ASYMPTOTIC BEHAVIOR OF AN ALLOCATION POLICY FOR REVENUE MANAGEMENT

WILLIAM L. COOPER

Department of Mechanical Engineering, University of Minnesota, 111 Church Street S.E., Minneapolis, Minnesota 55455, billcoop@me.umn.edu

(Received April 2000; revision received November 2000; accepted May 2001)

Revenue management has become an important tool in the airline, hotel, and rental car industries. We describe asymptotic properties of revenue management policies derived from the solution of a deterministic optimization problem. Our primary results state that, within a stochastic and dynamic framework, solutions arising out of a single well-known linear program can be used to generate allocation policies for which the normalized revenue converges in distribution to a constant upper bound on the optimal value. We also show similar asymptotic results for expected revenues. In addition, we describe counterintuitive behavior that can occur when allocations are updated during the booking process (updating allocations can lead to lower expected revenue). These results add to the understanding of allocation policies and help to make concrete the statement that simple policies from easy-to-solve formulations can be relatively effective, even when analyzed in the more realistic stochastic and dynamic framework.

1. INTRODUCTION

The practice of revenue management has become increasingly important in a number of areas, perhaps most notably the airline, hotel, and rental car industries. A common thread throughout these industries is the need to sell various products during a finite time horizon before the resources from which the products are made perish. For instance, in the airline context, such a resource may be a seat on a flight on a particular departure date. Often, a single resource can be combined with other resources to make different products, as is the case with networks of flights. In cases where there are many different types of resources and many different types of products, the question of how to most profitably control such a system can become quite complex. For recent surveys on revenue management, see Weatherford and Bodily (1992) or McGill and van Ryzin (1999).

Early revenue management models considered the stochastic and dynamic nature of the demand process in rather limited ways. The advantage of such approaches is that they yield easy-to-compute and easy-to-implement policies. For models with a single-leg flight, see (among many) Belobaba (1989), Curry (1990), or Brumelle and McGill (1993). Recent studies have considered more realistic models for the demand process. Such work on Markov-decision-process-type models of single-leg flights includes Stone and Diamond (1992), Lee and Hersh (1993), Lautenbacher and Stidham (1999), Liang (1999), Subramanian et al. (1999), and Zhao and Zheng (2001). If a single-leg flight is in question and the arrival process has independent increments, then natural monotonicity properties can be exploited to make the computation and storage of optimal policies possible. See You (1999) for some similar results regarding multileg flights. In all of these cases, the methods rely heavily on the fact that the demand processes are relatives of the Poisson process.

Although such single-leg methods can provide a good deal of insight, it is often the case that the real question is how to control a network of flight legs, from which many different itineraries can be constructed. Here, however, the so-called "curse of dimensionality" makes the computation and storage of optimal policies difficult or impossible for even moderate-sized problems. In these cases, it is often necessary to fall back on optimization methods that ignore much of the randomness in the demand; see Curry (1990), Glover et al. (1982), Simpson (1989), or Williamson (1992). Mathematical programming approaches for railway and rental car revenue management problems can be found in, respectively, Ciancimino et al. (1999) and Geraghty and Johnson (1997). For a comprehensive list of revenue management work, we refer the reader to the survey paper of McGill and van Ryzin (1999).

Mathematical programming approaches yield computable and "reasonable" policies, but typically disregard some seemingly key aspects of the problem. On the other hand, it is frequently the case that it is impossible to compute optimal policies when these factors are explicitly taken into account. Nevertheless, it is almost certainly the case that the "realistic" demand models offer the appropriate context within which to analyze proposed policies. Thus, we are led to a fundamental question: Are the policies that arise out of the simpler models good, when they are examined within the context of a more realistic stochastic and dynamic model? In a series of influential papers, Gallego and van Ryzin (1994, 1997) and Talluri and van Ryzin (1998) addressed this issue. They showed that for a general class of problems, certain simple policies do, in fact, perform well in comparison to optimal policies. In part, these

Subject classifications: Inventory, perishable items: revenue/yield management. Probability, stochastic model applications.
Area of review: Manufacturing, Service, and Supply-Chain Operations.

0030-364X/02/5004-0720 $05.00
1526-5463 electronic ISSN

results showed that when problems are scaled properly, the ratio of the expected revenue from their heuristics to the expected revenue from an optimal policy approaches one.

In this paper, we describe asymptotic properties, also under appropriate scaling and normalization, of revenue management policies derived from solutions to simple deterministic optimization problems. Our primary results state that, within a stochastic and dynamic framework, solutions arising out of a single linear program can be used to generate good policies. More precisely, we show that for a sequence of problems with increasingly large capacities and expected demands, the normalized sequence of revenues converges in distribution to a constant upper bound on the optimal value. In the simplest terms, this means that the distribution function of the normalized revenue "approaches" the point mass at the upper bound as the capacities and expected demands grow. We also show analogous results for the expected revenue. One further consequence is that the ratio of the expected revenue from the LP-based heuristic to that from an optimal policy converges to one as capacities and mean demands grow, as in the references above.

The linear program we study is a well-known classical revenue management formulation whose behavior, when viewed within the stochastic and dynamic framework, has not received much mathematical analysis in the literature. One contribution of this paper is the presentation of precise mathematical statements about the performance of the LP-based allocation policy within the context of a general stochastic model of the demand process. Our results should not be considered an endorsement of an LP-based allocation policy; rather, our objective is simply to conduct a formal analysis while providing general insights into the behavior of this type of acceptance/rejection rule.

