is represented by a nonnegative vector in the $p_{\max}$-dimensional real space satisfying (9)–(10). System (11) immediately generalizes Equation (4). We also note that it holds for all (n,r) with $n = 0$ because (10) implies that the first term is zero, and for $(n,r) = (Y_T, 0)$ because $\mathcal{S}$ does not contain a state with inventory level $Y_T + 1$, and consequently the second term is trivially zero. Thus, (11) incorporates (4) and (6) as well.

Formulation (3) involves maximizing the objective

$$E[R(0) - CI(R(0) < z)] = \sum_{(n,r)\in\bar{\mathcal{S}}} rP_{(n,r)}(0) - C\sum_{(n,r)\in\bar{\mathcal{S}}: r<z} P_{(n,r)}(0), \tag{12}$$

subject to evolution system (11) with the initial conditions (7)–(8). We will refer to this optimal control problem as PH.

Note that PH is a standard optimal control problem; see Vinter (2000). In fact, because the evolution system (11) is bilinear in $P_{(n,r)}$ and $\pi(p \mid t,n,r)$ (i.e., linear in the state and control variables separately, but jointly quadratic with constant coefficients), and the objective is linear in $P_{(n,r)}$, this formulation belongs to a special class of bilinear optimal control problems, long used in applied models in engineering (see Mohler 1973).

3. Existence of a Solution and Optimality Conditions

In this section, we prove the existence of an optimal solution to the risk-adjusted problem (in its optimal control form PH) and derive its optimality conditions.

First, we note that problem PH satisfies the conditions of the existence theorem in Lee and Markus (1967, Theorem 4, p. 259), which allows us to assert the existence of the optimal solution in the appropriate class of admissible controls. One of the theorem's requirements is that the image of the admissible control set at time t under the mapping given by the right-hand side of system (11) is a convex set. Note that the discrete-price controls $p(t,n,r)$ do not satisfy this requirement (image would be a discrete set), whereas the extended controls $\pi(\cdot \mid t,n,r)$ do. This is our primary reason for control embedding.

Proposition 1. *Problem PH has an optimal solution: There exists a control that maximizes* (12) *for the system governed by* (11) *with initial conditions* (7)–(8) *in the class of measurable functions* $\boldsymbol{\pi}(t) = (\pi(p \mid t,n,r), (n,r) \in \bar{\mathcal{S}})$ *subject to constraints* (9)–(10).

Corollary 1. *The solution of the stochastic control problem* (3), *under the finite discrete-price assumption in the class of randomized state-feedback controls, exists.*

A standard analysis technique for optimal control is the necessary optimality condition in the form of the Pontryagin maximum principle; see Vinter (2000). To apply the maximum principle to PH, we introduce the adjoint variables $\eta_{(n,r)}(t)$ for each of the evolution equations in (11). As often occurs in these types of problems, $\eta_{(n,r)}(t)$ can be interpreted as a conditional expectation of the objective term $R(0) - CI(R(0) < z)$ (see Lemma 1 below). We will denote the vector of all adjoint variables by $\boldsymbol{\eta}(\cdot)$, and the vectors of all state probabilities and policies by $\mathbf{P}(\cdot)$ and $\boldsymbol{\pi}(\cdot)$, respectively (as functions of time). The following proposition is proved in the online appendix. An electronic companion to this paper is available as part of the online version that can be found at http://or.journal.informs.org/.

Proposition 2. *For an optimal state trajectory* $\mathbf{P}(\cdot)$ *and control* $\boldsymbol{\pi}(\cdot)$ *in problem PH, there exists* $\boldsymbol{\eta}(\cdot)$ *satisfying the adjoint system of differential equations*

$$\frac{d\eta_{(n,r)}}{dt} = \sum_p \lambda(t,p)\pi(p \mid t,n,r)[\eta_{(n,r)}(t) - \eta_{(n-1,r+p)}(t)], \quad (n,r) \in \mathcal{S}, \quad (13)$$

$$\frac{d\eta_{(n,r)}}{dt} = 0, \quad (n,r) \in \bar{\mathcal{S}} \setminus \mathcal{S}, \quad (14)$$

with the terminal condition

$$\eta_{(n,r)}(0) = r - CI(r < z), \quad (n,r) \in \bar{\mathcal{S}}, \quad (15)$$

such that $\pi(\cdot \mid t,n,r)$ *delivers the maximum in the auxiliary problem*

$$\max \; P_{(n,r)}(t) \sum_p \lambda(t,p)[\eta_{(n-1,r+p)}(t) - \eta_{(n,r)}(t)] \cdot \pi(p \mid t,n,r), \quad (16)$$

$$\text{s.t.} \; \sum_p \pi(p \mid t,n,r) = 1 \quad (17)$$

for all $(n,r) \in \mathcal{S}$ *and for almost all* $t \in [T,0]$.

Remark 1. In (16), $P_{(n,r)}(t)$ is merely a coefficient in a linear expression to be maximized over $\pi(\cdot \mid t,n,r)$. Then, if $\pi(\cdot \mid t,n,r)$ delivers the maximum in (16) for some strictly positive value of $P_{(n,r)}(t)$, then that same $\pi(\cdot \mid t,n,r)$ would also deliver the maximum for any other $P_{(n,r)}(t)$, including $P_{(n,r)}(t) = 0$. We see that the initial value problem for state system (11) with initial conditions (7)–(8) and the terminal value problem for adjoint Equations (13)–(14) with terminal condition (15) separate in the following sense. If one wants to find some $\mathbf{P}(\cdot)$, $\boldsymbol{\pi}(\cdot)$, and $\boldsymbol{\eta}(\cdot)$ satisfying the optimality conditions, it is possible to do so by first finding a solution to the terminal value problem for the adjoint variables $\boldsymbol{\eta}(\cdot)$, where $\boldsymbol{\pi}(\cdot)$ delivers the maximum in (16) independently of $\mathbf{P}(\cdot)$. Second, one solves the initial value problem for $\mathbf{P}(\cdot)$ using $\boldsymbol{\pi}(\cdot)$ found in the first step. Because this procedure is very efficient, control problems where the state and adjoint equations separate have a reputation for being quite easy to solve in practice.

The following lemma (proved in the online appendix) provides an interpretation for $\eta_{(n,r)}(t)$:

Lemma 1. *Under any control* $\boldsymbol{\pi}(\cdot)$, *the solution* $\boldsymbol{\eta}(\cdot)$ *of the adjoint system* (13)–(14) *with the terminal condition* (15) *is such that*

$$\eta_{(n,r)}(t) = E[R(0) - CI(R(0) < z) \mid Y(t) = n, R(t) = r].$$

Remark 2. Lemma 1 implies that problem (16)–(17) has a very natural interpretation: For any inventory n and revenue r at any time t, the optimal decision is to maximize the expected rate of increment in our objective per unit of time.

Remark 3. The existence results imply that the value function $J(Y_T, T, C, z)$ of (3) is defined for all $Y_T \in \mathbb{N}$, $T < 0$, C, and z. We may also extend it to $Y_T = 0$ and $T = 0$ by $J(0,T,C,z) = J(Y_T,0,C,z) = -CI(0 < z)$. For any solution $\boldsymbol{\eta}(\cdot)$ of the adjoint system (13)–(14), Lemma 1 implies that $\eta_{(n,r)}(t) - r$ corresponds to the objective function of the problem of the form (3) with initial inventory n, time horizon $[t,0]$, and the revenue threshold $z - r$. The optimal value of such a problem is equal to $J(n,t,C,z-r)$. If under the optimal controls the state (n,r) has a nonzero probability at time t—that is $P_{(n,r)}(t) > 0$—then the adjoint variables satisfy $\eta_{(n,r)}(t) = r + J(n,t,C,z-r)$. For an arbitrary triple $\mathbf{P}(\cdot)$, $\boldsymbol{\pi}(\cdot)$, $\boldsymbol{\eta}(\cdot)$ satisfying the conditions of Proposition 2, this equality relation need not be satisfied in states of probability zero at time t. The reason is that the policies in these states do not affect the optimal value of the problem for a system whose initial state is $(Y_T, 0)$. This can also be seen from the fact that, when $P_{(n,r)}(t) = 0$, the optimal value in (16) is delivered by any $\pi(\cdot \mid t,n,r)$ satisfying (17). However, for any triple $\mathbf{P}(\cdot)$, $\boldsymbol{\pi}(\cdot)$, $\boldsymbol{\eta}(\cdot)$ that satisfies the conditions of Proposition 2, and for any state (n,r) at time t including those of probability zero, Lemma 1 implies that $\eta_{(n,r)}(t) \leqslant r + J(n,t,C,z-r)$. Otherwise, we have a contradiction with the optimality of $J(n,t,C,z-r)$.

The following corollary gives a generalization of Lemma 1 in case of arbitrary terminal conditions (such as situations with salvage value or disposal costs, to be discussed later).

Corollary 2. *Under any control* $\boldsymbol{\pi}(\cdot)$, *the solution* $\boldsymbol{\eta}(\cdot)$ *of the adjoint system* (13)–(14) *with arbitrary terminal conditions*

$$\eta_{(n,r)}(0) = F(n,r)$$

is such that

$$\eta_{(n,r)}(t) = E[F(Y(0), R(0)) \mid Y(t) = n, R(t) = r].$$

From the maximum principle, we know that the optimal control satisfies (16)–(17) for almost all t. If $P_{(n,r)}(t) > 0$, then by dividing the objective term (16) by $P_{(n,r)}(t)$, we get the negative of the right-hand side of the adjoint equation (13) at time t. Thus, when $P_{(n,r)}(t) > 0$, the value of the right-hand side in (13) is the same for all controls delivering the maximum in (16)–(17). We can now combine (16)–(17) with (13) to get the Hamilton-Jacobi-Bellman equation

$$\frac{d\eta_{(n,r)}}{dt} = \min_{\pi(\cdot \mid t,n,r)} \left\{ \sum_p \lambda(t,p)\pi(p \mid t,n,r) \cdot [\eta_{(n,r)}(t) - \eta_{(n-1,r+p)}(t)] \right\}. \tag{18}$$

For given model parameters Y_T, T, C, z, a system of equations consisting of (18) for all $(n,r) \in \mathcal{S}$ and $d\eta_{(n,r)}/dt = 0$ for $(n,r) \in \bar{\mathcal{S}} \setminus \mathcal{S}$ with the boundary condition (15) provides a sufficient optimality condition. A solution $\eta_{(n,r)}(t)$ to (18) will satisfy $\eta_{(n,r)}(t) = J(n,t,C,z-r) + r$ due to Remark 3.

Remark 4. We now show that the optimal controls in the original control space (i.e., the discrete price set) exist as well. In doing so, we also establish that the right-hand side of (18) coincides with

$$\min_p \lambda(t,p)[\eta_{(n,r)}(t) - \eta_{(n-1,r+p)}(t)]. \tag{19}$$

Note that the right-hand side of (18) is linear in controls $\pi(\cdot \mid t,n,r)$, and the set of feasible controls for given (t,n,r) has a polyhedral structure: $\pi(\cdot \mid t,n,r)$ are nonnegative and satisfy (9). Recall that the optimum in linear optimization problems is attained at a vertex that in our case, corresponds to $\pi(p \mid t,n,r) = 1$ for some p. Thus, if the term attaining the minimum in (19) is unique (e.g., if the terms $\lambda(t,p)[\eta_{(n,r)}(t) - \eta_{(n-1,r+p)}(t)]$ are all distinct for different p), then the minimum in (18) is uniquely attained by $\pi(p \mid t,n,r) = 1$ for some p. Then, the right-hand side of (18) reduces to (19), and the policy in the state (n,r) at time t is not randomized. If several terms attain the minimum in (19) with respect to p, then the minimum in (18) is not unique and may be attained by $\pi(\cdot \mid t,n,r)$ with several positive components (i.e., several vertices and their linear combinations attain the minimum). In particular, such a phenomenon occurs at switch points between different prices. However, we can modify any such $\pi(\cdot \mid t,n,r)$ by concentrating all probability on an appropriately chosen singleton set of prices without modifying the value of $\sum_p \lambda(t,p)\pi(p \mid t,n,r)[\eta_{(n,r)}(t) - \eta_{(n-1,r+p)}(t)]$ (and thus showing that this value is equal to (19)). This permits an optimal policy with a singleton support set in the form $p = p(t,n,r)$. Specifically, we can choose the maximum p for which (19) is attained to be that singleton. The same tie-breaking rule was also employed, for example, by Zhao and Zheng (2000). Henceforth, $p(t,n,r)$ will specifically refer to the maximum price attaining the minimum in (19), and we will refer to the corresponding policies as the "singleton price."

