Operations Research Letters 37 (2009) 273–279

Contents lists available at ScienceDirect

Operations Research Letters

journal homepage: www.elsevier.com/locate/orl

On the structural properties of a discrete-time single product revenue management problem

Seray Aydın [a], Yalçın Akçay [b], Fikri Karaesmen [a,*]

[a] Department of Industrial Engineering, Koç University, Istanbul, Turkey

[b] College of Business Administration and Economics, Koç University, Istanbul, Turkey

ARTICLE INFO

Article history:
Received 10 October 2008
Accepted 6 March 2009
Available online 17 March 2009

Keywords:
Revenue management
Stochastic dynamic programming
Optimal control policies

ABSTRACT

We consider a multi-period revenue management problem in which multiple classes of demand arrive over time for the common inventory. The demand classes are differentiated by their revenues and their arrival distributions. We investigate monotonicity properties of varying problem parameters on the optimal reward and the policy.

1. Introduction

We consider a finite horizon, single-product inventory control problem in which the decision-maker accepts or rejects customer requests coming from multiple demand classes. Customer requests can be for multiple units of the product (*batch* orders) and we allow *partial* fulfillment of demand for accepted requests. The optimal decisions depend on factors such as the available inventory, relative profitability of demand classes, projected volume and mix of future demand (distribution of future demand), and time-to-go till the end of the time horizon. Clearly, this is a typical *revenue management* problem, which has garnered great interest from researchers (see Talluri and van Ryzin [1] for a comprehensive survey of revenue management literature). Revenue management has also become a very powerful *managerial tool* to exploit the revenue-enhancement potential in many businesses (Cross [2], Smith et al. [3]). Revenue management empowers these businesses to effectively address the challenges of matching supply and demand. However, due to the many difficulties in successful implementation of revenue management systems, companies have not been able to *fully* realize the benefits from the cutting-edge tools that researchers have developed in the last few decades (Lahoti [4]). One of these obstacles is the *estimation* of parameters used in the underlying revenue management models, which govern the principles of how to allocate and reserve resources for high profit customers. Understanding the impact of each parameter on optimal admission policies is a key factor in successful revenue management practices, as it allows managers to perform *what-if* analysis when faced with changing parameter values (possibly because of estimation errors). This paper focuses on the structure of optimal admission policies in a well-established model of dynamic revenue management. In particular, the investigation focuses on the effects of perturbations of the problem parameters on the optimal admission policy and the optimal reward. The parameters of interest include arrival probabilities of different classes as well as their rewards. Such an investigation is crucial for designing admission policies which are *robust* to changes in the parameters.

Revenue management literature has gone a long way in establishing the structure of optimal policies and our results complements some of the existing results. Here, we only review work that is particularly related to our paper. The structure of the optimal admission policy for the basic dynamic revenue management is established in Lee and Hersh [5]. Notably, Lee and Hersh establish the optimality of a nested threshold-type admission policy. These results were later streamlined and generalized by Lautenbacher and Stidham [6]. Brumelle and Walczak [7] obtain further results in the challenging semi-Markov case. In the presence of inaccurate estimates of customer arrival distributions, Birbil et al. [8] illustrate, using simulations, the benefit of using robust optimization techniques in reducing expected revenue variability. Talluri and Van Ryzin [1] present a summary of the most important results in both the static and dynamic versions of the problem.

* Corresponding address: Department of Industrial Engineering, Koc University, Rumeli Feneri Yolu, 34450 Sariyer, Istanbul, Turkey. Tel.: +90 212 338 1718; fax: +90 212 338 1548.

E-mail address: fkaraesmen@ku.edu.tr (F. Karaesmen).

0167-6377/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.orl.2009.03.001

Among the rich literature investigating dynamic policies in revenue management, the work of Lautenbacher and Stidham [6] is of particular importance for us for two reasons. First, we borrow their model which is versatile and subsumes some of the well-established dynamic and static models in the literature. Second, Lautenbacher and Stidham emphasize the significance of focusing on the monotonicity of certain operators appearing in the value function of the dynamic program. This is also the approach we take for investigating the monotonicity properties related to changes in the problem parameters. Koole [9] presents an excellent overview of monotonicity results in Markov Decision Process with applications in queueing. In particular, for varying arrival probabilities, we adapt some of the recent ideas from Çil, Örmeci and Karaesmen [10] to the model of [6]. The main focus of [10] is on continuous-time infinite-horizon queueing-inventory models with stationary parameters. The discrete-time model with non-stationary parameters considered here poses additional challenges but it turns out that corresponding results can be obtained. We also obtain additional results pertaining to parameter effects that are particular to the discrete-time revenue management setting. To our knowledge, the only other paper that investigates related monotonicity issues in a revenue management context is Cooper and Gupta [11]. That paper investigates the effect of demand distributions on the expected optimal reward in a single-period setting.

Designing robust policies for revenue management when problem parameters are uncertain has recently attracted attention. For instance, Lan, Gao, Ball and Karaesmen [12] consider the case with limited demand information employing ideas from competitive analysis of on-line algorithms. We do not explicitly consider the challenging robust policy design problem in this paper but our results provide basic guidelines as to what sort of changes in the optimal policy are anticipated as problem parameters are varied within a given uncertainty set.