Although we deal with limiting behavior, our work also suggests that such simple policies may be good when mean demand and remaining capacities are both "large." In the context of airline booking, for decisions made far from the time of departure, this is typically the case: Demand has yet to materialize, and all seats are unsold. One potential implication is that one might use a crude LP-based policy far from departure, and then switch to a more detailed decision rule near the departure date. Subramanian et al. (1999) suggest that revenues are, in fact, most sensitive to what occurs close to the time of departure.

While our results are related to those of other authors mentioned above, we must highlight several differences. One distinguishing feature of the work presented here is the emphasis on the convergence in distribution of the normalized revenue. Most other work has focused solely on expected revenues (one exception being Feng and Xiao 1999, where variability effects are factored into the decision model). In addition, we allow for arrival processes other than the compound Poisson process or its discrete analog. The key assumption herein is that the arrival process, when normalized, converges in distribution to the original mean value, a vector of constants. By emphasizing this condition, we believe that we show precisely what is the mechanism underlying the asymptotic behavior (albeit in a different context) of expected values first demonstrated in the references above. This will (hopefully) help make the results "intuitive." Using this assumption as a starting point, we show that the normalized revenue converges in distribution, and we also obtain similar results for expected values.

We focus only on the problem of controlling the availability of products. Such availability problems can be viewed as special cases (when the arrival process is Poisson) of the pricing problems studied in Gallego and van Ryzin (1997). There, however, only a more general case (but requiring a controlled compound Poisson demand process) is mentioned explicitly. Here, we allow different types of demand processes, but we study only the classical formulation alluded to above, and thereby provide direct insights into its performance. To provide additional insights into the behavior of such LP-based allocation policies, we describe some possible effects of updating allocations throughout the booking period. Specifically, we provide a simple example that exhibits the counterintuitive property whereby updating the allocations can lead to lower expected revenues.

The remainder of the paper is organized as follows. Section 2 introduces the model, §3 presents the LP and describes the resulting policy, §4 provides the basic asymptotic results, §5 discusses the issue of updating the policy, and §6 provides a brief summary.

2. PROBLEM FORMULATION

In this section, we begin by introducing the model. We consider a situation where there are m resources, combinations of which can be used to create n possible products. Assume that there is a finite time horizon $(0, \tau]$ during which sales can be made. Over this period, customers arrive according to a stochastic process and request to buy products. Each arriving customer is assumed to belong to one of the n product classes, where a class j customer is defined to be someone who requests product j. Throughout, we will use interchangeably the terms demand, request, and arrival. We make the convenient modeling assumption (and one that has been used fruitfully in numerous research papers and applications) that arriving customers "belong" to classes. We assume that the revenue from selling product j is given by the fare f_j, and there is no variability of fares *within* a class. In practice, this is often not the case, but we feel that it is a fairly reasonable modeling assumption. In fact, we can model the problem this way by relabeling distinct fares within a class as different fare classes, so long as the number of different fares within each original class is finite. Other authors have directly modeled random fares (see Talluri and van Ryzin 1998).

We are given an m-vector (all vectors are column vectors, unless stated otherwise) of resources $c \equiv (c_i)$, and an $(m \times n)$-matrix $A \equiv (a_{ij})$. The entry $c_i > 0$ represents the amount of resource i present at the beginning of the time horizon ($t = 0$), and the entry $a_{ij} \geqslant 0$ represents the amount of resource i requested by a class j customer. These entries need not be zeroes and ones (or even integers). A policy is a rule for accepting or rejecting requests. Whenever a class j product is sold, we say that we have accepted a class j customer. In this case he pays his fare f_j, consumes his resource requirement, and thereafter there are fewer resources available to sell to future customers. (Note that we assume that there are no cancellations or no-shows.) On the other hand, we say that a customer's request is rejected if he is not sold his requested product. In this case, no revenue is realized and no resources are consumed.

We do not provide here specific details about the dynamics of the arrival process other than that requests arrive one at a time to the system over $(0, \tau]$ according to a random point process that has a finite number of points (almost surely); i.e., the overall arrival process is a simple point process. For now, note that we have not specified whether the model is in discrete or continuous time, nor have we assumed that it is a Poisson process or a discrete approximation thereof. This highlights the generality of the results and will serve to emphasize precisely what are the important features that lead to our asymptotic statements.

In a standard airline context, this model can represent demand for a network of flights that depart within a particular day. In this case, the resources might be leg-cabin combinations, and the products might be itinerary-fare class combinations. In this case, row i of A would have zeroes in the columns corresponding to fare class-itineraries that do not use leg-cabin i, and ones in the columns corresponding to itinerary-fare classes that do use leg-cabin i. Throughout, we will use the term "customer" in a general sense; for instance, a class j customer could actually be a group that must be either entirely accepted or entirely rejected. Group demand is modeled easily in our framework by having values in A other than zero and one. For an airline example, the vector c represents the number of seats in each leg-cabin throughout the network. In this context, time 0 is when the airline begins selling tickets, and time τ is the departure time of the flights. In applications, there may be extremely large numbers of itinerary-fare classes and leg-cabins. Dealing with this issue, though important, is somewhat tangential to this paper.