4. Structural Properties of Optimal Policies

In this section, we consider the structural properties of optimal policies. We also compare the obtained policies with those obtained for the risk-neutral model of Gallego and van Ryzin (1994). Recall that the latter model is a special case of our model with $z = 0$.

Next, we show that the value of the right-hand side of Equation (18) does not depend on r for $r \geqslant z$. This is a reality check for the model because the revenue-level requirement should not affect the policy after the required level has been reached. Consider a change of variables $\bar{\eta}_{(n,r)} = \eta_{(n,r)} - r$, which can be interpreted as the expected future increment in the objective value compared to its value in the current state (expected future revenues for $r \geqslant z$). For values of $r \geqslant z$, Equation (18) becomes

$$\frac{d\bar{\eta}_{(n,r)}}{dt} = \min_{\pi(p \mid t,n,r)} \left\{ \sum_p \lambda(t,p)\pi(p \mid t,n,r) \cdot [\bar{\eta}_{(n,r)}(t) - \bar{\eta}_{(n-1,r+p)}(t) - p] \right\},$$

and the terminal condition (15) transforms, for all n and $r \geqslant z$, to

$$\bar{\eta}_{(n,r)}(0) = 0.$$

We know that the solution $\bar{\boldsymbol{\eta}}$ exists and is uniquely determined by

$$\bar{\eta}_{(n,r)}(t) = \eta_{(n,r)}(t) - r = J(n,t,C,z-r)$$

for $r \geqslant z$. For a fixed level of inventory n', now consider any $r' \geqslant z$ and the subsystems of equations governing the collections of variables

$$\{\bar{\eta}_{(n,r)}\colon n \in \{0,1,\ldots,n'\}, r \in \{r',\ldots,r' + (n'-n)p_{\max}\}\}.$$

These subsystems of equations and the corresponding terminal conditions are identical for different $r' \geqslant z$. Thus, we have the following proposition that the optimal controls can be selected to be identical for all states with the same level of inventory and any revenue-to-date level $r \geqslant z$:

Proposition 3. *For all $r > r' \geqslant z$ and all t, n, $\bar{\eta}_{(n,r)}(t) = \bar{\eta}_{(n,r')}(t)$. Moreover,* $\operatorname{Arg\,min}_p \lambda(t,p)[\bar{\eta}_{(n,r)}(t) - \bar{\eta}_{(n-1,r+p)}(t) - p]$ *does not depend on r, and there exist corresponding optimal controls that do not depend on r: $\pi(p \mid t,n,r) = \pi(p \mid t,n,r')$ for all $r > r' \geqslant z$ and all t,n.*

This proposition formalizes the fact that a risk-adjusted optimal policy coincides with a risk-neutral one after the required revenue level has been reached, a conclusion that certainly applies to the singleton price controls as well. Let $p^0(t,n,r)$ be an optimal "singleton price" policy in the risk-neutral problem with the same initial inventory Y_T. The corresponding adjoint variables will be denoted by $\boldsymbol{\eta}^0(\cdot)$. As before (see Remark 4), $p^0(t,n,r)$ denotes the largest p that attains the minimum in $\min_p \lambda(t,p)[\eta^0_{(n,r)}(t) - \eta^0_{(n-1,r+p)}(t)]$. Using standard results on the monotonicity of risk-neutral policies, e.g., decrease of price in n at each fixed moment in time, we can conclude that after reaching the required level z, the risk-adjusted policy possesses the same monotonic properties. However, numerical experiments described in §7 show that the risk-adjusted optimal policy need not be monotonic in the states where $r < z$ (below the required level).

For values of $r < z$, the optimal price $p(t,n,r)$ will, in general, depend on r. The behavior of p near the critical threshold $r = z$ is of some interest. The relation between the optimal prices in the risk-adjusted and risk-neutral models near this threshold is given in the following proposition. The intuition behind this result is that when the current value of revenue is close to the required level z, it becomes more important to cross the threshold than to collect a high price. Because a higher demand rate is generally achieved by a price reduction, this situation creates a downward pressure on the current price. The proposition applies to the range of r no lower than one optimal risk-neutral price below the threshold. It is intuitively clear that the downward pressure on price may exist only within this band because, for lower values of r, it may be preferable to step over the threshold by means of higher prices. Such behavior of optimal prices is almost obvious if only one item is available for sale and is indeed observed in the numerical experiments.

Proposition 4. *For all (t,n,r) such that $z > r \geqslant z - p^0(t,n,r)$, we have $p(t,n,r) \leqslant p^0(t,n,r)$.*

Proof. Consider, for some (t,n,r) such that $z > r \geqslant z - p^0(t,n,r)$, the optimal prices in the risk-adjusted $p \equiv p(t,n,r)$ and risk-neutral $p^0 \equiv p^0(t,n,r)$ models, respectively. Suppose, on the contrary, that we have $p > p^0$. Because $p(t,n,r)$ and $p^0(t,n,r)$ attain the minimums of the right-hand sides in the corresponding differential equations for $\eta_{(n,r)}$ and $\eta^0_{(n,r)}$, and $p^0(t,n,r)$ is the largest p attaining its corresponding minimum, we conclude that

$$\lambda(t,p)[\eta^0_{(n,r)} - \eta^0_{(n-1,r+p)}] > \lambda(t,p^0)[\eta^0_{(n,r)} - \eta^0_{(n-1,r+p^0)}],$$

$$\lambda(t,p)[\eta_{(n,r)} - \eta_{(n-1,r+p)}] \leqslant \lambda(t,p^0)[\eta_{(n,r)} - \eta_{(n-1,r+p^0)}].$$

By rewriting these inequalities, we get

$$[\lambda(t,p) - \lambda(t,p^0)]\eta^0_{(n,r)} > \lambda(t,p)\eta^0_{(n-1,r+p)} - \lambda(t,p^0)\eta^0_{(n-1,r+p^0)}, \tag{20}$$

$$[\lambda(t,p) - \lambda(t,p^0)]\eta_{(n,r)} \leqslant \lambda(t,p)\eta_{(n-1,r+p)} - \lambda(t,p^0)\eta_{(n-1,r+p^0)}. \tag{21}$$

Because $r + p > r + p^0 \geqslant z$, we have $\eta_{(n-1,r+p)} = \eta^0_{(n-1,r+p)}$ and $\eta_{(n-1,r+p^0)} = \eta^0_{(n-1,r+p^0)}$. Therefore, the right-hand sides of inequalities (20) and (21) coincide. Subtracting the second inequality from the first, we get

$$[\lambda(t,p) - \lambda(t,p^0)](\eta^0_{(n,r)} - \eta_{(n,r)}) > 0.$$

Note that $\eta^0_{(n,r)} = r + J(n,t,C,-r) > r + J(n,t,C,z-r) = \eta_{(n,r)}$, and it follows that $\lambda(t,p) > \lambda(t,p^0)$. Due to the monotonicity of $\lambda(t,p)$ in p, we have $\lambda(t,p) \leqslant \lambda(t,p^0)$, a contradiction. Thus, $p^0 \geqslant p$ $(p(t,n,r) \leqslant p^0(t,n,r)$ in the original nonabbreviated notation). □

In the risk-neutral model of Gallego and van Ryzin (1994), the price usually increases immediately following a sale. Thus, a typical realization of the price process has upward jumps at the times of sales. In the case of the risk-adjusted model, however, the price may decrease following a sale. This happens because, as the above proposition suggests, the optimal prices must be lower than the risk-neutral prices when current revenues are slightly below the threshold z (within one optimal risk-neutral price below the threshold). Indeed, consider the following example of model parameters that lead to such a price drop. Let the set of allowed prices be $\{1,2,3\}$ and suppose that the demand is time homogeneous: $\lambda(t,1) = \lambda_1 > 0$, $\lambda(t,2) = \lambda_2 > 0$, and $\lambda(t,3) = 0$. Let $Y_T = 2$ and $z = 3$. At the end of the planning interval, we have $\eta_{(n,r)}(0) = r - CI(r < 3)$. Note that $\eta_{(n,r)}(t)$ will be very close to $\eta_{(n,r)}(0)$ as t approaches zero, due to its continuity. To illustrate a drop in price, we seek λ_1 and λ_2 such that $p = 1$ is selected for the state $(n,r) = (1,2)$, and $p = 2$ for the state $(n,r) = (2,0)$. That is, the optimal price drops from 2 to 1 following a transition from $(2,0)$ to $(1,2)$. At $t = 0$, this is enforced by inequalities

$$\lambda_1(\eta_{(1,2)}(0) - \eta_{(0,3)}(0)) < \lambda_2(\eta_{(1,2)}(0) - \eta_{(0,4)}(0)),$$
$$\lambda_1(\eta_{(2,0)}(0) - \eta_{(1,1)}(0)) > \lambda_2(\eta_{(2,0)}(0) - \eta_{(1,2)}(0)).$$

Note that the shutdown price $p = 3$ cannot be selected near $t = 0$ because $\eta_{(n,r)}(0) - \eta_{(n-1,r+p)}(0) < 0$ for all (n,r). Because these inequalities are strict, they will also enforce the drop in price for some range of t near 0. Substituting the expression for $\eta_{(n,r)}(0)$ in the first inequality, we get

$$\lambda_1(2 - CI(2<3) - 3 + CI(3<3)) < \lambda_2(2 - CI(2<3) - 4 + CI(4<3)),$$

which is equivalent to $\lambda_1(-1-C) < \lambda_2(-2-C)$ or

$$\lambda_1 > \lambda_2\left(1 + \frac{1}{1+C}\right).$$

Similarly, from the second inequality, we get

$$\lambda_1(0 - CI(0<3) - 1 + CI(1<3)) > \lambda_2(0 - CI(0<3) - 2 + CI(2<3))$$

or

$$\lambda_1 < 2\lambda_2.$$

Thus, whenever demand rates satisfy

$$\lambda_2\left(1 + \frac{1}{1+C}\right) < \lambda_1 < 2\lambda_2$$

and a sale happens near $t=0$ in state $(n,r)=(2,0)$, then as the system jumps to the state $(1,2)$, the optimal price drops from 2 to 1.

The following proposition summarizes intuitive monotonic properties of the value function with respect to t, n, and r.

PROPOSITION 5. *Let* $\boldsymbol{\eta}(\cdot)$ *be a solution of* (18). *Then, for all* t, n, r:

(a) $\eta_{(n,r)}(t) \leqslant \eta_{(n-1,\, r+p_{\max})}(t)$;
(b) $\eta_{(n,r)}(t)$ *is nonincreasing in* t;
(c) *for an optimal singleton price control* $p(t,n,r)$, $\lambda(t, p(t,n,r)) > \lambda(t, p_{\max}) = 0$ *implies that* $\eta_{(n,r)}(t) < \eta_{(n-1,\, r+p(t,n,r))}(t)$;
(d) $\eta_{(n,r)}(t) \leqslant \eta_{(n+1,r)}(t)$;
(e) $\eta_{(n,r)}(t) \leqslant \eta_{(n,r+1)}(t) - 1$.

PROOF. Part (a) follows from the following relaxation argument. Consider a state (n,r) at time t and a modified problem where one item is guaranteed to sell at price $p_{\max}$ at the end of the planning horizon. The optimal expected revenue of the modified problem in the given state is certainly greater than or equal to the one in the original problem. However, the state $(n-1, r+p_{\max})$ is equivalent to the initial state of the modified problem in terms of inventory and current level of revenue, which establishes part (a).

To establish part (b), we recall that by Remark 4, the right-hand side of (18) coincides with (19). From part (a), it follows that the value of (19) satisfies

$$\begin{aligned}\min_p \lambda(t,p)[\eta_{(n,r)}(t) - \eta_{(n-1,r+p)}(t)] \\ \leqslant \lambda(t,p_{\max})[\eta_{(n,r)}(t) - \eta_{(n-1,r+p_{\max})}(t)] \leqslant 0.\end{aligned}$$

Therefore, the right-hand side of (18) is nonpositive and $d\eta_{(n,r)}/dt \leqslant 0$. Part (b) follows.

To establish part (c), recall that $p(t,n,r)$ denotes the largest price attaining the minimum in

$$\min_p \lambda(t,p)[\eta_{(n,r)}(t) - \eta_{(n-1,r+p)}(t)].$$

If the value of this minimum is zero, then it must be that $p(t,n,r) = p_{\max}$. The condition $\lambda(t, p(t,n,r)) > \lambda(t, p_{\max}) = 0$ implies that $p(t,n,r) < p_{\max}$. Part (c) follows.

