CORC Technical Report TR-2004-01

Managing Flexible Products on a Network

G. Gallego* G. Iyengar† R. Phillips‡ A. Dubey§

Abstract

A *flexible product* is a menu of two or more alternatives products serving the same market. Purchasers of flexible products are assigned to one of the alternatives at a later date. Gallego and Phillips [9] show that capacitated suppliers, such as airlines and hotels, can potentially improve revenue by offering flexible products in addition to traditional specific products. In this paper, we extend the concept of flexible products to networks. We study the network revenue management problem with flexible products in two different settings: one where the demand for each product is independently and exogenously generated; and the other where the demand is driven by a consumer choice model. We show that in both these settings the optimal value of the stochastic optimization problem can be closely approximated by the optimal value of a deterministic linear program. In the independent demand case the corresponding linear program is of modest size. When the demand is driven by a customer choice model, the linear program has exponentially many variables; however, we show that for an important class of consumer choice models the linear program can be efficiently solved using column-generation. We report the findings of numerical experiments with a real airline subnetwork and show how the results vary as a function of the key inputs.

*IEOR Department, Columbia University. Research partially supported by NSF grant DMI-02-18104

†IEOR Department, Columbia University. Email: garud@ieor.columbia.edu. Research partially supported by NSF grants CCR-00-09972 and DMS-01-04282.

‡Nomis Solutions. Research partially supported by NSF grant DMI-02-18104

§School of Operations Research and Industrial Engineering, Cornell University

1 Background and introduction

Gallego and Phillips [9] define a *flexible product* as "a set of two or more alternatives serving the same market such that a purchaser of the flexible product is assigned one of the alternatives by the seller at a later date". This is in contrast to the more traditional notion of a *specific product* in which the customer purchases a known alternative. They illustrated flexible products with the example of an airline with three daily flights from New York's Kennedy airport (JFK) to San Francisco International (SFO). Customers can book seats on any one of the three flights. In addition, the airline also offers a flexible product "JFK → SFO" at a discount. Customers purchasing this flexible product are guaranteed a seat on one of the three flights; but the assignment to a specific flight is only finalized a day prior to departure. The airline has complete freedom in assigning the flexible passengers to any of the three flights. Since flexible products are offered at a discount, they have the potential to attract additional price-sensitive customers as well as improve overall capacity utilization.

The concept of flexible products is applicable to any industry selling fixed, perishable capacity when some of the products are considered substitutes by some customers. For example, a hotel operator with several properties in Manhattan could take bookings for a flexible "Manhattan lodging" product in addition to specific bookings for the Soho property, the Midtown property, etc. A traveler booking "Manhattan lodging" will be assigned to a specific property at a specified time before her arrival. Examples of industries that can employ flexible products include on-line and broadcast advertising, freight transportation, tour industries, and contract manufacturing. Some specific applications in these industries are discussed in Gallego and Phillips [9]. While the concept of flexible products has broad application, for consistency and ease of understanding this paper is couched in terms of passenger airlines.

Flexible products offer two main benefits:

(a) *Higher capacity utilization*: Since purchasers of flexible products are assigned to specific products after the demand uncertainty has been largely resolved, offering flexible products can improve capacity utilization.

(b) *Demand induction*: Since purchasers of flexible products are informed of their assignment well after the time of purchase, they are likely to be considered inferior products. Consequently, by offering flexible products at a discount the firm could attract additional customers without entirely cannibalizing the full-fare demand.

Gallego and Phillips [9] present a two-flight example in which offering a flexible product at a discount can generate significant additional revenue *provided* the booking limit for the flexible product is chosen appropriately. Under their consumer-choice model, the benefit from offering the flexible product depends on the discount – if the discount is too high, cannibalization of higher-fare specific products outweighs the benefits from flexibles; and if the discount is too low, there is inadequate demand for flexibles to provide significant benefits.

In this paper we extend the concept of flexible products to networks. In Section 2 we formulate the network revenue management problem with flexible products when the requests arrive according to independent Poisson processes. We show that the value of the stochastic optimization problem is closely approximated by the optimal value of a suitably defined linear program (LP). We show that the deterministic approximation remains valid for the more general setting when requests arrive according to independent point processes. In addition to showing that the value of the LP is asymptotically optimal in the limit of long planning horizon and large capacity, we also provide a lower bound for the rate of convergence (see Corollary 1). This refinement is new and also applies to revenue management problems without flexible products.

In Section 3 we study the case when the demand is driven by a consumer choice model and the flexible products are chosen endogenously. This is, to our knowledge, the first paper to use a choice model in a network context. We formulate a dynamic program (DP) for solving the stochastic problem and show that the value function DP can be closely approximated by the optimal solution of an appropriately constructed LP. (This LP formulation has begun receiving attention in the revenue management literature, e.g. see van Ryzin and Liu [16].) Although this LP has exponentially many variables, we show that for a broad class of consumer choice models, the optimal set of flexible products and the corresponding booking limits can be efficiently computed by a column-generation algorithm. We show that this deterministic approximation remain valid as long as the arrival processes are Poisson processes. In addition to establishing asymptotic optimality, we also show a lower bound on the rate of convergence.

Section 4 reports the results of a detailed numerical study of a small airline network. We investigate how the increase in revenue from offering flexible products depends on various system parameters such as time horizon, discount, outside competition, etc. A key insight from these numerical studies is that the benefits from offering flexible products is greatest when capacity is fairly closely matched with demand. In Section 5 we include some concluding remarks and discuss avenues of further research.

2 Network model with independent demand and flexible products

We consider a network consisting of m *resources* and n *specific products* that are combinations of one or more resources with $n \geq m$. We allow for the possibility that different products may use the same combination of resources (This would be the case with multiple fare classes for the same itinerary at a passenger airline.). The initial capacity is given by the vector $c \in \Re^m_+$. This capacity must be consumed over the sales horizon $[0, T]$ and any capacity unused at time T is worthless. Let $A = [a_{li}] \in \Re^{m \times n}$ denote the incidence matrix of the specific products, i.e. $a_{li} = 1$ if the i-th specific product consumes resource l and 0 otherwise[1]. The revenue vector of the specific products is given by $p \in \Re^n_+$.

The firm also offers f *flexible products*. Each flexible product k consists of a set $F_k \subset N = \{1, \ldots, n\}$ of specific products such that every customer purchasing the flexible product k must be assigned to one of the $f_k = |F_k|$ alternatives in the set F_k. The assignment to specific products occurs at time T, i.e. at the end of the sales period. Note that the firm does *not* have the option of reneging on a flexible product sale; it only has the flexibility to choose which specific product to assign. Let $F = \{1, \ldots, f\}$ denote the set of flexible products and let $r \in \Re^f_+$ denote the vector of unit revenues associated with the flexible products.

We consider a continuous time model. Requests for each product arrive according to independent Poisson processes. The rate of the Poisson process for the i-th specific product (resp. k-th flexible product) is given by λ_i, $i = 1, \ldots, n$ (resp. γ_k, $k = 1, \ldots, f$). While it is not conceptually difficult to allow the rate vectors λ and γ to be time-dependent, for ease of exposition we assume that the rates are constant. (We consider more general arrival processes in Section 2.2 and Section 2.4.) The firm has the option of accepting or rejecting the request of an arriving customer; however, it must ensure that every admitted request for a specific product is satisfied and every

[1] Although we develop the analysis in the case where each booking is for a single product, much of the analysis will extend to the case in which bookings are for multiple units of product.

request for a flexible product k is assigned to some specific product in the set F_k. Furthermore we assume that all customer commitments are firm, i.e. there are no cancellations or no-shows.

Let $v \in \mathcal{Z}_+^n$ be the number of specific product bookings that have been accepted. Then, at any time, the network can be described by the pair of vectors $(s, y) \in \mathcal{Z}_+^m \times \mathcal{Z}_+^f$, where $s \in \mathcal{Z}_+^m$ defined by $s = c - Av$ denotes the vector of resource capacity not committed to specific products and $y \in \mathcal{Z}_+^f$ denotes the vector of accepted requests for flexible products. The state (s, y) is said to be *feasible* if, and only if, the requests y can be satisfied by the residual capacity s. The following proposition characterizes the set $\mathcal{A}$ of feasible states.

Proposition 1 *Let $\mathcal{A}$ denote the set of feasible states for the network. Then $(s, y) \in \mathcal{A}$ if and only if $s \leq c$, and there exist vectors $z_k \in \mathcal{Z}_+^{f_k}$, $k = 1, \ldots, f$, satisfying*

$$\begin{array}{rcrcll} & & \sum_{k=1}^{f} B_k z_k & \leq & s, & \\ y_k & - & 1_k' z_k & = & 0, & k = 1, \ldots, f, \end{array} \tag{1}$$

where $1_k \in \Re^{f_k}$ is a vector of all ones and B_k denotes the sub-matrix of A obtained by selecting the columns corresponding to the indices in F_k, $k = 1, \ldots, f$.

Proof

Suppose (1) is satisfied. Assign the requests for the k-th flexible product according to the vector z_k, i.e. assign z_{ki} requests to the i-th specific product in the set F_k, $i = 1, \ldots, f_k$. The first constraint in (1) ensures that this assignment does not violate the capacity constraints.

Conversely, suppose the state $(s, y) \in \mathcal{A}$, i.e. the requests for the flexible products can be satisfied by the residual capacity s. Then it follows that for each k, there exists $z_k \in \mathcal{Z}_+^{f_k}$ such that $\sum_{k=1}^{f} B_k z_k \leq s$ with $1_k' z_k = y_k$.