Returning to the general case, the revenue management problem at hand is to determine a policy that maximizes the expected revenue from the sale of the products, subject to using at most c resources. Somewhat more formally, we seek a policy π that maximizes $E(f \cdot N^\pi)$, where f is the n-dimensional row vector $f \equiv (f_j)$ and N^π is the n-vector in which entry N_j^π is the number of accepted class j customers over $(0, \tau]$ when policy π is in use. We will also use the notation $R^\pi \equiv f \cdot N^\pi$ for the (random) amount of revenue when π is in use. Of course, a policy needs to satisfy (almost surely) $AN^\pi \leqslant c$ and $0 \leqslant N^\pi \leqslant D(\tau)$, where $D(\tau) \equiv (D_j(\tau))$ is the n-vector whose j-th entry is the total number of class j customer demands. We define the value v as

$$v \equiv \sup_{\pi \in \Pi} \{E(f \cdot N^\pi) : AN^\pi \leqslant c,\ 0 \leqslant N^\pi \leqslant D(\tau)\}. \tag{1}$$

Note that the constraints must be satisfied for almost every sample path of the demand process. Here, the set Π represents the set of allowable policies. Roughly speaking, this is the set of policies that make acceptance/rejection decisions at each time t based only upon information acquired up to t. An element π of Π maps (for each t) the history up to t to $\{0, 1\}^n$, where a one in the jth position means a class j arrival will be accepted if it arrives at time t, and a zero means a class j arrival will be rejected if it arrives at time t. We will use the notation $\pi(t)$ to denote the decision rule in effect at time t; note, however, that the rule can depend upon the "history of the process up to t." In addition, we do not allow availability to affect the demand process. Although this is a somewhat unrealistic assumption, it is one that has been made by numerous authors and is, in fact, implicit in many of the works cited in the Introduction.

Observe that we have not yet specified how a policy π and a demand realization combine to form a realization of N^π. To describe this interaction, let $D(\cdot)$ be the vector of arrival point processes, and let $D_j(t)$ be the number of class j requests in $(0, t]$; note that this is consistent with our earlier definition of $D_j(\tau)$. The notation $D(\cdot)$ is used to refer to the cumulative demand *process*; i.e., $D(\cdot) = \{D(t) : t \in (0, \tau]\}$. Later, we will employ similar notation when we want to emphasize something is a function (of time).

Under $\pi \in \Pi$, the number of accepted class j customers is then given by the random variable

$$N_j^\pi \equiv \int_{(0, \tau]} \pi_j(s) D_j(ds), \tag{2}$$

where $\pi_j(s)$ is the projection of π onto its jth component. For an informal explanation of the above, note that the increment $D_j(ds)$ is one when there is an arrival at time s and zero otherwise. Therefore, the expression on the right side of (2) simply counts the number of class j arrivals that occur at times when we are willing to accept class j arrivals. In a discrete-time situation, the integral above reduces to a summation. Because we have allowed general arrival processes, there may be no optimal Markovian policy; i.e., decisions may depend upon the entire history of the arrival process. We will not attempt to obtain optimal policies (or prove their existence); indeed, the dearth of implementable methods for finding optimal policies points out the importance of studying heuristics. We close this section by introducing the general notation $R^\pi(t)$ and $N_j^\pi(t)$, which represent, respectively, the revenue generated up to time t and the number of class j customers accepted up to time t, each under the arbitrary policy π. In other words $N_j^\pi(t)$ is defined as in (2), but with τ replaced by t, and $R^\pi(t) \equiv f \cdot N^\pi(t)$.

3. THE POLICY

In this section, we describe the linear program alluded to in the Introduction, and we show how to derive the heuristic from an optimal solution. In the following, $\mu < \infty$ is an n-vector where the jth entry represents the expected class j demand over $(0, \tau]$: $\mu_j \equiv ED_j(\tau)$. The linear program is given by

$$\max_x \{f \cdot x : Ax \leqslant c,\ 0 \leqslant x \leqslant \mu\}, \tag{3}$$

and we will denote by x^* an optimal solution. In case of multiple optima, we choose among them in an arbitrary manner. Also, define $\beta \equiv f \cdot x^*$ to be the optimal objective value of the linear program. The policy π^{LP} is defined by

$$\pi_j^{\mathrm{LP}}(t) \equiv 1(D_j(t) \leqslant x_j^*), \tag{4}$$

where $1(E)$ is the indicator function of the event E. Recall that $D_j(t)$ includes any class j arrival at time t; so (4) simply formalizes the policy "accept up to x_j^* customers in class j." When viewed within Gallego and van Ryzin's (1997) framework, this rule is closely related to what they call a make-to-stock policy. Assorted variations of the policy π^{LP} are quite well known to practitioners. For instance, rules of this type are discussed in Phillips et al. (1991), Williamson (1992), Phillips (1994), Geraghty and Johnson (1997), and Ciancimino et al. (1999), to name just a few. The policy π^{LP} requires solving an LP just once, at time $t = 0$, and then using the resulting allocation over the entire horizon $(0, \tau]$. In practice, it is common to solve such an LP at various points in time during $(0, \tau]$, modifying the LP each time to account for the fact that mean future demand and remaining capacity have changed. The derived policy may then be followed only up to the next re-solve point. We will revisit this issue in §5.