2. Model

The model that we employ was first introduced, to our knowledge, by Lautenbacher and Stidham [6]. Suppose that time is divided into decision periods such that at most one request is received in any given period but the customer can demand more than one unit of the product. Let K be the number of decision periods. Time is indexed by k in our model, where $k = K$ is the first period and $k = 1$ is the last period after which all inventories perish. There are n demand classes, with Class i offering to pay R_i, $i = 1, 2, \dots, n$, for a unit of the product. Assume that $R_1 \geq R_2 \geq R_3 \geq \cdots \geq R_n$, without loss of generality. Let p_{ibk} be the probability that a customer belonging to demand Class i (referred to as a class i customer) requests b units of inventory in period k, and p_{0k} be the probability that no customers arrive in period k. We assume B_i is an upper bound on the batch demand size for Class i customers. Note that $p_{0k} + \sum_{i=1}^{n}\sum_{b=1}^{B_i} p_{ibk} = 1$ for all $k = 1, 2, \dots, K$. As in [6], we assume that in each period a request can be partially fulfilled. This model is fairly versatile and is referred to as the "omnibus model" formulation in [6]. Taking the batch demand size to be at most one unit, we obtain the standard single-arrival dynamic model. Using uniformization and discretization, stationary or non-stationary Poisson arrivals with batch requests can be captured. In addition, the frequently employed static model which makes the assumption that lower class fares always arrive earlier than higher class fares can also be captured by choosing the arrival probabilities such that the lower class demands arrive at earlier periods than the higher class demands.

The decision-maker's problem of maximizing expected revenues over the entire finite time horizon can be modeled using a dynamic programming formulation. Let $v_k(x)$ be the *expected maximum revenue-to-go* in period k when there are x units of inventory are available. We can express $v_k(x)$ as

$$v_k(x) = \sum_{i=1}^{n}\sum_{b=1}^{B_i} p_{ibk}\left(\max_{\kappa_i\in\{0,1,\dots,\min(b,x)\}} \kappa_i R_i + v_{k-1}(x-\kappa_i)\right) + p_{0k} v_{k-1}(x) \tag{1}$$

with boundary conditions $v_k(0) = 0$ for all k and $v_0(x) = 0$ for all x. In the above formulation, κ_i is the inventory assigned to the Class i customer, requesting b units of the product. Note that κ_i is an integer between 0 and $\min(x, b)$. We can rewrite the value function in (1) as a combination of the fictitious and rationing *event operators* defined in [10]. The *batch rationing operator* $T_{b\_RT_i}$ determines the number of inventory units assigned to Class i customers and the *fictitious operator* T_{FIC} represents the fictitious event corresponding to no demand arrivals in the period. These two operators when applied on a function $f(x)$ yield $T_{b\_RT_i} f(x) = \max_{\kappa_i \leq \min\{x,b\}}\{\kappa_i R_i + f(x-\kappa_i)\}$ and $T_{FIC} f(x) = f(x)$, respectively. Hence, we have

$$v_k(x) = \sum_{i=1}^{n}\sum_{b=1}^{B_i} p_{ibk} T_{b\_RT_i} v_{k-1}(x) + p_{0k} T_{FIC} v_{k-1}(x). \tag{2}$$

3. Structural properties

In this section, we first describe a number of basic structural properties of the model in Section 2 that are well known in revenue management literature. We then present our results on the impact of varying two particular problem parameters – the arrival probabilities and the rewards. In this section, we also use numerical examples to illustrate interesting policy implications of our analytical results.

3.1. Preliminaries

To set up the stage, we first describe some properties that are well known [6,1]. We later use these preliminary results to establish our analysis in Sections 3.2 and 3.3.

Proposition 1. *$T_{b\_RT_i}$ and T_{FIC} event operators have the following properties:*
1. *If $f(x)$ is non-decreasing in x, then the $T_{b\_RT_i} f(x)$ and $T_{FIC} f(x)$ are also non-decreasing in x.*
2. *If $f(x)$ is concave in x, then the $T_{b\_RT_i} f(x)$ and $T_{FIC} f(x)$ are also concave in x.*

As in Lautenbacher and Stidham [6], the properties of the operators $T_{b\_RT_i}$ and T_{FIC} can be combined to yield the properties of the value function as summarized in the next proposition:

Proposition 2. *The maximum expected revenue-to-go function, $v_k(x)$ is*
1. *A non-decreasing function of the inventory level, x,*
2. *A non-decreasing function of the time remaining till the end of the finite time horizon, k,*
3. *A concave function of the inventory level, x.*