∎

The distribution variables $z_k \in \Re_+^{f_k}$ can be viewed as "second-stage" variables in a stochastic linear program in the sense that $(s, y) \in \mathcal{A}$ if, and only if, there exist second-stage variables satisfying (1). Since all requests must be satisfied, the vector y *cannot* be interpreted as a second-stage variable.

2.1 Dynamic programming formulation

We assume that the firm is risk-neutral and maximizes total expected revenue. Let $V(s, y, t)$ denote the maximum achievable revenue starting from the state (s, y) at time t. Then, $V(s, y, t)$ is characterized by the Hamilton-Jacobi-Bellman (HJB) recursion

$$\begin{aligned} -\frac{\partial V(s,y,t)}{\partial t} &= \sum_{i=1}^{n} \lambda_i \max\{p_i + V(s-A_i,y,t) - V_t(s,y,t), 0\} \\ &\quad + \sum_{k=1}^{f} \gamma_k \max\{r_k + V(s,y+e_k,t) - V(s,y,t), 0\}, \end{aligned} \tag{2}$$

where e_k denotes the k-th standard basis vector in $\Re^f$, and the boundary conditions are

$$\begin{aligned} V(s,y,T) &= 0, \qquad (s,y) \in \mathcal{A}, \\ V(s,y,t) &= -\infty, \quad (s,y) \notin \mathcal{A}, \quad t \in [0,T]. \end{aligned}$$

The goal of the recursion is to compute $V(c) = V(c,0,0)$. Even for small values of m, n and c, the size of the state space $\mathcal{A}$ is very large; consequently, the complexity of computing the solution of the HJB recursion (2) is prohibitive. In practice, therefore, one has to resort to heuristics to approximate $V(s,y,t)$. In the next section, we construct an asymptotically optimal policy that can be determined by solving a linear program.

2.2 Deterministic control problem

In this section we develop a deterministic approximation for the stochastic revenue optimization problem when requests arrive according to independent point processes – not necessarily Poisson. Let $V(c)$ denote the optimal value of the corresponding stochastic control problem. The recursion (2) characterizes $V(c)$ in the case when the arrival processes are Poisson. In the general stochastic setting considered in this section, even characterizing $V(c)$ is impossible!

The first step toward the deterministic approximation is to replace the stochastic demand by its the expected value, i.e. the demand is $d_x = \mathbf{E}[n_x(T)]$ and $d_y = \mathbf{E}[n_y(T)]$, where $n_x(T)$ and $n_y(T)$ denote the random number of requests for specific and flexible products respectively. Since the demand is deterministic, the firm's decision reduces to computing the pair of booking limits $(x,y) \in \mathcal{Z}_+^{n+f}$ for the specific and flexible products respectively.

The next step is to relax the integrality constraints on (x,y). Then Proposition 1 implies that the firm's decision problem reduces to the following LP.

$$\begin{array}{rllllllll} V^D(c) = & \max & p'x & + & r'y & & & & \\ & \text{subject to} & Ax & & & + & Bz & \leq & c, \\ & & & & y & - & Uz & = & 0, \\ & & x & & & & & \leq & d_x, \\ & & & & y & & & \leq & d_y, \\ & & x, & & y, & & z & \geq & 0 \end{array} \quad (3)$$

where

$$B = \begin{pmatrix} B_1 & B_2 & \dots & B_f \end{pmatrix},$$

B_k denotes the sub-matrix of A obtained by selecting the columns of A corresponding to the indices in F_k, and

$$U = \begin{pmatrix} 1_1' & 0 & \dots & 0 \\ 0 & 1_2' & \dots & 0 \\ 0 & 0 \dots & \dots & 1_f' \end{pmatrix}.$$

We call (3) the Network Linear Program with Flexibles (NLPF). The NLPF is a generalization of the LP approximation of the usual network revenue management problem without flexibles studied by Williamson [19], Phillips [14], and Talluri and van Ryzin [17], among others.

2.3 Bid-price heuristic for networks with flexible products

Before showing the key results on the convergence of a deterministic heuristic, we motivate a bid-price heuristic for controlling availabilities of both flexible and specific products on a network. We assume, without loss of generality, that the demand vectors $d_x > 0$ and $d_y > 0$. The dual of (3) is given by

$$\begin{array}{rlllllll} \min & c'\theta & + & d_x'\mu & + & d_y'\nu & & \\ \text{subject to} & A'\theta & + & \mu & & & \geq & p, \\ & B'\theta & & & + & U'\nu & \geq & U'r, \\ & \theta, & & \mu, & & \nu & \geq & 0. \end{array}$$

Since $\mu \geq \max\{p - A'\theta, 0\}$ and $d_x > 0$, it follows that $\mu = \max\{p - A'\theta, 0\}$. Thus, for a specific product i, $p_i > \theta' Ae_i$ implies by complementary slackness that $x_i = d_{xi}$, i.e. all incoming requests for the specific product i should be accepted. Similarly, $U'\nu \geq U'r - B'\theta$, and $\nu \geq 0$ together imply

1. Fix $\theta \in \Re^m_+$ (bid-price)

2. Let (s, y) denote the current state of the network.

3. Accept a request for the specific product j if $p_j > \theta' A_j$ and $(s - A_j, y) \in \mathcal{A}$.
   If request is accepted, set the state $(s, y) \leftarrow (s - A_j, y)$.

4. Accept a request for the flexible product k if $r_k > \min_{j \in F_k}\{\theta' A_j\}$ and $(s, y+e_k) \in \mathcal{A}$.
   If request is accepted, set the state $(s, y) \leftarrow (s, y + e_k)$.

Figure 1: Bid-price heuristic for network models with flexibles

$U'\nu \geq \max\{U'r - B'\theta, 0\}$. This constraint implies that for every specific product $i \in F_k$ we have $\nu_k \geq \max\{r_k - \theta' A_i, 0\}$. Consequently,

$$\nu_k \geq \max_{i \in F_k} \max\{r_k - \theta' A_i, 0\}$$

Now, since $d_{yk} > 0$ for all k, it follows that it is optimal to select ν_k as small as possible, so

$$\nu_k = \max\left\{r_k - \min_{j \in F_k}\{\theta' A e_j\}, 0\right\}.$$

Consequently, $r_k > \min_{j \in F_k}\{\theta' A e_j\}$ implies $\nu_k > 0$, which in turn, by complementary slackness, implies that $y_k = e_k' U z = d_{yk}$. This motivates the "bid-price" heuristic displayed in Figure 1. Note that when the heuristic accepts a request for the specific product j the necessary resources are immediately allocated, i.e. the residual capacity $s \leftarrow s - A_j$. However, when request for a flexible product k is accepted it is *not* immediately assigned to a specific product in F_k – the assignment is made at the end of the time horizon.

Bid price heuristics are a popular and effective approach to managing specific products on a network. Talluri and van Ryzin [17] have shown that a bid price heuristic is asymptotically optimal for an airline offering only specific products as the number of seats and time periods increases. A similar asymptotic optimality property holds for the bid-price heuristic for combined flexible and specific products.

2.4 Asymptotic optimality of the deterministic control problem

In this section, we show that all the standard results for LP approximation of the revenue management problem extends to NLPF with simple modifications. We refine the asymptotic optimality result by establishing a lower bound on the rate of convergence.

Proposition 2 *The deterministic value function $V^D(c)$ is an upper bound on the stochastic value function $V(c)$.*

Proof

Let (n_x^ρ, n_y^ρ) denote the random number of requests admitted by any feasible control policy ρ for the stochastic problem by time T. Then Proposition 1 implies that

$$\begin{array}{rcrcl} An_x^\rho & & + & Bn_z^\rho & \leq & c, \\ & - & n_y^\rho & + & Un_z^\rho & = & 0, \end{array}$$

for some $n_z^\rho \in \mathcal{Z}_+^{\sum_k f_k}$. Taking expectations, it follows that the vector $(\mathbf{E}[n_x^\rho], \mathbf{E}[n_y^\rho], \mathbf{E}[n_z^\rho])$ is feasible for NLPF. Consequently, it follows that the expected reward R_ρ of the policy ρ satisfies

$$R_\rho = \mathbf{E}\left[p' n_x^\rho + r' n_y^\rho\right] = p'\mathbf{E}[n_x^\rho] + r'\mathbf{E}[n_y^\rho] \leq V^D(c).$$

Since the policy ρ was arbitrary, the result follows.

∎

Next, we prove a lower bound on $V(c)$ in terms of the deterministic value function $V^D(c)$. Let

$$\chi_i(T) = \frac{\sqrt{\mathbf{Var}(n_{xi}(T))}}{\mathbf{E}[n_{xi}(T)]}, \quad i = 1, \ldots, n, \qquad \chi_k(T) = \frac{\sqrt{\mathbf{Var}(n_{yk}(T))}}{\mathbf{E}[n_{yk}(T)]}, \quad k = 1, \ldots, f,$$

denote the coefficients of variation of the number of requests for the various products. Let

$$\chi(T) = \max\left\{ \max_{1 \leq i \leq n}\{\chi_i(T)\}, \max_{1 \leq k \leq f}\{\chi_k(T)\}\right\}$$

denote the maximum coefficient of variation over all products.