Note that π^{LP} allocates space to the various classes. For this reason, policies of this type are also often called discrete allocation polices. Under a variety of modeling assumptions, it is known that optimal policies are generally not of this form. For a number of multiclass, single-leg models, it has been shown that there is an optimal "nested" policy. In simple terms, this means that seats that are available to lower-class demand are always available to higher-class demand. Although it is not always possible to provide an ordering of the demand classes from low to high within the network setting, one may (pathwise) improve the revenue from the policy π^{LP} by enforcing an appropriate nesting scheme for classes that can be ordered. It is also common to obtain other types of policies from (3). In practice, one such important method is not to create allocations at all, but to instead use leg dual values as "bid prices." (See Simpson 1989; Williamson 1992; Phillips 1994; and Talluri and van Ryzin 1998, 1999.)

Returning to the issues at hand, under π^{LP}, the number of accepted class j demands is represented by N_j^{LP}. For a real number r, define $\lfloor r \rfloor$ to be the largest integer that does not exceed r; i.e., $\lfloor r \rfloor \equiv \max\{i \in \mathbb{Z} : i \leqslant r\}$. It follows that

$$N_j^{\mathrm{LP}} = D_j(\tau) \wedge \lfloor x_j^* \rfloor, \tag{5}$$

where $a \wedge b \equiv \min\{a, b\}$. The random revenue generated by π^{LP} is given by $R^{\mathrm{LP}} \equiv f \cdot N^{\mathrm{LP}}$. The policy π^{LP} does not "look into the future" and does not cause us to exceed capacity. The relation (5) shows why π^{LP} can be quite bad for problems that have many fare classes with very small mean demands: For many j, $\lfloor x_j^* \rfloor$ may be forced to be zero by the demand constraint of the LP, thereby causing us to leave too many resources unsold. In practice, this shortcoming can be corrected by aggregation methods. Similar problems occur when expected demand is low relative to capacity. These can also, to a degree, be overcome through various modifications to π^{LP}. Such issues are not addressed in this paper, because they do not occur in the asymptotic cases we study. Finally, note that the optimal objective value β is *not* the expected revenue from using policy π^{LP}.

4. ASYMPTOTIC RESULTS

We begin this section by introducing a sequence of problems indexed by k. For the kth problem, we replace our initial capacity vector by kc, and we assume the arrival process is given by the n-vector $D^{(k)}(\cdot)$. For each problem, we assume (as before) that, with probability one, there is (a) a finite number of arrivals and (b) at most one arrival at any point in time. We also will assume that $ED^{(k)}(\tau) = k\mu$. All notation from the original problem carries over to the kth problem with a parenthetical k added as a superscript. For instance, $v^{(k)}$ is the value of the kth problem; see (1). For each k we maintain the same resource-consumption matrix A and fare vector f. The kth linear program is given by

$$\max_x \{f \cdot x : Ax \leqslant kc,\ 0 \leqslant x \leqslant k\mu\}. \tag{6}$$

If x^* is our optimal solution from the original linear program, then kx^* is an optimal solution for the kth linear program. So, the kth LP has optimal objective value $k\beta$. In cases where there are multiple optimal solutions to the kth problem, we adopt the convention that we use kx^* to create our allocations.

Hereafter, we shall refer to quantities associated with the kth problem as "scaled." Scaled quantities that are subsequently divided by k will be called "normalized."

Recall that given a sequence of random variables $\{X_k\}$, we say that X_k *converges in distribution* to the random variable X (written $X_k \stackrel{\mathcal{D}}{\to} X$) if $\lim_{k\to\infty} P_k(X_k \in S) = P(X \in S)$ for every X-continuity set S. Here, P_k gives the distribution of X_k, and P gives the distribution of X. In words, this means that the distribution of X_k becomes close to that of X when k is large. For more on this, other equivalent definitions, and extensive coverage of convergence in distribution, see Billingsley (1968), Fristedt and Gray (1997), or other probability texts. As mentioned in the Introduction, our results follow from the assumption that the vector of the normalized number of arrivals converges in distribution to the vector of the original expected number of arrivals:

$$k^{-1} D^{(k)}(\tau) \stackrel{\mathcal{D}}{\to} \mu. \tag{7}$$

We believe that by emphasizing this condition, we show precisely what is the mechanism underlying the asymptotic expected-value behavior of the type first demonstrated in Gallego and van Ryzin (1994, 1997) and Talluri and van Ryzin (1998). Using (7) as a starting point, we show that the normalized revenue converges in distribution to $f \cdot x^*$. From this, it then follows that the ratio of the expected revenue from the allocation policy to that from an optimal policy approaches one.

Before proceeding, we will briefly describe a few classes of arrival processes that satisfy the condition (7). Perhaps the most important example occurs when the original demand process is a (possibly nonhomogeneous) Poisson process with rate function $\lambda(\cdot)$; i.e., $ED(t) = \int_{(0,t]} \lambda(s)\,ds$ and $\mu = \int_{(0,\tau]} \lambda(s)\,ds$. Then (7) will be satisfied if, on the kth problem, we make the demand process a Poisson process with rate function $\lambda^{(k)}(\cdot) \equiv k\lambda(\cdot)$. Such a construction is described in Gallego and van Ryzin (1997).