Concavity of the value function $v_k(x)$ means that marginal value of an inventory is non-increasing with the current inventory level, x. This has an important implication on the structure of the optimal policy. Let ℓ^*_{ik} be defined as follows: $\ell^*_{ik} = \max\{x : v_k(x)-v_k(x-1) > R_i\}$. More explicitly, ℓ^*_{ik} is the maximum possible inventory on hand such that if the current inventory on hand, x, is less than or equal to ℓ^*_{ik}, it is optimal to reject the whole Class i batch. Similarly, if the current inventory level, x, is greater than or equal to $\ell^*_{ik}+1$, it is optimal to satisfy Class i demand until either the inventory level drops down to ℓ^*_{ik} or the whole batch is satisfied. Here, ℓ^*_{ik} is the optimal threshold value for Class i demand such that the optimal policy will reject the whole Class i batch if $x < \ell^*_{ik}$, partially satisfy (i.e. satisfy up to inventory level ℓ^*_{ik}) the demand if $\ell^*_{ik} < x < \ell^*_{ik} + b$, and satisfy the entire batch if $x \geq \ell^*_{ik} + b$. Therefore, a threshold policy is the optimal policy in our model. It is obvious that if the reward of a Class i customer is higher than the reward of a Class-j customer, then the optimal threshold value of class-i will be lower than that of class-j. Let us summarize these well-known results (see [6,1]). First, it can be shown that a Class 1 demand is always accepted: $\ell^*_{1k} = 0$, for all $k = 1, 2, \ldots, K$. Second, the thresholds have a nested structure: $\ell^*_{1k} \leq \ell^*_{2k} \leq \cdots \leq \ell^*_{nk}$ for all $k = 1, 2, \ldots, K$.

3.2. Effects of varying arrival probabilities

Let us first explain how we vary arrival probabilities in our analysis. An increase by ε in any arrival probability p_{ibk}, leads to the new arrival probability $p_{ibk}+\varepsilon$ and causes a corresponding decrease in the fictitious event probability; that is the new fictitious event probability becomes $p_{0k} - \varepsilon$. We assume that ε is small enough that both $p_{ibk} + \varepsilon$ and $p_{0k} - \varepsilon$ stay in the interval [0, 1].

We first consider the effect of an increase in p_{ibt} on $v_k(x)$. Note that increasing p_{ibt} and thereby decreasing p_{0t} has the effect of increasing the probability of a controlled event (admission) and decreasing the probability of an uncontrolled event (fictitious).

Proposition 3. *$v_k(x)$ is a non-decreasing function of p_{ibt} $\forall k = 1, 2, \ldots, K$, where $1 \leq t \leq K$.*

We omit the proof of Proposition 3 which is straightforward using a sample-path argument or induction. The proposition establishes that the optimal expected reward is non-decreasing in the arrival probabilities p_{ibt}. Next, we consider the effects of varying the arrival probabilities on the optimal policy determined by the thresholds ℓ^*_{ik}. Noting that the thresholds ℓ^*_{ik} are determined by the difference $v_k(x) - v_k(x-1)$, we focus on the monotonicity properties of these differences with respect to arrival probabilities. To clarify the comparison, let us make the dependence on the perturbed parameter explicit by introducing $v'_k(x, p_{jbt})$ which is defined as the value of $v_k(x)$ for a parameter value of p_{jbt}.

Proposition 4. *$v_k(x)$ is a supermodular function of p_{jbt} and x, i.e. $v'_k(x, p_{jbt}+\varepsilon)-v'_k(x-1, p_{jbt}+\varepsilon) \geq v'_k(x, p_{jbt})-v'_k(x-1, p_{jbt})$ $\forall k = 1, 2, .K$, where $1 \leq t \leq K$ and $0 \leq \varepsilon \leq p_{0t}$.*

Proof. Consider two systems, system 1 and system 2. All model parameters of these two systems, as well as their demand distributions are identical except for some period t, where $1 \leq t \leq K$. In the tth period, the arrival probability of a particular Class j customer with a batch demand of size $\tilde{b}$ units is given by $p_{j\tilde{b}t}$ in system 1, whereas the likelihood of the same event in system 2 is given by $p_{j\tilde{b}t}+\varepsilon$. Let $v_k(x)$ be the optimal value function of system 1 in period k and $v^\varepsilon_k(x)$ be the optimal value function of system 2. From the definition of supermodularity, we need to show $v^\varepsilon_k(x) - v^\varepsilon_k(x-1) \geq v_k(x) - v_k(x-1)$. Let us define the marginal value function $\Delta f = f(x) - f(x-1)$. Hence, the above expression can be written as $\Delta v^\varepsilon_k(x) \geq \Delta v_k(x)$. For $k = 0, 1, \ldots, t-1$, supermodularity holds trivially since $v_k(x) = v^\varepsilon_k(x)$. Hence, we next verify $\Delta v^\varepsilon_t(x) \geq \Delta v_t(x)$, i.e., $k = t$, which can be written as $\Delta T_{\tilde{b}\_RT_j} v^\varepsilon_{t-1}(x) - \Delta T_{FIC} v^\varepsilon_{t-1}(x) \geq 0$, or equivalently as:

$$\max_{\kappa_j \leq \min\{x,\tilde{b}\}} \{\kappa_j R_j + v^\varepsilon_{t-1}(x-\kappa_j)\} - \max_{\kappa_j \leq \min\{x-1,\tilde{b}\}} \{\kappa_j R_j + v^\varepsilon_{t-1}(x-1-\kappa_j)\} \geq v^\varepsilon_{t-1}(x) - v^\varepsilon_{t-1}(x-1). \tag{3}$$

Let κ^*_{jx} be the optimal number of units of Class j demand filled out of a batch of size $\tilde{b}$ in period t when the inventory level is x. Due to the concavity of $v^\varepsilon_{t-1}(x)$ in x, there are only three possibilities for the pair $(\kappa^*_{j(x-1)}, \kappa^*_{jx})$: (0,0), $(\tilde{b}, \tilde{b})$, or $(\kappa^*_{jx} - 1, \kappa^*_{jx})$. For the cases (0,0) and $(\tilde{b}, \tilde{b})$, the desired equality holds by concavity. For the case $(\kappa^*_{j(x-1)}, \kappa^*_{jx})$, the inequality becomes $R_j \geq v^\varepsilon_{t-1}(x) - v^\varepsilon_{t-1}(x-1)$ which is true if any admission is to take place. This proves statement (3).