The proof of (d) follows immediately from

$$\begin{aligned}\eta_{(n,r)}(t) &= r + J(n,t,C,z) \leqslant r + J(n+1,t,C,z)\\ &= \eta_{(n+1,r)}(t),\end{aligned}$$

with the inequality in the middle obtained from a relaxation argument. Similarly, part (e) is immediately obtained from

$$\begin{aligned}\eta_{(n,r)}(t) &= r + J(n,t,C,z-r) \leqslant r + J(n,t,C,z-(r+1))\\ &= \eta_{(n,r+1)}(t) - 1. \quad \square\end{aligned}$$

5. Generalization

5.1. General Terminal Value

The approach of reducing the stochastic optimization problem (3) to a deterministic optimal control problem can be applied to a much wider class of problems. The existence results and optimality conditions derived for (3) can be easily obtained for an optimization problem of the form

$$\max E[F(Y(0), R(0))]$$

with an arbitrary finite function $F(n,r)$. The only difference in the optimality conditions is that terminal condition (15) needs to be replaced with a more general

$$\eta_{(n,r)}(0) = F(n,r).$$

For example, a particularly interesting extension is an introduction of disposal costs or salvage values for unsold items into the model. For each item, let the disposal cost be $D<0$ (or, salvage value, $D>0$). Then, an immediate generalization of (3) is

$$\max E[R(0) + DY(0)] - CP(R(0) + DY(0) < z).$$

For a nonrandom D, the corresponding $F(n,r)$ is of the form $F(n,r) = r + Dn - CI(r+Dn<z)$. If D is stochastic, we use the slightly more general $F(n,r) = E_D[r + Dn - CI(r+Dn<z)]$.

5.2. Extended Risk-Neutral Horizon

Another generalization of our formulation is to consider a planning horizon extending beyond zero, the time at which we assess the loss probability, into a later time period in which we only maximize expected revenue. This problem can be written as

$$\max E[R(T')] - CP[R(0) < z], \tag{22}$$

where $T' > 0$.

Note that this objective function can be rewritten as $E[R(T') - R(0)] + E[R(0)] - CP[R(0) < z]$. The value $E[R(T') - R(0)]$ is the amount of revenues earned after time 0; that is, when risk no longer affects the decision and an optimal policy can depend on the current inventory only. Let the value of the optimal expected revenues $E[R(T') - R(0)]$ on $[0, T']$, given the inventory n at zero, be $\tilde{J}(n, T')$. Then, we can solve problem (22) as the problem of maximizing $E[F(Y(0), R(0))]$, where $F(n,r) = \tilde{J}(n,T') + r - CI(r<z)$.

5.3. Variance Risk Measure

We also note that the traditional mean-variance approach to incorporation of risk, while it can be analyzed by similar techniques, does not fit into the framework of this paper. Indeed, the mean-variance formulation would be to

$$\max E[R(0)] - C\mathrm{Var}(R(0)).$$

This objective can be rewritten as

$$E[R(0)] - C(E[R^2(0)] - E[R(0)]^2)$$
$$= \sum_{(n,r)\in\bar{\mathcal{S}}} rP_{(n,r)}(0) - C\Bigg[\sum_{(n,r)\in\bar{\mathcal{S}}} r^2 P_{(n,r)}(0) - \Bigg(\sum_{(n,r)\in\bar{\mathcal{S}}} rP_{(n,r)}(0)\Bigg)^2\Bigg],$$

which is quadratic in $\mathbf{P}(0)$. Therefore, the boundary conditions for the adjoint variables $\boldsymbol{\eta}(\cdot)$ will depend on $\mathbf{P}(0)$, and the evolution and adjoint equations governing $\mathbf{P}(\cdot)$ and $\boldsymbol{\eta}(\cdot)$ will no longer separate in the sense of Remark 1. Thus, the mean-variance approach appears to be applicable only to limited-size dynamic pricing problems.

5.4. Multiproduct Case

Note also that it is easy to extend our model to the case of m perishable products if the minimal price increments for all products have the same value, and each customer can only purchase a single item of a single product. Indeed, the combined revenue value will still be represented by a single scalar r in the state description. The state description will also contain current inventory levels n_i for each product $i = 1, \ldots, m$. The feasible controls will be given by a discrete vector of prices $(p_1, \ldots, p_m) \in \{0, \ldots, p_{\max}\}^m$ for each product. It can be shown that the same analytic techniques as the ones used in this paper are applicable. Of course, the computational complexity will increase as the number of products m gets larger. However, this issue is also present in a multiproduct variation of the model of Gallego and van Ryzin (1994). There is no inherent difference in the treatment of a revenue-level requirement in 1- and m-product situations.

5.5. Moving Revenue Target Generalization

Our analysis of problem (3) under the discrete price assumption leads to a Hamilton-Jacobi-Bellman (HJB) equation of the form familiar from the stochastic control literature. However, it is not obvious how to handle the problem in the form (1)–(2) directly by standard stochastic control techniques, and, in this section, we describe a practically relevant generalization of model (1)–(2) that does not have an immediate solution through the standard method but can be approached with the techniques of this paper.

Consider a company that has a *moving revenue target* $z(t)$. The target could represent cash flow requirements that result from the need to cover ongoing operational costs. Due to the stochastic nature of sales, it is impossible to guarantee that the target is met with certainty; however, management may have some probability $\pi(t)$ with which it wants to exceed the revenue target $z(t)$ at time t. This naturally leads to a problem of the form

$$\max \; E[R(0)] \tag{23}$$

$$\text{subject to } \; P[R(t) > z(t)] \geqslant \pi(t) \quad \text{for almost all } t \in [T, 0]. \tag{24}$$

We will assume that $z(t)$ and $\pi(t)$ are continuous functions of t. We use $R(t) > z(t)$ in this formulation instead of $R(t) \geqslant z(t)$ for technical reasons to ensure upper semicontinuity of the function $\pi(t) - P[R(t) > z(t)]$, a requirement for known optimal control results. This formulation can be cast as an optimal control problem with a state constraint. Indeed, using probability vector $\mathbf{P}(\cdot)$ in the same way as before, objective (23) can be expressed as

$$E[R(0)] = \sum_{(n,r)\in\bar{\mathcal{S}}} rP_{(n,r)}(0), \tag{25}$$

while constraint (24) can be rewritten as a *state constraint*

$$P(R(t) > z(t)) = \sum_{(n,r)\in\bar{\mathcal{S}}} P_{(n,r)}(t) I(r > z(t)) \geqslant \pi(t) \quad \text{for almost all } t \in [T, 0]. \tag{26}$$

Therefore, we obtain an optimal control problem with the objective to maximize the functional (25) of the probability vector $\mathbf{P}(\cdot)$, whose evolution is described by system (11) with the initial conditions (7)–(8), and whose value is restricted by the state constraint (26). We will refer to this problem as PG.

This formulation satisfies the conditions of Theorem 5.2.1 of Clarke (1990, p. 211). This includes upper semicontinuity of the function

$$g(t, \mathbf{P}) = \pi(t) - \sum_{(n,r)\in\bar{\mathcal{S}}} P_{(n,r)} I(r > z(t))$$

specifying the state constraint as $g(t, \mathbf{P}(t)) \leqslant 0$ because in our problem, $z(t)$ is continuous, $\mathbf{P}$ is guaranteed to be nonnegative, and, therefore, each term of the form $-P_{n,r} I(r > z(t))$ is upper semicontinuous. Whereas Theorem 5.2.1 of Clarke (1990) is very general and applies to the case of nonsmooth problem data (even including nonsmooth dynamics of the system), our problem is much simpler because all functions, except for one used in the state constraint, are smooth. Thus, the main difference in optimality conditions for this formulation as compared to Proposition 2 will be due to the state constraint:

Proposition 6. *For an optimal state trajectory* $\mathbf{P}(\cdot)$ *and controls* $\boldsymbol{\pi}(\cdot)$ *in problem* PG, *there exist a measurable function* $\boldsymbol{\gamma}(\cdot)$ *from* $[T, 0]$ *to* $|\bar{\mathcal{S}}|$*-dimensional real space, a nonnegative measure* μ *on* $[T, 0]$, *a constant* $\zeta \in \{0, 1\}$, *and an adjoint trajectory* $\boldsymbol{\eta}(\cdot)$ *such that*:

(i) $\boldsymbol{\eta}(\cdot)$ *satisfies the adjoint system of differential equations*

$$\begin{aligned}\frac{d\eta_{(n,r)}}{dt} = \sum_p \lambda(t,p)\pi(p\mid t,n,r) \\ \cdot\bigg[\eta_{(n,r)}(t) + \int_{[T,t)} \gamma_{(n,r)}(s)\,\mu(ds) - \eta_{(n-1,r+p)}(t) \\ - \int_{[T,t)} \gamma_{(n-1,r+p)}(s)\mu(ds)\bigg], \quad (n,r)\in\mathcal{S}, \end{aligned} \tag{27}$$

$$\frac{d\eta_{(n,r)}}{dt} = 0, \quad (n,r)\in\bar{\mathcal{S}}\backslash\mathcal{S}; \tag{28}$$

(ii) $\pi(\cdot\mid t,n,r)$ *delivers the maximum in the auxiliary problem*

$$\begin{aligned}\max\ P_{(n,r)}(t)\sum_p \lambda(t,p)\bigg[\eta_{(n-1,r+p)}(t) + \int_{[T,t)} \gamma_{(n,r)}(s)\,\mu(ds) \\ - \eta_{(n,r)}(t) - \int_{[T,t)} \gamma_{(n-1,r+p)}(s)\,\mu(ds)\bigg]\pi(p\mid t,n,r)\end{aligned} \tag{29}$$

$$\text{s.t.}\ \sum_p \pi(p\mid t,n,r) = 1 \tag{30}$$

for all $(n,r)\in\mathcal{S}$ *and for almost all* $t\in[T,0]$;

(iii) *for almost all* $t\in[T,0]$ *with respect to measure* μ, *we have* $\gamma_{(n,r)}(t) = -I(r > z(t))$ *for* $r\neq z(t)$ *and* $\gamma_{(n,r)}(t)\in[-1,0]$ *for* $r = z(t)$, *and measure* μ *is supported on the set*

$$\bigg\{t\in[T,0]\colon \sum_{(n,r)\in\bar{\mathcal{S}}} P_{(n,r)}(t)I(r>z(t)) = \pi(t)\bigg\};$$

(iv) *the following transversality conditions hold*:

$$\eta_{(n,r)}(0) + \int_{[T,0]} \gamma_{(n,r)}(s)\,\mu(ds) = \zeta r, \quad (n,r)\in\bar{\mathcal{S}}; \tag{31}$$

(v) $\boldsymbol{\eta}(\cdot)$, μ, *and* ζ *are not simultaneously equal to zero.*

As one can see, the main complexity for the analysis of these optimality conditions lies in the presence of measure μ, which plays the role of a Lagrange multiplier for state constraint (26).

6. Optimal Policy Approximation Techniques

The optimal dynamic pricing policy for the risk-adjusted problem may be hard to compute because of the size of the state space and the large number of steps required for time discretization. Such a policy may also be hard to apply in practice if frequent price changes are not possible. Thus, a decision maker may want to have an assortment of methods that provide approximate solutions of a simpler structure.

A particularly simple type of policy is the *fixed-price policy* in which the price remains constant during the entire planning horizon. It is also very easy to find an optimal policy in this restricted class. Indeed, under fixed-price p, the total number of customer arrivals is a Poisson random variable N'_p with the rate

$$\Lambda(p) = \int_T^0 \lambda(t,p)\,dt. \tag{32}$$

The expected value of the number of sales N_p can be expressed as

$$\begin{aligned}E[N_p] &= E[\min\{Y_T, N'_p\}] = \sum_{n=1}^{Y_T} nP(N'_p = n) + Y_T P(N'_p > Y_T) \\ &= \sum_{n=1}^{Y_T} n\frac{\Lambda^n(p)e^{-\Lambda(p)}}{n!} + Y_T P(N'_p > Y_T) \\ &= \Lambda(p)\sum_{n=0}^{Y_T-1} \frac{\Lambda^n(p)e^{-\Lambda(p)}}{n!} + Y_T P(N'_p > Y_T) \\ &= \Lambda(p)P(N'_p < Y_T) + Y_T P(N'_p > Y_T),\end{aligned}$$

and the hybrid objective value as

$$\begin{aligned}&E[R(0)] - CP(R(0) < z) \\ &\quad = p(\Lambda(p)P(N'_p < Y_T) + Y_T P(N'_p > Y_T)) \\ &\qquad - CP(\min\{Y_T, N'_p\} < z/p).\end{aligned}$$

For each given value of p, it is easy to evaluate $\Lambda(p)$ using numerical integration techniques. The objective function is also easy to evaluate because it is expressed via the cumulative distribution function (and its complement) of the Poisson random variable with a given rate $\Lambda(p)$. We will denote the optimal fixed-price policy selected according to this method as FP.