Lemma 1 *$V(c) \geq \left(1 - 1.89\chi^{\frac{2}{3}}(T)\right)V^D(c)$ for all T such that $\chi^2(T) \leq \frac{1}{2}$.*

Proof

Let (x^*, y^*, z^*) denote the optimal solution of the NLPF (3). Let $\bar{\rho}$ denote a control policy that admits at most $(1-\epsilon)x^*$ specific requests and $(1-\epsilon)y^*$ flexible requests; and distributes the flexible requests according to $(1-\epsilon)z^*$ for some $\epsilon \in (0, 1)$. The policy $\bar{\rho}$ is clearly feasible and the expected revenue $R_{\bar{\rho}}$ is given by

$$R_{\bar{\rho}} = \mathbf{E}\left[\sum_{i=1}^{n} p_i \min\{n_{xi}(T), (1-\epsilon)x_i^*\} + \sum_{k=1}^{f} r_k \min\{n_{yk}(T), (1-\epsilon)y_k^*\}\right].$$

A generic term in the expression for $R_{\bar{\rho}}$ can be bounded as follows.

$$\begin{aligned}
\mathbf{E}[\min\{n_{xi}, (1-\epsilon)x_i^*\}] &\geq (1-\epsilon)x_i^*\mathbf{P}\big(n_{xi} \geq (1-\epsilon)x_i^*\big), & (4)\\
&\geq (1-\epsilon)x_i^*\mathbf{P}\big(n_{xi} \geq (1-\epsilon)d_{xi}\big), & (5)\\
&= (1-\epsilon)x_i^*\Big(1 - \frac{\chi_i^2(T)}{\chi_i^2(T)+\epsilon^2}\Big), & (6)\\
&\geq (1-\epsilon)x_i^*\Big(1 - \frac{\chi^2(T)}{\epsilon^2}\Big), & (7)
\end{aligned}$$

where (4) holds because $n_x \geq 0$, (5) follows from the fact that $x^* \leq d_x$, (6) follows from Marshall's inequality and (7) follows from the definition of $\chi^2(T)$. The generic lower bound (7) implies that

$$R_{\bar{\rho}} \geq (1-\epsilon)\Big(1 - \frac{\chi^2(T)}{\epsilon^2}\Big)(p'x^* + r'y^*) = (1-\epsilon)\Big(1 - \frac{\chi^2(T)}{\epsilon^2}\Big)V^D(c) \geq \Big(1 - \epsilon - \frac{\chi^2(T)}{\epsilon^2}\Big)V^D(c).$$

Maximizing the lower bound over ϵ yields $\epsilon^* = \big(2\chi^2(T)\big)^{\frac{1}{3}}$. Thus, for all such T such that $\epsilon^* < 1$, i.e. $2\chi^2(T) < 1$,

$$V(c) \geq R_{\bar{\rho}} \geq \Big(1 - (2^{\frac{1}{3}} + 2^{-\frac{2}{3}})\chi^{\frac{2}{3}}(T)\Big)V^D(c).$$

The result follows by explicitly computing the constants.

∎

Corollary 1 *Fix $\zeta > 1$ and consider the scaled stochastic problem in initial capacity ζc and time horizon ζT. Then, for all $\zeta > 0$ such that $\chi^2(\zeta T) \leq \frac{1}{2}$, we have*

$$V(\zeta c) \geq \zeta V^D(c)\Big(1 - 1.89\chi^{\frac{2}{3}}(\zeta T)\Big).$$

Proof

Let (x^*, y^*, z^*) denote the optimal solution of the NLPF (3). Then it is easy to check that $\zeta(x^*, y^*, z^*)$ is optimal for the NLPF corresponding to the stochastic problem with initial capacity ζc and time horizon ζT. Consequently the result follows from Lemma 1.

∎

When requests for the specific and flexible products arrive according to independent Poisson processes, the maximum coefficient of variation $\chi(T) = \frac{1}{\sqrt{\lambda_{\min} T}}$, where $\lambda_{\min} = \min\{\min_{1\leq i\leq n}\{\lambda_i\}, \min_{1\leq k\leq f}\{\gamma_k\}\}$. Corollary 1 implies the following convergence rate

$$V(\zeta c) \geq \zeta V^D(c)\Big(1 - \frac{1.89}{(\zeta\lambda_{\min}T)^{\frac{1}{3}}}\Big).$$

Lemma 1 and Corollary 1 refine the previously known asymptotic optimality results by providing a lower bound on the rate of convergence.

2.5 Flexible products and capacity utilization

In this section we discuss a simple example with two parallel flights that highlights the role of flexible products in correcting short-term capacity mismatch. In this section, the term *revenue* will refer to the value of the deterministic value function $V^D(c)$.

Consider a small airline network described as follows.

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad c = \begin{bmatrix} 100 \\ 120 \end{bmatrix}, \quad d = \lambda T = \begin{bmatrix} 75 \\ 300 \end{bmatrix}, \quad p = \begin{bmatrix} 600 \\ 400 \end{bmatrix}. \tag{8}$$

Thus, the network consists of two parallel flights: product 1 is the higher quality product and product 2 is the lower quality product. There is a capacity mismatch in the sense that product 1 is under-subscribed whereas product 2 is over-subscribed. The firm typically has access to several different controls to correct this situation, e.g., decrease the price of product 1 to induce more demand, reduce the capacity of flight 1, increase the price of product 2 to earn more revenue, etc. However, time scales over which these controls apply are often very different. Consider a situation where the average demand is matched to capacity on the time scales over which these controls apply; and the mismatch detailed in (8) is a *local* mismatch, i.e. a random perturbation when viewed from the longer time scale.

Flexible products provide an additional control for the firm to smooth out these local perturbations, i.e. shape local demand to increase revenue. Suppose the firm offers the flexible product $F = \{1, 2\}$ at a discount $\beta \in [0, 1)$, i.e. $r = \beta \min\{p_1, p_2\}$. Suppose at a discount β, requests for each of the two specific products switch to the flexible product with probability

$$\kappa(\beta) = \max\{e^{-a_1\beta} - a_0, 0\}.$$

In our computations, $a_0 = 0.4$ and $a_1 = -\frac{\ln(a_0)}{0.8}$ leading to the $\kappa(\beta)$ curve shown in Figure 2. The switching probability $\kappa(\beta) = 0$ for $\beta \in [0.8, 1)$, i.e. no requests switch unless the discount is at least 20% and $\kappa(0) = 0.6$, i.e. only 60% of the requests switch to the flexible product even when it is free. Although we assume that the switching rate is the same for both the specific products, this is not necessary and the analysis provided below easily extends to the more general case. In this

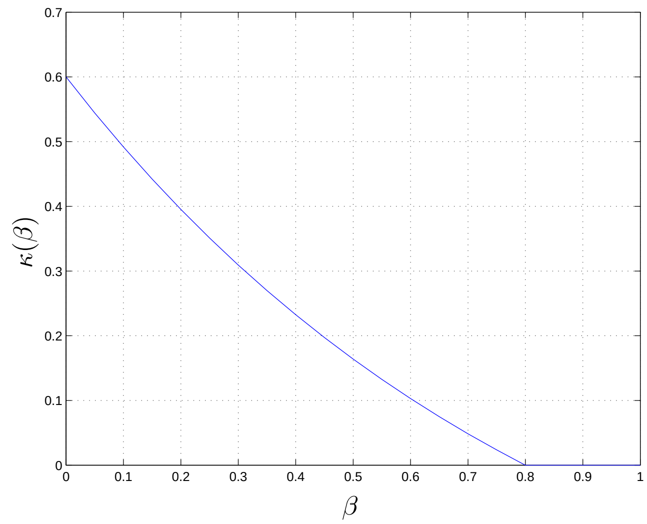

Figure 2: Switching probability $\kappa(\beta)$

example we assume that there is no demand induction; therefore, the computed revenue will be a lower bound.

Figure 3 plots the relative increase in revenue as a function of β. The maximum revenue enhancement is 4.5% at $\beta = 0.635$ and the revenue is higher only when $\beta \in (0.5, 0.8)$. For small values of β – that is high discounts for the flexible product – the increase in capacity utilization does not compensate for the decrease in revenue per request, i.e. the flexible products cannibalize the higher fare products.

Figure 4 plots the revenue enhancement from offering a flexible product when the demand is $d = [300, 75]'$, i.e. demand for product 1 is higher than capacity. Not surprisingly the revenue enhancement is higher in this case. The maximum enhancement is 8.03% and occurs at $\beta = 0.525$; and the revenue is enhanced for $\beta \in (0.33, 0.8)$.

In the case where the demand $d = [75, 300]'$, it is obvious that the intent behind offering the flexible product is to route some of the excess demand from product 2 to product 1. However, by offering the flexible product to customers of *both* the products, the firm is unnecessarily subsidizing customers requesting product 1. An alternative is to offer the flexible product *only* to customers requesting product 2. Figure 5 plots the revenue enhancement when the flexible product is only offered to customers requesting product 2. In this case, revenue enhancement occurs for all values of β that induce a switch-over, i.e. $\beta \in (0, 0.8)$. The maximum enhancement still occurs at $\beta = 0.635$;

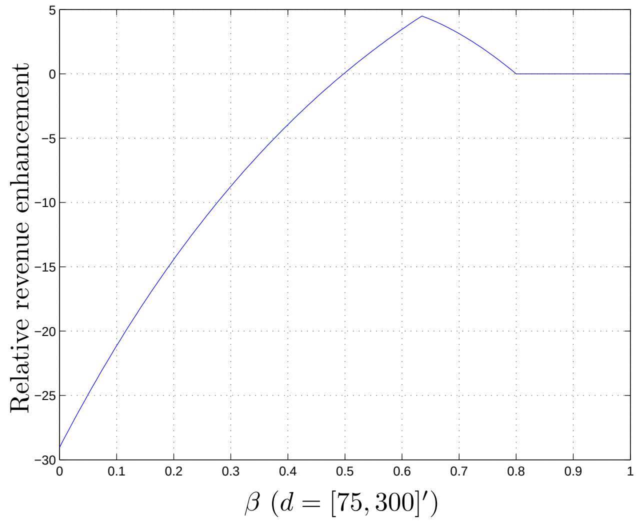

Figure 3: Revenue enhancement when flexible offered to customers of both products

however, now it is 6.82%. Note that since product 1 is the higher quality product, one may be able to induce demand for the flexible product without offering any discount!