Now, let $k = t+1$. The statement in the proposition can be stated as

$$\sum_{\substack{i=1\\ i\neq j}}^{n} \sum_{b=1}^{B_i} p_{ib(t+1)} \Delta T_{b\_RT_i} v^\varepsilon_t(x) + \sum_{\substack{b=1\\ b\neq\tilde{b}}}^{B_j} p_{jb(t+1)} \Delta T_{b\_RT_j} v^\varepsilon_t(x) + p_{j\tilde{b}(t+1)} \Delta T_{\tilde{b}\_RT_j} v^\varepsilon_t(x) + p_{0(t+1)} \Delta T_{FIC} v^\varepsilon_t(x)$$
$$\geq \sum_{\substack{i=1\\ i\neq j}}^{n} \sum_{b=1}^{B_i} p_{ib(t+1)} \Delta T_{b\_RT_i} v_t(x) + \sum_{\substack{b=1\\ b\neq\tilde{b}}}^{B_j} p_{jb(t+1)} \Delta T_{b\_RT_j} v_t(x) + p_{j\tilde{b}(t+1)} \Delta T_{\tilde{b}\_RT_j} v_t(x) + p_{0(t+1)} \Delta T_{FIC} v_t(x).$$

Since supermodularity holds in period t, we know that $\Delta v^\varepsilon_t(x) \geq \Delta v_t(x)$. Hence, the above inequality would hold if $\Delta T_{b\_RT_i} v^\varepsilon_t(x) \geq \Delta T_{b\_RT_i} v_t(x)$ for all $i = 1, \ldots, n$. Based on the definition of the batch rationing operator, this expression is equivalent to

$$\max_{\kappa_i \leq \min\{x-1,b\}} \{\kappa_i R_i + v_t(x-1-\kappa_i)\} + \max_{\kappa_i \leq \min\{x,b\}} \{\kappa_i R_i + v^\varepsilon_t(x-\kappa_i)\}$$
$$\geq \max_{\kappa_i \leq \min\{x-1,b\}} \{\kappa_i R_i + v^\varepsilon_t(x-1-\kappa_i)\} + \max_{\kappa_i \leq \min\{x,b\}} \{\kappa_i R_i + v_t(x-\kappa_i x)\}.$$

Table 1
Optimal threshold levels for Example 1.

t	1	2	3	4	5	6	7	8	9	10
ℓ^*_{2t}(min)	0	1	1	2	2	2	3	3	3	4
ℓ^*_{2t}	*1*	*1*	*2*	*2*	*3*	*3*	*4*	*4*	*5*	*5*
ℓ^*_{2t}(max)	1	2	2	3	4	4	5	6	6	7

Let κ_{ix} be the optimal number of units of inventory allocated to Class i demand in system 1, and κ^ε_{ix} be the optimal number of units of inventory allocated to Class-i demand in system 2, with x units of available inventory in period $t+1$ in both systems. Consequently,

$$\kappa_{i(x-1)}R_i + v_t(x-1-\kappa_{i(x-1)}) + \kappa^\varepsilon_{ix}R_i + v^\varepsilon_t(x-\kappa^\varepsilon_{ix}) \geq \kappa^\varepsilon_{i(x-1)}R_i + v^\varepsilon_t(x-1-\kappa^\varepsilon_{i(x-1)}) + \kappa_{ix}R_i + v_t(x-\kappa_{ix}). \tag{4}$$

Next, we prove the validity of the above inequality by considering all possible values for κ_{ix} and κ^ε_{ix}. First note that κ_{ix} and $\kappa_{i(x-1)}$ can differ at most by 1 unit due to the concavity of the value function $v_t(x)$. Further, if $\kappa_{ix} = \kappa_{i(x-1)}$, then it should be true that either $\kappa_{ix} = \kappa_{i(x-1)} = 0$ or $\kappa_{ix} = \kappa_{i(x-1)} = b$ (same property holds for κ^ε_{ix}). Also, due to the optimality of κ_{ix} and κ^ε_{ix}, and our hypothesis in period t, we have

$$\begin{aligned}
R_i \geq v^\varepsilon_t(x) - v^\varepsilon_t(x-1) &\geq v_t(x) - v_t(x-1)\\
R_i \geq v^\varepsilon_t(x-1) - v^\varepsilon_t(x-2) &\geq v_t(x-1) - v_t(x-2)\\
&\cdots\\
R_i \geq v^\varepsilon_t(x-\kappa^\varepsilon_{ix}+1) - v^\varepsilon_t(x-\kappa^\varepsilon_{ix}) &\geq v_t(x-\kappa^\varepsilon_{ix}+1) - v_t(x-\kappa^\varepsilon_{ix}).
\end{aligned}$$

Hence, in the first system, the optimal number of units of inventory allocated to Class i demand in period $t+1$ with x units of available inventory, κ_{ix}, is at least κ^ε_{ix}, i.e. $\kappa^\varepsilon_{ix} \leq \kappa_{ix}$. Now for any two integers w_1 and w_2, such that $0 \leq w_1 \leq b-1$ and $w_1 \leq w_2 \leq b-1$, consider the following cases