While the computation of the FP policy is easy, it is hard to get analytic results directly. On the other hand, it is possible to approximate the total demand random variable N'_p using the normal distribution: $N'_p \approx \Lambda(p) + X\sqrt{\Lambda(p)}$, where X is $N(0,1)$. In the following, we consider the uncapacitated case. The analysis can be generalized by considering prices $p > z/Y_T$ (otherwise, reaching the revenue target is impossible). Using $\Phi(x)$ and $\phi(x)$ to denote the standard normal cumulative distribution and density functions, respectively, and $\rho(p) = p\Lambda(p)$ to denote the expected revenues in the uncapacitated case, we can approximate the hybrid objective by $\rho(p) - C\Phi(f(p))$, where

$$f(p) = \sqrt{\Lambda(p)}\bigg(\frac{z}{\rho(p)} - 1\bigg).$$

Note that $\Lambda(p)$ is typically assumed to be decreasing where it is positive, whereas $\rho(p)$ is assumed to be unimodal and to satisfy the regularity condition $\lim_{p\to\infty}\rho(p) = 0$. To simplify the analysis, we additionally assume that $\rho(p)$ has the unique maximizer $\tilde{p}$, and $\rho(p)$ is strictly increasing (decreasing) for $p < \tilde{p}$ ($p > \tilde{p}$). Using first-order optimality conditions, we immediately get the following proposition (the proof is in the online appendix):

PROPOSITION 7. *Suppose that $\Lambda(p)$ is defined and differentiable for all $p \in \mathbb{R}_+$. Then, the price p^* that maximizes the normal approximation $\rho(p) - C\Phi(f(p))$ to the hybrid objective satisfies the equation*

$$r'(p) - C\phi(f(p))\left(\frac{\Lambda'(p)}{2\sqrt{\Lambda(p)}}\left(\frac{z}{\rho(p)} - 1\right) - \frac{zr'(p)\sqrt{\Lambda(p)}}{r^2(p)}\right) = 0. \qquad (33)$$

Moreover, if $z \geqslant \rho(\tilde{p})$, then $p^* \geqslant \tilde{p}$. If $z < \rho(\tilde{p})$, there exists $\bar{p}_1 < \tilde{p} < \bar{p}_2$ such that $\rho(\bar{p}_1) = \rho(\bar{p}_2) = z$, and either $p^* > \bar{p}_2$ or $\bar{p}_1 < p^* < \tilde{p}$.

An improvement upon policy FP is possible if we partition the entire planning horizon into k equal intervals and use a fixed price on each of them. We define a *myopic policy with k prices* (denoted as Mk) as the policy that uses k prices selected in the beginning of each interval as the optimal fixed prices from that point until the end of the planning horizon. We call this policy myopic because it ignores the possibility of future price changes. Such a policy maintains the advantages of the FP policy: ease of computation and implementation. It will be suboptimal compared to the optimal state feedback policy with k fixed prices selected at the beginning of each interval (we will denote such a policy as Ok). Because the policy Ok has to be computed using dynamic programming, it will be considerably harder to compute if the number of different prices $p_{\max}$ is large (resulting in a very large state space of the size $Y_T p_{\max}$). On the other hand, the policy Ok will lead to significantly simpler price paths than a fully dynamic pricing policy, and although lacking the transparency of the FP or Mk policies, may still be reasonably easy to implement in practice.

Experimental comparisons of the approximate policies mentioned above to the true optimum is presented in the next section, along with other numerical experiments.

7. Numerical Experiments

In this section, we present a series of numerical experiments with our model. In these experiments, we utilize a discrete-time approximation of Equation (18), as described in the online appendix. The online appendix also shows how we can organize calculations so that the optimal solution for a range of required level of revenues z from zero up to a specified maximum value can be obtained simultaneously. We consider this feature to be very important in practice because it saves computation time, lets a user of the model observe the effect of different values for z, and facilitates the selection of an appropriate value.

Section 7.1 gives an example of the optimal policy in the case of a stationary demand process. We will see that a managerially significant improvement in the probability of reaching a required level of revenue can be achieved at some moderate expense in terms of the expected value of revenue. However, we will also see that the effects of the risk-adjusted objective compared to the risk-neutral case decrease as the initial inventory increases. The effects, however, may also depend on the structure of demand. Thus, §7.2 describes an experiment in which we compare the following types of demand: one stationary and two nonstationary ones with generally increasing/decreasing willingness of customers to pay. This comparison is done based on the families of efficient frontiers for different values of z. We see that better results can be achieved when willingness to pay is increasing in time. Sections 7.1 and 7.2 also provide comparisons of the approximate policies with the optimal one. Finally, §7.3 presents a larger-scale example with 250 items and 25,000 time steps and shows that a managerially significant improvement in the probability of reaching the required level of revenue can still be attained.

7.1. The Behavior of the Model in the Case of a Stationary Demand Process

We first study the effects of introducing a risk term into a numerical example presented in §3.4 of Gallego and van Ryzin (1994). In that example, the demand function is stationary of the form

$$\lambda(t, p) = \bar{\lambda}\exp(-ap). \qquad (34)$$

The value of $\bar{\lambda}$ is selected so that $\bar{\lambda}\exp(-1)|T| = 10$. The initial inventory ranges from 1 to 20. The above choices correspond to those of the Gallego and van Ryzin (1994) example. We have also specified $a = 0.1$, $p_{\max} = 100$, $C = 100$, and $z = 200$. The value of $p_{\max} = 100$ signifies that 101 discrete price levels were used: $\mathcal{P} = \{0, 1, \ldots, 100\}$. The discretization step size was $\delta t = 10^{-3}|T|$, resulting in $K = 1{,}000$ time intervals. The computational procedure was implemented in `C++` and compiled with `g++` version 3.3.3 on a 3.4 GHz P4 workstation with 1 Gb of RAM running Linux. The computation of the optimal policy for this example took 2.54 seconds.

The optimal expected revenues for each $Y_T \in \{1, \ldots, 20\}$ and $z \in \{0, \ldots, 200\}$ can be obtained as $\eta^1_{(Y_T, r)} - r$, where $r = 200 - z$ (the quantity $\eta^1_{(n, r)}$ denotes the "expected revenues" component of the adjoint variable $\eta_{(n, r)}$ under an optimal pricing policy; see the online appendix for computational details). Note that all of this sensitivity information is obtained in a single run. These values are shown in Table 1. Naturally, the revenues tend to drop with a decrease in Y_T. In addition, we see that for many values of Y_T, the revenues decrease, then increase slightly, drop again, and finally, increase as z varies. That is, the optimal $E[R(0)]$ is higher when z slightly exceeds $E[R(0)]$, with smaller expected revenues for z-values immediately above and below. Such behavior of the optimal expected revenues can be explained by the fact that one does not need to sacrifice as much in expected revenues to achieve a given increase in $P(R(0) \geqslant z)$ when z is near the center

Table 1. The optimal expected revenues for several levels of initial inventory Y_T and desired levels of revenue z: stationary demand example of Gallego and van Ryzin (1994).

	Y_T									
z	1	3	5	7	9	11	13	15	17	19
200	24.0	54.3	73.0	84.9	92.2	96.4	98.4	99.3	99.8	99.9
190	24.0	54.3	73.0	84.9	92.2	96.2	98.3	99.3	99.7	99.8
180	24.0	54.3	73.0	84.9	92.1	96.0	98.1	99.3	99.6	99.7
170	24.0	54.3	73.0	84.9	91.7	95.7	98.1	99.2	99.5	99.6
160	24.0	54.3	73.0	84.6	91.2	95.6	98.1	99.0	99.3	99.4
150	24.0	54.3	73.0	83.8	90.9	95.6	97.9	98.8	99.1	99.2
140	24.0	54.3	72.8	83.0	91.0	95.5	97.8	98.7	99.0	99.1
130	24.0	54.3	71.5	82.9	91.1	95.1	97.5	98.4	98.9	99.0
120	24.0	54.3	70.3	83.2	90.8	95.4	97.6	98.5	98.8	99.0
110	24.0	54.3	70.1	83.4	90.2	95.0	97.3	98.5	98.8	99.0
100	24.0	52.0	70.8	82.9	89.9	94.8	97.6	98.4	98.9	99.1
90	24.0	50.0	71.2	82.0	90.5	95.1	97.5	98.5	99.0	99.1
80	24.0	50.4	70.5	82.0	90.5	95.5	97.5	98.6	99.1	99.2
70	24.0	51.9	69.7	83.1	91.0	95.6	97.7	98.8	99.2	99.3
60	23.9	52.0	70.2	83.5	91.5	95.8	98.0	98.9	99.3	99.5
50	18.2	51.0	71.6	84.1	91.8	96.0	98.2	99.2	99.5	99.6
40	17.5	51.8	72.2	84.6	92.0	96.2	98.4	99.3	99.7	99.8
30	22.3	53.2	72.7	84.8	92.1	96.3	98.5	99.4	99.8	99.9
20	21.7	53.9	72.9	84.9	92.2	96.4	98.5	99.5	99.8	100.0
10	22.6	54.2	73.0	84.9	92.2	96.4	98.5	99.5	99.9	100.0
0	24.0	54.3	73.0	84.9	92.2	96.4	98.5	99.5	99.9	100.0

of the distribution of $R(0)$ (which roughly corresponds to $E[R(0)]$). As z deviates from the center of the distribution, a more prominent decrease in $E[R(0)]$ results from optimizing the risk-adjusted objective (3). As the deviation of z from the center becomes more extreme, it is difficult to achieve a significant increase in $P(R(0) \geqslant z)$. Therefore, the term $-CP(R(0) < z)$ will not significantly affect the policy, and, as a consequence, the revenues. In Table 2, we also present the difference in the expected revenues under the risk-neutral expected revenue maximization versus the

Table 2. The percentage difference in the optimal expected revenues in the risk-neutral and the risk-adjusted models relative to the risk-neutral case for each level of initial inventory Y_T and desired level of revenue z.

	Y_T									
z	1	3	5	7	9	11	13	15	17	19
200	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.2	0.1	0.1
190	0.0	0.0	0.0	0.0	0.0	0.2	0.2	0.2	0.2	0.2
180	0.0	0.0	0.0	0.0	0.1	0.4	0.4	0.2	0.3	0.3
170	0.0	0.0	0.0	0.0	0.5	0.7	0.4	0.3	0.4	0.4
160	0.0	0.0	0.0	0.4	1.1	0.8	0.4	0.5	0.6	0.6
150	0.0	0.0	0.0	1.3	1.4	0.8	0.6	0.7	0.8	0.8
140	0.0	0.0	0.3	2.3	1.3	0.9	0.7	0.8	0.9	0.9
130	0.0	0.0	2.1	2.4	1.2	1.4	1.0	1.1	1.0	1.0
120	0.0	0.0	3.8	2.0	1.5	1.0	0.9	1.0	1.1	1.0
110	0.0	0.0	4.1	1.8	2.2	1.5	1.2	1.0	1.1	1.0
100	0.0	4.4	3.1	2.4	2.6	1.7	0.9	1.1	1.0	0.9
90	0.0	8.6	2.5	3.5	1.9	1.4	1.0	1.0	0.9	0.9
80	0.0	7.7	3.5	3.5	1.9	0.9	1.0	0.9	0.8	0.8
70	0.0	4.6	4.7	2.2	1.3	0.8	0.8	0.7	0.7	0.7
60	0.4	4.4	4.0	1.7	0.8	0.6	0.5	0.6	0.6	0.5
50	31.9	6.5	2.0	1.0	0.4	0.4	0.3	0.3	0.4	0.4
40	37.1	4.8	1.1	0.4	0.2	0.2	0.1	0.2	0.2	0.2
30	7.6	2.1	0.4	0.1	0.1	0.1	0.0	0.1	0.1	0.1
20	10.6	0.7	0.1	0.0	0.0	0.0	0.0	0.0	0.1	0.0
10	6.2	0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Table 3. The difference in the optimal $P(R(0) \geqslant z)$ in the risk-neutral and the risk-adjusted models for each level of initial inventory Y_T and desired level of revenue z.