We next consider the case of a general continuous-time arrival process $D(\cdot)$ satisfying the conditions of §2, for which $ED_j(\tau) = \mu_j < \infty$ and $Var(D_j(\tau)) < \infty$ for all j, and the expected number of arrivals at each time point is zero. We can then obtain a sequence of demand processes by defining, for each k, the stochastic process $D^{(k)}(\cdot) = \widetilde{D}^{1,k}(\cdot) + \cdots + \widetilde{D}^{k,k}(\cdot)$, where $\{\widetilde{D}^{i,k}(\cdot)\colon i = 1, \ldots, k\}$ are k i.i.d. "copies" of the process $D(\cdot)$. It is known that processes so constructed satisfy (7). Furthermore, if we assume (a) that $\{\widetilde{D}^i(\cdot)\colon i = 1, 2, \ldots\}$ is an i.i.d. sequence of demand processes and (b) that $\widetilde{D}^{i,k}(\cdot) = \widetilde{D}^i(\cdot)$ for all $k \geqslant i$; then $k^{-1}D^{(k)}(\tau) = k^{-1}(\widetilde{D}^1(\tau) + \cdots + \widetilde{D}^k(\tau)) \to \mu$ with probability one by the strong law of large numbers. Condition (7) then holds because almost-sure convergence implies convergence in distribution. The assumption that the expected number of arrivals at each time point is zero implies that in the original (nonscaled) problem, for any particular instant in continuous time, the probability of having an arrival is zero. In turn, this implies that with probability one, there are not multiple arrivals at any time point in the kth problem (see Franken et al. 1982, Theorem 1.3.10).

A similar technique will work for discrete-time processes with the minor adjustment that we must introduce, at each step, a successively finer time grid so that we do not allow for the possibility of multiple arrivals in one period. Specifically, we replace each time point with k independent copies of itself.

Note that we require only that the random variables $k^{-1}D^{(k)}(\tau)$ converge in distribution, but we do not make this requirement of the demand *process*. However, constructing a sequence of problems in the manner described above has the desirable property that it preserves the shape of the expected buildup of demand. For instance, if in the original problem the expected number of demands is much higher just prior to departure than far from departure, then this will again be the case in the kth problem. There are other ways to pick $D^{(k)}(\cdot)$ to fit within the required framework, and we need not preserve the shape of demand in this way for the results to hold. However, it certainly seems desirable if we want to maintain the connection with the original problem.

LEMMA 1. *Suppose that X and $\{X_k\}$ are $\mathbb{R}^n$-valued random variables such that $X_k \xrightarrow{\mathcal{D}} X$, and suppose that the functions $h_k : \mathbb{R}^n \to \mathbb{R}$ converge uniformly on compact sets to a continuous function $h : \mathbb{R}^n \to \mathbb{R}$. Then $h_k(X_k) \xrightarrow{\mathcal{D}} h(X)$.*

A proof can be found in Billingsley (1968, p. 34). We now are now ready to present our first main result.

PROPOSITION 1. *Suppose that* (7) *holds*; *i.e., suppose that the normalized number of arrivals converges in distribution to the original expected number of arrivals. Then*
(i) $k^{-1}R^{\mathrm{LP}(k)} \xrightarrow{\mathcal{D}} \beta$ *as* $k \to \infty$, *and*
(ii) $k^{-1}ER^{\mathrm{LP}(k)} \to \beta$ *as* $k \to \infty$,
where β is a random variable that satisfies $P(\beta = f \cdot x^*) = 1$.

PROOF. We prove (i) first. Applying (5) to the kth problem, we see that $N_j^{\mathrm{LP}(k)} = D_j^{(k)}(\tau) \wedge \lfloor kx_j^* \rfloor$. Therefore,

$$\begin{aligned} k^{-1}R^{\mathrm{LP}(k)} &= k^{-1}\sum_{j=1}^{n} f_j\big(D_j^{(k)}(\tau) \wedge \lfloor kx_j^*\rfloor\big) \\ &= \sum_{j=1}^{n} f_j\big(k^{-1}D_j^{(k)}(\tau) \wedge k^{-1}\lfloor kx_j^*\rfloor\big) \\ &\xrightarrow{\mathcal{D}} \sum_{j=1}^{n} f_j\big(\mu_j \wedge x_j^*\big) = \sum_{j=1}^{n} f_j x_j^*. \end{aligned}$$

(Recall that f_j is the fare paid by class j; the quantities in parentheses above are *not* arguments of a function f_j.) The convergence in distribution is justified by Lemma 1, because the functions $g_k : \mathbb{R}^n \to \mathbb{R}$ defined by $g_k(y) \equiv \sum_{j=1}^n f_j(y_j \wedge k^{-1}\lfloor kx_j^*\rfloor)$ converge uniformly on compact sets to g defined by $g(y) \equiv \sum_{j=1}^n f_j(y_j \wedge x_j^*)$, which is a continuous function. The random variables $k^{-1}R^{\mathrm{LP}(k)}$ are uniformly bounded (by the number β); therefore, (ii) follows from (i). □

Examination of the proof above suggests a general method for proving similar results for other types of policies. For a particular policy π under consideration, one need only demonstrate that $k^{-1}N^{\pi(k)} \xrightarrow{\mathcal{D}} x^*$.