In this section, *revenue* refers to the deterministic value function $V^D(c)$. However, the asymptotic optimality result (see Lemma 1 and Corollary 1) implies that for a sufficiently long time horizon T and large capacity c, similar conclusions will hold for the stochastic value function $V(c)$.

2.6 Flexible products and capacity allocation under demand uncertainty

In this section we investigate the value of flexible products in mitigating the effects of demand uncertainty. We consider the deterministic control problem on a network with

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad p = \begin{bmatrix} 600 \\ 600 \end{bmatrix}.$$

In this section we assume that the demand for the two products is uncertain. (This is analogous to assumption that in the stochastic model the arrival rate is uncertain.) Concretely, the demand takes values

$$d_0 = \begin{bmatrix} 300 \\ 150 \end{bmatrix}, \quad d_1 = d_0 + \begin{bmatrix} 100 \\ -50 \end{bmatrix}, \quad d_2 = d_0 + \begin{bmatrix} -100 \\ 50 \end{bmatrix}.$$

We consider two different models for demand uncertainty:

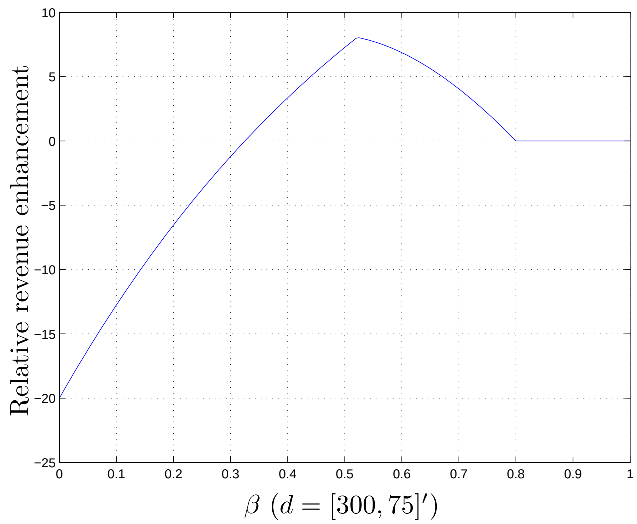

Figure 4: Revenue enhancement when flexible offered to customers of both products

(i) Bayesian model: In this model, we assume that there is a prior distribution π on the set $\mathcal{D} = \{d_0, d_1, d_2\}$. The distribution π is known to the firm; however, the actual realization is not known. In our computations, we set π to the uniform distribution $\pi = \frac{1}{3}[1, 1, 1]'$.

(ii) Robust model: In this model, we do not make any distributional assumptions. The only information available to the firm is that the demand vector $d \in \text{conv}(\mathcal{D})$.

In this section we assume that the company can choose the level of capacity to offer for each product. We assume that the cost of capacity c is linear and the cost vector $q = [450, 450]'$. This is a gross simplification – in any practical setting, there are fixed costs and often capacity can only be purchased in multiples of a fixed quantity. However, we believe that the simple setting of linear costs allows one to concentrate on studying the benefits of offering flexible products.

2.6.1 Capacity allocation in the Bayesian setting

We assume that the firm has to commit to a certain capacity *before* observing demand realization; however, the firm makes the decision of how much of this demand to admit *after* observing the realized demand. This assumption is consistent with the deterministic control problem (15) where the allocation x is chosen *after* observing the demand. Thus, the capacity allocation problem in

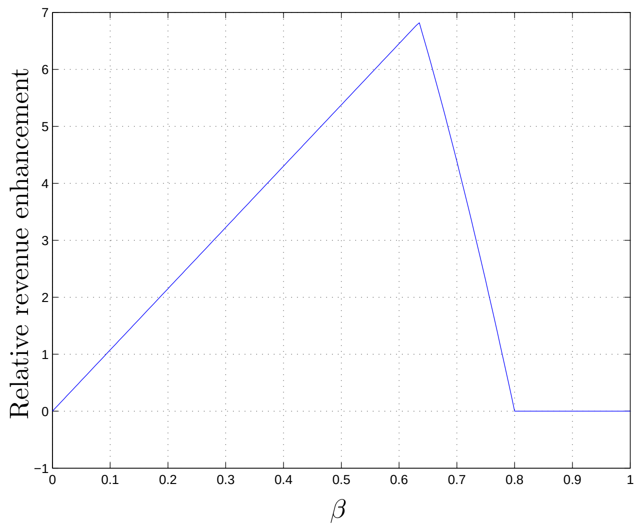

Figure 5: Revenue enhancement when flexible offered only to customers of product 2

this setting is given by

$$\begin{array}{llllll} \max & -q'c + \sum_{l=1}^{3} \pi_l (p'x^{(l)}) & & & & \\ \text{subject to} & -c \;+\; Ax^{(l)} & \leq & 0, & l = 1,2,3, & \\ & \quad\quad\quad x^{(l)} & \leq & d^l, & l = 1,2,3, & \\ & c, \quad\quad\; x^{(l)} & \geq & 0, & l = 1,2,3, & \end{array} \tag{9}$$

where the decision vectors are the installed capacity, $c \in \Re^2_+$, and the vector of admitted requests, $x^{(l)} \in \Re^2_+$, in the three scenarios $l = 1,2,3$.

We now introduce the flexible product $F = \{1,2\}$ at a discount β, i.e. $r = \beta \min\{p_1, p_2\}$. The capacity allocation problem is now given by

$$\begin{array}{lllll} \max & -q'c + \sum_{l=1}^{3} \pi_l (p'x^{(l)} + r'z^{(l)}) & & & \\ \text{subject to} & -c \;+\; Ax^{(l)} \;+\; Az^{(l)} & \leq & 0, & l = 1,2,3, \\ & \quad\quad\quad x^{(l)} & \leq & (1-\kappa(\beta))d^l, & l = 1,2,3, \\ & \quad\quad\quad z_1^{(l)} + z_2^{(l)} & \leq & \kappa(\beta)(d_1^l + d_2^l), & l = 1,2,3, \\ & c, \quad x^{(l)}, \quad z^{(l)} & \geq & 0, & l = 1,2,3. \end{array} \tag{10}$$

Figure 6 plots the relative increase in revenue as a function of the discount β. The maximum revenue enhancement is 4.5% and it occurs at $\beta = 0.635$. Figure 7 plots the installed capacity as a function of β. Since the flexible products are able to smooth out the uncertainty in the demand by moving demand to the underutilized resource, the firm is able to install a higher capacity and

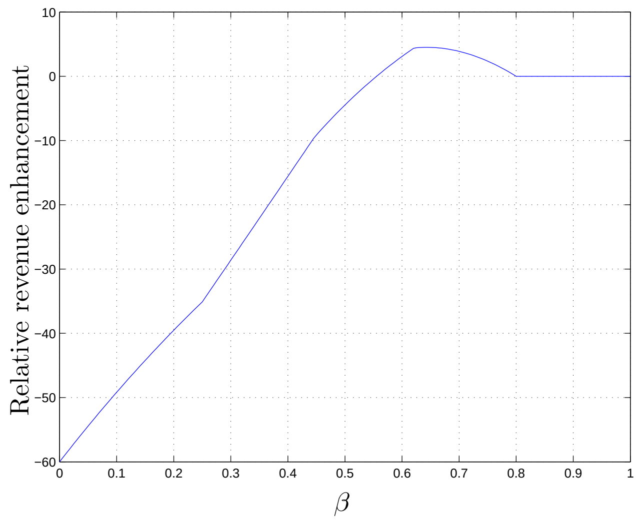

Figure 6: Revenue enhancement as function of β

generate a higher revenue by offering flexible products when β is moderately large. As β decreases, the installed capacity increases; however, the increase in accepted requests is not able to compensate for the loss in revenue per customer and the total revenue fall. Finally, at very small β most of demand switches to flexibles and the small revenue generated from the flexibles does not warrant adding extra capacity; consequently, at very small β the installed capacity starts declining.

2.6.2 Capacity allocation in the robust setting

The setting is that of the previous section, namely the firm has to commit to a certain capacity *before* observing demand realization; however, the firm makes the decision of how much of this demand to admit *after* observing the realized demand. The capacity allocation in the robust setting is given by

$$\max \quad -q'c + \min_{d \in \mathrm{conv}(\mathcal{D})} \{\max \{r'x : Ax \leq c, x \leq d\}\}. \tag{11}$$

The robust setting can be thought of as a game between the firm and nature: the firm chooses the capacity c, the nature chooses $d \in \mathrm{conv}(\mathcal{D})$, and then the firm gets to choose the allocation x. For details on the robust optimization framework, see [4].