	Y_T									
z	1	3	5	7	9	11	13	15	17	19
200	0.000	0.000	0.000	0.000	0.000	0.001	0.003	0.005	0.005	0.003
190	0.000	0.000	0.000	0.000	0.000	0.003	0.008	0.009	0.007	0.005
180	0.000	0.000	0.000	0.000	0.002	0.011	0.017	0.015	0.011	0.008
170	0.000	0.000	0.000	0.000	0.011	0.026	0.028	0.021	0.017	0.012
160	0.000	0.000	0.000	0.005	0.032	0.045	0.038	0.028	0.023	0.016
150	0.000	0.000	0.000	0.024	0.061	0.062	0.047	0.040	0.028	0.022
140	0.000	0.000	0.002	0.063	0.091	0.076	0.056	0.050	0.032	0.026
130	0.000	0.000	0.027	0.112	0.111	0.085	0.071	0.055	0.033	0.030
120	0.000	0.000	0.083	0.148	0.122	0.089	0.080	0.053	0.033	0.032
110	0.000	0.000	0.161	0.164	0.128	0.100	0.088	0.046	0.032	0.031
100	0.000	0.036	0.212	0.169	0.125	0.103	0.081	0.039	0.030	0.029
90	0.000	0.131	0.217	0.164	0.116	0.099	0.067	0.031	0.026	0.026
80	0.000	0.268	0.211	0.142	0.101	0.080	0.049	0.024	0.022	0.022
70	0.000	0.299	0.187	0.108	0.080	0.056	0.032	0.018	0.017	0.018
60	0.001	0.256	0.136	0.076	0.053	0.033	0.019	0.012	0.012	0.012
50	0.120	0.218	0.083	0.045	0.028	0.015	0.009	0.006	0.007	0.007
40	0.378	0.138	0.044	0.022	0.011	0.005	0.003	0.002	0.003	0.003
30	0.375	0.060	0.019	0.007	0.003	0.001	0.001	0.001	0.001	0.001
20	0.195	0.025	0.005	0.001	0.000	0.000	0.000	0.000	0.000	0.000
10	0.080	0.003	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
0	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

maximization of the risk-adjusted objective (3). The difference becomes smaller for larger values of Y_T because the uncertainty in revenues relative to the expected value and, consequently, the effects of the risk term on policies, decrease.

Table 3 shows the effect of the risk-adjusted objective in terms of the increase in $P(R(0) \geqslant z)$ versus the risk-neutral case. As one can see, a significant increase in this probability can be achieved for various values of z (especially for the smaller Y_Ts). Unfortunately, the achieved increase in $P(R(0) \geqslant z)$ becomes smaller with an increase in the initial inventory. Table 4 provides a summary of these comparisons for the value of initial inventory $Y_T = 10$. One can see that the practically significant improvements (above 5%) in $P(R(0) \geqslant z)$ are obtained in the wide range of z from 60 to 160. Smaller improvements can also be important if they correspond to the low-range percentiles of the distribution of $R(0)$, such as $z = 30, \ldots, 60$. For example, the probability of failing to reach $z = 50$ has decreased from 4.79% to 1.99% (corresponding to an increase in $P(R(0) \geqslant z)$ from 0.9521 to 0.9801).

We also examine the features of the optimal price $p(T, n, r)$ at time T (beginning of the planning horizon). Of course, in the original problem formulation, the initial value of the revenue is zero, but due to our simultaneous sensitivity analysis, we have the optimal price available for different initial levels of revenue as well. The price exhibits strikingly different behavior from the price optimal under the risk-neutral maximization. The price profile is shown in Figure 1. As suggested by Proposition 4, the optimal price is below the one for the risk-neutral case for values of initial revenue r immediately below z. However, we see that the prices for even smaller revenues r can become significantly higher and subsequently drop. This behavior is intuitively reasonable because even a small possibility of reaching the critical revenue threshold z can encourage a price increase. However, as the chance of reaching

Table 4. The summary of the optimal $E[R(0)]$ and $P(R(0) \geqslant z)$ for different values of z, and corresponding changes $\Delta E[R(0)]$ and $\Delta P(R(0) \geqslant z)$ compared to the risk-neutral case ($C = 0$) for $Y_T = 10$.

z	$E[R(0)]$	$\Delta E[R(0)]$	$P(R(0) \geqslant z)$	$\Delta P(R(0) \geqslant z)$
200	93.4	−1.2	0.0024	0.0024
190	92.2	−2.4	0.0070	0.0070
180	91.3	−3.3	0.0155	0.0155
170	90.8	−3.8	0.0302	0.0299
160	90.6	−4.0	0.0543	0.0510
150	90.7	−3.9	0.0921	0.0750
140	91.0	−3.6	0.1476	0.0949
130	90.8	−3.8	0.2247	0.1083
120	90.1	−4.5	0.3240	0.1183
110	88.7	−5.9	0.4422	0.1293
100	86.6	−8.0	0.5706	0.1337
90	85.0	−9.6	0.6961	0.1286
80	87.6	−7.0	0.8042	0.1117
70	87.7	−6.9	0.8912	0.0875
60	89.6	−5.0	0.9485	0.0558
50	91.5	−3.1	0.9801	0.0280
40	93.1	−1.5	0.9938	0.0107
30	94.1	−0.5	0.9985	0.0030
20	94.5	−0.1	0.9997	0.0005
10	94.6	0.0	1.0000	0.0001
0	94.6	0.0	1.0000	0.0000

Figure 1. The initial optimal price $p(T, n, r)$ for each n, r.

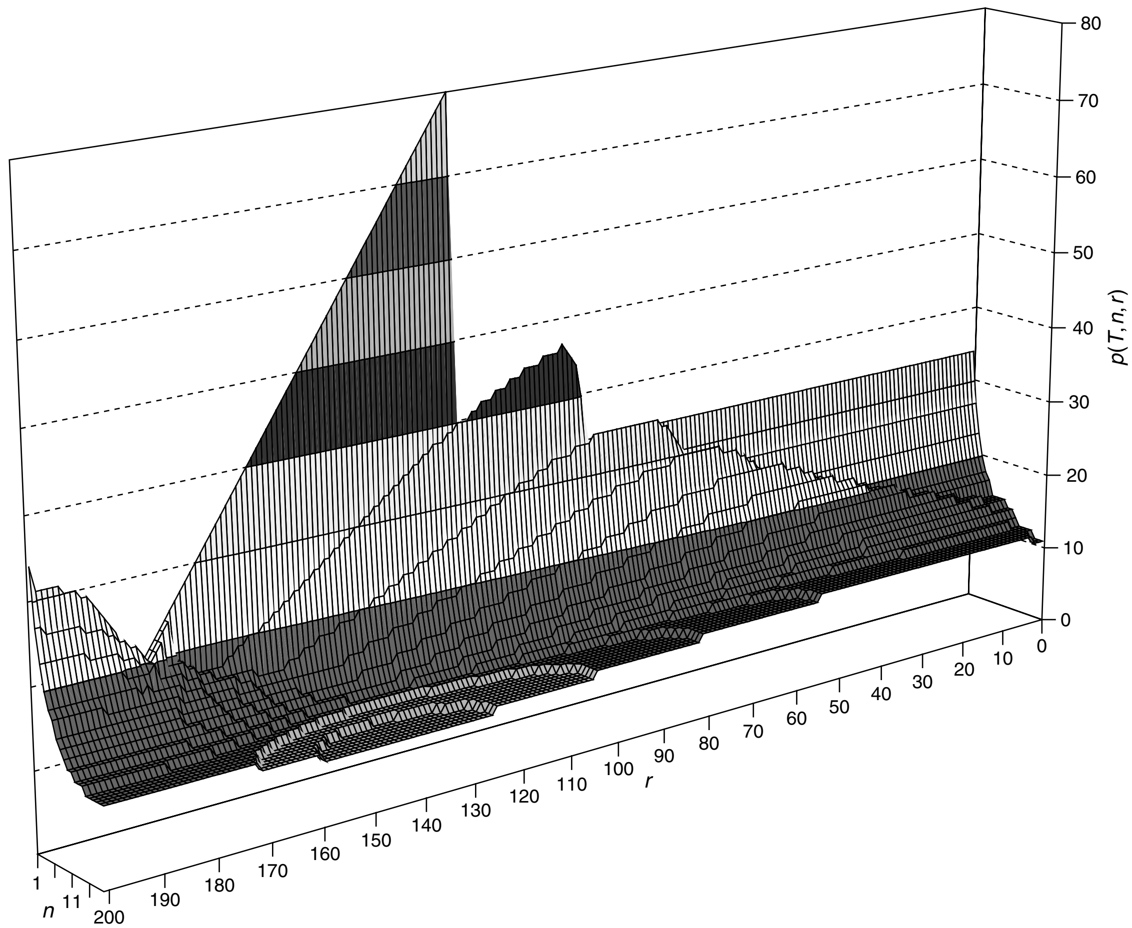

z becomes smaller, maximizing expected revenue becomes more important, which results in a lower price. Another important observation is that even under the exponential demand model employed in this experiment, $p(T, n, r)$ is not necessarily decreasing in n for a fixed $r < z$. One possible explanation is that having more items on hand increases the chance to reach the level z, thus driving the price higher under some scenarios.

Finally, in Table 5, we present the performance of the approximate policies discussed in §6 applied to this example. Along with the fixed price, we consider myopic and optimal policies with 2, 4, 10, and 20 prices (denoted as FP, M2, M4, M10, M20, O2, O4, O10, and O20, respectively). The entries in the table represent the maximum gap between the optimal value of the hybrid objective and its approximation measured as the percentage of the optimal risk-neutral revenues (which give an upper bound for the optimal hybrid objective function) over all values of $z = 0, \ldots, 200$. Generally, the performance of every approximate method improves with the initial inventory Y_T. This is natural because uncertainty in revenues decreases with Y_T. The optimal fixed-price policy, FP, results in a much larger gap (at least 9%) than Mk or Ok policies (below 4% and and 1% for $k = 20$ and larger values of Y_T). Increasing the number of prices in the myopic policy usually improves the gap. For some values, it actually increases the gap because the myopic policy does not take into account likely future price changes. The performance of Ok policies improves monotonically and is much better than that of Mk policies with the same number of prices.

7.2. The Behavior of the Model in the Case of Nonstationary Distribution of the Reservation Price

Our second experiment is related to the observation that improvements in $P(R(0) \geqslant z)$ tended to get smaller with an increase in Y_T under the exponential stationary demand structure. This, however, is strongly related to the property of stationarity of the reservation price distribution. To recover a reservation price distribution from the rate of the sales process, let $\bar{\lambda}(t) = \lambda(t, 0)$ be the (finite) maximum of $\lambda(t, p)$ with respect to p at each $t \in [T, 0]$. Then, the quantity $\lambda(t, p)/\bar{\lambda}(t)$ can be interpreted as the probability that a customer arriving according to a Poisson process with rate $\bar{\lambda}(t)$ buys a product priced at p (see, for example, Gallego and van Ryzin 1994 and Zhao and Zheng 2000). This interpretation is consistent with the view that a customer compares a directly unobservable quantity, his or her

Table 5. The maximum gap between the approximate and optimal hybrid objective values as a percentage of the optimal risk-neutral revenues for different values of Y_T: the stationary demand example of Gallego and van Ryzin (1994).

