M&SOM MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Vol. 10, No. 2, Spring 2008, pp. 288–310
ISSN 1523-4614 | EISSN 1526-5498 | 08 | 1002 | 0288
informs®
DOI 10.1287/msom.1070.0169

On the Choice-Based Linear Programming Model for Network Revenue Management

Qian Liu
Industrial Engineering and Logistics Management Department, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong, qianliu@ust.hk

Garrett van Ryzin
Graduate School of Business, Columbia University, New York, New York 10027, gjv1@columbia.edu

Gallego et al. [Gallego, G., G. Iyengar, R. Phillips, A. Dubey. 2004. Managing flexible products on a network. CORC Technical Report TR-2004-01, Department of Industrial Engineering and Operations Research, Columbia University, New York.] recently proposed a choice-based deterministic linear programming model (CDLP) for network revenue management (RM) that parallels the widely used deterministic linear programming (DLP) model. While they focused on analyzing "flexible products"—a situation in which the provider has the flexibility of using a collection of products (e.g., different flight times and/or itineraries) to serve the same market demand (e.g., an origin-destination connection)—their approach has broader implications for understanding choice-based RM on a network. In this paper, we explore the implications in detail. Specifically, we characterize optimal offer sets (sets of available network products) by extending to the network case a notion of "efficiency" developed by Talluri and van Ryzin [Talluri, K. T., G. J. van Ryzin. 2004. Revenue management under a general discrete choice model of consumer behavior. *Management Sci.* **50** 15–33.] for the single-leg, choice-based RM problem. We show that, asymptotically, as demand and capacity are scaled up, only these efficient sets are used in an optimal policy. This analysis suggests that efficiency is a potentially useful approach for identifying "good" offer sets on networks, as it is in the case of single-leg problems. Second, we propose a practical decomposition heuristic for converting the static CDLP solution into a dynamic control policy. The heuristic is quite similar to the familiar displacement-adjusted virtual nesting (DAVN) approximation used in traditional network RM, and it significantly improves on the performance of the static LP solution. We illustrate the heuristic on several numerical examples.

Key words: network revenue management; choice behavior; multinomial logit choice model; dynamic programming; linear programming
History: Received: October 21, 2004; accepted: February 16, 2007. Published online in *Articles in Advance* December 17, 2007.

1. Introduction and Literature Review

Among both practitioners and researchers, there is growing interest in modeling customer choice behavior in revenue management problems. This interest stems partly from a long-standing dissatisfaction with the limitations of the traditional independent demand model.[1] Indeed, heuristic corrections to the independent demand model to account for "buy-up" and "buy-down" effects date back to the earliest work on capacity control problems (see Belobaba 1987a, b). More-rigorous modeling of such a phenomenon has been a persistent goal of RM research, and customer choice models provide a natural solution.

In addition, the emergence of low-cost airlines offering simplified, undifferentiated fare structures has rekindled interest in modeling customer choice behavior. In this new environment, there are few fare restrictions, and often, fare products differ only in terms of price. Customers appear to be quite willing to exploit the resulting purchase flexibility; for example, industry colleagues have commented that in markets with unrestricted fares, often the only observed bookings are in the lowest available fare

[1] The independent demand model is one in which demand for products is assumed to be mutually statistically independent and also unaffected by the availability controls applied by the seller; see Talluri and van Ryzin (2004b) for a detailed discussion of the independent demand model.

class. In such environments, the assumption that each fare product is purchased exclusively by a distinct customer segment—an acceptable (if not entirely satisfying) practical compromise in the past—has now moved into the realm of pure modeling absurdity. These industry trends have accelerated interest in customer choice modeling as an alternative foundation on which to base RM models and systems.

Several researchers have looked at approximate analysis of customer choice behavior for single-leg RM problems. As mentioned, Belobaba (1987a, b) proposed a correction to the expected marginal seat revenue (EMSR) heuristics that introduces a probability of buying a higher fare when a low fare is closed. (See also Belobaba and Weatherford 1996.) Phillips (1994) proposed a state-contingent approach to revenue management that adjusts controls based on forecasts that depend on the controls in effect (the system state) at any point in time. Talluri and van Ryzin (2004a) provided an exact analysis of the optimal control policy for a single-leg model of RM under a general discrete choice model of demand. A key result is that the optimal policy can be characterized in terms of an ordered sequence of "efficient" offer sets. These sets are efficient in the sense that they provide the most favorable trade-off between expected revenue and expected capacity consumption. One of the contributions of our work is to generalize this analysis of efficient offer sets to the network case.

Network RM problems are discussed extensively in Talluri and van Ryzin (2004b, Chapter 3). Most work to date on network problems is based on the independent demand model. However, even in the independent demand case, the resulting optimization problems are difficult due to the high dimensionality of the state space. Hence, research in the area has focused on various approximate methods, mainly math programming and decomposition approaches.

The most popular math programming approximation is the DLP model, first introduced by Simpson (1989) and later investigated by Williamson (1988, 1992), Talluri and van Ryzin (1999a, b), and Bertsimas and Popescu (2003). In this model, demand for each product is treated as a deterministic quantity equal to the mean forecasted demand. A linear program is then solved to find the optimal mix of demand to accept given the capacity constraints on each leg of the network. The primal solution to the DLP model is typically not used directly; rather, the optimal dual prices are used to construct bid price controls.[2] Alternatively, the dual prices are used in a displacement-adjusted virtual nesting (DAVN) scheme or dynamic programming (DP) decomposition scheme, in which product revenue values at each leg of the network are adjusted by subtracting an estimate of the displacement cost on other legs of the network (the dual prices), and then modified leg-level problems are formulated using these displacement-adjusted revenue values. This decomposes the complex network problem into a collection of computationally tractable leg-level problems. Again, see Talluri and van Ryzin (2004b, Chapter 3) for an in-depth discussion of such decomposition approaches.

The earliest work on choice behavior in networks is the passenger origin and destination simulator (PODS) studies of Belobaba and Hopperstad (1999). This work has focused on understanding the revenue management implications of passenger choice behavior on traditional RM methods (primarily, methods based on the independent demand model). An interesting body of applied research on choice modeling in networks is the work of Andersson (1989, 1998), and Algers and Besser (2001), who report a research and development effort at SAS to apply logit choice models to estimate buy-up and recapture factors at one of SAS's hubs. Zhang and Cooper (2005) analyze choice among parallel flights in the same market (e.g., different departure times between the same origin-destination pair). The model assumes that customers choose among the same fare class on different flights, but not among fare classes (e.g., customer segments are still effectively separated by the fare-class restrictions). They develop bounds and approximations to the resulting dynamic program. The shortcoming of this work is that the approximation approach is specialized to the parallel flight case. van Ryzin and Vulcano (2004) apply a simulation-based optimization approach to compute optimal parameters of a virtual

[2] A bid price control sets threshold prices for each leg in the network. A product request is then accepted (i.e., the product is open for sale) if its revenue exceeds the sum of the bid prices of the legs required to satisfy the request. See Talluri and van Ryzin (1999a, b), as well as Talluri and van Ryzin (2004b, Chapter 3), for a discussion and analysis of bid price controls.

nesting control policy under a general model of network choice behavior. Although their demand model is quite general, their approach is restricted to a specific virtual nesting control scheme.

As mentioned, our work is directly motivated by the work of Gallego et al. (2004). While nominally addressing a problem with flexible fare products, their model is quite general and has important implications for choice-based network RM. Their model can be considered to be the natural choice-based analog of the DLP model of traditional network RM. We make this connection precise below, but it is important to recognize. Our objective is to explore the implications of their work for choice-based RM, relate it to the existing body of knowledge on network RM problems, and extend their ideas where possible.

Finally, we note that other models and methods addressing choice behavior issues are dynamic pricing models such as those studied by Bitran et al. (1998), Feng and Gallego (1995, 2000), Feng and Xiao (2000a, b, c), Gallego and van Ryzin (1994, 1997), and Maglaras and Meissner (2006). However, most models assume that only one product is sold at one price at any point in time. Customers then face a binary choice of whether or not to buy. However, in the airline case in particular, many fare products are offered simultaneously, and customers choose among them based on their preferences for price, and also nonprice factors such as refundability and whether they can meet various purchase restrictions (e.g., Saturday night stay). In the network case, customers also choose among different routings and flight departure times. The exception is the network models of Gallego and van Ryzin (1997), which allow for multiple products to be sold at the same time, although again, the control in their analysis is the price of each product rather than its availability. Maglaras and Meissner (2006) explore a unified framework for the multiproduct dynamic pricing problem and the capacity control problem by demand aggregation.

2. Overview

The remainder of the paper is organized as follows. We begin in §3 by formulating a general model of revenue management under customer choice behavior. The problem is to find the optimal set of products to offer at any point in time (the optimal *offer set*) based on the current remaining capacities and remaining time. We then give the corresponding CDLP for this problem as proposed by Gallego et al. (2004), which can be considered as a deterministic approximation of the original stochastic problem.

We then establish in §4 that the solution to the CDLP model is in fact asymptotically optimal for the stochastic network choice problem because the demand and capacity are scaled up proportionally (this scaling is made precise below). This result is not surprising because it parallels the behavior of the DLP solution in the independent demand case. (See Cooper 2002, Gallego and van Ryzin 1997, Talluri and van Ryzin 1999a.) Still, the result is reassuring and shows that CDLP shares some of the desirable theoretical behavior of the DLP.

More important, the asymptotic optimality of the CDLP provides useful insights about optimal offer sets. In particular, in §5 we use the CDLP and our asymptotic analysis to extend the notion of efficient sets introduced by Talluri and van Ryzin (2004a). The notion of efficiency here is an efficient trade-off between expected revenue and the expected vector of consumption rates. It is a single-output, multiple-input notion of efficiency, analogous to notions of efficiency from data envelopment analysis (Cooper et al. 1978).

In §6, we explore practical applications of the CDLP model. We first briefly review the column-generation strategy for solving the model proposed by Gallego et al. (2004), which provides a general framework for efficiently computing the solution to what is otherwise an exponentially large math program. We then develop a decomposition heuristic for using the dual prices of the CDLP analogous to DAVN-like decomposition schemes developed for the traditional independent demand model. We specialize both procedures to the case in which customers are assumed to be divided into segments, each of which has a disjoint consideration set of products. We show that the computation and approximation methods are quite efficient in this special case. In §7, we test our decomposition heuristic using some numerical examples. The results show that the decomposition heuristic produces significant improvements in average revenue compared to a direct application of the CDLP solution. Moreover, its performance is much less sensitive to the frequency of reoptimizations, which is an important property given that solving the CDLP for large networks can be quite computationally complex. Finally, conclusions are provided in §8.

3. Model

We begin by formulating a general statement of the network RM problem under customer choice behavior. The network has m legs and provides n products (a product is defined by an itinerary and fare class combination). The initial capacities are denoted $c = (c_1, \ldots, c_m)$. The set of products is denoted by $N = \{1, \ldots, n\}$. Each product j has an associated revenue r_j. Define the incidence matrix $A = [a_{ij}]_{m \times n}$. If leg i is used by product j, $a_{ij} = 1$; otherwise, $a_{ij} = 0$. The jth column vector A_j is the incidence vector for product j, and the ith row vector A^i is the incidence vector for leg i. With slight abuse of notations for convenience, we use the notation $i \in A_j$ to denote that leg i is used by product j, and $j \in A^i$ to denote that product j uses leg i.

We consider discrete periods indexed forward in time by t. We assume at most one arrival of a customer within each period; that is, time is divided sufficiently finely so that the probability of more than one request is negligible. The probability of an arrival in each period is denoted by λ. While it is not difficult to allow these arrival probabilities to depend on time t, to simplify the exposition we assume the arrival probability is constant over time.

The firm's control is the set of products it makes available at each point in time. We call this subset $S \subseteq N$ of open (available) products the firm's *offer set*. Given an offer set S, an arriving customer chooses product j in S with probability $P_j(S)$, where $P_j(S) = 0$ if $j \notin S$. $P_0(S)$ denotes the no-purchase probability, and by total probability $\sum_{j \in S} P_j(S) + P_0(S) = 1$. As a practical matter, these probabilities would most likely be derived from a parametric discrete choice model of the type discussed in §6.3, although here we consider them as generic choice probabilities. Again, one can easily allow these probabilities to vary over time, but to simplify the notation and exposition we omit this generalization in what follows.[3]

[3] For the time-inhomogeneous case, the general approach is to break the planning horizon into subintervals such that the choice probabilities and arrival rates are constant over each interval. Then one can apply the analysis developed below to each subinterval.