The function

$$f(d) = \max \left\{ r'x : Ax \leq c, 0 \leq x \leq d \right\},$$

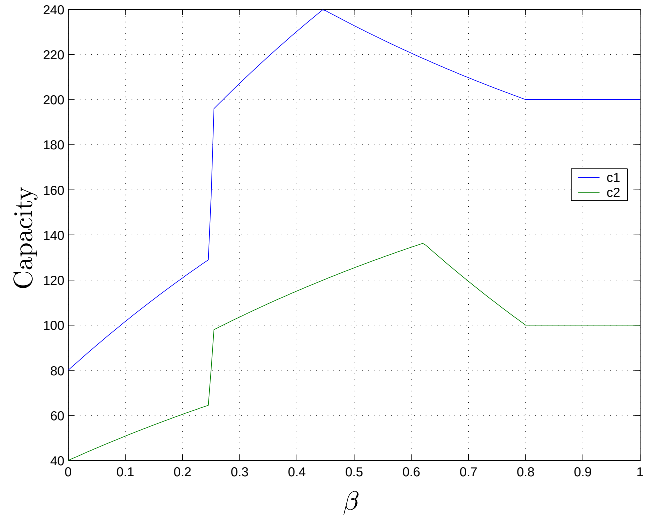

Figure 7: Installed capacity as function of β

is concave in d. Therefore,

$$\min_{d\in \mathrm{conv}(\mathcal{D})}\{f(d)\} = \min_{d\in\mathcal{D}}\{f(d)\},$$

and (11) can be simplified as follows.

$$\begin{array}{rlllll} \max & -q'c + v & & & & \\ \text{subject to} & -c & + \; Ax^{(l)} & & \leq & 0, \quad l = 1,2,3, \\ & & - \; p'x^{(l)} & + \; v & \leq & 0, \\ & & x^{(l)} & & \leq & d^l, \quad l = 1,2,3, \\ & c, & x^{(l)} & & \geq & 0, \quad l = 1,2,3, \end{array} \tag{12}$$

The robust capacity allocation problem when the firm offers the flexible product with $F = \{1, 2\}$ at a discount β, i.e. $r = \beta \min\{p_1, p_2\}$, is given by

$$\begin{array}{rllllll} \max & -q'c + v & & & & & \\ \text{subject to} & -c & + \; Ax^{(l)} & + \; Az^{(l)} & & \leq & 0, \quad l = 1,2,3, \\ & & - \; p'x^{(l)}) & - \; r(z_1^{(l)} + z_2^{(l)}) & + \; v & \leq & 0, \\ & & x^{(l)} & & & \leq & (1-\kappa(\beta))d^l, \quad l = 1,2,3, \\ & & & z_1^{(l)} + z_2^{(l)} & & \leq & \kappa(\beta)(d_1^l + d_2^l), \quad l = 1,2,3, \\ & c, & x^{(l)}, & z^{(l)} & & \geq & 0, \quad l = 1,2,3. \end{array} \tag{13}$$

Figure 8 plots the relative increase in revenue as a function of the discount β. The maximum

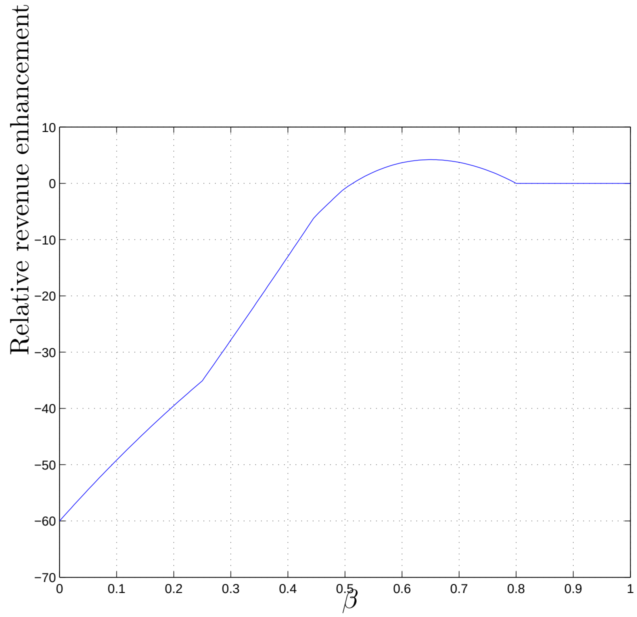

Figure 8: Revenue enhancement as function of β

revenue enhancement is 4.5% and it occurs at $\beta = 0.635$. Figure 9 plots the installed capacity as a function of β.

3 Network choice model with flexible products

The model in Section 2 assumed that the arrival rates of requests for different products were independent of the set of products being offered. However, it is well known that potential customers who find their most desired product unavailable often switch to another product [18]. To model this situation we consider a model in which demand rates depend upon which products are available. At any $t \in [0, T]$ the requests for specific (resp. flexible) products arrive according to independent Poisson processes with a rate vector $\lambda(S) \in \Re^n_+$ (resp. $\gamma(S) \in \Re^f_+$) that depends on the set $S \subseteq N \cup F$ of products (both flexible and specific) being offered. We assume that $S = \emptyset$ is always possible and that $\lambda(\emptyset) = 0$ and $\gamma(\emptyset) = 0$. Note that the arrival rates are assumed to be time independent. This is an important assumption and is further discussed in Section 6. As before, the state of the network is given by the pair of vectors $(s, y) \in \mathcal{Z}^m_+ \times \mathcal{Z}^f_+$, where s denotes the residual capacity and y denote the vector of accepted requests for flexible products. The firm's control in this model is the choice of the set of products S to offer at any given time $t \in S$. Unlike the previous section, the firm *must* accept all requests for the set of products S open at $t \in [0, T]$.

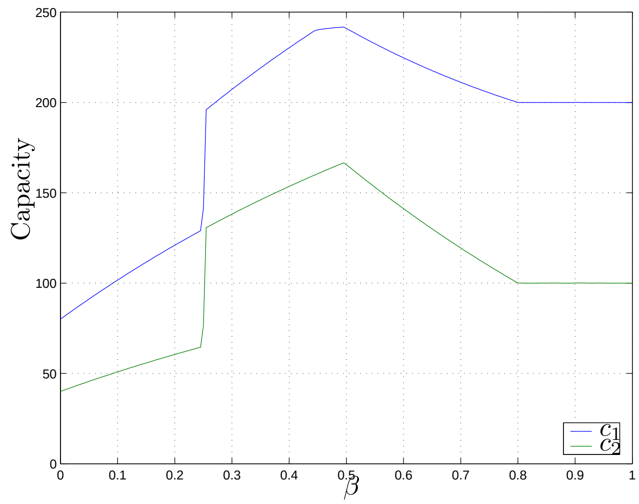

Figure 9: Installed capacity as function of β

3.1 Dynamic programming formulation

Let $V(s,y,t)$ denote the maximum achievable revenue starting from the state (s,y) at time t. Then $V(s,y,t)$ is characterized by the HJB recursion

$$\begin{aligned} -\frac{\partial V(s,y,t)}{\partial t} &= \max_{S\subseteq N\cup F}\Big\{\sum_{i=1}^{n}\lambda_i(S)\big(p_i+V(s-A_i,y,t)-V_t(s,y,t)\big) \\ &\qquad + \sum_{k=1}^{f}\gamma_k(S)\big(r_k+V(s,y+e_k,t)-V(s,y,t)\big)\Big\}, \qquad (14) \end{aligned}$$

and the boundary conditions

$$\begin{aligned} V(s,y,T) &= 0, \qquad (s,y)\in\mathcal{A}, \\ V(s,y,t) &= -\infty, \quad (s,y)\notin\mathcal{A}, \quad t\in[0,T]. \end{aligned}$$

The goal of the recursion is to compute $V(c)=V(c,0,0)$. Note that the *only* control in (14) is the choice of the set of products $S\subseteq N\cup F$; once the set is chosen the firm *must* accept all arriving requests. For most practical networks, the complexity of computing $V(c)$ is prohibitive, and one has to resort to heuristics. In the next section, we construct an asymptotically optimal policy by solving a deterministic control problem. However, unlike in Section 2, we are *not* able to extend this result to the case when the requests arrive according to a stationary point process.

3.2 Deterministic control problem with customer choice

In this section we construct a deterministic approximation for the stochastic control problem. The deterministic approximation is obtained by assuming that the requests arrive according to a deterministic rate. Since the rates are assumed to be time independent, the firm's problem reduces to maximizing revenue by selecting a collection $S_l \subseteq N \cup F$, $l = 1, \ldots, L$, of subsets of products to offer and corresponding *lengths* of time $t(S_l)$, $l = 1 \ldots, L$, for which to offer these sets, i.e. the decision problem reduces to the following LP

$$
\begin{array}{rclrcrcl}
V^D(c) & = & \max & \sum_S \big(p'\lambda(S) + r'\gamma(S)\big)t(S) & & & & \\
& & \text{subject to} & A\Big(\sum_S \lambda(S)t(S)\Big) & + & Bz & \leq & c, \\
& & & \sum_S \gamma(S)t(S) & - & Uz & = & 0, \\
& & & \sum_S t(S) & & & \leq & T, \\
& & & t(S) & & & \geq & 0, \quad \forall S \subseteq N \cup F, \\
& & & & & z & \geq & 0.
\end{array} \tag{15}
$$

The decision variables in (15) are the times $\{t(S) : S \subseteq N \cup F\}$ and the composition variables z, i.e. the total number of variables is $2^{n+f} + \sum_{k=1}^{f} f_k - 1$. Since (15) has $m+f+1$ constraints, there exists an optimal solution with no more than $m + f + 1$ non-zero variables. Moreover, suppose the k-th flexible product is offered for any period. Then at least one of the components of the z_k vector must be strictly positive. On the other hand if the k-th flexible product is never offered then the row corresponding to k in the constraint $\sum_S \gamma(S)t(S) - Uz = 0$ is effectively redundant and can be dropped without affecting the solution. Thus, we have the following result.