LEMMA 2. $ER^{\mathrm{LP}(k)} \leqslant v^{(k)} \leqslant k\beta$ *for all* $k \geqslant 1$, *where* $v^{(k)}$ *is the value of the kth problem.*

Variants of the above have appeared in several papers, including Gallego and van Ryzin (1997) and Chen et al. (1998). The first inequality follows from the definition of $v^{(k)}$. The second inequality follows from the observation that for any policy, we have $N^{\pi(k)} \leqslant D^{(k)}(\tau)$ and $AN^{\pi(k)} \leqslant kc$, and therefore $EN^{\pi(k)} \leqslant k\mu$ and $A(EN^{\pi(k)}) \leqslant kc$. Consequently, $EN^{\pi(k)}$ is feasible for the LP, so the second inequality holds.

In the following, we see that LP-based allocations are asymptotically good in comparison to any other policy. In conjunction with Proposition 1, this can be seen as the primary result of this paper. Below, the revenues are scaled, but *not* normalized. However, one should note that division by $v^{(k)}$ has, in a sense, taken the normalization role previously held by division by k. Statement (ii) below is the precise analog of the asymptotic results of the Gallego and van Ryzin and Talluri and van Ryzin papers. Observe that (ii) is indeed a consequence of (7), supporting our earlier assertion that (7) helps explain why such results are true.

PROPOSITION 2. *Suppose that* (7) *holds. Then,*
(i) $R^{\mathrm{LP}(k)}/v^{(k)} \xrightarrow{\mathscr{D}} 1$ *as* $k \to \infty$, *and*
(ii) $ER^{\mathrm{LP}(k)}/v^{(k)} \to 1$ *as* $k \to \infty$.

PROOF. To prove (ii), note that Lemma 2 implies that

$$\frac{ER^{\mathrm{LP}(k)}}{k\beta} \leqslant \frac{ER^{\mathrm{LP}(k)}}{v^{(k)}} \leqslant 1.$$

Statement (ii) now follows from the preceding display by part (ii) of Proposition 1. To show (i), note that

$$\frac{R^{\mathrm{LP}(k)}}{v^{(k)}} = \frac{R^{\mathrm{LP}(k)}}{k\beta} \cdot \frac{k\beta}{ER^{\mathrm{LP}(k)}} \cdot \frac{ER^{\mathrm{LP}(k)}}{v^{(k)}}.$$

Part (ii) of Proposition 1 and part (ii) of this proposition together imply that the sequence of numbers $a_k \equiv (k\beta/ER^{\mathrm{LP}(k)}) \cdot (ER^{\mathrm{LP}(k)}/v^{(k)}) \to 1$ as $k \to \infty$. Statement (i) now follows by Proposition 1 and Lemma 1. □

Note that (ii) does not guarantee that $|ER^{\mathrm{LP}(k)} - v^{(k)}| \to 0$. Because $v^{(k)}$ is proportional to k, it is possible that the difference increases without bound. It can be shown (see Gallego and van Ryzin 1997) that for Poisson cases, the difference grows proportionally to at most $\sqrt{k}$. In addition, it is possible to construct examples to show that, in general, no sharper rate can be obtained. However, under the additional assumptions that demand is Poisson and that $x_j^* < \mu_j$, it can be proved that $\limsup_{k\to\infty} |ER^{\mathrm{LP}(k)} - v^{(k)}| \leqslant \sum_{j=1}^{n} f_j$. For details, see Cooper (2000), which includes material from an earlier version of this paper.

One natural question is whether or not there are other simple policies with provably good asymptotic behavior. If $A\mu \leqslant c$, then it is possible to prove the precise analogs of Propositions 1 and 2 for the first-come–first-served (FCFS) policy. It is not too surprising that FCFS is fairly good in such cases, because the condition $A\mu \leqslant c$ implies that, on average, the total amount of demand is not enough to exceed capacity. When $A\mu \leqslant c$ does not hold, then it is possible to construct examples that show, in a sense, that FCFS is asymptotically "very bad." Although these results pertaining to FCFS are rather intuitive, they do require considerable notational buildup. Consequently, we again direct the interested reader to Cooper (2000).

5. IS RE-SOLVING BETTER?

The primary focus of this paper has been the policy π^{LP}, which is derived by solving a single linear program to create an allocation that is used over the entire booking horizon. In this section we ask the question: Is it better (on average) to update the allocations during the booking period? In §3 we raised the issue that, in practice, such policies are indeed typically updated. One standard way that this is accomplished is by "re-solving" (later in the booking period) a modified LP that takes into account adjustments in remaining capacity and expected remaining demand. More specifically, the capacity is decremented by the space already taken up by booked customers, and the expected demand is replaced by the conditional expected future demand, given the history of the process up to the re-solve time. The allocation obtained from the solution to this problem is then used until the end of the horizon τ (or the next re-solve time). It seems plausible that one can obtain an improvement over π^{LP} by updating allocations in this fashion, because these updates allow one to take into account deviations away from the expected demand buildup.

In this section, we show by example that policies based upon re-solving can, in fact, yield lower expected revenues than does π^{LP}. This is perhaps a counterintuitive result, because one might expect that a policy that responds to observed information would be better. Note that it is simple to come up with sample paths for which some particular policy will outperform another seemingly "better" policy, where "better" refers to yielding higher expected revenue. So, it should not be surprising if re-solving does worse on some sample paths. Here, however, we emphasize that the *expected* revenue yielded by re-solving can be strictly worse than that given by following π^{LP}.