Figure 1 **Network and Products for the Running Example**

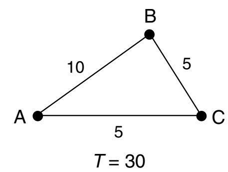

Product	Origin-destination	Class	Fare
1	A → C	H	1,200
2	A → B → C	H	800
3	A → B	H	600
4	A → C	L	800
5	A → B → C	L	500
6	A → B	L	300

The state of the network is the available capacity denoted by a vector $x = (x_1, \ldots, x_m) \geq 0$. We assume the firm is risk neutral and seeks to maximize expected revenues. The firm's decision problem, then, is to find a state-dependent policy for choosing an offer set S at each time t that maximizes their total expected revenues.

To illustrate the model and analysis, the following running example is used throughout the paper.

Running Example. An airline network consists of three cities and three flights with leg capacities as shown in Figure 1. Each flight offers high- and low-fare classes, so a total of six products are offered. The descriptions of each product are shown in Figure 1 as well. Customers are divided into three segments according to their fare-class preferences and origin-destination markets. Specifically, Segment 1 includes customers who fly from city A to B and are willing to buy in both fare classes. Segments 2 and 3 consist of customers who fly from city A to C but differ in terms of their class preferences; Segment 2 consists of business customers who will only buy class H, whereas Segment 3 is composed of leisure customers who will only buy class L. Table 1 describes precisely the characteristics of each segment. The second column of Table 1 gives the probability of an arriving customer belonging to a given segment; the third column specifies the consideration set of each segment; the last column gives the vector of preference "weights" for each product and no-purchase preference "weight" (the last component in the vector). The choice probabilities are determined from this vector of weights

Table 1 **Segments and Their Characteristics for the Running Example**

Segment	Probability	Consideration set	Preference vector
1	0.2	{3, 6}	(5, 8, 2)
2	0.3	{1, 2}	(10, 5, 5)
3	0.5	{4, 5}	(5, 10, 10)

using the multinomial logit (MNL) choice model and are discussed in §6.3, but, in essence, the probability that a customer buys a product is given by the ratio of its preference weight to the total weight of all alternatives, which includes all the open products in the customer's consideration set together with the no-purchase alternative.

Six products result in 2^6 possible offer sets. Table A1 in the appendix lists each offer set and the associated choice probabilities.

3.1. Dynamic Programming Formulation

This decision problem can be formulated as a DP. Let the value function, denoted $V_t(x)$, be defined as the optimal expected revenue obtainable from time t through to the terminal time T given that the vector of remaining capacities at time t is x. The Bellman equation is then

$$\begin{aligned} V_t(x) &= \max_{S\subseteq N}\left\{\sum_{j\in S}\lambda P_j(S)(r_j+V_{t+1}(x-A_j)) \right.\\ &\qquad\left. +(\lambda P_0(S)+1-\lambda)V_{t+1}(x)\right\}\\ &= \max_{S\subseteq N}\left\{\sum_{j\in S}\lambda P_j(S)(r_j-(V_{t+1}(x)-V_{t+1}(x-A_j)))\right\}\\ &\quad +V_{t+1}(x), \end{aligned} \tag{1}$$

where the second equation follows from the fact that $\sum_{j\in S}P_j(S)+P_0(S)=1$. We do not allow overbooking, that is, $x-A_j\geq 0$, for all $j\in S$. The boundary conditions are

$$V_t(0)=0,\quad t=1,\ldots,T,$$
$$V_{T+1}(x)=0,\quad \forall\, x\geq 0.$$

3.2. CDLP Approximation

Unfortunately, the above DP model is not solvable for most realistic networks because of the large dimensionality of the state space (e.g., typical numbers for even a modest-sized airline network are capacities on the order of 100 with $m=50$ flights, leading to 10^{100} states). Hence, as mentioned above, the only practical approach is to try to approximate the decision problem. One popular approach is to use a deterministic approximation, in which stochastic quantities are replaced by their mean (expected) values and capacity and demand are assumed to be continuous. This reduces the problem to a math program, which is normally much easier to solve. This is the approach taken by Gallego et al. (2004).

Specifically, their CDLP is formulated as follows: Let S denote the firm's offer set as before. If set S is offered and a customer arrives, a deterministic quantity $P_j(S)$ of product j is sold. (This is simply the mean value of demand for product j when S is offered.) Let $R(S)$ denote the expected revenue generated from an arriving customer when S is offered. That is,

$$R(S)=\sum_{j\in S}r_jP_j(S).$$

We will treat $R(S)$ as a deterministic revenue in the CDLP. Similarly, let $Q_i(S)$ denote the probability of using a unit of capacity on leg i, $i=1,\ldots,m$, given that we offer set S, and denote the vector of capacity consumption probabilities $Q(S)=(Q_1(S),\ldots,Q_m(S))^\top$. Then

$$Q(S)=AP(S),$$

where $P(S)=(P_1(S),\ldots,P_n(S))^\top$ is the vector of purchase probabilities. Again, for the CDLP, we will treat $Q(S)$ as a deterministic consumption rate vector. Likewise, the arrival probability λ is treated as deterministic, measuring the rate of customers arriving in each period.

Suppose the sequence of offer sets is $\{S(1),\ldots,S(T)\}$. Let the number of periods during which the subset S is offered be denoted by $t(S)$. Since demand is deterministic and the choice probabilities are time homogeneous, it is irrelevant during which periods we offer set S; rather, only the total number of periods in which S is offered matters. Indeed, any permutation of a given sequence of sets will produce the same revenue and consume the same vector of capacities. Hence, it is sufficient to view $t(S)$ as the firm's decision variables. We make one further relaxation and allow the variables $t(S)$ to be continuous as well. Effectively, this amounts to assuming we can use a set S for some fraction of a period (e.g., a value $t(S)=2.1$ would say that we offer S for two whole periods and 1/10 of another period). The objective is to find the total time $t(S)$ to offer each set S such that we maximize the firm's revenue. This leads to the following linear program:[4]

$$
\begin{aligned}
V^{\mathrm{CDLP}} = \max\ & \sum_{S\subseteq N} \lambda R(S)t(S)\\
\text{s.t.}\ & \sum_{S\subseteq N} \lambda Q(S)t(S) \le c\\
& \sum_{S\subseteq N} t(S) \le T\\
& t(S) \ge 0, \quad \forall S \subseteq N.
\end{aligned} \tag{2}
$$

A few features of this LP are worth noting. First, it has an exponential number of variables, since with n products there are 2^n possible offer sets S and, hence, 2^n corresponding variables $t(S)$. Enumerating all these variables is not feasible for practical problems. Nevertheless, as shown by Gallego et al. (2004) and as discussed further in §6 below, column-generation techniques can be used to try to solve this LP efficiently. Also, as mentioned, an optimal solution to the CDLP specifies the total time (but not sequence in which) each set should be offered. This ambiguity about the order in which sets are offered creates problems when applying the CDLP primal solution to the original stochastic problem. Again, we discuss this issue further in §§6 and 7. Finally, since there are $m+1$ constraints in the LP, there are at most $m+1$ basic variables with positive time values. Hence, even though the number of variables is astronomically large, only at most $m+1$ sets end up being used in the optimal solution. This observation again motivates the use of column-generation techniques.

For reference, the dual of the CDLP is

$$
\begin{aligned}
\min\ & \pi^\top c + T\sigma\\
\text{s.t.}\ & \lambda\pi^\top Q(S) + \sigma \ge \lambda R(S), \quad \forall S \subseteq N\\
& \pi \ge 0\\
& \sigma \ge 0,
\end{aligned} \tag{3}
$$

where π is the vector of dual prices associated with the leg capacity constraints (the first constraint in (2)), and σ is the dual price associated with the time constraint (the second constraint in (2)). Intuitively, the optimal dual prices π provide an estimate of the marginal value of capacity on each leg of the network, and the optimal dual value σ provides an estimate of the marginal value (i.e., opportunity cost) of time.

Returning to our running example, there are 64 possible offer sets, and since there are only three legs, at most four sets are offered with positive time values. The capacity on legs AB, AC, and BC is $(10, 5, 5)$. The number of periods is 30. The CDLP solution is then $t(S = \{1, 2, 3\}) = 16.35$, $t(S = \{1, 2, 3, 4\}) = 2.48$, $t(S = \{1, 2, 3, 5\}) = 10.30$, and $t(S = \{1, 2, 3, 4, 5\}) = 0.87$. That is, products 1, 2, and 3 are offered for the entire time; products 4 and 5 are offered for 3.35 and 11.17 periods, respectively; and product 6 should never be offered.

4. Asymptotic Optimality of the CDLP

One well-known property of the traditional DLP is that its solution is asymptotically optimal for the original stochastic RM problem. (See, for example, Cooper 2002.) Here, we show that this same property holds for the CDLP.

To do so, we restate the stochastic RM problem as follows: At the beginning of the selling horizon, the firm has the initial set of capacities $c = (c_1, \ldots, c_m)$. It controls the availability of products through a control policy μ, which maps states of the system to control actions (offer sets). The offer set chosen under policy μ at time t is denoted by $S_\mu(t \mid \mathcal{F}_t)$, where $\mathcal{F}_t$ denotes the history of the system up to time t. For simplicity, we abbreviate $S_\mu(t \mid \mathcal{F}_t)$ as $S_\mu(t)$ (the dependence on $\mathcal{F}_t$ being implicit from here on). The quantities of product sold during time t when policy μ is used are denoted by the n-dimensional random vector $N(S_\mu(t))$, where $N_j(S_\mu(t)) = 1$ indicates a sale of product j and $N_j(S_\mu(t)) = 0$ indicates no sale of product j.

We denote by $\mathcal{M}$ the class of all *admissible policies*. That is, those policies are nonanticipating (i.e., the control at time t depends only on the history of the process up to time t, $\mathcal{F}_t$) and satisfy

$$
\sum_{t=1}^{T} AN(S_\mu(t)) \le c \quad \text{(a.s.)},
$$

[4] This LP is formulated for the initial capacity c and the entire horizon T. When the CDLP is resolved periodically, however, the right-hand side capacity becomes the remaining capacity vector, x, and the right-hand side time becomes the remaining time t.

so they do not sell more capacity than is available. With this setup, the DP formulation (1) can be written more abstractly as the following stochastic control problem:

$$V^* = \max_{\mu \in \mathcal{M}} \; E\left[\sum_{t=1}^{T} r^\top N(S_\mu(t))\right]$$
$$\text{s.t.} \; \sum_{t=1}^{T} AN(S_\mu(t)) \le c \quad \text{(a.s.)}$$
$$S_\mu(t) \subseteq N, \quad \forall\, t = 1, \ldots, T, \tag{4}$$

where V^* denotes the optimal expected revenue over the entire time horizon when applying the optimal policy μ^*.

We first show that the optimal objective function value of the CDLP (2) provides an upper bound on the optimal expected revenue in the stochastic problem (4); that is,

Proposition 1. $V^* \le V^{\text{CDLP}}$.

Proof. Let $S_{\mu^*}(t)$, $t = 1, \ldots, T$ denote the sequence of optimal controls to the stochastic control problem (4). Then, because μ^* is an admissible policy, pathwise we must have

$$\sum_{t=1}^{T} AN(S_{\mu^*}(t)) \le c.$$

Therefore, in expectation,

$$\sum_{t=1}^{T} AE[N(S_{\mu^*}(t))] \le c.$$

Next, note that

$$\sum_{t=1}^{T} E[N_j(S_{\mu^*}(t))] = \sum_{S \subseteq N} \lambda P_j(S) t_{\mu^*}(S), \quad j = 1, \ldots, n,$$

where

$$t_{\mu^*}(S) = E\left[\sum_{t=1}^{T} 1\{S_{\mu^*}(t, \omega) = S\}\right],$$

is the expected time we offer set S under policy μ^*. (The argument ω emphasizes that the control actions depend on the sample path realization of demand.) That is, the expected quantity of sales of product j is simply the expected rate of sales of j given offer set S multiplied by the expected time each set S is used, summed over all sets S. (This follows from Wald's equation.) Thus, we have

$$\sum_{S \subseteq N} \lambda AP(S) t_{\mu^*}(S) = \sum_{S \subseteq N} \lambda Q(S) t_{\mu^*}(S) \le c,$$

where $Q(S) = AP(S)$ as before. Hence, $t_{\mu^*}(S)$, $S \subseteq N$ is a feasible solution to CDLP (2). Further, note that the expected revenue earned from the offer sets $S_{\mu^*}(t)$ in the stochastic control problem (4) is

$$V^* = \sum_{t=1}^{T} r^\top E[N(S_{\mu^*}(t))] = \sum_{S \subseteq N} \lambda r^\top P(S) t_{\mu^*}(S),$$

which is exactly the objective function value of the feasible solution $t_{\mu^*}(S)$, $S \subseteq N$ to the CDLP (2). Since the objective function from this feasible solution is no more than the optimal CDLP revenue, we therefore have $V^* \le V^{\text{CDLP}}$ as claimed. □

We next show that this upper bound from the CDLP is asymptotically tight as both capacity and demand are scaled up proportionately. Specifically, consider a sequence of problems, both stochastic and deterministic, indexed by θ, in which the initial capacities and the number of periods are increased by a factor of θ to θc and θT, respectively. We call these the *θ-scaled problems*. The case $\theta = 1$ corresponds to the original (unscaled) problems (4) and (2) above. Let V_θ^* denote the optimal expected revenue in the θ-scaled stochastic problem and V_θ^{CDLP} denote the optimal revenue in the θ-scaled CDLP problem. We will show:

Proposition 2.

$$\lim_{\theta \to \infty} \frac{1}{\theta} V_\theta^* = \lim_{\theta \to \infty} \frac{1}{\theta} V_\theta^{\text{CDLP}} = V^{\text{CDLP}}.$$

Proof. Let $t^*(S)$, $S \subseteq N$ denote the optimal solution to the unscaled CDLP problem (2). Then, it is not hard to see that $\theta t^*(S)$, $S \subseteq N$ is an optimal solution for the θ-scaled CDLP with optimal value θV^{CDLP}. Hence, the second limit above is straightforward.