Case	$(\kappa_{ix}, \kappa_{i(x-1)}, \kappa^\varepsilon_{ix}, \kappa^\varepsilon_{i(x-1)})$	Supermodularity Inequality
1	$(0, 0, 0, 0)$	$v_t(x-1) + v^\varepsilon_t(x) \geq v_t(x) + v^\varepsilon_t(x-1)$
2	$(w_1+1, w_1, 0, 0)$	$v^\varepsilon_t(x) - v^\varepsilon_t(x-1) \geq R_i$
3	$(b, b, 0, 0)$	$v^\varepsilon_t(x) - v^\varepsilon_t(x-1) \geq v_t(x-b) - v_t(x-1-b)$
4	(w_1+1, w_1, w_2+1, w_2)	$R_i \geq R_i$
5	(b, b, w_2+1, w_2)	$R_i \geq v_t(x-b) - v_t(x-1-b)$
6	(b, b, b, b)	$v^\varepsilon_t(x-b)+v_t(x-1-b) \geq v_k(x-b)+v^\varepsilon_k(x-1-b)$

Cases 1 and 6 are true due to the supermodularity of $v_k(x)$ in period k and x. Case 2 is satisfied since no Class i demand is filled in the second system. In Case 3, the left hand side of the inequality is greater than or equal to than R_i, whereas the left hand side is less than or equal to R_i, hence is true. Case 4 trivially holds. In Case 5, the inequality is true since all type-i demand is filled in the first system. As a result, $v_{t+1}(x)$ is supermodular with respect in x and $p_{ib(t+1)}$. Clearly, the supermodularity property is also valid for any $k > t+1$. ■

Let us discuss the implications of Proposition 4. Since the admission thresholds ℓ^*_{ik} are determined by the difference $v_k(x) - v_k(x-1)$ and further this difference is shown to be non-decreasing in the probability of arrival p_{ibt}, we conclude that increasing p_{ibt} causes the admission thresholds ℓ^*_{ik} to be non-decreasing for all i and for all $k > t$.

Now consider a ten-period ($K = 10$) problem scenario with *two* classes of arrivals with unit demands ($B_1 = B_2 = 1$), which are stationary over time and bring respective rewards of $R_1 = 3$ and $R_2 = 1$. The arrival probability of *more valuable* customer class is equal to 0.2, whereas the arrival probability of *less valuable* customer class is equal to 0.6, i.e., $p_1 = 0.2$, $p_2 = 0.6$. The initial inventory level is 10. We refer to this particular scenario as the BASE CASE, and consider its variations of to illustrate implications of Proposition 4.

Example 1. Suppose in the BASE CASE, we let the arrival probability of Class 1 customers assume values between 0.1 and 0.3. Proposition 4 establishes that the optimal threshold ℓ^*_{2t} in each period is non-decreasing in p_1. This implies that when $p_1 \in [0.1, 0.3]$, ℓ^*_{2t} takes its lowest value when $p_1 = 0.1$ and its highest value when $p_1 = 0.3$. Table 1 reports these upper and lower bounds for the optimal threshold, denoted by ℓ^*_{2t}(min) and ℓ^*_{2t}(max) respectively, as well as the optimal threshold for the BASE CASE, ℓ^*_{2t}. We conclude, for instance, that despite the uncertainty in p_1, the optimal threshold for period 10 is between 4 and 7.

In Example 1, the uncertain parameter p_1 pertains to the *higher* reward customer class. The intuition that is exhibited in Table 1 is that increasing the arrival probability (or availability of demand) for this class should lead to increased protection of the inventory from other classes, thereby resulting in higher thresholds. Nevertheless, Proposition 4 establishes a much stronger result, stating that increasing the arrival probability of *any* demand class leads to higher levels of protection for *all* classes. To observe this interesting result numerically, we vary p_2, the arrival probability of the less valuable class, in the next example.

Example 2. Suppose in the BASE CASE we let the arrival probability of Class 2 customers range between 0.5 and 0.7. Similar to our notation in Example 1, ℓ^*_{2t}(min) corresponds to the case with $p_2 = 0.5$ and ℓ^*_{2t}(max) corresponds to the case with $p_2 = 0.7$. Table 2 reports these results, along with the optimal thresholds for the BASE CASE, ℓ^*_{2t}. The increase in p_2 leads to a similar effect on the threshold levels as the increase in p_1. However, the intuition is somewhat different. The marginal value of an additional inventory is increasing when demand availability is higher. This may make the marginal inventory too valuable for a Class 2 demand. Therefore, additional protection for the marginal inventory is needed which is achieved by increasing the threshold.

Next, we investigate the second order properties of the value function $v_k(x)$ as a function of the arrival probability p_{ibk}.

Proposition 5. *$v_k(x)$ is neither concave nor convex in p_{ibk}.*

Table 2
Optimal threshold levels for Example 2.

t	1	2	3	4	5	6	7	8	9	10
ℓ_{2t}^* (min)	1	1	2	2	2	3	3	4	4	5
ℓ_{2t}^*	*1*	*1*	*2*	*2*	*3*	*3*	*4*	*4*	*5*	*5*
ℓ_{2t}^* (max)	1	1	2	3	3	4	4	5	6	6

Proposition 5, which we state here without a proof (but can easily be verified with a counterexample), establishes that, in general, the optimal reward does not have nice second-order properties despite being non-decreasing in p_{ibt}. The next proposition establishes the supermodularity of the value function in k and x.