Proposition 3 *There exists an optimal solution of the LP (15) with $t(S) > 0$ for at most $m + 1$ subsets $S \subseteq N \cup F$.*

Since the number of sets in the optimal collection $\mathcal{S}^*$ is relatively small, namely $m + 1$, column generation can be an efficient algorithm for solving (15). However, as we shall see, the effectiveness of column generation depends upon the consumer choice model defining $\gamma(S)$ and $\lambda(S)$.

3.2.1 Column generation algorithm

A column generation algorithm for solving an LP works with a subset of variables, i.e. a subset of columns of the constraint matrix, and then suitably augments this collection one column, i.e. one

1. Select an initial collection $\mathcal{S}^{(0)}$ of subsets of $N \cup F$. Set $q \leftarrow 0$.

2. Solve the restricted primal LP (16) corresponding to $\mathcal{S}^{(q)}$. Let $(t^{(q)}, z^{(q)})$ denote the optimal solution of the restricted primal and let $(u^{(q)}, v^{(q)}, \beta^{(q)})$ denote the corresponding optimal dual vector.

3. Using the explicit construction in Lemma 2 compute the "minimum set" S corresponding to $(u^{(q)}, v^{(q)}, \beta^{(q)})$. Let
$$l(S) = \frac{(u^{(q)})' A\lambda(S) + (v^{(q)})'\gamma(S) + \beta}{p'\lambda(S) + r'\gamma(S)},$$
denote the "length" of set S.

4. If $l(S) \geq 1$, the solution of the current restricted primal problem is optimal for the full primal LP (15). Stop.

5. Else, set $q \leftarrow q + 1$, $\mathcal{S}^{(q)} \leftarrow \mathcal{S}^{(q-1)} \cup S$. Return to step 2.

Figure 10: Column generation algorithm

variable, at a time to compute an optimal solution. A column generation algorithm for solving (15) is to choose a collection $\mathcal{S}^{(0)}$ of subsets of $N \cup F$ and solve the restricted LP

$$\begin{array}{rrcrcl} \max & \sum_{S \in \mathcal{S}^{(0)}} \big(p'\lambda(S) + r'\gamma(S)\big) t(S) & & & & \\ \text{subject to} & A\Big(\sum_{S \in \mathcal{S}^{(0)}} \lambda(S) t(S)\Big) & + & Bz & \leq & c, \\ & \sum_{S \in \mathcal{S}^{(0)}} \gamma(S) t(S) & - & Uz & = & 0, \\ & \sum_{S \in \mathcal{S}^{(0)}} t(S) & & & \leq & T, \\ & t(S) & & & \geq & 0, \quad \forall S \in \mathcal{S}^{(0)}, \\ & & & z & \geq & 0. \end{array} \tag{16}$$

Notice that we have kept all the z variables in the restricted LP: this is not necessary – z_k must be included only if the k-th flexible product is contained in some set $S \in \mathcal{S}^{(0)}$.

Let $(t^{(0)}, z^{(0)})$ denote the optimal solution of (16) and let $(u^{(0)}, v^{(0)}, \beta^{(0)})$ denote the optimal dual solution, i.e. the reduced cost at the optimal solution. The solution $(t^{(0)}, z^{(0)})$ is optimal for the full LP (15) provided the "cost" $(u^{(0)})' A\lambda(S) + (v^{(0)})'\gamma(S) + \beta^{(0)}$ of the columns $S \notin \mathcal{S}^{(0)}$ is more than its contribution to revenue $p'\lambda(S) + r'\gamma(S)$, i.e.

$$\min_{S \subseteq N \cup F} \left\{ \frac{(u^{(0)})' A \lambda(S) + (v^{(0)})' \gamma(S) + \beta^{(0)}}{p'\lambda(S) + r'\gamma(S)} \right\} \geq 1, \qquad \forall S \notin \mathcal{S}^{(0)}. \tag{17}$$

Suppose there exists a set S that violates (17). Then the column generation algorithm augments the collection $\mathcal{S}^{(0)}$ with the set that achieves the minimum in (17) and repeats the process.

Column generation is successful only if (17) can be efficiently solved for all values of (u, v, β). Clearly, this is impossible for arbitrary assignments of the demand rates $\{\lambda(S), \gamma(S)\}, S \subseteq N \cup F$. However, in practical applications, we can restrict ourselves to rates $\{\lambda(S), \gamma(S)\}$ that are given by some rational consumer choice models. In this paper we consider the class of consumer choice models defined below.

Definition 1 *Let $g, \tilde{g} \in \Re^n_+$ and $h, \tilde{h} \in \Re^f_+$ be given fixed positive vectors. Then the demand vectors $\lambda(S)$ and $\gamma(S)$ are defined as follows:*

$$\lambda_j(S) = \begin{cases} \frac{g_j}{\tilde{g}(S \cap N) + \tilde{h}(S \cap F) + 1}, & j \in S \cap N, \\ 0, & \text{otherwise}, \end{cases} \tag{18}$$

$$\gamma_k(S) = \begin{cases} \frac{h_k}{\tilde{g}(S \cap N) + \tilde{h}(S \cap F) + 1}, & k \in S \cap F, \\ 0, & \text{otherwise}, \end{cases} \tag{19}$$

where $a(S) = \sum_{j \in S} a_j$.

Both independent demands and the Multinomial Logit (MNL) model as well as other "attraction models" such as the probit are included in the class of customer choice model in Definition 1[2].

The following lemma provides an explicit solution for the column generation problem (17) corresponding to the choice model in Definition 1.

Lemma 2 *Let $u \in \Re^n_+$, $v \in \Re^f_+$, $\beta \geq 0$ and assume that consumer choices are given by a model that belongs to the class defined in Definition 1. Define*

$$x_j = \frac{u' A_j g_j + \beta \tilde{g}_j}{g_j p_j}, \quad j \in N, \tag{20}$$

[2]Ben-Akiva and Lerman [3] show how the Multinomial Logit model implies the choice probabilities given in Definition 1. Talluri and Van Ryzin [18] apply the Multinomial Logit Model to a single-leg case. Anderson, Palma and Thisse [1] provide an overview of attraction models.

where A_j is the j-th column of the capacity utilization matrix A; and

$$y_k = \frac{v_k h_k + \beta \tilde{h}_k}{h_k r_k}, \quad k \in F. \tag{21}$$

Let $\{w_l : l = 1, \ldots, n+f\}$ denote the values $\{x_j : j \in N\}$ and $\{y_k : k \in F\}$ arranged in increasing order. Let $S_{w_l} = \big\{j \in N : x_j \leq w_l\big\} \cup \big\{k \in F : y_k \leq w_l\big\}$ and let l^* denote the largest index l such that

$$\sum_{j \in S_{w_l} \cap N} g_j p_j \big(x_j - w_l\big) + \sum_{k \in S_{w_l} \cap F} h_k r_k \big(y_k - w_l\big) + \beta \geq 0.$$

Then

$$\min_{S \subseteq N \cup F} \left\{ \frac{u' A \lambda(S) + v' \gamma(S) + \beta}{p' \lambda(S) + r' \gamma(S)} \right\} = \frac{u' A \lambda(S_{w_{l^*}}) + v' \gamma(S_{w_{l^*}}) + \beta}{p' \lambda(S_{w_{l^*}}) + r' \gamma(S_{w_{l^*}})},$$

i.e. $S_{w_{l^*}}$ is an optimal set.