To show the first limit, we construct an admissible policy μ for the θ-scaled stochastic problem as follows: Offer each set S a deterministic amount of time $\theta t^*(S)$. The order in which the sets are offered is arbitrary. Let $D(S, t)$ denote the random vector of product demand given that set S is offered for t units of time, so the vector of demand generated by offering S for $\theta t^*(S)$ units of time is $D(S, \theta t^*(S))$.

Under our policy μ, however, we will not accept all demand generated by offering S. Rather, we limit the demand accepted to the mean demand $\theta t^*(S)\lambda P(S)$. That is, the amount of demand accepted when S is offered is $N_\mu(S) = \min\{D(S, \theta t^*(S)), \theta t^*(S)\lambda P(S)\}$. Since $\theta t^*(S), S \subseteq N$ is a feasible solution for the θ-scaled CDLP, we know that

$$\sum_{S\subseteq N} \theta t^*(S)\lambda A P(S) \leq \theta c$$

and, hence, it follows that

$$\sum_{S\subseteq N} A N_\mu(S) = \sum_{S\subseteq N} A \min\{D(S, \theta t^*(S)), \theta t^*(S)\lambda P(S)\} \leq \theta c \quad \text{(a.s.)}$$

as well. Therefore, μ is an admissible policy. Moreover, the sample path revenue from policy μ is

$$\sum_{S\subseteq N} r^\top N_\mu(S) = \sum_{S\subseteq N} r^\top \min\{D(S, \theta t^*(S)), \theta t^*(S)\lambda P(S)\}.$$

Dividing both sides above by θ and letting $\theta \to \infty$, we find

$$\begin{aligned}
&\lim_{\theta\to\infty} \frac{1}{\theta} \sum_{S\subseteq N} r^\top N_\mu(S) \\
&= \lim_{\theta\to\infty} \frac{1}{\theta} \sum_{S\subseteq N} r^\top \min\{D(S, \theta t^*(S)), \theta t^*(S)\lambda P(S)\} \\
&= \lim_{\theta\to\infty} \sum_{S\subseteq N} r^\top \min\left\{\frac{1}{\theta} D(S, \theta t^*(S)), t^*(S)\lambda P(S)\right\} \\
&= \sum_{S\subseteq N} r^\top \min\{t^*(S)\lambda P(S), t^*(S)\lambda P(S)\} \\
&= \sum_{S\subseteq N} t^*(S)\lambda r^\top P(S) \\
&= V^{\text{CDLP}},
\end{aligned}$$

where the third equality above follows from the law of large numbers and the continuity of the minimum function. Thus, the scaled revenue of the policy μ converges almost surely to the upper bound V^{CDLP}. That the expected revenue of the policy μ also converges follows from the dominated convergence theorem, since the demand for each product j is trivially bounded above by θT (the total number of possible customer arrivals). This proves the result. □

This result suggests that offering sets for a deterministic amount of time as directed by the CDLP primal solution should be a good approximate policy for the stochastic choice-based RM problem when capacities and demand volumes are large. Still, such asymptotic analyses are known to be quite crude, and given that small percentage differences in performance in RM problems are economically quite significant, one must treat such performance guarantees with a fair degree of caution. In §7 below, we test the quality of the static CDLP solution and the corresponding upper bound on some numerical examples to give a more practical sense of their performance.

5. Efficient Sets

A key concept in the exact analysis of the choice-based single-leg RM model in Talluri and van Ryzin (2004a) is that of an "efficient" offer set. Intuitively, efficient sets are those offering the most favorable trade-off between expected revenue and expected capacity consumption. Specifically, as above, let $R(S)$ denote the expected revenue and $Q(S)$ (a scalar in the single-leg case) denote the probability of selling a unit (the consumption rate) using offer set S. A set T is said to be *inefficient* if a mixture of other offer sets can be used to generate more revenue for the same (or lower) consumption rate. That is, there exists a set of convex weights $\alpha(S), S \subseteq N$ satisfying $\sum_S \alpha(S) = 1$ and $\alpha(S) \geq 0, \forall S \subseteq N$ such that

$$R(T) < \sum_S \alpha(S) R(S) \tag{5}$$

$$Q(T) \geq \sum_S \alpha(S) Q(S). \tag{6}$$

If no such mixture exists, set T is said to be *efficient*. When $R(S)$ and $Q(S)$ are plotted on a scatter plot, efficient sets are those that lie on the efficient frontier in $(Q(S), R(S))$ space (refer to in Talluri and van Ryzin 2004a, Figure 1, p. 19).

As shown by Talluri and van Ryzin (2004a), efficient sets are important because they are the only ones used in the optimal solution to the single-led RM problem. Eliminating inefficient sets can reduce the complexity of the problem. Moreover, how these efficient sets are used provides interesting insights. Specifically, they can be ordered such that both $R(S)$ and $Q(S)$ are increasing, and sets are used in this order as a function of capacity and time. That is, at any point in time, with more capacity remaining, higher sets in the ordering are used; and for any given capacity, with more time remaining, lower sets in the ordering are used. This monotonicity has important implications for the optimality of nested booking limit/protection-level policies as well.

Is there an analogous notion for networks? The answer is somewhat mixed. While there is a natural extension of the definition of efficient sets to networks, it cannot be used to eliminate sets for the exact stochastic network DP in general. Still, we show next that efficient sets are the only sets used in an optimal solution to CDLP and, by our asymptotic analysis above, this can be used to argue that, asymptotically, they are the only ones used in the stochastic DP as well. This provides some evidence (albeit not irrefutable proof) that efficient sets are good ones to use. Moreover, in some special cases, like the parallel flight leg problem studied by Zhang and Cooper (2005), one can show that efficient sets are the only optimal ones to use for the exact stochastic DP.

5.1. Characterization of the Efficient Sets

We begin with a precise definition of efficient offer sets in networks, which is the natural extension of the single-leg definition. Specifically, a set T is said to be *inefficient* if there exists a set of convex weights $\alpha(S)$, $\sum_S \alpha(S) = 1$, $\alpha(S) \geq 0, \forall S \subseteq N$ such that (5) and (6) hold, where $Q(T)$ is now a vector of consumption rates. That is, a set T is inefficient if a mixture of other sets produces strictly greater expected revenue and consumes no more capacity (in expectation) on each of the m legs of the network. If no such weights exist, T is said to be *efficient*. Alternatively, efficient sets can be characterized as follows (See the appendix for a proof.)

Proposition 3. *A set T is efficient if and only if, for some $\pi = (\pi_1, \ldots, \pi_m)^\top \geq 0$, T is the optimal solution to*

$$\max_S \{R(S) - \pi^\top Q(S)\}. \tag{7}$$

In other words, to generate all efficient sets, let π range over all the nonnegative m-vectors and collect the solutions to the above problem; the union of all such solutions are the efficient sets. This result is analogous in Talluri and van Ryzin (2004a, Proposition 1).

To illustrate, for our running example the third and fourth columns of Table A1 in the appendix describe expected revenues and consumption rates on each leg; the last column in the table shows which sets are efficient. Among all 64 possible offer sets, there are 16 efficient sets.

5.2. The Optimality of Efficient Sets

What's significant about efficient sets in the network case? One important fact is that only efficient sets are used in the solution to the CDLP. Specifically,

Proposition 4. *If $t^*(T) > 0$ is the optimal solution to the CDLP* (2), *then T is an efficient set.*

Proof. Note that the reduced cost of a column S in the CDLP (2) is

$$\lambda R(S) - \lambda \pi^\top Q(S) - \sigma,$$

where π and σ are the dual variables from (3). If $t^*(T) > 0$, then set T is part of the optimal basis and has a reduced cost of zero at the optimal solution, π^* and σ^*. However, since π^* and σ^* are dual feasible, by (3) they must satisfy

$$\lambda R(S) - \lambda \pi^{*\top} Q(S) - \sigma^* \leq 0, \quad \forall S \subseteq N.$$

Therefore, we have

$$R(S) - \pi^{*\top} Q(S) \leq R(T) - \pi^{*\top} Q(T) \quad \forall S \subseteq N,$$

so T maximizes $R(S) - \pi^{*\top} Q(S)$. Hence, by Proposition 3, T must be efficient. □

Thus, efficiency is a defining characteristic of the optimal sets produced by the CDLP. Unfortunately, one cannot make the same claim for the exact DP (1). The essential reason is that the displacement costs $V_{t-1}(x) - V_{t-1}(x - A_j)$ in the inner optimization of (1) are not generally additive in the components of A_j. To see why this matters, suppose the value function was additive in the sense that for each product j, we could express the displacement cost as

$$V_{t-1}(x) - V_{t-1}(x - A_j) = \Delta V_{t-1}^\top(x) A_j, \tag{8}$$

where $\Delta V_{t-1}^\top(x)$ is an m-dimensional vector. Heuristically, one can think of this vector as being a gradient of the value function. Were this true, the inner optimization of (1) could be written

$$\max_{S \subseteq N} \{\lambda(R(S) - \Delta V_{t-1}^\top(x) Q(S))\}.$$

In this case, Proposition 3 would imply that only efficient sets would be chosen. However, in general this decomposition of the value function is not possible, and hence we cannot guarantee that only efficient sets are optimal. This phenomenon is essentially a generalization of that investigated in Talluri and van Ryzin (1999a), who showed that bid-price policies in general are not optimal under the standard linear programming model because of a similar nonadditivity of the value function.

One special case in which this additivity property holds, however, is when each product uses only a single leg, such as the parallel flight problem studied by Zhang and Cooper (2005). In this case, we can define $\Delta V_{t-1}^i(x) = V_{t-1}(x) - V_{t-1}(x - e_i)$, where e_i denotes the ith unit vector (the vector with a one in the ith component and zeros elsewhere). If each product j uses only a single leg, then each A_j is a unit vector as well and (8) always holds. Therefore, for problems like the parallel flight choice problem, one can indeed assert that only efficient sets are optimal for the exact DP.

Another fact about efficient sets is that in the same asymptotic scaling used in Proposition 2, there is no loss in optimality from restricting a policy to using only efficient sets. This is immediate from the proof of Proposition 2, where we constructed an asymptotically optimal policy that used only the optimal sets from the CDLP, which by Proposition 4 are always efficient. Hence, in this sense, it is asymptotically optimal to always use efficient sets. This provides another piece of theoretical evidence that efficient sets are good to use and again parallels the results of Talluri and van Ryzin (1999a), who showed that bid-price policies, while not optimal in general, are asymptotically optimal in the same sense under the standard independent demand model.

5.3. Partial Ordering of Efficient Sets

As we saw above, in the single-leg case there is an ordering of the efficient sets such that an efficient set with higher purchase probability has higher expected revenue. For the network choice model, however, sets are only partially ordered.

Proposition 5. *Suppose S_1 and S_2 are two efficient sets. If $Q(S_2) \geq Q(S_1)$, then $R(S_2) \geq R(S_1)$ as well.*

Proof. Since S_2 is an efficient set, it is the optimal solutions to (7) for some nonnegative vectors π. That is,

$$R(S_2) - \pi^\top Q(S_2) \geq R(S) - \pi^\top Q(S), \quad \forall S \subseteq N.$$

In particular, taking $S = S_1$, this implies

$$R(S_2) - R(S_1) \geq \pi^\top (Q(S_2) - Q(S_1)).$$

Now, if $R(S_2) < R(S_1)$, then the above implies

$$\pi^\top (Q(S_2) - Q(S_1)) < 0.$$

However, since $\pi \geq 0$, this contradicts the fact that $Q(S_2) - Q(S_1) \geq 0$, so we must have $R(S_2) \geq R(S_1)$ as claimed. □

This says that the efficient sets are partially ordered, in the sense that a set S_2 that produces a vector of purchase probabilities higher than a given set S_1 in all components also must have higher expected revenues. Although this ordering is essential to understanding the way in which optimal sets are used in the single-leg model, the partial ordering in the network case appears (so far, to us) to be considerably less useful. Still, the generalization is worth noting.