	Approximate policy (%)								
Y_T	FP	M2	M4	M10	M20	O2	O4	O10	O20
1	24.51	19.87	20.18	20.43	22.93	13.29	5.31	1.56	0.43
2	23.41	21.48	18.77	19.71	19.89	10.90	4.44	2.29	1.78
3	20.77	17.69	16.14	13.82	13.60	8.04	5.04	2.21	1.31
4	18.31	15.76	12.42	10.77	9.87	8.03	4.45	2.04	1.17
5	16.05	14.99	10.31	8.37	7.40	7.92	4.31	2.05	1.14
6	14.63	11.94	7.74	6.11	5.23	7.88	4.27	2.05	1.12
7	12.73	10.11	6.76	4.22	3.67	7.09	4.08	2.03	1.12
8	11.93	7.80	5.30	3.83	3.23	6.46	3.88	1.94	1.11
9	10.82	8.64	5.03	4.02	3.37	6.12	3.70	1.82	1.04
10	9.25	6.80	4.90	3.66	3.14	6.01	3.39	1.72	0.95
11	9.10	6.66	5.51	3.84	3.24	5.48	3.18	1.55	0.84
12	10.49	7.53	5.43	4.01	3.53	5.37	2.93	1.47	0.76
13	11.33	8.40	5.55	3.86	3.36	5.41	2.77	1.34	0.69
14	12.01	8.90	5.78	3.99	3.54	5.44	2.66	1.24	0.60
15	12.43	9.21	5.99	4.07	3.61	5.45	2.57	1.16	0.55
16	12.27	8.29	5.48	4.12	3.68	5.21	2.53	1.11	0.52
17	12.38	8.41	5.54	4.17	3.73	5.07	2.45	1.07	0.51
18	12.47	8.47	5.57	4.19	3.76	5.00	2.40	1.04	0.50
19	12.52	8.50	5.59	4.20	3.75	4.96	2.37	1.02	0.49
20	12.54	8.52	5.62	4.21	3.76	4.94	2.36	1.02	0.49

reservation price, to p, and makes a purchase if the reservation price is higher than p. Accordingly, for each t, a probability distribution function of the form $F_t(p) = 1 - \lambda(t,p)/\bar{\lambda}(t)$ can be thought of as a distribution of reservation price. Gallego and van Ryzin (1994) considered a case in which $F_t(p)$ is *stationary*, i.e., independent of t, while Zhao and Zheng (2000) studied a general, nonstationary case. In their analysis of the stationary case, Gallego and van Ryzin (1994) showed that the fixed-price heuristic is asymptotically optimal. This suggests that it may be harder to achieve significant improvements in $P(R(0) \geqslant z)$ by a dynamic (versus fixed) pricing policy for larger initial inventories. As well, the fixed-pricing policy is itself not useful for improving $P(R(0) \geqslant z)$ when the inventory is large. Thus, there is little hope of improving $P(R(0) \geqslant z)$ when the reservation price distribution is stationary and inventory is large. In contrast, Zhao and Zheng (2000) present an example in which the reservation price distribution is not stationary, and, accordingly, the dynamic pricing strategy is extremely relevant even for large inventories. Thus, we will compare the stationary model form used previously:

$$\lambda_0(t,p) = \bar{\lambda}\exp\{-a_0 p\},$$

with two nonstationary models:

$$\lambda_1(t,p) = \bar{\lambda}\exp\{-(a_1 + bt)p\},$$

$$\lambda_2(t,p) = \bar{\lambda}\exp\{-(a_2 - bt)p\},$$

where $b > 0$. The parameters were chosen so that $a_0 = 0.2$, $a_1 = 0.3$, $a_2 = 0.1$, and $b = 0.2/|T|$. The function $\lambda_0(t,p)$ is a standard exponential demand function. The second function, $\lambda_1(t,p)$, corresponds to a situation when the value of each item deteriorates with time and customers are willing to pay less and less for it; and the third one, $\lambda_2(t,\, p)$, exhibits the opposite behavior. We used $Y_T = 50$, $z = 300$, $p_{\max} = 100$, and $\bar{\lambda} = \exp(1)Y_T/|T|$. The calculations were done for 5,000 time steps for each of the values of C from the set

$$\{0.005, 5.0, 50.0, 125.0, 250.0, 375.0, 500.0, 5{,}000.0\}$$

(total computation time was 384 seconds for each model). This resulted in an approximation to the efficient frontier in the space of $E[R(0)]$, $P(R(0) \geqslant z)$. Again, note that the simultaneous sensitivity analysis feature of our approach provided us with the optimal $E[R(0)]$, and $P(R(0) \geqslant z)$ for all values of z in the range from 0 to 300. Therefore, the efficient frontiers were simultaneously computed for $0 \leqslant z \leqslant 300$. Because it is impractical to present these curves for each value of z, we only plot them for values $z \in \{200, 210, \ldots, 290, 300\}$ in Figure 2. The efficient frontiers of the most favorable shape correspond to the "increasing value" demand structure because they show a possibility of significant improvement in $P(R(0) \geqslant z)$ without sacrificing much of the expected revenue. One can see this in the length of the "shoulders" in the top portion of the curve (for each respective value of z in all three figures) as well as the sharpness of the curve's drop from the shoulders. In contrast, the case of "decaying value" demand structure appears to be the most unfavorable because the maximum possible improvement in probability is smaller and requires

Figure 2. The family of efficient frontiers for $Y_T = 50$: stationary $\lambda_0(t, p)$, "decaying value" $\lambda_1(t, p)$, and "increasing value" $\lambda_2(t, p)$ demand structures.

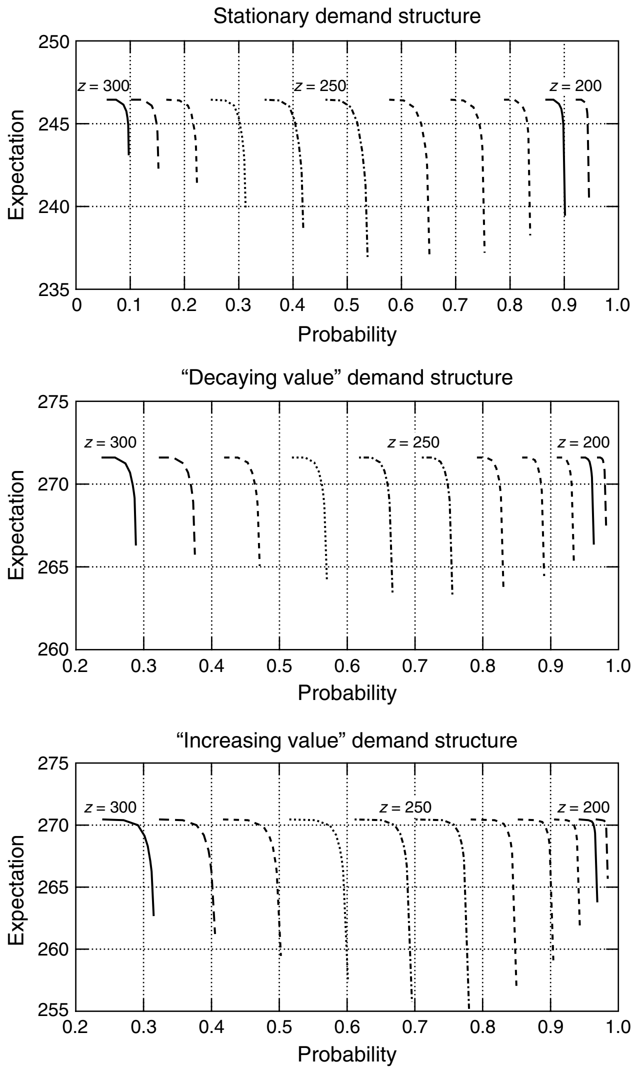

a significant sacrifice in expected revenue. Thus, one can see that the presence of nonhomogeneity of demand and its type may significantly change the effects from introduction of a risk term, even for larger inventories.

Finally, we examine the performance of the approximate methods as a function of cost C of failure to reach the desired level of revenues and the type of nonhomogeneity of demand structure. Table 6 shows the maximum gap (over z) for each demand structure, each level of C, and the approximations FP (fixed price), M10 (myopic with 10 prices), and O10 (optimal with 10 prices). Generally, the difficulty of the problem increases with C for each of these approximations. The FP policy performs much better for the stationary demand structure relative to "decreasing" or "increasing value" structures. M10 does not improve over FP significantly when nonstationarity is present, whereas there is a noticeable improvement for the stationary demand structure. O10 significantly improves upon both FP and M10 for either demand structure. This shows that successful optimization of the hybrid objective generally requires some form of feedback, but the feedback has to be optimal when the reservation price distribution is nonstationary. A smaller gap for the "decaying value" structure in the case of the O10 policy is explained by limited space for improvement in probability to reach the desired level z compared to other demand structures.

7.3. The Behavior of the Model for Large Values of Initial Inventory

It is interesting to examine the behavior of the model for much larger values of initial inventory. In this experiment, we studied the performance of our model on the example with $Y_T = 250$. We used a linear demand function, with the reservation price distribution changing in time in a manner similar to $\lambda_2(t, p)$:

$$\lambda_3(t, p) = \bar{\lambda}(1.0 - (a_2 - bt)p)^+,$$

where $\bar{\lambda} = \exp(1)Y_T/|T|$, $a_2 = 0.1$, $b = 0.2/|T|$ (the same as in the previous experiment). These parameters result in a shutdown price of at most 10 because $1.0 \leqslant 10(a_2 - bt)$ for all $t \in [T, 0]$. We note that although we also performed a similar large inventory experiment with the exponential demand function, we saw little improvement in $P(R(0) \geqslant z)$. Other model parameters were $z = 1{,}500$ and $C = 5{,}000$. Again, we note that the value $z = 1{,}500$ is simply the upper range of z values considered in the experiment, and the results for all smaller z are obtained simultaneously due to the setup of the numerical procedure. The discrete-time approximation utilized $K = 25{,}000$ time steps (total computation time was 2,265 seconds). We first show, in Figure 3, the probability mass function of $R(0)$ before the risk term is incorporated into the model. The mean, which is equal to 847.35, and the median of 852 are indicated by dashed and solid lines, respectively. Note a skew of the distribution to the left. This left skewness means that one has a somewhat higher probability of getting low revenues even though most of the realizations of the revenue distribution belong to the area with higher revenue values. Figure 4 superimposes the probability mass functions $P(R(0) = r)$ of the total revenue in the risk-neutral (solid line) and risk-adjusted with $z = 800$ (dashed line) problems. These distributions are computed by solving the initial value problem for state system (11) under the risk-neutral and risk-adjusted optimal controls, respectively. The probability that the total revenue is below $z = 800$ (that is, the area under the curve to the left of $z = 800$) reduces significantly under the risk-adjusted policy. It is achieved by shifting the bulk of the total revenue distribution to the left. The mean of the total revenues decreases from 847.35 to 824.94, and the standard deviation from 36.99 to 24.56. Figure 5 shows an (absolute) improvement in $P(R(0) \geqslant z)$

Table 6. The maximum gap between the approximate and optimal hybrid objective values as a percentage of the optimal risk-neutral revenues for different values of C: stationary $\lambda_0(t,p)$, "decaying value" $\lambda_1(t,p)$, and "increasing value" $\lambda_2(t,p)$ demand structures.

	Approximate policy								
	FP (%)			M10 (%)			O10 (%)		
C	$\lambda_0(t,p)$	$\lambda_1(t,p)$	$\lambda_2(t,p)$	$\lambda_0(t,p)$	$\lambda_1(t,p)$	$\lambda_2(t,p)$	$\lambda_0(t,p)$	$\lambda_1(t,p)$	$\lambda_2(t,p)$
0.005	0.97	5.34	4.95	0.20	3.68	4.03	0.18	0.20	0.26
5	1.01	5.40	5.00	0.25	3.80	4.18	0.19	0.21	0.27
50	1.24	5.88	5.51	0.38	4.68	4.90	0.32	0.28	0.37
125	4.29	6.69	6.46	0.52	6.20	5.59	0.53	0.37	0.57
250	4.58	8.14	8.07	1.19	8.53	6.88	1.02	0.50	0.93
375	4.97	13.02	13.01	1.56	10.08	8.06	1.20	0.66	1.28
500	5.36	13.63	13.76	1.77	11.12	9.38	1.57	0.88	1.67
5,000	40.72	73.32	82.75	12.94	55.43	62.64	11.42	6.41	12.81

relative to the risk-neutral model for various values of z and a relative decrease in revenue with respect to the base value of 847.35 (computed with $C = 0$). The improvement in $P(R(0) \geqslant z)$ can be significant, reaching close to 9% for intermediate values of $z = 800, \ldots, 900$. We also note that, for example, a company with a target level of revenue $z = 817$ facing a 20% loss probability could reduce this to 12% with relatively little (2.3%) loss in expected revenues. In general, the overall cost of improvement in probability in terms of expected revenue is not very high. On the other hand, as seen from the graph, the decrease in revenue is always less than 3%.

8. Conclusions and Directions for Future Research

This paper presents a new dynamic pricing model that permits control of both expected revenues and the risk that total revenues will fall below a desired minimum. This model would permit decision makers in a number of emerging application areas for revenue management to seek pricing policies that balance maximization of expected revenues against risk of poor performance. We show that this model falls within a general class of continuous-time optimal control problems, demonstrate the existence of optimal solutions under reasonable assumptions, provide optimality conditions, and explore some of the structural properties of solutions. We also present an efficient computational procedure that allows evaluation of this model on problems of realistic size. Our approach also provides useful and numerically tractable analytical tools that can aid in decision making.

The future research related to this model may include the following topics:

(1) embedding this model into a game-theoretic framework that models the effects of risk in a competitive setting;

(2) studying the effects of price guarantees (Levin et al. 2007) on the behavior of our model;

(3) extending the model to include such factors as multiple product types, returns, exchanges, and cancellations; and

(4) incorporating loss-probability risk into general inventory management systems including variable and stochastic costs, and stochastic interest rates.