To our knowledge, there are no prior formal mathematical results that compare the performance of revenue management heuristics based upon the frequency with which they are updated. In light of this, we feel that although the example we will present is "unrealistic," it does serve to emphasize one key fact: Frequent updating/re-solving is *not necessarily* better. Of course, it is reasonable to expect that in many cases, frequent updating will give better performance than infrequent (or no) updating. In addition, there are significant practical benefits to updating, such as the incorporation of user input, the adjustment of estimates of model parameters, and the use of overbooking techniques to deal with cancellations and no-shows.

To present our example, we first must state more formally what is meant by a *re-solving policy*. We begin by specifying a re-solve time t'. The re-solving policy, denoted π^{RS}, is as follows. Solve the linear program (3) at time 0, and follow π^{LP} up to and including time t'. Then, at time t' solve the updated LP:

$$\max_y \{f \cdot y : Ay \leqslant c - AN^{\mathrm{LP}}(t'),\quad 0 \leqslant y \leqslant E(D(\tau) - D(t') \mid H_{t'})\}, \tag{8}$$

where $H_{t'}$ is the history of the process up to time t'. We let y^* denote an optimal solution of (8). Observe that the vector $c - AN^{\mathrm{LP}}(t')$ is the amount of capacity remaining unsold at time $t = t'$. Similarly, $E(D(\tau) - D(t') | H_{t'})$ is the expected

Table 1. Effects of re-solving at Time 1 on bookings in $(1, 2]$.

$D_1(1)$	y_1^*	y_2^*	$N_1^{RS} - N_1^{RS}(1)$	$N_2^{RS} - N_2^{RS}(1)$	$N_1^{LP} - N_1^{LP}(1)$	$N_2^{LP} - N_2^{LP}(1)$
0	1	1	$(D_1(2) - 0) \wedge 1$	$(D_2(2) - D_2(1)) \wedge 1$	$(D_1(2) - 0) \wedge 2$	0
1	1	0	$(D_1(2) - 1) \wedge 1$	0	$(D_1(2) - 1) \wedge 1$	0
$\geqslant 2$	0	0	0	0	0	0

future demand from time t' onward (not inclusive), conditional upon the process up to t'. When demand is Poisson, this expression reduces to $E(D(\tau) - D(t'))$. After time t', the policy π^{RS} follows the allocation given by y^*; i.e., accept a class j request at time $t \in (t', \tau]$ if and only if $D_j(t) - D_j(t') \leqslant y_j^*$. Formally, π^{RS} is given by

$$\pi^{RS}(t) \equiv \begin{cases} \pi_j^{LP}(t) & t \leqslant t' \\ 1(D_j(t) - D_j(t') \leqslant y_j^*) & t > t'. \end{cases}$$

We are now ready to present the example that shows that π^{RS} can yield strictly worse expected revenue than π^{LP}. In the following example, it is possible to obtain an optimal policy; see Stone and Diamond (1992) or Liang (1999). However, by focusing on a simple example we are able to demonstrate clearly the possible drawbacks of updating suboptimal policies. For more complicated problems, it is not possible to compute optimal policies, so how to best apply heuristics becomes particularly important.

Example 1. Consider a single-leg problem with two fare classes: Class 1 and Class 2, who pay $f_1 = 10$ and $f_2 = 2$, respectively. Suppose $A = (1, 1)$, $c = 2$, $\tau = 2$, and $t' = 1$. Assume also that the demand processes for Class 1 and Class 2 are independent homogeneous Poisson processes, each with rate 1. By inspection, we see that $x_1^* = 2$, $x_2^* = 0$ is an optimal solution to (3).

For any policy π, we can decompose R^π into the sum of revenues prior to Time 1 and revenues after Time 1 as follows: $R^\pi = R^\pi(1) + (R^\pi - R^\pi(1))$. Therefore, we can compute the expected revenue from π by

$$ER^\pi = E(R^\pi(1)) + E(E(R^\pi - R^\pi(1) \mid H_1)). \tag{9}$$

Table 1 describes the number of accepted bookings during the time period $(1, 2]$ for both π^{LP} and π^{RS}, under different possible scenarios for Class 1 demand prior to Time 1. Note also that no Class 2 customers are accepted by either π^{LP} or π^{RS} during the time period $(0, 1]$. In light of Table 1, Equation (9), and the fact that $N^{LP}(1) = N^{RS}(1)$, one can verify the claim that $ER^{RS} < ER^{LP}$ by checking that

$$E\big(10(D_1(2) \wedge 1) + 2((D_2(2) - D_2(1)) \wedge 1) \mid D_1(1) = 0\big) < E\big(10(D_1(2) \wedge 2) \mid D_1(1) = 0\big).$$

This is equivalent to verifying that

$$E(10(Q_1 \wedge 1) + 2(Q_2 \wedge 1)) < E(10(Q_1 \wedge 2)), \tag{10}$$

where Q_1 and Q_2 are independent Poisson random variables, each with mean 1. Direct calculation shows that the left-hand side of (10) is 7.59, and the right-hand side of (10) is 8.96, thereby proving the claim. It is also possible to obtain explicit algebraic conditions on general fares f_1 and f_2, so that (10) holds.