6. Computation and Decomposition Approximations

In this section we look at how to solve the CDLP efficiently and how to use its outputs to construct a more accurate decomposition approximation of the stochastic DP (1). Both are important steps in making the CDLP a viable model in practice.

6.1. Solving CDLP by Column Generation

As already noted, the CDLP (2) is exponentially large; it has 2^n primal variables, corresponding to all possible subsets of the set of network products N. For even modest-sized networks, this makes direct application of the model impractical. However, as Gallego et al. (2004) point out, one can use column-generation techniques to attempt to overcome this complexity. Roughly speaking, we start with a limited number of columns (subsets) and solve a reduced LP using only these columns. Using the resulting dual solution, we then check to see if any columns left out of the problem have a positive reduced cost relative to these dual prices. If so, a positive reduced-cost column is added and the LP is resolved. If there are no such positive reduced-cost columns, then the current solution is optimal. In this manner, columns are generated as needed as the problem is solved; the hope is that only a modest number of columns needs to be generated before we reach optimality.

Consider the primal LP (2). The reduced primal problem is identical except that we initially consider only a limited number of subsets (columns), denoted $\mathcal{N} = \{S_1, S_2, \ldots, S_k\}$. Therefore, the reduced problem is

$$
\begin{aligned}
V^{\mathrm{CDLP}}(\mathcal{N}) = \max \quad & \sum_{S\in\mathcal{N}} \lambda R(S) t(S) \\
\text{s.t.} \quad & \sum_{S\in\mathcal{N}} \lambda Q(S) t(S) \le c \\
& \sum_{S\in\mathcal{N}} t(S) \le T \\
& t(S) \ge 0, \quad \forall\, S \in \mathcal{N}. \qquad (9)
\end{aligned}
$$

Let π and σ be the dual prices for the first and second constraints, respectively, of this reduced problem. We want to check to see if these values are feasible for the problem (3). To do so, we must determine if any sets S not contained in our collection $\mathcal{N}$ have a positive reduced cost. This is achieved by solving the following column-generation subproblem:

$$
\begin{aligned}
&\max_{S\subseteq N} \{\lambda R(S) - \lambda \pi^\top Q(S) - \sigma\} \\
&\quad = \lambda \max_{S\subseteq N} \{R(S) - \pi^\top Q(S)\} - \sigma. \qquad (10)
\end{aligned}
$$

If the optimal value for this problem is nonpositive, then π and σ are dual feasible, and our current solution to the reduced problem is in fact optimal for the original problem (2). Otherwise, the optimal solution S^* corresponds to a subset with positive reduced cost, i.e., $\lambda R(S^*) - \lambda\pi^\top Q(S^*) - \sigma > 0$. Therefore, if we add S^* to the collection of columns $\mathcal{N}$ and resolve (9), the optimal objective function value will increase. Along with the new solution, we get a new set of dual prices, which we then use to check again for another positive reduced-cost column, etc. Again, the hope here is that a relatively modest number of columns are generated before an optimal solution is reached, although in the worst case it is possible that an exponential number of columns are required.

The main practical difficulty in this approach is solving the column-generation subproblem (10). In general, this could be an NP-hard problem (see Bront et al. 2006 for detailed discussion), although as we show in the following, some special cases can be solved quite efficiently. Even if it is an NP-hard problem, one can attempt to solve it approximately to identify a positive reduced-cost column (if not the most positive reduced-cost column). Still, determining which classes of choice models lead to efficient column-generation procedures and how to approximately solve the subproblem in complex cases are largely open questions worthy of additional research. (See, for example, Bront et al. 2006.)

6.2. Decomposition Approximation Method

How should one use the output of the CDLP model in actual applications? One approach, of course, is to apply the primal solution directly. This would involve offering a collection of (at most) $m+1$ subsets, each for a fixed amount of time, as given by the solution $t^*(S)$, $S \subseteq N$. However, this approach has a few problems. First, as mentioned, the order in which the sets are offered is not specified, although one can try various heuristic approaches to ordering the offer sets. The main difficulty, though, is that the solution is static and does not adjust to changes in demand and remaining time and capacity. One can of course attempt to resolve the model frequently, but, as previously suggested, this is likely to be quite computationally complex. Also, one is still left with the problem of deciding which of the $m+1$ subsets from the basic solution to use at each point in time once the model is resolved.

In this section, we develop an alternative approach for using the CDLP solution that overcomes these problems. It produces a dynamic policy that prescribes a unique offer set to use as a function of the vector of remaining capacities and remaining time. In this way, the policy dynamically adjusts the offer set to changing network conditions. Moreover, the approach does not rely on frequently reoptimizing the CDLP and can even be applied by only solving the CDLP once up-front. The approach is motivated by the decomposition ideas used in traditional network RM and can be viewed as the choice-based equivalent of methods such as DAVN and DP decomposition (see Talluri and van Ryzin 2004b, Chapter 3, for a discussion of DAVN and DP decomposition).

The main idea of the approach is to decompose the network DP (1) into a collection of leg-level DPs, each of which is only one dimensional and therefore easy to solve exactly. To do this, the decomposition uses the optimal solution to (3) to approximate the marginal value of capacity elsewhere in the network. Again, this idea is exactly analogous to using dual prices from the DLP to compute displacement-adjusted revenues in the traditional network RM case.

Specifically, let $\pi^* = (\pi_1^*, \ldots, \pi_m^*)^\top$ denote the optimal dual prices from (3) corresponding to the m leg capacity constraints. Consider approximating the problem at a given leg i. We approximate the network value function at this leg by

$$V_t(x) \approx V_t^i(x_i) + \sum_{l \neq i} \pi_l^* x_l, \tag{11}$$

where $V_t^i(x_i)$ is a dynamic (time-dependent) and nonlinear approximation of the value of the capacity on leg i and $\pi_l^* x_l$ are static (time-independent) and linear approximations of the value of capacity elsewhere in the network. For a given product j that uses leg i ($i \in A_j$), (11) yields the following estimate of the opportunity cost of selling product j:

$$V_t(x) - V_t(x - A_j) \approx \Delta V_t^i(x_i) + \sum_{l \in A_j, l \neq i} \pi_l^*,$$

where $\Delta V_t^i(x_i) = V_t^i(x_i) - V_t^i(x_i - 1)$. Then, using the approximation (11) in the DP recursion (1) we obtain

$$V_t^i(x_i) = \max_{S \subseteq N} \left\{ \sum_{j \in S} \lambda P_j(S) \left(r_j - \sum_{l \in A_j\, l \neq i} \pi_l^* - \Delta V_{t+1}^i(x_i) \right) \right\} + V_{t+1}^i(x_i), \tag{12}$$

with boundary conditions

$$V_{T+1}^i(x_i) = 0, \quad \forall x_i \geq 0,$$
$$V_t^i(0) = 0, \quad \forall t = 1, \ldots, T.$$

This is now a one-dimensional DP.[5] It can be solved efficiently, provided that the inner problem of optimizing over subsets at each stage is not too computationally complex. Note that this approximation effectively accounts for the network effects by replacing each product revenue r_j by the *pseudorevenue* (also called the *displacement-adjusted revenue* $r_j - \sum_{l \in A_j\, l \neq i} \pi_l^*$). This is precisely how DAVN and DP approximations are formed in traditional RM. Also note that the inner optimization problem is of exactly the same form as the column-generation problem (10), and so again its complexity will depend on the choice model. (This is discussed further below.)

We repeat this approximation at each leg i of the network, yielding a set of one-dimensional value functions $V_t^i(x_i)$, $i = 1, \ldots, m$. These leg-level value functions can then be combined to form a dynamic approximation of the network value function, denoted $\hat{V}_t(x)$; that is,

$$V_t(x) \approx \hat{V}_t(x) \equiv \sum_{i=1}^{m} V_t^i(x_i).$$

Given this additive approximation, we then select a set dynamically at each time by solving

$$\max_{S \subseteq N} \left\{ \sum_{j \in S} \lambda P_j(S)(r_j - \Delta \hat{V}_{t+1}^\top(x) A_j) \right\}, \tag{13}$$

where $\Delta \hat{V}_{t+1}(x) = (\Delta V_{t+1}^1(x_1), \ldots, \Delta V_{t+1}^m(x_m))^\top$ is the vector of approximate displacement costs. Again, this problem is of the same form as the column-generation subproblem (10). We give a specific example next.

6.3. Application to the MNL Choice Model with Disjoint Segments

Consider the special case of the MNL choice model. (See Ben-Akiva and Lehrman 1985 for a description of the MNL model.) We assume each customer is interested only in a subset of the entire product set. We call this set the customer's *consideration set*. Customers choose among the available products in their consideration set according to the MNL model. We assume there are L segments of customers and that each segment l has a distinct consideration set C_l. Let $s_l = |C_l|$ denote the number of products in consideration set C_l. We further assume that the consideration sets are disjoint, so that $C_l \cap C_k = \varnothing$ if $l \neq k$. Let λ_l denote the probability that an arrival is from segment l, so $\sum_{l=1}^{L} \lambda_l = \lambda$ is the total arrival probability. We call this the *MNL with disjoint consideration sets* model. We focus on this model because of its computational efficiency. In particular, under the MNL model with disjoint segments, the column-generation subproblem (10) (and hence the various inner optimization problems in the decomposition heuristic) can be separated by segment, and each segment subproblem can be solved by a simple ranking procedure.

[5] Here, by one-dimensional, we are referring to the capacity dimension only; time is always an implicit dimension of the DP.

(See §6.3.1 for a detailed discussion.) However, we emphasize that our decomposition heuristic applies to general choice models; the only difference is the complexity of solving the combinatorial subset selection subproblems. Indeed, Bront et al. (2006) have recently extended our decomposition heuristic to a MNL model with overlapping segments by exploiting integer programming techniques to solve the combinatorial subproblems.

For this model, it is convenient to define a binary vector y_l that indicates which products in consideration set C_l are offered. Each component y_{lj} of y_l corresponds to the availability of product j and is defined by

$$y_{lj} = \begin{cases} 1 & \text{if product } j \text{ is available, } j \in C_l; \\ 0 & \text{if product } j \text{ is unvailable, } j \in C_l. \end{cases}$$

Similarly, we let $P_{lj}(y_l)$ denote the probability of a sale of product j to a customer of segment l given y_l. Under the MNL choice model, the choice probability is defined by a preference vector, denoted v_l, that indicates the preference "weight" for each product contained in C_l. This vector, together with the no-purchase preference value, v_{l0}, determines a customer's choice probabilities as follows:

$$P_{lj}(y_l) = \frac{v_{lj} y_{lj}}{\sum_{i \in C_l} v_{li} y_{li} + v_{l0}}.$$

The following proposition shows why the MNL model is computationally efficient. This same result is shown by Gallego et al. (2004), although for completeness we provide an alternative proof of it in the appendix.

Proposition 6. *Consider the optimization problem*

$$\max_{y_l \in \{0,1\}^{s_l}} \frac{\sum_{j \in C_l} w_j v_{lj} y_{lj}}{\sum_{j \in C_l} v_{lj} y_{lj} + v_{l0}}. \tag{14}$$

Rank the values w_j in a decreasing order; that is, $w_{[1]} \geq \cdots \geq w_{[i]} \geq \cdots \geq w_{[s_l]}$. Then there is a critical value k^*, $1 \leq k^* \leq s_l$, such that the optimal solution to the above problem is given by

$$y^*_{lj} = \begin{cases} 1 & \text{if } w_j \geq w_{[k^*]}, \\ 0 & \text{if } w_j < w_{[k^*]}. \end{cases}$$

This property is useful because, again, both the column-generation subproblem (10) in the CDLP and the set selection problem (13) in the decomposition heuristic are of this form. Proposition 6 shows that both problems can be solved by an efficient ranking procedure. We next discuss each problem in turn.