Figure 3. The probability mass function $P(R(0) = r)$ of the total revenue in the risk-neutral problem for linear "increasing value" demand structure $\lambda_3(t,p)$ and $Y_T = 250$.

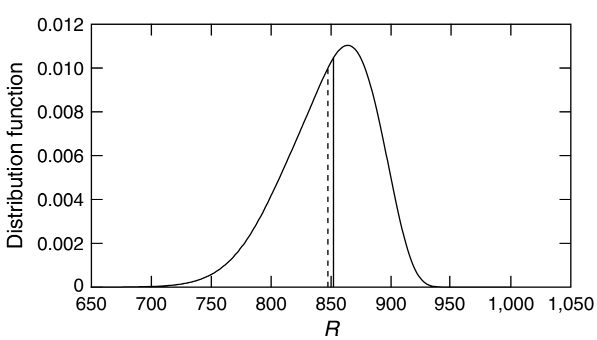

Note. The mean and the median are indicated by dashed and solid lines, respectively.

Figure 4. The probability mass functions $P(R(0) = r)$ of the total revenue in the risk-neutral (solid line) and risk-adjusted with $z = 800$ (dashed line) models for linear "increasing value" demand structure $\lambda_3(t,p)$ and $Y_T = 250$.

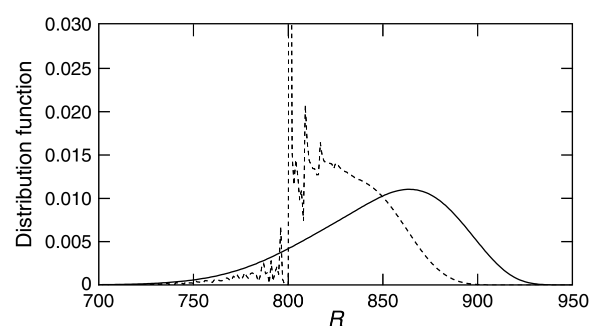

Figure 5. The improvement in $P(R(0) \geqslant z)$, and the decrease in $E[R(0)]$ for linear "increasing value" demand structure $\lambda_3(t, p)$ and $Y_T = 250$.

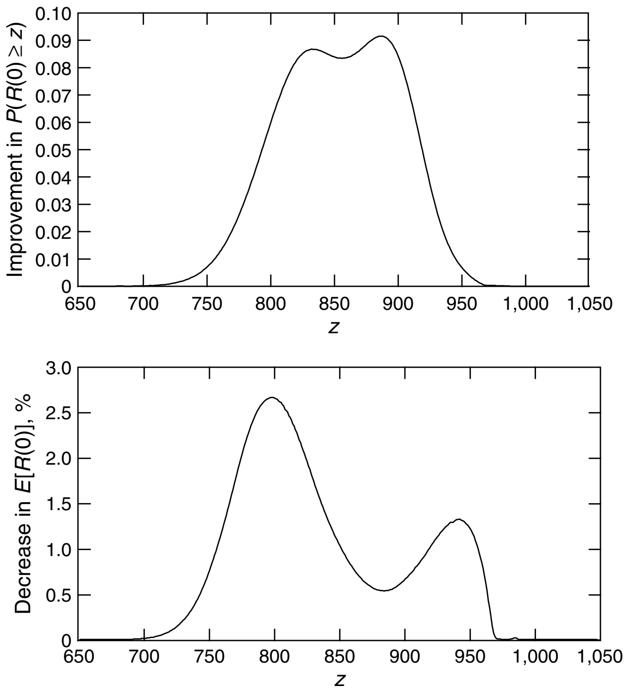

9. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at http://or.journal.informs.org/.

References

Agrawal, V., S. Seshadri. 2000. Impact of uncertainty and risk aversion on price and order quantity in the newsvendor problem. *Manufacturing Service Oper. Management* **2**(4) 410–423.

Belobaba, P. 1987. Airline yield management: An overview of seat inventory control. *Transportation Sci.* **21** 63–73.

Bitran, G., R. Caldentey. 2003. An overview of pricing models for revenue management. *Manufacturing Service Oper. Management* **5**(3) 203–229.

Chan, L., Z. Shen, D. Simchi-Levi, J. Swann. 2004. Coordination of pricing and inventory decisions: A survey and classification. D. Simchi-Levi, D. Wu, Z. Shen, eds. *Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era*. Kluwer Academic Publishers, Boston, MA, 335–392.

Chen, F., A. Federgruen. 2000. Mean-variance analysis of basic inventory models. Working paper, Columbia University, New York.

Chen, X., M. Sim, D. Simchi-Levi, P. Sun. 2006. Risk aversion in inventory management. *Oper. Res.* Forthcoming.

Clarke, F. 1990. *Optimization and Nonsmooth Analysis. Classics in Applied Mathematics*, Vol. 5, 2nd ed. Society for Industrial and Applied Mathematics, Philadelphia, PA.

Eeckhoudt, L., C. Gollier, H. Schlesinger. 1995. The risk-averse (and prudent) newsboy. *Management Sci.* **41**(5) 786–794.

Elmaghraby, W., P. Keskinocak. 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices and future directions. *Management Sci.* **49**(10) 1287–1309.

Feng, Y., B. Xiao. 1999. Maximizing revenues of perishable assets with a risk factor. *Oper. Res.* **47**(2) 337–341.

Feng, Y., B. Xiao. 2000a. A continuous time yield management model with multiple prices and reversible price changes. *Management Sci.* **46**(5) 644–657.

Feng, Y., B. Xiao. 2000b. Optimal policies of yield management with multiple predetermined prices. *Oper. Res.* **48**(2) 332–343.

Feng, Y., B. Xiao. 2004. A combined risk and revenue analysis for managing perishable products. Working paper, Chinese University of Hong Kong, Hong Kong, China.

Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* **40**(8) 999–1020.

Gīhman, I., A. Skorohod. 1979. *Controlled Stochastic Processes*. Springer-Verlag, New York.

Gnedenko, B. 1967. *The Theory of Probability.* Chelsea Publishing Co., New York.

Kahneman, D., A. Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* **47**(2) 263–291.

Lee, E., L. Markus. 1967. *Foundations of Optimal Control Theory*. John Wiley & Sons Inc., New York.

Levin, Y., J. McGill, M. Nediak. 2007. Price guarantees in revenue management and dynamic pricing. *Oper. Res.* **55**(1) 75–97.

Lim, A., G. Shanthikumar. 2007. Relative entropy, exponential utility, and robust dynamic pricing. *Oper. Res.* **55**(2) 198–214.

Martínez-de Albéniz, V., D. Simchi-Levi. 2006. Mean-variance trade-offs in supply contracts. *Naval Res. Logist.* **53**(7) 603–616.

McGill, J., G. van Ryzin. 1999. Revenue management: Research overview and prospects. *Transportation Sci.* **33**(2) 233–256.

Mohler, R. 1973. *Bilinear Control Processes with Applications to Engineering, Ecology, and Medicine. Mathematics in Science and Engineering*, Vol. 106. Academic Press, New York/London, UK.

Phillips, R. 2005. *Pricing and Revenue Optimization*. Stanford University Press, Stanford, CA.

Phillips, R. 2006. Personal communication. Nomis Solutions, CA.

Secomandi, N., K. Abbott, T. Atan, E. A. Boyd. 2002. From revenue management concepts to software systems. *Interfaces* **32**(2) 1–11.

Simon, H. 1956. Rational choice and the structure of the environment. *Psych. Rev.* **63**(2) 129–138.

Talluri, K., G. van Ryzin. 2004. *The Theory and Practice of Revenue Management*. Kluwer, Boston, MA.

Vinter, R. 2000. *Optimal Control. Systems & Control: Foundations & Applications*. Birkhäuser, Boston, MA.

Weatherford, L., S. Bodily. 1992. A taxonomy and research overview of perishable-asset revenue management: Yield management, overbooking, and pricing. *Oper. Res.* **40** 831–844.

Yano, C., S. Gilbert. 2004. Coordinated pricing and production/procurement decisions: A review. A. Chakravarty, J. Eliashberg, eds. *Managing Business Interfaces: Marketing, Engineering Issues and Manufacturing Perspectives*. Kluwer Academic Publishers, Norwell, MA, 65–104.

Zhao, W., Y. Zheng. 2000. Optimal dynamic pricing for perishable assets with nonhomogeneous demand. *Management Sci.* **46**(3) 375–388.

OPERATIONS RESEARCH
DOI 10.1287/opre.1070.0438ec
pp. ec1–ec4

e-companion

ONLY AVAILABLE IN ELECTRONIC FORM

Electronic Companion—“Risk in Revenue Management and Dynamic Pricing” by Yuri Levin, Jeff McGill, and Mikhail Nediak, *Operations Research* 2007, DOI 10.1287/opre.1070.0438.

1. Online Appendix: Mathematical Proofs

To establish notation, we first briefly present a standard statement of the well-known Pontryagin maximum principle for a minimization problem.

1.1. Pontryagin Maximum Principle

For a generic optimal control problem on the fixed interval $[T, 0]$

$$\min \; g(\mathbf{x}(0)) \tag{EC1}$$

$$\text{s.t.} \; \frac{d\mathbf{x}}{dt} = f(t, \mathbf{x}(t), \mathbf{u}(t)), \tag{EC2}$$

$$\mathbf{x}(T) = \mathbf{x}_0, \tag{EC3}$$

$$\mathbf{u}(t) \in U(t), \tag{EC4}$$

one introduces *adjoint variables* $\mathbf{y}$ for equations (EC2), and a *Hamiltonian* function $H(t, \mathbf{x}, \mathbf{y}, \mathbf{u}) = \mathbf{y}^T f(t, \mathbf{x}, \mathbf{u})$. The maximum principle claims that for an optimal $\mathbf{x}(\cdot)$, $\mathbf{u}(\cdot)$, there exists $\mathbf{y}(\cdot)$ and a constant $\zeta \in \{0, 1\}$ not both identically equal to zero such that

- $\mathbf{y}(\cdot)$ satisfies the *adjoint system* of differential equations

$$\frac{d\mathbf{y}}{dt} = -\frac{\partial H}{\partial \mathbf{x}}(t, \mathbf{x}(t), \mathbf{y}(t), \mathbf{u}(t))$$

with the boundary condition

$$\mathbf{y}(0) = -\zeta \frac{\partial g}{\partial \mathbf{x}}(\mathbf{x}(0))$$

- $\mathbf{u}(t)$ delivers the maximum of $H(t, \mathbf{x}(t), \mathbf{y}(t), \mathbf{u})$ with respect to $\mathbf{u} \in U(t)$ almost everywhere on $[T, 0]$.

1.2. Proof of Proposition 2

The above proposition is formulated for a minimization problem, so one needs appropriate adjustments to be able to use the maximum principle. We simply select

$$g(\mathbf{P}) = -\sum_{(n,r)\in\bar{\mathcal{S}}} rP_{(n,r)} + C \sum_{(n,r)\in\bar{\mathcal{S}}:\, r<z} P_{(n,r)}.$$

To apply the maximum principle to our problem, we introduce the adjoint variables $\eta_{(n,r)}(t)$ for each of the evolution equations in (11). The Hamiltonian function has the form

$$H(t, \mathbf{P}, \boldsymbol{\eta}, \boldsymbol{\pi}) = \sum_{(n,r)\in\bar{\mathcal{S}}} \eta_{(n,r)} \left[-\sum_{p} \lambda(t, p)\pi(p \mid \cdot, n, r)P_{(n,r)} + \sum_{p:\,(n+1,r-p)\in\mathcal{S}} \lambda(t, p)\pi(p \mid \cdot, n+1, r-p)P_{(n+1,r-p)} \right].$$

Note that H is linear in $\boldsymbol{\pi}$. The term corresponding to $\pi(p \mid \cdot, n, r)$ in H is of the form

$$\lambda(t, p)P_{(n,r)}[\eta_{(n-1,r+p)} - \eta_{(n,r)}]\pi(p \mid \cdot, n, r).$$

Moreover, we can rewrite the Hamiltonian as

$$H(t, \mathbf{P}, \boldsymbol{\eta}, \boldsymbol{\pi}) = \sum_{(n,r)\in\mathcal{S}} P_{(n,r)} \sum_{p} \lambda(t, p)[\eta_{(n-1,r+p)} - \eta_{(n,r)}]\pi(p \mid \cdot, n, r),$$

and observe that the problem of maximizing H with respect to $\boldsymbol{\pi}$ for a given $\boldsymbol{\eta}$ and $\mathbf{P}$ separates into completely independent linear subproblems of the form (16)–(17) for each $(n, r) \in \mathcal{S}$. The resulting adjoint system has the form (13)–(14) and the terminal condition is

$$\eta_{(n,r)}(0) = \zeta[r - CI(r < z)], \quad (n, r) \in \overline{\mathcal{S}}.$$

We observe that $\zeta = 0$ would imply $\boldsymbol{\eta} = 0$, which would contradict the requirement of the maximum principle that both ζ and $\boldsymbol{\eta}$ are not identically zero. We can select, without loss of generality, $\zeta = 1$ and get (15).