Observe that in the example, the poor performance of π^{RS} is not caused by rounding effects that may be introduced when translating from an LP solution to an allocation. It is simple to construct examples where the discrete nature of demand can cause re-solving the LP to behave poorly. For instance, re-solving at a time when all conditional expected remaining demands are strictly below one will result in an updated allocation that is identically zero. Clearly, however, such behavior would be corrected in practice. The example above demonstrates that, absent such effects, counterintuitive behavior is still possible. Also, one might argue that this example displays a drawback peculiar to discrete allocation policies. Although there may be some merit to this assertion, the example does point out the fact that, when choosing within a class of suboptimal policies, taking more information into account need not always be better. Furthermore, it should point out that, from a mathematical perspective, the effects of updating and re-solving are not yet fully understood.

6. CONCLUSIONS

In this paper we have described a general stochastic and dynamic framework for network revenue management problems. Working within this framework, we have presented limit theorems that help formalize the notion that, although optimal policies can be difficult or impossible to compute, relatively simple and easy-to-compute policies can be good. Our results are similar in spirit to earlier work of Gallego and van Ryzin (1994, 1997) and Talluri and van Ryzin (1998). We have studied different modes of convergence, presented relatively simple proofs, and maintained a specific focus on LP-based allocations, thereby adding to the overall understanding of such policies. In addition, we have shown by example that a standard method for updating allocations can lead to worse performance. This, in turn, suggests a future line of research that examines when and how to update heuristic policies.

ACKNOWLEDGMENTS

The author is grateful to the referees and the area editor. Their valuable suggestions and insights helped improve the presentation of this paper.

REFERENCES

Belobaba, P. P. 1989. Application of probabilistic decision model to airline seat inventory control. *Oper. Res.* **37** 183–197.

Billingsley, P. 1968. *Convergence of Probability Measures*. John Wiley and Sons, New York.

Brumelle, S. L., J. I. McGill. 1990. Airline seat allocation with multiple nested fare classes. *Oper. Res.* **41** 127–137.

Chen, V. C. P., D. Günther, E. L. Johnson. 1998. A Markov decision problem based approach to the airline YM problem. Working paper, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA.

Ciancimino, A., G. Inzerillo, S. Lucidi, L. Palagi. 1999. A mathematical programming approach for the solution of the railway yield management problem. *Transportation Sci.* **33** 168–181.

Cooper, W. L. 2000. Asymptotic behavior of some revenue management policies: Examples and counterexamples. Technical report, Department of Mechanical Engineering, University of Minnesota, Minneapolis, MN.

Curry, R. E. 1990. Optimal airline seat allocation with fare classes nested by origins and destinations. *Transportation Sci.* **24** 193–204.

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

Franken, P., D. König, U. Arndt, V. Schmidt. 1982. *Queues and Point Processes*. John Wiley and Sons, New York.

Fristedt, B., L. Gray. 1997. *A Modern Approach to Probability Theory*. Birkhäuser, Boston, MA.

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

——, ——. 1997. A multiproduct dynamic pricing problem and its applications to network yield management. *Oper. Res.* **45** 24–41.

Geraghty, M. K., E. Johnson. 1997. Revenue management saves National Car Rental. *Interfaces* **27** 107–127.

Glover, F., R. Glover, J. Lorenzo, C. McMillan. 1982. The passenger mix problem in scheduled airlines. *Interfaces* **12** 73–79.

Lautenbacher, C. J., S. Stidham. 1999. The underlying Markov decision process in the single-leg airline yield management problem. *Transportation Sci.* **33** 136–146.

Lee, T. C., M. Hersh. 1993. A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Sci.* **27** 252–265.

Liang, Y. 1999. Solution to the continuous time dynamic yield management model. *Transportation Sci.* **33** 117–123.

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

Phillips, R. L. 1994. A marginal value approach to airline origin and destination revenue management. J. Henry, P. Yvon, eds. *Proceedings of the 16th Conference on System Modeling and Optimization*. Springer-Verlag, New York.

——, D. W. Boyd, T. A. Grossman. 1991. An algorithm for calculating consistent itinerary flows. *Transportation Sci.* **25** 225–239.

Simpson, R. W. 1989. Using network flow techniques to find shadow prices for market and seat inventory control. MIT Flight Transportation Laboratory Memorandom M89-1, MIT, Cambridge, MA.

Stone, R., M. Diamond. 1992. Optimal inventory control for a single flight leg. Technical report, Operations Research Division, Northwest Airlines, Minneapolis, MN.

Subramanian, J., S. Stidham, C. J. Lautenbacher. 1999. Airline yield management with overbooking, cancellations, and no-shows. *Transportation Sci.* **33** 147–168.

Talluri, K., G. van Ryzin. 1998. An analysis of bid-price controls for network revenue management. *Management Sci.* **44** 1577–1593.

——, ——. 1999. A randomized linear programming method for computing network bid prices. *Transportation Sci.* **33** 207–216.

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

Williamson, E. L. 1992. Airline network seat control. Ph.D. dissertation, MIT, Cambridge, MA.

You, P.-S. 1999. Dynamic pricing in airline seat management for flights with multiple flight legs. *Transportation Sci.* **33** 192–206.

Zhao, W., Y.-S. Zheng. 2001. A dynamic model for airline seat allocation with passenger diversion and no-shows. *Transportation Sci.* **35** 80–98.