6.3.1. Solving the CDLP for the MNL with Disjoint Segments. Under the MNL with disjoint segments model, the column-generation subproblem (10) separates by segment; that is, each segment l chooses products from its consideration set C_l, and hence (10) reduces to

$$\max_{S \subseteq N} \{\lambda R(S) - \lambda \pi^\top Q(S)\} = \sum_{l=1}^{L} \lambda_l \max_{y_l \in \{0,1\}^{s_l}} \frac{\sum_{j \in C_l} (r_j - \pi^\top A_j) v_{lj} y_{lj}}{\sum_{j \in C_l} v_{lj} y_{lj} + v_{l0}}.$$

Each term on the right-hand side above is of precisely the form required by Proposition 6, where $w_j = r_j - \pi^\top A_j$. Therefore, we can find the optimal solution by simply ranking the products in each consideration set by their displacement-adjusted revenue values $r_j - \pi^\top A_j$. By Proposition 6, then, the optimal offer set consists of the k^* highest-ranked products for some k^*. Therefore, the optimization problem reduces to checking at most $s_l = |C_l|$ possibilities. Hence, this subproblem is quite efficient to solve.

6.3.2. Decomposition Heuristic for MNL with Disjoint Segments. The decomposition heuristic under the MNL with disjoint segments model is also efficient computationally. Indeed, the leg-level DPs (12) become

$$V_t^i(x_i) = \sum_{l=1}^{L} \lambda_l \max_{y_l \in \{0,1\}^{s_l}} \left\{ \left(\sum_{j \in C_l} \left(r_j - \Delta V_{t+1}^i(x_i) - \sum_{h \in A_j, h \neq i} \pi_h^* \right) v_{lj} y_{lj} \right) \cdot \left(\sum_{j \in C_l} v_{lj} y_{lj} + v_{l0} \right)^{-1} \right\} + V_{t+1}^i(x_i)$$

with the boundary conditions

$$V_{T+1}^i(x_i) = 0, \quad \forall x_i \geq 0,$$
$$V_t^i(0) = 0, \quad \forall t = 1, \ldots, T.$$

Table 2 **Dynamic Offer Sets for the Running Example at Period** $t = 20$

Remaining capacity	Offer set	Remaining capacity	Offer set
(4, 4, 4)	{1, 2, 3, 4, 5}	(1, 3, 2)	{1, 4}
(4, 3, 3)	{1, 2, 3, 4, 5}	(3, 1, 1)	{3}
(3, 3, 4)	{1, 2, 3, 4, 5}	(1, 3, 1)	{1, 4}
(3, 4, 3)	{1, 3, 4}	(2, 2, 2)	{1, 3}
(4, 3, 2)	{1, 3, 4}	(3, 0, 1)	{3}
(4, 2, 3)	{1, 2, 3, 5}	(1, 3, 0)	{1, 4}
(3, 3, 3)	{1, 3, 4}	(0, 3, 1)	{1, 4}
(3, 2, 2)	{1, 3}	(3, 1, 0)	{3}
(2, 3, 3)	{1, 3, 4}	(2, 2, 1)	{1, 3}
(2, 2, 3)	{1, 3}	(1, 2, 2)	{1}
(2, 3, 2)	{1, 3, 4}	(2, 1, 2)	{3}
(3, 2, 1)	{1, 3}	(0, 2, 1)	{1}
(3, 1, 2)	{2, 3}	(2, 0, 1)	{3}
(2, 3, 1)	{1, 3, 4}	(1, 2, 0)	{1}

Again, the inner optimization above is of the form required by Proposition 6, where $w_j = r_j - \Delta V_{t+1}^i(x_i) - \sum_{h\in A_j\, h\neq i} \pi_h^*$. Hence, the optimization at each stage in this DP can be solved by simply ranking products by these weights and checking which cutoff value k^* yields the highest objective function value. Note that the weights again have the interpretation of displacement-adjusted revenues—the difference is that leg i has a dynamic displacement cost of $\Delta V_{t+1}^i(x_i)$, whereas all other legs have static displacement costs given by the dual prices π.

Once these individual leg-level DPs are solved, sets are selected dynamically in real time by solving (13). Again, this is an easy problem because (13) is also of the form required by Proposition 6; hence, a simple ranking procedure can be used to identify the optimal offer set.

Returning to our running example, we illustrate how the optimal offer sets are dynamically generated based on the CDLP outputs: The dual solution obtained with the starting leg capacities is $(0, 800, 500)$. Substituting this into (12), we compute each leg-level value function. The optimal offer sets at each time t are then determined by solving (13). Table 2 shows the resulting offer sets for a range of the residual capacities at time $t = 20$.

7. Numerical Examples

To test the relative performance of our decomposition heuristic, we conducted numerical experiments using two example networks. One is a collection of parallel flights; the other is a small hub-and-spoke network. In each case, customers were assumed to follow the behavior of the MNL with disjoint segments model. The parameters of the MNL choice models were varied to simulate different degrees of customer willingness to switch among alternatives in the consideration sets. In addition, different load factors were simulated by changing the network capacities.

The following policies were tested:

INDEP. This policy implements the DP decomposition policy based on the independent demand model. We first solved the DLP, where mean demand was computed assuming that all products are offered. That is, the probability of a request for product j was taken as $\lambda_j = \lambda P_j(N)$. The mean demand for product j in the DLP is then $T\lambda_j$. Once the DLP model was solved, we used the dual prices in a dynamic programming decomposition scheme, as described in Talluri and van Ryzin (2004b, Chapter 3). This scheme is in fact equivalent to the decomposition scheme of §6.2, but with $P_j(S) = \lambda_j$ for all j and S (the independent demand model assumption). This policy serves as a benchmark to evaluate the benefits of explicitly using a choice-based modeling in the optimization procedure.

CDLP. This policy implements the static CDLP solution. The optimal primal solution to the deterministic LP (2) gives the total time to offer each set. Because the sequence in which the sets are offered is ambiguous, we considered two ways to construct this sequence. The first is that sets were offered according to their indexes in the solutions to (2), that is, sets were indexed as they were generated and then simply offered according to this index order. In the second approach, we randomly generated the sequence to offer.[6] Our tests showed that neither approach dominated the other. Hence, we only report results for the first method in the tables below.

[6] To make these two procedures more precise, in our running example the efficient sets generated (in order) were $S_1 = \{1, 2, 3\}$, $S_2 = \{1, 2, 3, 4\}$, $S_3 = \{1, 2, 3, 5\}$, and $S_4 = \{1, 2, 3, 4, 5\}$ with associated CDLP solution $t_1^* = 16.35$, $t_2^* = 2.48$, $t_3^* = 10.30$, and $t_4^* = 0.87$. Hence, in the first approach, S_1 is offered for $t_1^* = 16.35$ units of time, then S_2 for $t_2^* = 2.48$ units of time, etc. In the second approach, we randomly permuted the indices and obtained the sequence $\{3, 2, 1, 4\}$; hence, S_3 was first offered for $t_3^* = 10.30$ units of time, then S_2 was offered for $t_2^* = 2.48$ units of time, etc.

DCOMP. This is our choice-based decomposition heuristic, as described in §6.2.

We also computed an upper bound on the optimal expected revenue based on the CDLP value (2). We used this upper-bound to bound the suboptimality gap of our DCOMP method. In the tables below, we denote the upper-bound revenue by UB REV, and the revenues produced by the DCOMP, CDLP, and INDEP methods, by DCOMP REV, CDLP REV, and INDEP REV, respectively. We denote the percentage gap between UB and DCOMP by %OPT-GAP. The revenue improvements of DCOMP relative to CDLP and INDEP are denoted %DCOMP-CDLP and %DCOMP-INDEP, respectively.

We also tested the effect of reoptimizing the policies periodically throughout the simulated booking process. In the no-reoptimization case, the deterministic LPs are solved only once, and the dual variables associated with the capacity constraints are used to estimate the marginal capacity for the entire time horizon. In the reoptimization procedure, we divided the total time horizon into five equal-sized periods and resolved the problem at the beginning of each period, using the remaining time, capacity, and estimated demand-to-come as input parameters to the policies.

Simulations were run under various load factors by scaling all leg capacities by a common factor α. Five different load factors were tested. We also simulated various degrees of choice behavior by varying the no-purchase preference value vector v_0. A zero no-purchase value vector describes a perfect substitution case in which customers are perfectly willing to substitute among the products in their consideration sets. As v_0 increases, the probability of customers not purchasing rather than substituting increases. When v_0 is very large, customer behavior approaches the independent demand model.

In the parallel flight example, the number of simulation runs was fixed at 20,000, whereas in the small network, the number of runs was fixed at 3,000. We used $T = 300$ periods in the parallel flight example and $T = 1{,}000$ periods in the small network example. With these simulation sample sizes, the revenue estimates had relative errors in the range 0.1%–0.6% with 99% confidence.

7.1. Simulation Example 1: Parallel Flights

This set of simulations is based on a network with three parallel flights with two fare classes on each flight (high (H) and low (L)), producing six products in total. We assumed there were two customer segments corresponding to the two fare classes (H and L). That is, we assumed customers were willing to choose among flight times within a class, but were not willing to purchase up or down from their preferred fare classes. This is essentially the model studied by Zhang and Cooper (2005).

The arrival probabilities for each segment were $(\lambda_H, \lambda_L) = (0.2, 0.3)$. The initial capacities for each flight were $c = (30, 50, 40)$. For class H, the fares for each flight were $r_H = (800, 1{,}000, 600)$ and the preference vector was $v_H = (5, 10, 1)$. For class L, the revenue values for each flight were $r_L = (400, 500, 300)$ and the preference vector was $v_L = (5, 1, 10)$. We simulated different capacity scaling factors α ranging from 0.6 to 1.4 and for various no-purchase vectors $v_0 = (v_{oH}, v_{0L})$, where v_{0H} (respectively, v_{0L}) is again interpreted as the "weight" that class H (respectively, L) places on the no-purchase option. The simulation results with and without reoptimization are summarized in Tables 3 and 4, respectively.

Note that the DCOMP heuristic produces consistent revenue gains over the CDLP policy, especially in the tightly constrained ($\alpha = 0.6$ and $\alpha = 0.8$) capacity cases. The gains in these cases are on the order of 1%–5%, which is quite significant. Also note the extremely large gains (up to 13% improvement in revenue) relative to the independent demand model policy (INDEP). The improvements are largest in the cases where $v_0 = (0, 0)$, that is, when customers are perfectly willing to substitute among the products in their consideration sets. This is intuitive and illustrates the potential improvement in revenue from explicitly considering choice behavior when customers are willing to substitute.

Table 4 shows that the improvements from using DCOMP over CDLP are not as large when the parameters of the policies are reoptimized. However, the gains are still significant, especially in the high load factor case ($\alpha = 0.6$), where improvements are on the order of 1%–2%.

Table 5 reports each policy's gain in revenue from reoptimization. It is clear from this table that the

Table 3 Simulation Results for Parallel Flights Without Reoptimization

α	v_0	UB REV	DCOMP REV	CDLP REV	INDEP REV	%OPT-GAP	%DCOMP-CDLP	%DCOMP-INDEP
0.6	(0, 0)	55,200	53,356	51,051	51,841	−3.34	4.51	2.92
	(1, 5)	53,400	51,866	49,186	51,264	−2.87	5.45	1.17
	(5, 10)	50,400	48,396	46,476	48,023	−3.98	4.13	0.78
	(10, 20)	45,139	43,132	41,673	42,835	−4.45	3.50	0.69
0.8	(0, 0)	67,200	64,626	63,066	64,393	−3.83	2.47	0.36
	(1, 5)	65,600	63,189	61,139	62,461	−3.67	3.35	1.17
	(5, 10)	59,446	57,122	55,572	55,231	−3.91	2.79	3.42
	(10, 20)	47,431	46,621	46,137	46,316	−1.71	1.05	0.66
1.0	(0, 0)	78,000	75,176	74,141	72,129	−3.62	1.40	4.22
	(1, 5)	76,000	73,622	71,512	68,793	−3.13	2.95	7.02
	(5, 10)	60,731	60,222	59,293	58,908	−0.84	1.57	2.23
	(10, 20)	47,442	47,339	47,251	47,248	−0.22	0.19	0.19
1.2	(0, 0)	88,800	87,082	84,511	77,193	−1.93	3.04	12.81
	(1, 5)	78,117	77,534	75,750	72,574	−0.75	2.36	6.83
	(5, 10)	61,039	60,845	60,483	60,361	−0.32	0.60	0.80
	(10, 20)	47,442	47,435	47,401	47,372	−0.02	0.07	0.13
1.4	(0, 0)	93,200	92,762	89,498	81,533	−0.47	3.65	13.77
	(1, 5)	78,117	78,038	77,882	74,392	−0.10	0.20	4.90
	(5, 10)	61,038	60,993	60,916	60,660	−0.07	0.13	0.55
	(10, 20)	47,442	47,441	47,408	47,380	−0.00	0.07	0.13