1.3. Proof of Lemma 1

We first assert that the following function is constant along the trajectories of $(\mathbf{P}, \boldsymbol{\eta})$ satisfying the system formed by (11) and (13)–(14), and *arbitrary* boundary conditions (that is, not only (7)–(8) and (15)):

$$G(\mathbf{P}, \boldsymbol{\eta}) = \sum_{(n,r)\in\overline{\mathcal{S}}} P_{(n,r)}\eta_{(n,r)}.$$

Indeed, due to bilinearity of H in $(\mathbf{P}, \boldsymbol{\eta})$, we have

$$H(t, \mathbf{P}, \boldsymbol{\eta}, \boldsymbol{\pi}) = \sum_{n,r} \eta_{(n,r)} \frac{\partial H}{\partial \eta_{(n,r)}} = \sum_{n,r} P_{(n,r)} \frac{\partial H}{\partial P_{(n,r)}}.$$

By construction of H, and the definition of the adjoint system, we have

$$\frac{dP_{(n,r)}}{dt} = \frac{\partial H}{\partial \eta_{(n,r)}} \qquad \text{and} \qquad \frac{d\eta_{(n,r)}}{dt} = -\frac{\partial H}{\partial P_{(n,r)}}.$$

Therefore,

$$\begin{aligned}\frac{dG(\mathbf{P}(t), \boldsymbol{\eta}(t))}{dt} &= \sum_{n,r}\left[\frac{dP_{(n,r)}}{dt}\eta_{(n,r)} + P_{(n,r)}\frac{d\eta_{(n,r)}}{dt}\right] = \sum_{n,r}\left[\frac{\partial H}{\partial \eta_{(n,r)}}\eta_{(n,r)} - P_{(n,r)}\frac{\partial H}{\partial P_{(n,r)}}\right] \\ &= H(t, \mathbf{P}, \boldsymbol{\eta}, \boldsymbol{\pi}) - H(t, \mathbf{P}, \boldsymbol{\eta}, \boldsymbol{\pi}) = 0.\end{aligned}$$

We conclude that $G(\mathbf{P}(t), \boldsymbol{\eta}(t))$ is constant as a function of t. Consider now $G(\mathbf{P}(t), \boldsymbol{\eta}(t))$ evaluated for $\boldsymbol{\eta}(t)$ satisfying terminal conditions (15), and $\mathbf{P}(t)$ satisfying initial conditions

$$\begin{gathered}P_{(n',r')}(t') = 1, \\ P_{(n,r)}(t') = 0, \quad n \neq n', \quad r \neq r',\end{gathered}$$

corresponding to the system being in state (n', r') at time t' with probability one. Then

$$\begin{aligned}\eta_{(n',r')}(t') &= G(P(t'), \eta(t')) = G(P(T), \eta(T)) = \sum_{n,r}[r - CI(r < z)]P_{(n,r)}(T) \\ &= E[R(T) - CI(R(T) < z) \mid Y(t') = n', R(t') = r'].\end{aligned}$$

1.4. Proof of Proposition 7

The optimality condition (33) can be rewritten as

$$r'(p)\left(1 + \frac{C\phi(f(p))z\sqrt{\Lambda(p)}}{r^2(p)}\right) = \frac{C\phi(f(p))\Lambda'(p)}{2\sqrt{\Lambda(p)}}\left(\frac{z}{\rho(p)} - 1\right).$$

Note that if $z \geqslant \rho(\tilde{p})$ then the right-hand-side of this equation is nonpositive for all p. The left-hand-side is nonpositive for $p \geqslant \tilde{p}$. Thus, the optimal solution p^* that solves the equation must satisfy $p^* \geqslant \tilde{p}$. On the other hand, if $z < \rho(\tilde{p})$ then, by continuity of $\rho(p)$ and its strict increasing (decreasing) behaviour for $p < \tilde{p}$ $(p > \tilde{p})$, there exist $\bar{p}_1$ and $\bar{p}_2$ such $z = \rho(\bar{p}_1) = \rho(\bar{p}_2)$ and $\bar{p}_1 < \tilde{p} < \bar{p}_2$. Note that the right-hand-side is positive for $\bar{p}_1 < p < \bar{p}_2$ and negative for $p > \bar{p}_2$ and $p < \bar{p}_1$. The left-hand-side is positive for $p < \tilde{p}$ and negative for $p > \tilde{p}$. Thus, we either have $p^* > \bar{p}_2$ or $\bar{p}_1 < p^* < \tilde{p}$.

2. Online Appendix: Computational Issues

In our experiments, we utilize a discrete-time approximation of equation (18) obtained as follows. Replace $d\eta_{(n,r)}/dt$ in (18) with a first-order approximation $(\eta_{(n,r)}(t)-\eta_{(n,r)}(t-\delta t))/\delta t$, and let $\pi(p\mid\cdot,n,r)$ be fixed on the interval $[t-\delta t,t)$. As we have already noted in Remark 4, restricting the measure $\pi(\cdot\mid t,n,r)$ to be concentrated on a singleton support set does not affect the optimality. We will then have

$$\frac{\eta_{(n,r)}(t)-\eta_{(n,r)}(t-\delta t)}{\delta t}=\min_{p\in\mathcal{P}}\{\lambda(t,p)[\eta_{(n,r)}(t)-\eta_{(n-1,r+p)}(t)]\},$$

or, equivalently,

$$\eta_{(n,r)}(t-\delta t)=\eta_{(n,r)}(t)+\delta t\max_{p\in\mathcal{P}}\lambda(t,p)[\eta_{(n-1,r+p)}(t)-\eta_{(n,r)}(t)]. \tag{EC5}$$

The entire planning interval $[T,0]$ can be split into K equal intervals of length $\delta t=|T|/K$. A discrete-time approximation to the optimal policy is provided by p attaining the maximum in (EC5). The evaluation starts at $t=0$, where the values for $\boldsymbol{\eta}(t)$ are given by the boundary conditions (15), and proceeds down to $t=T$ in steps of length δt. At each time step, we start with $n=1$ and increase it up to $n=Y_T$, while, for each n, the index r can run in either direction. Due to the independence of the optimal policy from r for $r\geqslant z$, the highest value of r that needs to be considered in equation (EC5) is $z-1$ (assume that z is integer). An optimal policy for all $r\geqslant z$ can be found from

$$\eta^0_n(t-\delta t)=\eta^0_n(t)+\delta t\max_{p\in\mathcal{P}}\lambda(t,p)[\eta^0_{n-1}(t)+p-\eta^0_n(t)], \tag{EC6}$$

which is obtained by substituting η^0_n+r for $\eta_{(n,r)}$ in (EC5).

Consider the equation (EC5) where we substitute the (discrete-time) optimal price $p(t-\delta t,n,r)$:

$$\eta_{(n,r)}(t-\delta t)=\eta_{(n,r)}(t)+\delta t\lambda(t,p(t-\delta t,n,r))[\eta_{(n-1,r+p(t-\delta t,n,r))}(t)-\eta_{(n,r)}(t)].$$

This equation is linear in $\boldsymbol{\eta}(\cdot)$, and one can decompose it into two linearly independent components $\boldsymbol{\eta}^1(\cdot)$ and $\boldsymbol{\eta}^2(\cdot)$ such that

$$\eta_{(n,r)}(t)=\eta^1_{(n,r)}(t)-C(1-\eta^2_{(n,r)}(t)),\quad (n,r)\in\mathcal{S}. \tag{EC7}$$

Note that, $\eta^1_{(n,r)}(t)$ represents the "return" component of the adjoint variable $\eta_{(n,r)}(t)$, while $\eta^2_{(n,r)}(t)$ represent the "risk" component. They satisfy equations of the same form

$$\eta^i_{(n,r)}(t-\delta t)=\eta^i_{(n,r)}(t)+\delta t\lambda(t,p(t-\delta t,n,r))[\eta^i_{(n-1,r+p(t-\delta t,n,r))}(t)-\eta^i_{(n,r)}(t)],\quad i=1,2, \tag{EC8}$$

but different boundary conditions

$$\eta^1_{(n,r)}(0)=r,$$
$$\eta^2_{(n,r)}(0)=I(r\geqslant z),$$

for all $(n,r)\in\mathcal{S}$. Corollary 2 establishes that the continuous-time versions of $\eta^1_{(n,r)}(t)$ and $\eta^2_{(n,r)}(t)$ are equal to $E[R(0)\mid Y(t)=n,R(t)=r]$ and $P(R(0)\geqslant z\mid Y(t)=n,R(t)=r)$, respectively. Thus, the calculations can be organized as follows. We keep a separate record of $\boldsymbol{\eta}^1(t)$ and $\boldsymbol{\eta}^2(t)$. Once we know $\boldsymbol{\eta}^1(t)$ and $\boldsymbol{\eta}^2(t)$, we reconstruct $\boldsymbol{\eta}(t)$ using (EC7). After that, we calculate the (discrete-time) optimal $p(t-\delta t,n,r)$ as the one delivering the maximum in (EC5), and, finally, we compute $\boldsymbol{\eta}^1(t-\delta t)$ and $\boldsymbol{\eta}^2(t-\delta t)$. The advantage of organizing our computation in this way is that, at the end of this procedure, we will have computed separate discrete-time approximations to $E[R(0)]$ and $P(R(0)\geqslant z)$ in the form of $\eta^1_{(Y_T,0)}(T)$ and $\eta^2_{(Y_T,0)}(T)$, respectively.

Another important advantage of this decomposition is that it permits a "simultaneous sensitivity analysis" on the values of the desired level of revenue z and the initial inventory Y. Note that, for continuous time versions of $\boldsymbol{\eta}^1(\cdot)$ and $\boldsymbol{\eta}^2(\cdot)$,

$$\eta^1_{(n,r)}(T)-C(1-\eta^2_{(n,r)}(T))=\eta_{(n,r)}(T)=J(n,T,C,z-r)+r,$$

$$\begin{aligned}\eta^1_{(n,r)}(T)&=E[R(0)\mid Y(T)=n,R(T)=r]\\&=E[R(0)\mid Y(T)=n,R(T)=0]+r,\end{aligned}$$

$$\begin{aligned}\eta^2_{(n,r)}(T) &= E[I(R(0) \geqslant z) \mid Y(T) = n, R(T) = r] \\ &= E[I(R(0) \geqslant z - r) \mid Y(T) = n, R(T) = 0].\end{aligned}$$

In other words, $\eta^1_{(n,r)}(T) - r$ and $\eta^2_{(n,r)}(T)$ represent the *optimal* $E[R(0)]$ and $P(R(0) \geqslant z')$ for the risk adjusted problem (3) with the desired level of revenue $z' = z - r$ and the initial inventory n. To maximize the benefits of this decomposition one may also compute $\eta^1_{(n,r)}(T)$ and $\eta^2_{(n,r)}(T)$ for (n, r) outside the set $\mathcal{S}$ of the original problem. For example, the states (Y_T, r) with $r > 0$ may be of interest, since they correspond to a smaller desired level of revenue $z' = z - r$ and the same initial inventory Y_T.

Finally, observe that even in the risk-neutral case when $C = 0$, knowing $\eta^2_{(n,r)}(T)$ is very useful. Since a pricing policy for the case of $C = 0$ will be selected independently of r, the values of

$$\eta^2_{(n,r)}(T) = P(R(0) \geqslant z - r \mid Y(T) = Y_T, R(T) = 0), \quad r = 0, \ldots, z$$

are computed under identical policies, and can be used to recover the probability mass function of $R(0)$. Indeed, for all r',

$$\begin{aligned}P(R(0) = r' \mid Y(T) = n, R(T) = 0) &= P(R(0) \geqslant r' \mid Y(T) = n, R(T) = 0) - P(R(0) \geqslant r' + 1 \mid Y(T) = n, R(T) = 0) \\ &= \eta_{(n, z-r')}(T) - \eta_{(n, z-r'-1)}(T).\end{aligned}$$

We thus have the distribution of $R(0)$ in the risk-neutral model of Gallego and van Ryzin (1994) for every level of inventory from 1 to Y_T.