Proposition 6. *$v_k(x)$ is a supermodular function of x and k.*

Proof. We would like to show that $v_k(x) - v_k(x-1) \geq v_{k-1}(x) - v_{k-1}(x-1)$. Let $\Delta v_k^m(x)$ be defined as the *marginal value per unit* of inventory when m units out of x units of available inventory is allocated to demand in period k, i.e.,

$$\Delta v_k^m(x) = \frac{v_k(x) - v_k(x-m)}{m}. \tag{5}$$

Based on this definition, we can express $\Delta v_k(x) = v_k(x) - v_k(x-1)$ as follows:

$$\begin{aligned}\Delta v_k(x) = &\sum_{i=1}^{n}\sum_{b=1}^{B_i} p_{ibk}\left(\max_{\kappa_i \leq \min(b,x)} \kappa_i R_i + v_{k-1}(x-\kappa_i)\right) + p_{0k} v_{k-1}(x)\\ &- \sum_{i=1}^{n}\sum_{b=1}^{B_i} p_{ibk}\left(\max_{\kappa_i \leq \min(b,x-1)} \kappa_i R_i + v_{k-1}(x-\kappa_i-1)\right) + p_{0k} v_{k-1}(x-1).\end{aligned}$$

Substituting $p_{0k} = 1 - \sum_{i=1}^{n}\sum_{b=1}^{B_i} p_{ibk}$ into the above equation, we can simplify $\Delta v_k(x)$ as

$$\Delta v_k(x) = \Delta v_{k-1}(x) + \sum_{i=1}^{n}\sum_{b=1}^{B_i} p_{ibk}\left\{\max_{\kappa_i \leq \min(b,x)} \kappa_i R_i - \kappa_i \Delta v_{k-1}^{\kappa_i}(x) - \max_{\kappa_i \leq \min(b,x-1)} \kappa_i R_i - \kappa_i \Delta v_{k-1}^{\kappa_i}(x-1)\right\}. \tag{6}$$

Note that, in order to show the supermodularity of $v_k(x)$ it suffices to show in Eq. (6) that

$$\sum_{i=1}^{n}\sum_{b=1}^{B_i} p_{ibk}\left\{\max_{\kappa_i \leq \min(b,x)} \kappa_i R_i - \kappa_i \Delta v_{k-1}^{\kappa_i}(x) - \max_{\kappa_i \leq \min(b,x-1)} \kappa_i R_i - \kappa_i \Delta v_{k-1}^{\kappa_i}(x-1)\right\} \geq 0. \tag{7}$$

Due to concavity of $v_k(x)$ as a function of x, we know that

$$v_{k-1}(x) - v_{k-1}(x-\kappa_i) \leq v_{k-1}(x-1) - v_{k-1}(x-\kappa_i-1) \Rightarrow \Delta v_{k-1}^{\kappa_i}(x) \leq \Delta v_{k-1}^{\kappa_i}(x-1). \tag{8}$$

Let κ_i^* be the optimal number of units of inventory assigned to Class i in period k with x units of inventory available. Similarly, let $\tilde{\kappa}_i^*$ be the optimal number of units of inventory assigned to Class i in period k with $x-1$ units of inventory available. Then, we can rewrite (7) as follows

$$\sum_{i=1}^{n}\sum_{b=1}^{B_i} p_{ibk}\left\{[\kappa_i^* R_i - \kappa_i^* \Delta v_{k-1}^{\kappa_i^*}(x)] - [\tilde{\kappa}_i^* R_i - \tilde{\kappa}_i^* \Delta v_{k-1}^{\tilde{\kappa}_i^*}(x-1)]\right\} \geq 0. \tag{9}$$

Since we already know that $v_k(x)$ is concave in x, it should be true that either $\kappa_i^* = \tilde{\kappa}_i^*$ or $\kappa_i^* = \tilde{\kappa}_i^* + 1$ for all $i = 1, \ldots, n$. If $\kappa_i^* = \tilde{\kappa}_i^*$, then

$$[\kappa_i^* R_i - \kappa_i^* \Delta v_{k-1}^{\kappa_i^*}(x)] - [\tilde{\kappa}_i^* R_i - \tilde{\kappa}_i^* \Delta v_{k-1}^{\tilde{\kappa}_i^*}(x-1)] = \kappa_i^*(\Delta v_{k-1}^{\kappa_i^*}(x-1) - \Delta v_{k-1}^{\kappa_i^*}(x))$$

which is nonnegative due to (8). On the other hand, if $\kappa_i^* = \tilde{\kappa}_i^* + 1$, we have

$$[\kappa_i^* R_i - \kappa_i^* \Delta v_{k-1}^{\kappa_i^*}(x)] - [\tilde{\kappa}_i^* R_i - \tilde{\kappa}_i^* \Delta v_{k-1}^{\tilde{\kappa}_i^*}(x-1)] = R_i + (\kappa_i^* - 1)\Delta v_{k-1}^{\kappa_i^*-1}(x-1) - \kappa_i^* \Delta v_{k-1}^{\kappa_i^*}(x).$$