<u>Proof</u>

Suppose $\lambda(S)$ and $\gamma(S)$ are as in (18) and (19) respectively. Then

$$\frac{u' A \lambda(S) + v' \gamma(S) + \beta}{p' \lambda(S) + r' \gamma(S)} = \frac{\Big(\sum_{j \in S \cap N} g_j p_j x_j + \sum_{k \in S \cap F} h_k r_k y_k\Big) + \beta}{\Big(\sum_{j \in S \cap N} g_j p_j + \sum_{k \in S \cap F} h_k r_k\Big)}.$$

Thus,

$$\min_{S \subseteq N \cup F} \left\{ \frac{u' A \lambda(S) + v' \gamma(S) + \beta}{p' \lambda(S) + r' \gamma(S)} \right\} \geq \zeta \tag{22}$$

if and only if

$$\Big(\sum_{j \in S \cap N} g_j p_j \big(x_j - \zeta\big) + \sum_{k \in S \cap F} h_k r_k \big(y_k - \zeta\big) \Big) + \beta \geq 0, \quad \forall S \subseteq N \cup F. \tag{23}$$

Let ζ^* denote the optimal value of the optimization problem in (22). Since $\beta \geq 0$, (23) implies that

$$\zeta^* \geq w_1. \tag{24}$$

Let S^* denote the set that achieves the minimum value ζ^*. We will show by contradiction that $S^* \subseteq S_{\zeta^*} = \big\{j \in N : x_j \leq \zeta^*\big\} \cup \big\{j \in F : y_j \leq \zeta^*\big\}$. Eq. (24) implies that $S_{\zeta^*} \neq \emptyset$. Suppose $S^* \not\subseteq S_{\zeta^*}$. Then it follows that

$$\begin{aligned} & \Big(\sum_{j \in S_{\zeta^*} \cap N} g_j p_j \big(x_j - \zeta^*\big) + \sum_{k \in S_{\zeta^*} \cap F} h_k r_k \big(y_k - \zeta^*\big) \Big) + \beta \\ < & \ \Big(\sum_{j \in S^* \cap N} g_j p_j \big(x_j - \zeta^*\big) + \sum_{k \in S^* \cap F} h_k r_k \big(y_k - \zeta^*\big) \Big) + \beta \\ = & \ 0. \end{aligned}$$

Rearranging terms, the above inequality implies that

$$\frac{\left(\sum_{j\in S_{\zeta^*}\cap N} g_j p_j x_j + \sum_{k\in S_{\zeta^*}\cap F} h_k r_k y_k\right) + \beta}{\left(\sum_{j\in S_{\zeta^*}\cap N} g_j p_j + \sum_{k\in S_{\zeta^*}\cap F} h_k r_k\right)} < \zeta^*.$$

A contradiction. Moreover,

$$\begin{aligned}
&\frac{u'A\lambda(S_{\zeta^*}) + v'\gamma(S_{\zeta^*}) + \beta}{p'\lambda(S_{\zeta^*}) + r'\gamma(S_{\zeta^*})} \\
&= \frac{\left(\sum_{j\in S^*\cap N} g_j p_j x_j + \sum_{k\in S^*\cap F} h_k r_k y_k + \beta\right) + \left(\sum_{j\in (S_{\zeta^*}\setminus S^*)\cap N} g_j p_j x_j + \sum_{k\in (S_{\zeta^*}\setminus S^*)\cap F} h_k r_k y_k\right)}{\left(\sum_{j\in S_{\zeta^*}\cap N} g_j p_j + \sum_{k\in S_{\zeta^*}\cap F} h_k r_k\right)}, \\
&= \frac{\zeta^*\left(\sum_{j\in S^*\cap N} g_j p_j + \sum_{k\in S^*\cap F} h_k r_k\right) + \left(\sum_{j\in (S_{\zeta^*}\setminus S^*)\cap N} g_j p_j x_j + \sum_{k\in (S_{\zeta^*}\setminus S^*)\cap F} h_k r_k y_k\right)}{\left(\sum_{j\in S_{\zeta^*}\cap N} g_j p_j + \sum_{k\in S_{\zeta^*}\cap F} h_k r_k\right)}, \quad (25) \\
&\leq \frac{\zeta^*\left(\sum_{j\in S^*\cap N} g_j p_j + \sum_{k\in S^*\cap F} h_k r_k\right) + \zeta^*\left(\sum_{j\in (S_{\zeta^*}\setminus S^*)\cap N} g_j p_j + \sum_{k\in (S_{\zeta^*}\setminus S^*)\cap F} h_k r_k\right)}{\left(\sum_{j\in S_{\zeta^*}\cap N} g_j p_j + \sum_{k\in S_{\zeta^*}\cap F} h_k r_k\right)}, \quad (26) \\
&\leq \zeta^*,
\end{aligned}$$

where (25) follows from the fact that $\left(u'A\lambda(S^*)+v'\gamma(S^*)+\beta\right) = \zeta^*\left(\sum_{j\in S^*\cap N} g_j p_j + \sum_{k\in S^*\cap F} h_k r_k\right)$, and (26) follows from the fact that $x_j \leq \zeta^*$ and $y_k \leq \zeta^*$for all $x_j, y_k \in S_{\zeta^*}$. Thus, the set S_{ζ^*} is also an optimal set. Therefore, one can restrict attention to sets of the form

$$S_\zeta = \left\{j \in N : x_j \leq \zeta\right\} \cup \left\{k \in F : y_k \leq \zeta\right\}$$

for $\zeta \geq w_1$. The result now follows by recognizing that the sets S_ζ are identical for $w_l \leq \zeta < w_{l+1}$, $l = 1, \ldots, n+f$.

∎

The column generation problem can be solved by a simple sorting because the function

$$g(\mu,\nu) = \frac{\left(\sum_{i\in N} g_i p_i x_i \mu_i + \sum_{k\in F} h_k r_k y_k \nu_k\right) + \beta}{\left(\sum_{i\in N} g_i p_i \mu_i + \sum_{k\in F} h_k r_k \nu_k\right)},$$

is quasi-concave on the set $(\mu,\nu) \in [0,1]^{n+f}$. (This has been implicitly proved above.) Therefore, the optimal value is achieved at the boundary. See [16] for an alternate proof of Lemma 2 that exploits this fact. The characterization of the "minimum set" in Lemma 2 leads to the column generation algorithm in Figure 10.

The optimal basis for the restricted primal corresponding to $\mathcal{S}^{(q)}$ is feasible for the restricted primal corresponding to $\mathcal{S}^{(q+1)}$. Therefore, it can be used as the starting basis. Typically, the re-optimization only requires one to four pivot steps.

3.3 Asymptotic optimality of the deterministic control

In this section we compare the value $V(c)$ achievable in the stochastic problem to the value $V^D(c)$ achievable in the deterministic control problem. The results in this section are motivated by the work of van Ryzin and Liu [16].

Proposition 4 *The deterministic value function $V^D(c)$ is an upper bound for the stochastic value function $V(c)$.*

Proof

Fix any feasible control policy ρ for the stochastic problem. Let $(n_x^\rho, n_y^\rho, n_z^\rho)$ denote the random number of specific and flexible requests admitted by the policy ρ, and the corresponding distribution variables. Since the policy ρ is feasible, we must have

$$An_x^\rho - Bn_z^\rho \leq c, \quad n_y^\rho = Un_z^\rho, \quad \text{a.s.}$$

Thus, we have that

$$A\mathbf{E}[n_x^\rho] - B\mathbf{E}[n_z^\rho] \leq c, \quad \mathbf{E}[n_y^\rho] = U\mathbf{E}[n_z^\rho]. \tag{27}$$

The model dynamics and Wald's Lemma imply that

$$\begin{aligned} \mathbf{E}[n_x^\rho] &= \textstyle\sum_{S\subseteq N\cup F} \lambda(S)\mathbf{E}[t^\rho(S)], \\ \mathbf{E}[n_y^\rho] &= \textstyle\sum_{S\subseteq N\cup F} \gamma(S)\mathbf{E}[t^\rho(S)], \end{aligned} \tag{28}$$

where $t^\rho(S)$, $S \subseteq N \cup F$, denotes the random time for which the policy ρ offers the set S. From (27) and (28) it follows that $\big(\{\mathbf{E}[t^\rho(S)] : S \subseteq N \cup F\}, \mathbf{E}[z^\rho]\big)$ is feasible for the LP (15). Therefore, the expected revenue R_ρ of the policy ρ satisfies

$$R_\rho = \mathbf{E}\big[p'n_x^\rho + r'n_y^\rho\big] = \sum_{S\subseteq N\cup F} (p'\lambda(S) + r'\gamma(S))\mathbf{E}[t^\rho(S)] \leq V^D(c).$$

Since the policy ρ is arbitrary, the result follows.

∎

Next, we prove a lower bound for the stochastic value function $V(c)$. Let $n_x(S,t)$ (resp. $n_y(S,t)$)) denote the random number of requests for specific (resp. flexible) products. Let $\chi_i(S,t)$ (resp. $\chi_k(S,t)$) denote the coefficient of variation of the i-th specific (resp. k-th flexible) product with the set S is offered for t units of time. Recall that in this section we restrict ourselves to the special case where the requests arrive according to a Poisson process. Therefore,

$$\chi_i(S,t) = \frac{\sqrt{\mathbf{Var}(n_{xi}(S,t))}}{\mathbf{E}[n_{xi}(t)]} = \frac{1}{\sqrt{\lambda_i(S)t}}, \quad i = 1,\ldots,n,$$
$$\chi_k(S,t) = \frac{\sqrt{\mathbf{Var}(n_{yk}(t))}}{\mathbf{E}[n_{yk}(t)]} = \frac{1}{\sqrt{\gamma_k(S)t}}, \quad k = 1,\ldots,f.$$

Let

$$\chi(S,t) = \max\Big\{ \max_{1\le i\le n}\{\chi_i(S,t)\}, \max_{1\le k\le f}\{\chi_k(S,t)\}\Big\}.$$

Lemma 3 *Let $(\{t(S_l^*)\}_{l=1,\ldots,L}, z^*)$ denote the optimal solution of the LP (15). Suppose $\chi(t^*) = \max_{1\le l\le L}\{\chi(S_l^*,t_l^*)\} < \frac{1}{\sqrt{n+f}}$. Then $V^D(c) \ge V(c) \ge V^D(c)\Big(1 - 1.89\big((n+f)\chi^2(t^*)\big)^{\frac{1}{3}}\Big)$.*

Proof

Define a control policy $\bar\rho$ as follows:

(a) Offer the sets $S_1^*,\ldots,S_L^*$ in sequence.