CDLP policy improves significantly when reoptimized periodically; the DCOMP policy performance, in contrast, does not change significantly when it is reoptimized. The DCOMP method, therefore, appears to be more robust in the sense that its performance is less sensitive to the frequency of reoptimization. This behavior is likely to be an advantage in practical applications because the CDLP is quite computationally complex to solve, and, therefore, one would like to avoid frequent reoptimizations if possible. Notice

Table 4 Simulation Results for Parallel Flights with Reoptimization

α	v_0	UB REV	DCOMP REV	CDLP REV	INDEP REV	%OPT-GAP	%DCOMP-CDLP	%DCOMP-INDEP
0.6	(0, 0)	55,200	53,555	52,925	51,831	−2.98	1.19	3.33
	(1, 5)	53,400	52,288	51,305	51,267	−2.08	1.92	1.99
	(5, 10)	50,400	48,584	47,827	48,019	−3.60	1.58	1.18
	(10, 20)	45,139	43,283	42,870	42,810	−4.11	0.96	1.11
0.8	(0, 0)	67,200	64,855	63,185	64,394	−3.49	2.64	0.72
	(1, 5)	65,600	64,079	63,333	62,459	−2.32	1.18	2.59
	(5, 10)	59,446	57,231	56,847	55,233	−3.73	0.68	3.62
	(10, 20)	47,431	46,588	46,525	46,277	−1.78	0.14	0.67
1.0	(0, 0)	78,000	76,195	76,021	72,132	−2.31	0.23	5.63
	(1, 5)	76,000	73,738	73,118	68,767	−2.98	0.85	7.23
	(5, 10)	60,731	60,235	60,056	58,901	−0.82	0.30	2.26
	(10, 20)	47,442	47,302	47,321	47,233	−0.29	−0.04	0.15
1.2	(0, 0)	88,800	87,203	86,139	77,213	−1.80	1.24	12.94
	(1, 5)	78,117	77,510	77,092	72,553	−0.78	0.54	6.83
	(5, 10)	61,038	60,840	60,765	60,357	−0.32	0.12	0.80
	(10, 20)	47,442	47,403	47,440	47,367	−0.08	−0.08	0.08
1.4	(0, 0)	93,200	92,769	91,451	81,506	−0.46	1.44	13.82
	(1, 5)	78,117	78,008	77,941	74,410	−0.14	0.09	4.84
	(5, 10)	61,039	60,993	60,977	60,666	−0.07	0.03	0.54
	(10, 20)	47,442	47,408	47,447	47,373	−0.07	−0.08	0.07

Table 5 Percentage Gains in Revenue from Reoptimizing for Parallel Flights

α	v_0	% Gain-DCOMP	% Gain-CDLP	% Gain-INDEP
0.6	(0, 0)	0.37	3.67	−0.02
	(1, 5)	0.81	4.31	0.01
	(5, 10)	0.39	2.91	−0.01
	(10, 20)	0.35	2.87	−0.06
0.8	(0, 0)	0.35	0.19	0.00
	(1, 5)	1.41	3.59	−0.00
	(5, 10)	0.19	2.29	0.00
	(10, 20)	−0.07	0.84	−0.08
1.0	(0, 0)	1.36	2.54	0.00
	(1, 5)	0.16	2.25	−0.04
	(5, 10)	0.02	1.29	−0.01
	(10, 20)	−0.08	0.15	−0.03
1.2	(0, 0)	0.14	1.93	0.03
	(1, 5)	−0.03	1.77	−0.03
	(5, 10)	−0.01	0.47	−0.01
	(10, 20)	−0.07	0.08	−0.01
1.4	(0, 0)	0.01	2.18	−0.03
	(1, 5)	−0.04	0.08	0.02
	(5, 10)	0	0.10	0.01
	(10, 20)	−0.07	0.08	−0.01

that the INDEP method does not always benefit from reoptimization; in some cases, the expected revenue when reoptimizing is a bit worse than without reoptimizing. This phenomenon also occurs under the traditional independent demand model. (See Cooper 2002, Secomandi 2005 for detailed discussions.)

7.2. Simulation Example 2: Small Network

This example considers a small airline network consisting of seven flight legs, with two fare classes on each leg and a total of 22 products (including local and connecting itineraries). Figure 2 shows the network and Table 6 describes the products. The product set is segmented into 10 disjoint consideration sets, corresponding to 10 customer segments. Segments are defined in terms of their fare class preference (business or leisure) and their origin-destination market. Table 7 describes the 10 customer segments and their consideration sets and preferences. The simulation results for this network are shown in Tables 8 and 9. The first component of v_0 is the no-purchase "weight" for business customers (segments 1, 3, 5, 7, 9), and the second component is the no-purchase "weight" for leisure customers (segments 2, 4, 6, 8, 10). Again, as these values increase it makes the corresponding segment of customers less likely to substitute.

Table 6 Product Descriptions and Revenues

Product	Legs	Class	Revenue	Products	Legs	Class	Revenue
1	1	1	1,000	12	1	2	500
2	2	1	400	13	2	2	200
3	3	1	400	14	3	2	200
4	4	1	300	15	4	2	150
5	5	1	300	16	5	2	150
6	6	1	500	17	6	2	250
7	7	1	500	18	7	2	250
8	{2, 4}	1	600	19	{2, 4}	2	300
9	{3, 5}	1	600	20	{3, 5}	2	300
10	{2, 6}	1	700	21	{2, 6}	2	350
11	{3, 7}	1	700	22	{3, 7}	2	350

Without reoptimization, the DCOMP heuristic generally achieves a significant gain over the CDLP solution, although there are a few outlier cases where it does worse. We discuss the reason for this behavior

Figure 2 Two-Fare-Class Airline Network

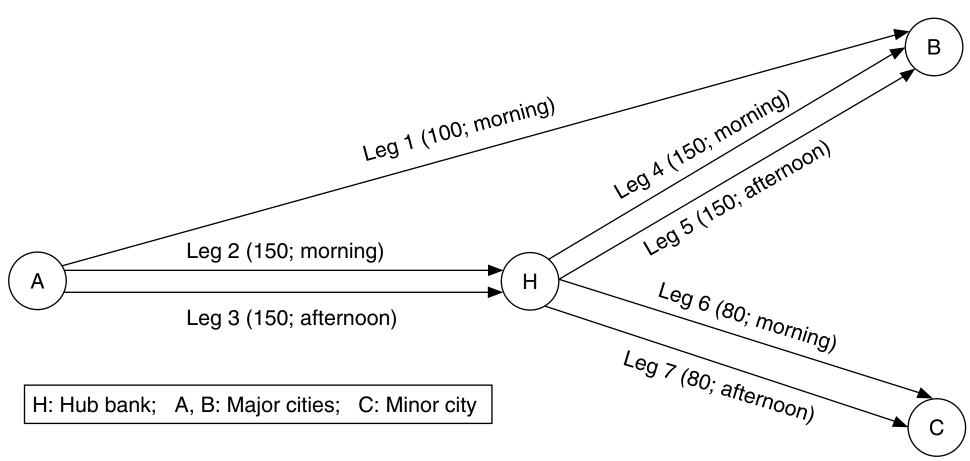

Table 7 Segments and Their Characteristics

Segment	Probability	Description	Consideration set	Preference value
1	0.08	AB Business	{1, 8, 9}	10, 5, 5
2	0.20	AB Leisure	{19, 20, 12}	10, 10, 5
3	0.05	AH Business	{2, 3}	10, 10
4	0.20	AH Leisure	{13, 14}	10, 10
5	0.10	HB Business	{4, 5}	10, 10
6	0.15	HB Leisure	{15, 16}	10, 5
7	0.02	HC Business	{6, 7}	10, 5
8	0.05	HC Leisure	{17, 18}	10, 10
9	0.02	AC Business	{10, 11}	10, 5
10	0.04	AC Leisure	{21, 22}	10, 10

below, but it is essentially due to the existence of multiple dual solutions to CDLP. When the policies are reoptimized, the relative improvement of the DCOMP method over the CDLP is smaller, but still significant, especially in the cases of tight capacity ($\alpha = 0.6$ and $\alpha = 0.8$). Again, the gains over the INDEP policy are very large in general. However, in one case ($\alpha = 0.6$, $v_0 = (0, 0)$), the INDEP method produced a larger revenue than the DCOMP method. This behavior, however, can again be attributed to the existence of multiple dual solutions in CDLP, as discussed below.

As in the previous example, there is little benefit to reoptimizing our DCOMP procedure, although reoptimizing improves the CDLP method significantly. This is illustrated in Table 10. Again, this illustrates the robustness and computational advantages of the DCOMP method.

7.3. Multiple Dual Solutions to CDLP

As noted, our DCOMP approach does worse than the CDLP method in some extreme cases—for instance, Example 2 in the cases $\alpha = 0.6$, $v_0 = (0, 0)$, and $\alpha = 1.2$, $v_0 = (0, 0)$ without reoptimization. This is due to the existence of multiple dual solutions to the CDLP (2). Multiple dual solutions result in different estimates of marginal capacities in the DCOMP policy. In our column-generation procedure, we chose the first product in each segment's consideration set as the initial offer set of the reduced LP. In the extreme case when $v_0 = (0, 0)$, we observed multiple dual solutions, and the particular choice of offer sets used to initialize the column-generation procedure has a significant impact on the dual prices produced, which in turn impacts the DCOMP policy.

To illustrate this effect, consider the case of $\alpha = 1.2$ and $v_0 = (0, 0)$ in Example 2 without reoptimization.

Table 8 Simulation Results for a Small Network Without Reoptimization

α	v_0	UB REV	DCOMP REV	CDLP REV	INDEP REV	%OPT-GAP	%DCOMP-CDLP	%DCOMP-INDEP
0.6	(0, 0)	186,400	172,818	176,354	178,290	−7.29	−2.00	−3.07
	(1, 5)	181,835	179,385	171,280	173,571	−1.35	4.73	3.35
	(5, 10)	166,017	163,643	157,621	160,035	−1.43	3.82	2.26
	(10, 20)	149,798	146,630	142,711	142,317	−2.11	2.75	3.03
0.8	(0, 0)	227,200	221,834	216,994	212,027	−2.36	2.23	4.63
	(1, 5)	216,062	213,813	207,014	202,771	−1.04	3.28	5.45
	(5, 10)	194,500	192,152	186,535	181,935	−1.21	3.01	5.62
	(10, 20)	165,560	163,900	162,599	160,055	−1.00	0.80	2.40
1.0	(0, 0)	256,000	252,135	246,612	235,550	−1.51	2.24	7.04
	(1, 5)	244,110	241,308	235,028	221,803	−1.15	2.67	8.79
	(5, 10)	213,833	212,413	210,266	199,603	−0.66	1.02	6.42
	(10, 20)	171,071	170,696	170,073	170,300	−0.22	0.37	0.23
1.2	(0, 0)	284,000	272,051	274,436	253,963	−4.21	−0.87	7.12
	(1, 5)	267,429	264,421	260,775	239,212	−1.12	1.40	10.54
	(5, 10)	217,738	217,722	217,452	212,920	−0.01	0.12	2.26
	(10, 20)	171,071	171,008	170,983	171,047	−0.04	0.01	−0.02
1.4	(0, 0)	309,000	306,862	300,673	271,258	−0.69	2.06	13.13
	(1, 5)	269,588	269,458	268,883	253,337	−0.05	0.21	6.36
	(5, 10)	217,738	217,731	217,729	215,630	−0.00	0.00	0.97
	(10, 20)	171,071	171,008	170,983	171,047	−0.04	0.01	−0.02

Table 9 **Simulation Results for a Small Network with Reoptimization**

α	v_0	UB REV	DCOMP REV	CDLP REV	INDEP REV	%OPT-GAP	%DCOMP-CDLP	%DCOMP-INDEP
0.6	(0, 0)	186,400	181,450	181,526	178,673	−2.66	−0.04	1.55
	(1, 5)	181,835	179,408	177,098	173,989	−1.33	1.30	3.11
	(5, 10)	166,017	163,679	161,833	160,444	−1.41	1.14	2.02
	(10, 20)	149,798	146,964	146,183	142,226	−1.89	0.53	3.33
0.8	(0, 0)	227,200	221,929	221,050	212,511	−2.32	0.40	4.43
	(1, 5)	216,062	213,836	211,590	202,577	−1.03	1.06	5.56
	(5, 10)	194,500	192,307	190,867	181,571	−1.13	0.75	5.91
	(10, 20)	165,560	164,160	163,849	159,772	−0.85	0.19	2.75
1.0	(0, 0)	256,000	252,301	251,522	235,247	−1.45	0.31	7.25
	(1, 5)	244,110	241,430	240,023	221,642	−1.10	0.59	8.93
	(5, 10)	213,833	212,502	211,846	199,222	−0.62	0.31	6.67
	(10, 20)	171,071	170,549	170,578	170,178	−0.30	−0.02	0.22
1.2	(0, 0)	284,000	280,816	279,625	253,978	−1.12	0.43	10.57
	(1, 5)	267,429	264,920	263,639	239,207	−0.94	0.49	10.75
	(5, 10)	217,738	217,443	217,449	212,544	−0.14	−0.00	2.30
	(10, 20)	171,071	170,949	170,948	170,826	−0.08	0.00	0.07
1.4	(0, 0)	309,000	306,741	305,074	271,498	−0.73	0.55	12.98
	(1, 5)	269,588	269,351	269,141	253,304	−0.09	0.08	6.34
	(5, 10)	217,738	217,590	217,529	215,267	−0.07	0.03	1.08
	(10, 20)	171,071	170,951	170,949	170,826	−0.07	0.00	0.07