We can simplify $(\kappa_i^* - 1)\Delta v_{k-1}^{\kappa_i^*-1}(x-1) - \kappa_i^* \Delta v_{k-1}^{\kappa_i^*}(x)$, using the definition in (5) and some simple manipulations as $(\kappa_i^* - 1)\Delta v_{k-1}^{\kappa_i^*-1}(x-1) - \kappa_i^* \Delta v_{k-1}^{\kappa_i^*}(x) = v_{k-1}(x-1) - v_{k-1}(x)$. Therefore, $R_i + (\kappa_i^* - 1)\Delta v_{k-1}^{\kappa_i^*-1}(x-1) - \kappa_i^* \Delta v_{k-1}^{\kappa_i^*}(x) = R_i + v_{k-1}(x-1) - v_{k-1}(x)$. Clearly, $R_i + v_{k-1}(x-1) - v_{k-1}(x)$ is nonnegative since at least one unit of inventory is assigned to Class i demand, i.e., $\kappa_i^* > 0$. As a result, the inequality in (9) is satisfied, which implies supermodularity of $v_k(x)$ in k and x. ■

By Proposition 6, the marginal value of an additional inventory is non-decreasing in the time remaining. In terms of the thresholds, the immediate conclusion is that ℓ_{ik}^* are non-decreasing in k for all i. Our numerical results for Examples 1 and 2, shown in Tables 1 and 2, clearly illustrate this fact.

Table 3
Optimal threshold levels for Example 3.

t	1	2	3	4	5	6	7	8	9	10
ℓ^*_{2t} (min)	0	0	1	1	1	2	2	2	3	3
ℓ^*_{2t}	*1*	*1*	*2*	*2*	*3*	*3*	*4*	*4*	*5*	*5*
ℓ^*_{2t} (min)	1	2	3	4	5	6	7	8	9	10

3.3. Effects of varying rewards

In this section, we focus on perturbations of the reward parameters R_i. An increase by ε to the reward of Class i has a similar interpretation to that of the variations we made for arrival probabilities in Section 3.2. The new reward is given by $R_i + \varepsilon$ with the assumption that $R_i + \varepsilon < R_{i-1}$ (without loss of generality). It is obvious that the optimal reward is non-decreasing in R_i which can be shown by a simple sample path argument. The next proposition establishes a related second order property.

Proposition 7. *$v_k(x)$ is a convex function of R_i.*

Proof. We only provide a sketch of the proof here. Using the standard linear programming (LP) formulation of the optimality equation (see [13] for details), it can be seen that R_i appears as a coefficient on the right hand side of an LP. Since the optimal value of a linear program is a convex function of its right-hand-side coefficients [14], we conclude that $v_k(x)$ is convex in R_i, for all $i = 1, \ldots, n$. ■

Proposition 7 establishes that the expected optimal reward is increasing and convex in R_i. Next, we focus on the effects of increasing R_i on the optimal thresholds.

Proposition 8. *$v_k(x)$ is*

1. *supermodular with respect to R_1 and x,*
2. *submodular with respect to R_n and x (as long as $R_n < R_{n-1}$).*

Proof. We present the proof of part (1). The proof of part (2) is similar. As before consider two systems, system 1 and system 2. All model parameters of these two systems are identical except the reward of a particular Class 1 customer. In system 1, the Class 1 reward is R_1, whereas the reward of the same class of customer in system 2 is given by $R_1 + \varepsilon$. Let $v_k(x)$ be the optimal value function of system 1 in period k and $v^\varepsilon_k(x)$ be the optimal value function of system 2. For $k = 0$, supermodularity holds trivially since $v_0(x) = v^\varepsilon_0(x) = 0\ \forall x$. Assume that for $k = t - 1$, $\Delta v^\varepsilon_{t-1}(x) \geq \Delta v_{t-1}(x)$ is true. Hence, we next need to verify for $k = t$.

The operators $T_{b_R T_i}$ $(i = 2 \ldots n)$ and T_{FIC} are not directly affected by an increase in R_1 and were already shown to preserve supermodularity in the proof of property 4. We only need to verify that supermodularity holds for $T_{b_R T_1}$ corresponding to Class 1. For this class, it was already shown that it is optimal to accept the entire batch, if sufficient inventory is available and as much of the batch as possible otherwise. In short,

$$\Delta v^\varepsilon_t(x) = v^\varepsilon_{t-1}(\max(x - b, 0)) + \max(x, b)(R_1 + \varepsilon) - v^\varepsilon_{t-1}(\max(x - b - 1, 0)) + \max(x - 1, b)(R_1 + \varepsilon)$$

and

$$\Delta v_t(x) = v_{t-1}(\max(x - b, 0)) + \max(x, b)R_1 - v_{t-1}(\max(x - b - 1, 0)) + \max(x - 1, b)R_1.$$

Using the induction assumption, and since $\varepsilon > 0$, we directly obtain: $\Delta v^\varepsilon_t(x) - \Delta v_t(x) \geq 0$. ■

Proposition 8 establishes that the optimal thresholds ℓ^*_{ik} are non-decreasing in R_1 and non-increasing in R_n for all i and for all k. Our next example illustrates how structural properties we proved in our paper can be employed in a setting that exhibits uncertainty in *multiple* problem parameters.