We tested two options for initializing the column-generation procedure: (1) the first product in each segment's offer set is used to initialize the reduced LP, and (2) the second product in each segment's offer set is used to initialize the reduced LP. These choices lead to different optimal dual solutions in the CDLP. Although the CDLP optimal values do not differ much in those two cases, the revenues produced by the DCOMP heuristic are quite different, as can be seen from Table 11. Indeed, Table 11 shows that with the proper choice of the initial subsets, the DCOMP method still generates larger revenue than the CDLP policy. The negative gains of the DCOMP policy observed in Tables 8 and 9, therefore, can be attributed to this multiple dual solution effect.[7] Although one would like the method to be less sensitive to this sort of occurrence, it is an inherent weaknesses of any method that uses dual rather than primal information. From our (albeit limited) experience, however, it appears that this multiple-dual-solution effect is only significant in extreme cases,

[7] Note that this phenomenon is more likely to have a significant effect when we solve the CDLP once and use a single dual solution for the entire time horizon. In the case of reoptimizing, we are using new dual solutions at each resolving point, which makes it less likely that a single case of multiple dual solutions will negatively impact the DCOMP heuristic. Of course, resolving frequently is not guaranteed to improve performance, as shown by Cooper (2002) and Secomandi (2005). Still, in our example, when the CDLP is resolved over time, its performance clearly improves.

Table 10 **The Percentage Gain in Revenue from Reoptimizing for a Small Network**

α	v_0	% Gain-DCOMP	% Gain-CDLP	% Gain-INDEP
0.6	(0, 0)	4.99	2.93	0.21
	(1, 5)	0.01	3.40	0.24
	(5, 10)	0.02	2.67	0.26
	(10, 20)	0.23	2.43	−0.06
0.8	(0, 0)	0.04	1.87	0.23
	(1, 5)	0.01	2.21	−0.10
	(5, 10)	0.08	2.32	−0.20
	(10, 20)	0.16	0.77	−0.18
1.0	(0, 0)	0.07	1.99	−0.13
	(1, 5)	0.05	2.13	−0.07
	(5, 10)	0.04	0.75	−0.19
	(10, 20)	−0.09	0.30	−0.07
1.2	(0, 0)	3.22	1.89	0.01
	(1, 5)	0.19	1.10	−0.00
	(5, 10)	−0.13	−0.00	−0.18
	(10, 20)	−0.03	−0.02	−0.13
1.4	(0, 0)	−0.04	1.46	0.09
	(1, 5)	−0.04	0.10	−0.01
	(5, 10)	−0.06	−0.09	−0.17
	(10, 20)	−0.24	−0.02	−0.13

Table 11 Revenues from Different Settings of the Initial Consideration Sets for $\alpha = 1.2$ and $v_0 = (0, 0)$ Without Reoptimization

Initial prod. selected	UB REV	DCOMP REV	CDLP REV	INDEP REV	%OPT-GAP	%DCOMP-CDLP	%DCOMP-INDEP
First	284,000	272,051	274,436	253,963	−4.21	−0.87	7.12
Second	284,000	279,756	273,843	254,023	−1.49	2.16	10.13

like the $v_0 = (0, 0)$ cases here, although more experience with the method is needed to make any strong claims in this regard.

8. Conclusions

The CDLP of Gallego et al. (2004) is the natural analog of the traditional deterministic LP, which is widely used in revenue management practice. Here, we have extended their analysis of this model, showing that its performance—as in the traditional independent demand case—is asymptotically optimal as capacity and demand are scaled up proportionately. We used these results to extend the concept of efficient sets developed by Talluri and van Ryzin (2004a) to the network case. Efficiency is a potentially useful concept for identifying good offer sets and thereby reducing the complexity of choice-based problems, although more work is needed on this issue. We also developed a heuristic that uses the dual information from the CDLP to decompose the network DP into a collection of leg-level DPs. It is the choice-based analog of traditional network decomposition methods such as DAVN and DP decomposition. The method has several attractive features: It is efficient to compute, recommends a unique offer set that changes dynamically in response to changes in remaining capacity and time, improves on the naive CDLP policy significantly—especially in high-load factor cases—and is much less sensitive to the frequency of reoptimization. Overall, the heuristic appears to be a viable practical method for using and enhancing the CDLP model of Gallego et al. (2004). Although both the CDLP and decomposition heuristic are computationally complex in general, we showed that under the MNL with disjoint segments model, both problems can be solved efficiently. Although this is a restrictive model, the recent work of Bront et al. (2006) indicates that the approach is viable for broader classes of choice models.

Acknowledgments

The authors thank Gustavo Vulcano for numerous detailed comments on earlier drafts of this work, especially regarding the computational examples. The authors also thank the anonymous referees and associate editor for their constructive feedback.

Appendix

Proof of Proposition 3. We first show the "only if" part. Define $\bar{R}$: $[0, 1]^m \to R$ as

$$\bar{R}(q) = \max_{\alpha}\left\{\sum_S \alpha(S)R(S):\ \sum_S \alpha(S)Q(S) \le q,\ \sum_S \alpha(S) = 1,\ \alpha(S) \ge 0, \forall S \subseteq N\right\}.$$

Then $\bar{R}(q)$ is increasing and concave in q. (See Theorem 5.1, p. 213, of Bertsimas and Tsitsiklis 1997.)

For any efficient set T, T satisfies $\bar{R}(q)$ for some q, which is concave in q. There exists a supporting hyperplane (u, v), $u \in R$, $v \in R^m$ such that

$$R(T) = v^\top Q(T) + u,$$
$$R(S) \le v^\top Q(S) + u, \quad S \subseteq N.$$

Therefore,

$$R(T) - v^\top Q(T) \ge R(S) - v^\top Q(S).$$

It is left to show that $v \ge 0$.

$$\sum_T \alpha(T)R(T) = \sum_T \alpha(T)(v^\top Q(T) + u) \le v^\top q + u.$$

Since $\sum_T \alpha(T)R(T) = \bar{R}(q)$ is increasing in q, it must be $v \ge 0$.

Next, we show the "if" part. Suppose T is the optimal solution to (7) for some $\pi = (\pi_1, \ldots, \pi_m) \ge 0$; then,

$$R(T) - \pi^\top Q(T) \ge R(S) - \pi^\top Q(S), \quad S \subseteq N,$$
$$(R(T) - R(S)) - \pi^\top (Q(T) - Q(S)) \ge 0. \tag{15}$$

Multiplying each inequality (15) with $\alpha(S)$, and summing those inequalities, we have

$$\left(R(T) - \sum_S \alpha(S)R(S)\right) - \pi^\top \left(Q(T) - \sum_S \alpha(S)Q(S)\right) \ge 0. \tag{16}$$

Table A1 Enumeration of All Offer Sets and the Associated Choice Probabilities, Consumption Probabilities, and Expected Revenues for the Running Example

S	$P_1(S)$	$P_2(S)$	$P_3(S)$	$P_4(S)$	$P_5(S)$	$P_6(S)$	$P_0(S)$	$R(S)$	$Q_1(S)$	$Q_2(S)$	$Q_3(S)$	Efficient?
{∅}	0	0	0	0	0	0	1	0	0	0	0	Y
{1}	0.20	0	0	0	0	0	0.80	240	0	0.20	0	Y
{2}	0	0.150	0	0	0	0	0.85	120	0.15	0	0.15	Y
{3}	0	0	0.14	0	0	0	0.86	84	0.14	0	0	Y
{4}	0	0	0	0.17	0	0	0.83	136	0	0.17	0	N
{5}	0	0	0	0	0.25	0	0.75	125	0.25	0	0.25	N
{6}	0	0	0	0	0	0.16	0.84	48	0.16	0	0	N
{1, 2}	0.15	0.075	0	0	0	0	0.775	240	0.075	0.15	0.075	Y
{1, 3}	0.20	0	0.14	0	0	0	0.66	324	0.14	0.20	0	Y
{1, 4}	0.20	0	0	0.17	0	0	0.63	376	0	0.37	0	Y
{1, 5}	0.20	0	0	0	0.25	0	0.55	365	0.25	0.20	0.25	N
{1, 6}	0.20	0	0	0	0	0.16	0.64	288	0.16	0.20	0	N
{2, 3}	0	0.150	0.14	0	0	0	0.71	204	0.29	0	0.15	Y
{2, 4}	0	0.150	0	0.17	0	0	0.68	256	0.15	0.17	0.15	N
{2, 5}	0	0.150	0	0	0.25	0	0.60	245	0.40	0	0.40	N
{2, 6}	0	0.150	0	0	0	0.16	0.69	168	0.31	0	0.15	N
{3, 4}	0	0	0.14	0.17	0	0	0.69	220	0.14	0.17	0	N
{3, 5}	0	0	0.14	0	0.25	0	0.61	209	0.39	0	0.25	N
{3, 6}	0	0	0.07	0	0	0.11	0.82	75	0.18	0	0	N
{4, 5}	0	0	0	0.10	0.20	0	0.70	180	0.20	0.10	0.20	N
{4, 6}	0	0	0	0.17	0	0.16	0.67	184	0.16	0.17	0	N
{5, 6}	0	0	0	0	0.25	0.16	0.59	173	0.41	0	0.25	N
{1, 2, 3}	0.15	0.075	0.14	0	0	0	0.635	324	0.215	0.15	0.075	Y
{1, 2, 4}	0.15	0.075	0	0.17	0	0	0.605	376	0.075	0.32	0.075	N
{1, 2, 5}	0.15	0.075	0	0	0.25	0	0.525	365	0.325	0.15	0.325	N
{1, 2, 6}	0.15	0.075	0	0	0	0.16	0.615	288	0.235	0.15	0.075	N
{1, 3, 4}	0.20	0	0.14	0.17	0	0	0.49	460	0.14	0.37	0	Y
{1, 3, 5}	0.20	0	0.14	0	0.25	0	0.41	449	0.39	0.20	0.25	N
{1, 3, 6}	0.20	0	0.07	0	0	0.11	0.62	315	0.18	0.20	0	N
{1, 4, 5}	0.20	0	0	0.10	0.20	0	0.50	420	0.20	0.30	0.20	N
{1, 4, 6}	0.20	0	0	0.17	0	0.16	0.47	424	0.16	0.37	0	N
{1, 5, 6}	0.20	0	0	0	0.25	0.16	0.39	413	0.41	0.20	0.25	N
{2, 3, 4}	0	0.15	0.14	0.17	0	0	0.54	340	0.29	0.17	0.15	N
{2, 3, 5}	0	0.15	0.14	0	0.25	0	0.46	329	0.54	0	0.40	Y
{2, 3, 6}	0	0.15	0.07	0	0	0.11	0.67	195	0.33	0	0.15	N
{2, 4, 5}	0	0.15	0	0.10	0.20	0	0.55	300	0.35	0.10	0.35	N
{2, 4, 6}	0	0.15	0	0.17	0	0.16	0.52	304	0.31	0.17	0.15	N
{2, 5, 6}	0	0.15	0	0	0.25	0.16	0.44	293	0.56	0	0.40	N
{3, 4, 5}	0	0	0.14	0.10	0.20	0	0.56	264	0.34	0.1	0.20	N
{3, 4, 6}	0	0	0.07	0.17	0	0.11	0.65	211	0.18	0.17	0	N
{3, 5, 6}	0	0	0.07	0	0.25	0.11	0.57	200	0.43	0	0.25	N
{4, 5, 6}	0	0	0	0.10	0.20	0.16	0.54	228	0.36	0.10	0.20	N
{1, 2, 3, 4}	0.15	0.075	0.14	0.17	0	0	0.465	460	0.215	0.32	0.075	Y
{1, 2, 3, 5}	0.15	0.075	0.14	0	0.25	0	0.385	449	0.465	0.15	0.325	Y
{1, 2, 3, 6}	0.15	0.075	0.07	0	0	0.11	0.595	315	0.255	0.15	0.075	N
{1, 2, 4, 5}	0.15	0.075	0	0.10	0.20	0	0.475	420	0.275	0.25	0.275	N
{1, 2, 4, 6}	0.15	0.075	0	0.17	0	0.16	0.445	424	0.235	0.32	0.075	N
{1, 2, 5, 6}	0.15	0.075	0	0	0.25	0.16	0.365	413	0.485	0.15	0.325	N
{1, 3, 4, 5}	0.20	0	0.14	0.10	0.20	0	0.36	504	0.34	0.30	0.20	Y
{1, 3, 4, 6}	0.20	0	0.07	0.17	0	0.11	0.45	451	0.18	0.37	0	N
{1, 3, 5, 6}	0.20	0	0.07	0	0.25	0.11	0.37	440	0.43	0.20	0.25	N
{1, 4, 5, 6}	0.20	0	0	0.10	0.20	0.16	0.34	468	0.36	0.30	0.20	N
{2, 3, 4, 5}	0	0.15	0.14	0.10	0.20	0	0.41	384	0.49	0.10	0.35	N
{2, 3, 4, 6}	0	0.15	0.07	0.17	0	0.11	0.50	331	0.33	0.17	0.15	N
{2, 3, 5, 6}	0	0.15	0.07	0	0.25	0.11	0.42	320	0.58	0	0.40	N
{2, 4, 5, 6}	0	0.15	0	0.10	0.20	0.16	0.39	348	0.51	0.10	0.35	N
{3, 4, 5, 6}	0	0	0.07	0.10	0.20	0.11	0.52	255	0.38	0.10	0.20	N
{1, 2, 3, 4, 5}	0.15	0.075	0.14	0.10	0.20	0	0.335	504	0.415	0.25	0.275	Y
{1, 2, 3, 4, 6}	0.15	0.075	0.07	0.17	0	0.11	0.425	451	0.255	0.32	0.075	N
{1, 2, 3, 5, 6}	0.15	0.075	0.07	0	0.25	0.11	0.345	440	0.505	0.15	0.325	Y
{1, 2, 4, 5, 6}	0.15	0.075	0	0.10	0.20	0.16	0.315	468	0.435	0.25	0.275	N
{1, 3, 4, 5, 6}	0.20	0	0.07	0.10	0.20	0.11	0.32	495	0.38	0.30	0.20	N
{2, 3, 4, 5, 6}	0	0.15	0.07	0.10	0.20	0.11	0.37	375	0.53	0.10	0.35	N
{1, 2, 3, 4, 5, 6}	0.15	0.075	0.07	0.10	0.20	0.11	0.295	495	0.455	0.25	0.275	N