Example 3. Suppose that in our BASE CASE, the arrival probabilities for the two demand classes in each period, p_1 and p_2, are uncertain but lie in the following intervals: $p_1 \in [0.1, 0.3]$ and $p_2 \in [0.5, 0.7]$. $R_2 = 1$ but R_1 is anticipated to vary between 2 and 4, i.e., $R_1 \in [2, 4]$. By Propositions 4 and 8, the optimal thresholds for Class 2, ℓ^*_{2t} are monotone in p_1, p_2 and R_1. This implies that, as long as these parameters lie in the above uncertainty sets, the optimal thresholds will lie in intervals corresponding to the extreme values of the parameter uncertainty sets. Therefore, in order to find the minimum values of the thresholds, denoted by ℓ^*_{2t} (min), it is sufficient to solve a single dynamic program with parameters $p_1 = 0.1, p_2 = 0.5$, and $R_1 = 2$. Similarly, to find the maximum values of the thresholds, denoted by ℓ^*_{2t} (max), it suffices to solve the problem with $p_1 = 0.3, p_2 = 0.7$, and $R_1 = 4$. Table 3 reports these two threshold levels, and the thresholds for the BASE CASE.

4. Extensions and discussion

Certain extensions to the model are straightforward. For instance, all of the properties go through if non-zero salvage values are assumed at the end of the horizon as long as the salvage value function satisfies the required properties for induction (i.e. concavity in the inventory level). Holding costs can also be handled in a straightforward manner for most properties as long as the holding cost function is increasing and convex. On the other hand, since the holding cost function applies over time, time related properties such as part 2 of Propositions 2 and 6 fail to hold.

The partial admission (or batch splitting) assumption which is frequently made in the revenue management literature is critical to most of the results in the paper. It appears that complete admission (no batch splitting) extensions are difficult even for some of the basic properties. First, the complete admission operator does not preserve concavity which is crucial for most of the result in this paper. In addition, Cil, Ormeci and Karaesmen [15] show that even a weaker result, the optimality of threshold policies, is only guaranteed under very restrictive assumptions (i.e. constant and identical batch sizes for all classes). Only a few properties that do not rely on concavity such as Propositions 3 and 7 continue to hold under the complete batch admission assumption.

Finally, a recent paper by Armony, Plambeck and Seshadri [16] shows that anticipated monotonicity results may not hold when customers renege from a queueing system. This implies that in our context monotonicity results under order cancelations may not be true.

Acknowledgements

This research was partially supported by TUBITAK and the TUBA-GEBIP programme. F. Karaesmen is grateful to the Dept. of Ind. Eng. and Man. Sci. of Northwestern University where part of this research was done.

References

[1] K.T. Talluri, G.J. Van Ryzin, The Theory and Practice of Revenue Management, Springer, 2004.
[2] R.G. Cross, Revenue Management: Hard-Core Tactics for Market Domination, Broadway, 1996.
[3] B.C. Smith, J.F. Leimkuhler, R.M. Darrow, Yield management at American airlines, Interfaces 22 (1) (1992) 8–31.
[4] A. Lahoti, Why CEOs should care about revenue management: How to minimize the implementation pains and maximize the benefits, OR. OR/MS Today, February 2002.
[5] T.C. Lee, M. Hersh, A model for dynamic airline seat inventory control with multiple seat bookings, Transportation Science 27 (1993) 252–265.
[6] C.J. Lautenbacher, S. Stidham, The underlying Markov decision process in the single-leg airline yield-management problem, Transportation Science 33 (1999) 136–146.
[7] S. Brumelle, D. Walczak, Dynamic airline revenue management with multiple semi-Markov demand, Operations Research 51 (1) (2003) 137–148.
[8] S.I. Birbil, J.B.G. Frenk, J.A.S. Gromicho, S. Zhang, The role of robust optimization in single-leg airline revenue management, Management Science 55 (1) (2009) 148–163.
[9] G. Koole, Monotonicity in Markov Reward and Decision Chains: Theory and Applications, Now Publishers Inc, 2007.
[10] E.B. Çil, E.L. Örmeci, F. Karaesmen, Effects of system parameters on the optimal policy structure in a class of queueing control problems, Queueing Systems, 2009, (in press).
[11] D. Gupta, W.L. Cooper, Stochastic comparisons in production yield management, Operations Research 53 (2) (2005) 377–384.
[12] Y. Lan, H. Gao, M. Ball, I. Karaesmen, Revenue management with limited demand information, Management Science 54 (9) (2008) 1594–1609.
[13] S.M. Ross, Introduction to Stochastic Dynamic Programming, Academic Press, 1995.
[14] M.S. Bazaraa, J.J. Jarvis, H.D. Sherali, Linear Programming and Network Flows, John Wiley & Sons, 1990.
[15] E.B. Çil, E.L. Örmeci, F. Karaesmen, Structural results on a batch acceptance problem for capacitated queues, Mathematical Methods of Operations Research 66 (2) (2007) 263–274.
[16] M. Armony, E. Plambeck, S. Seshadri, Sensitivity of optimal capacity to customer impatience in an unobservable M/M/S Queue (Why You Shouldn't Shout at the DMV), Manufacturing & Service Operations Management 11 (2009) 19–32.