If T is not an efficient set, then there exists $\alpha(S) \geq 0$, $\sum_S \alpha(S) = 1$ such that

$$R(T) < \sum_S \alpha(S)R(S),$$

$$Q(T) \geq \sum_S \alpha(S)Q(S).$$

Since $\pi \geq 0$,

$$\left(R(T) - \sum_S \alpha(S)R(S)\right) - \pi^\top\left(Q(T) - \sum_S \alpha(S)Q(S)\right) < 0,$$

which contradicts (16). □

Proof of Proposition 6. Denote $v_{lj} \cdot y_{lj} = z_{lj}$; then, we claim that

$$\max_{z_{lj}\in[0,\, v_{lj}]^{s_l}} \frac{\sum_{l\in C_l} w_j z_{lj}}{\sum_{l\in C_l} z_{lj} + v_{l0}} \qquad (17)$$

has the same optimal value as the optimization problem (14).

To simplify notation, we suppress the subscript l in (17) and define $w^T = (w_1, \ldots, w_s)$, $\mathscr{D} = [0, v_1] \times \cdots \times [0, v_s]$, $e^T = (1, \ldots, 1)$, and $z \in \mathscr{D}$. Then the objective function of (17) can be expressed as:

$$h(z) = \frac{h_1(z)}{h_2(z)} = \frac{w^T z}{e^T z + v_0}.$$

Note that both $h_1(z)$ and $h_2(z)$ are linear; in addition, $h_2(z)$ is always positive if $v_0 \neq 0$, and we define $h(0) := 0$ when $v_0 = 0$. Therefore, from results in Avriel (1976, p. 156) on quasiconvexity properties of the ratio of linear functions, we conclude that $h(z)$ is quasiconvex. Then, using the fact that the maximizer of a quasiconvex function on a closed convex set is achieved at boundary points, the optimal solution to (17) is either 0 or v_{lj}, which is equivalent to $y_{lj} = 0$ or $y_{lj} = 1$.

Without loss of generality, assume $w_1 \geq w_2 \geq \cdots \geq w_s$. We next show that the optimal solution to

$$\max_{z_i\in[0,\, v_i]} \frac{\sum_{i=1}^s w_i z_i}{\sum_{i=1}^s z_i + v_0}$$

is given by

$$z_i^* = \begin{cases} v_i & \text{if } i \leq i^* \\ 0 & \text{if } i > i^* \end{cases}$$

where $1 \leq i^* \leq s$.

Suppose the optimal solution $z_i^*, i = 1, \ldots, s$ does not have this property. Then there exists $1 \leq i < j \leq s$, such that $w_i > w_j$ and $z_i^* = 0$, $z_j^* = v_j$. Let $z_i' = \min\{v_i, v_j\}$, $z_j' = v_j - \min\{v_i, v_j\}$ and $z_l' = z_l^*$, $l = 1, \ldots, s$, but $l \neq i, j$. It is easy to check that z_i' is a feasible solution.

$$\begin{aligned}\sum_{l=1}^s w_l z_l' &= \sum_{l=1,\, l\neq i,\, j}^s w_l z_l' + w_i z_i' + w_j z_j' \\ &= \sum_{l=1,\, l\neq i,\, j}^s w_l z_l^* + w_i \min\{v_i, v_j\} \\ &\quad + w_j(v_j - \min\{v_i, v_j\}) \\ &= \sum_{l=1}^s w_l z_l^* + (w_i - w_j)\min\{v_i, v_j\} \\ &> \sum_{l=1}^s w_l z_l^*.\end{aligned}$$

The inequality holds because $w_i > w_j$ and $v_i, v_j > 0$. Also,

$$\sum_{l=1}^s z_l' = \sum_{l=1,\, l\neq i,\, j}^s z_l' + \min\{v_i, v_j\} + v_j - \min\{v_i, v_j\} = \sum_{l=1}^s z_l^*.$$

Thus,

$$\frac{\sum_{l=1}^s w_l z_l'}{\sum_{l=1}^s z_l' + v_0} > \frac{\sum_{l=1}^s w_l z_l^*}{\sum_{l=1}^s z_l^* + v_0}.$$

However, this contradicts the fact that z_i^* is the optimal solution. Therefore, the optimal solution z_i^* must satisfy

$$z_i^* = \begin{cases} v_i & \text{if } i \leq i^* \\ 0 & \text{if } i > i^*. \end{cases}$$

where i^* is defined by

$$i^* = \begin{cases} m & \dfrac{\sum_{l=1}^m w_l v_l}{\sum_{l=1}^m v_l + v_0} > \dfrac{\sum_{l=1}^{m-1} w_l v_l}{\sum_{l=1}^{m-1} v_l + v_0}; \\ \min\left\{1 \leq i \leq m-1 \;\middle|\; \dfrac{\sum_{l=1}^i w_l v_l}{\sum_{l=1}^i v_l + v_0} > \dfrac{\sum_{l=1}^{i+1} w_l v_l}{\sum_{l=1}^{i+1} v_l + v_0}\right\} \\ \quad \text{otherwise.} \end{cases}$$ □

References

Algers, S., M. Besser. 2001. Modeling choice of flight and booking class: A study using stated preference and revealed preference data. *Internat. J. Services Tech. Management* **2** 28–45.

Andersson, S. E. 1989. Operational planning in airline business—Can science improve efficiency? Experiences from SAS. *Eur. J. Oper. Res.* **43** 3–12.

Andersson, S. E. 1998. Passenger choice analysis for seat capacity control: A pilot project in Scandinavian Airlines. *Internat. Trans. Oper. Res.* **5** 471–486.

Avriel, M. 1976. *Nonlinear Programming: Analysis and Methods*. Prentice Hall, Englewood Cliffs, NJ.

Belobaba, P. P. 1987a. Air travel demand and airline seat inventory management. Ph.D. thesis, Flight Transportation Laboratory, MIT, Cambridge, MA.

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

Belobaba, P. P., C. Hopperstad. 1999. Boeing/MIT simulation study: PODS results update. *1999 AGIFORS Reservations and Yield Management Study Group Sympos. Proc.*, London, UK.

Belobaba, P. P., L. R. Weatherford. 1996. Comparing decision rules that incorporate customer diversion in perishable asset revenue management situations. *Decision Sci.* **27** 343–363.

Ben-Akiva, M., S. Lerman. 1985. *Discrete-Choice Analysis: Theory and Application to Travel Demand*. MIT Press, Cambridge, MA.

Bertsimas, D. J., I. Popescu. 2003. Revenue management in a dynamic network environment. *Transportation Sci.* **37** 257–277.

Bertsimas, D. J., J. N. Tsitsiklis. 1997. *Introduction to Linear Optimization*. Athena Scientific, Belmont, MA.

Bitran, G. R., R. Caldentey, S. V. Mondschein. 1998. Coordinating clearance markdown sales of seasonal products in retail chains. *Oper. Res.* **46** 609–624.

Bront, J. J. M., I. M. Diaz, J. M. Bront, G. Vulcano. 2006. A column generation algorithm for the choice-based linear programming model for network revenue management. Working paper, Computer Science Department, School of Science, University of Buenos Aires, Buenos Aires, Argentina.

Cooper, W. L. 2002. Asymptotic behavior of an allocation policy for revenue management. *Oper. Res.* **50** 720–727.

Cooper, W. W., A. Charnes, E. Rhodes. 1978. Measuring the efficiency of decision making units. *Eur. J. Oper. Res.* **2** 429–444.

Feng, Y., G. Gallego. 1995. Optimal starting times for end-of-season sales and optimal stopping times for promotional fares. *Management Sci.* **41** 1371–1391.

Feng, Y., G. Gallego. 2000. Perishable asset revenue management with Markovian time dependent demand intensities. *Management Sci.* **46** 941–956.

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

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

Feng, Y., B. Xiao. 2000c. Revenue management with two market segments and reserved capacity for priority customers. *Adv. Appl. Probab.* **32** 800–823.

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

Gallego, G., G. J. van Ryzin. 1997. A multi-product dynamic pricing problem and its applications to network yield management. *Oper. Res.* **45** 24–41.

Gallego, G., G. Iyengar, R. Phillips, A. Dubey. 2004. Managing flexible products on a network. CORC Technical Report TR-2004-01, Department of Industrial Engineering and Operations Research, Columbia University, New York. http://www.corc.ieor.columbia.edu/reports/techreports/.

Maglaras, C., J. Meissner. 2006. Dynamic pricing strategies for multiproduct revenue management problems. *Manufacturing Service Oper. Management* **8** 136–148.

Phillips, R. L. 1994. State-contingent airline yield management. Presentation in Session TC33.4, INFORMS Detroit, Detroit, MI.

Secomandi, N. 2005. An analysis of the control-algorithm re-solving issue in inventory and revenue management. Tepper Working Paper, Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA.

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

Talluri, K. T., G. J. van Ryzin. 1999a. An analysis of bid-price controls for network revenue management. *Management Sci.* **44** 1577–1593.

Talluri, K. T., G. J. van Ryzin. 1999b. A randomized linear programming method for computing network bid prices. *Transportation Sci.* **33** 207–216.

Talluri, K. T., G. J. van Ryzin. 2004a. Revenue management under a general discrete choice model of consumer behavior. *Management Sci.* **50** 15–33.

Talluri, K. T., G. J. van Ryzin. 2004b. *The Theory and Practice of Revenue Management*. Kluwer, New York.

van Ryzin, G. J., G. Vulcano. 2004. Computing virtual nesting controls for revenue management under consumer choice behavior. Technical report, Graduate School of Business, Columbia University, New York. http://www2.gsb.columbia.edu/divisions/dro/.

Williamson, E. L. 1988. Comparison of optimization techniques for origin-destination seat inventory control. Master's thesis, Flight Transportation Laboratory, MIT, Cambridge, MA.

Williamson, E. L. 1992. Airline network seat inventory control: Methodologies and revenue impacts. PhD thesis, Flight Transportation Laboratory, MIT, Cambridge, MA.

Zhang, D., W. L. Cooper. 2005. Revenue management for parallel flights with customer choice behavior. *Oper. Res.* **53** 415–431.