(b) Offer the set S_l^* for a random time

$$\tau_l = \min\Big\{t(S_l^*), \inf\big\{t : n_x(S_l^*,t) \ge \lfloor\lambda(S_l^*)t_l^*\rfloor, n_y(S_l^*,t) \ge \lfloor\gamma(S_l^*)t_l^*\rfloor\big\}\Big\}.$$

The choice of the time τ_l ensures that the policy $\bar\rho$ does not admit more requests than it can satisfy. Therefore, the policy $\bar\rho$ is feasible. Note that one can easily obtain a superior policy by continuing to offer the final set S_L^* until either time or capacity runs out.

Fix $\epsilon \in (0,1)$. From the definition of τ_l, we have that

$$\begin{aligned}
\mathbf{P}(\tau_l \ge (1-\epsilon)t_l^*) &= \mathbf{P}\Big(n_x\big(S_l^*,(1-\epsilon)t_l^*\big) < \lambda(S_l^*)t_l^*, n_y\big(S_l^*,(1-\epsilon)t_l^*\big) < \gamma(S_l^*)t_l^*\Big), \\
&\ge 1 - \Big(\sum_{i\in S_l^*\cap N} \mathbf{P}\big(n_{xi}(S_l^*,(1-\epsilon)t_l^*) \ge \lambda_i(S_l^*)t_l^*\big) \\
&\qquad + \sum_{k\in S_l^*\cap F} \mathbf{P}\big(n_{yk}(S_l^*,(1-\epsilon)t_l^*) \ge \gamma_k(S_l^*)t_l^*\big)\Big). \qquad (29)
\end{aligned}$$

Marshall's inequality bounds a generic term in (29) as follows.

$$\mathbf{P}\Big(n_{xi}(S_l^*,(1-\epsilon)t_l^*) \ge \lambda_i(S_l^*)t_l^*\Big) \le \frac{\chi_i^2(S^*,(1-\epsilon)t^*)}{\epsilon^2 + \chi_i^2(S^*,(1-\epsilon)t^*)} \le \frac{\chi^2(S_l^*,t_l^*)}{\epsilon^2} \le \frac{\chi^2(t^*)}{\epsilon^2},$$

where the second inequality follows from the fact that $\chi(S,t)$ is monotonically increasing in t; and the last inequality follows from the definition of $\chi(t^*)$. Substituting this bound in (29) we get

$$P\big(\tau_l \geq (1-\epsilon)t_l^*\big) \geq 1 - \frac{(n+f)\chi^2(t^*)}{\epsilon^2}. \tag{30}$$

Since τ_l is a bounded stopping time and $n_{xi}(S_l^*, t)$ is Poisson process with intensity $\lambda_i(S_l^*)$, the PASTA property implies that

$$\begin{aligned} \mathbf{E}[n_{xi}(S_l^*, \tau_l)] &= \lambda_i(S^*)\mathbf{E}[\tau_l], \\ &\geq \lambda_i(S_l^*)t_l^*(1-\epsilon)\mathbf{P}(\tau_l \geq (1-\epsilon)\tau_l), \\ &\geq \lambda_i(S_l^*)t_l^*(1-\epsilon)\Big(1 - \frac{(n+f)\chi^2(t^*)}{\epsilon^2}\Big). \end{aligned}$$

Therefore, the expected reward $R_{\bar{\rho}}$ of the policy $\bar{\rho}$ satisfies

$$\begin{aligned} R_{\bar{\rho}} &= \sum_{l=1}^{L}\Big(\sum_{i \in S_l^* \cap N} p_i \mathbf{E}[n_{xi}(S_l^*, \tau_l)] + \sum_{k \in S_l^* \cap F} r_k \mathbf{E}[n_{yk}(S_l^*, \tau_l)]\Big), \\ &\geq (1-\epsilon)\Big(1 - \frac{(n+f)\chi^2(t^*)}{\epsilon^2}\Big)V^D(c). \end{aligned}$$

Maximizing over ϵ, we get $\epsilon^* = \big(2(n+f)\chi^2(t^*)\big)^{\frac{1}{3}}$. The result follows by substituting ϵ^* in the bound above.

∎

Since the proof of Lemma 3 uses the PASTA property, relaxing the Poisson assumption appears to be difficult. An asymptotic result could perhaps be established when the arrival process is a renewal process. The following Corollary establishes an asymptotic optimality result.

Corollary 2 *Suppose $\chi(t^*) = \max_{1\leq l\leq L}\{\chi(S_l^*, t_l^*)\}$. Fix $\zeta > 1$ and consider the scaled stochastic problem in initial capacity ζc and time horizon ζT. Then, for all $\zeta > (n+f)\chi^2(t^*)$, the stochastic value function*

$$V(\zeta c) \geq \zeta V^D(c)\Big(1 - 1.89\Big(\frac{(n+f)\chi^2(t^*)}{\zeta}\Big)^{\frac{1}{3}}\Big).$$

The proof of this result is identical to that of Corollary 1.

Proposition 3 and Corollary 2 together argue that the deterministic value function $V^D(c)$ closely approximates the stochastic value function $V(c)$, at least in the limit of large T and large capacity c. However, it is not clear how to translate the optimal solution of the deterministic control problem into a good control policy for the stochastic model. One possibility is to use the control policy $\bar{\rho}$ constructed in the proof of Lemma 3 without deflating the times by $(1-\epsilon)$. Although the expected value $R_{\bar{\rho}} \approx V^D(c)$, its performance on a given sample path may not be satisfactory. In any case, there is the issue of selecting the optimal permutation of the sets. (Our preliminary numerical experiments suggest that the performance does not change much with the order of the sets.) Another possibility is to subdivide the interval $[0,T]$ and re-optimize at the end of the interval (or when the capacity allocated to the previous interval runs out). In a yet another approach Van Ryzin and Liu [16] use optimal dual vectors to decompose the n-dimensional HJB recursion into n 1-dimensional HJB recursions. The control policy is constructed by solving the 1-dimensional recursions to optimality. Van Ryzin and Liu [16] report that the performance of the control policy constructed from periodic re-optimization is close to (although, worse than) the policy constructed from the optimal dual variables.

4 Results for an Airline Subnetwork

In this section we report the results of our computational experiments with a subset of American Airlines flights from JFK, LGA, STL, ORD to SFO. The details of the flights in our model are shown in Table 1. There were $n = 26$ specific products consisting of all the direct flights shown in Table 1 and all connections from NYC to SFO satisfying time constraints. There were $f = 5$ flexible products consisting of the O-D pairs: NYC→SFO, NYC→STL, NYC→ORD, ORD→SFO, STL→SFO. The set F_k for each flexible product was set equal to all the flights that served the corresponding O-D pair.

Each O-D pair was assigned equal utility w. The vectors $g, \tilde{g} \in \Re^n_+$ were set to

$$g_i = \exp\left(w - \alpha p_i - \delta t_i\right), \quad \tilde{g}_i = \frac{g_j}{20}, \quad i = 1, \ldots, n, \tag{31}$$

where p_i is the fare for flight i, t_i is the duration of the trip, $\alpha = 0.004$ and $\delta = 0.05$ (the constants α and β were chosen to make the terms in the exponent comparable). For itineraries involving a connection, the flight time was set as the total travel time and the fare was set to be the sum of the fares on the legs.

The fare r_k of the k-th flexible product was set to $r_k = 0.75 \min_{i \in F_k}\{p_i\}$, i.e. the flexible products were sold at a 25% discount with respect to the least expensive specific product in F_k.

	Origin	Destination	Departure Time	Arrival Time	Capacity	Fare
1.	LGA	ORD	0600	0717	172	$194
2.	LGA	ORD	0700	0832	172	194
3.	LGA	ORD	0830	1000	172	194
4.	ORD	SFO	0933	1116	134	287
5.	ORD	SFO	1151	1325	134	567
6.	ORD	SFO	1347	1518	134	567
7.	JFK	SFO	0730	1049	160	567
8.	JFK	SFO	1545	1712	160	567
9.	LGA	STL	0610	0750	176	121
10.	LGA	STL	0750	0940	172	121
11.	JFK	STL	0920	1111	172	121
12.	STL	SFO	1608	1829	176	278
13.	STL	SFO	1940	2213	176	278

Table 1: Flight Subnetwork

The choice model coefficients h and $\tilde{h}$ were set to

$$h_k = \exp\left(\eta w - \alpha r_k - \delta \max_{j \in F_k}\{t_j\}\right), \quad \tilde{h}_k = \frac{h_k}{20}, \quad k = 1, \ldots, f. \tag{32}$$

with $\eta = 0.9$.

Let $R_f(T)$ (resp $R_{nf}(T)$) denote the deterministic value function $V^D(c)$, i.e. the optimal value of (15), when flexible products are offered (resp. are not offered). Note that $R_f(T)$ and $R_{nf}(T)$ both refer to the *total* expected revenue and not the expected revenue per unit time. The relative revenue enhancement $G_f(T)$ from offering flexible products is defined as follows

$$G_f(T) = \Big(\frac{R_f(T) - R_{nf}(T)}{R_{nf}(T)}\Big) \times 100.$$

In this section, we investigate how the gain $G_f(T)$ depends on the various system parameters, such as the time horizon T, the O-D utility w, the discount offered on the flexibles, etc. In the sequel, the LP (15) will be referred to as Flex-LP and the LP of the same form without flexibles as NoFlex-LP. Thus, $R_f(T)$ (resp. $R_{nf}(T)$) is the optimal solution of the Flex-LP (resp. NoFlex-LP).

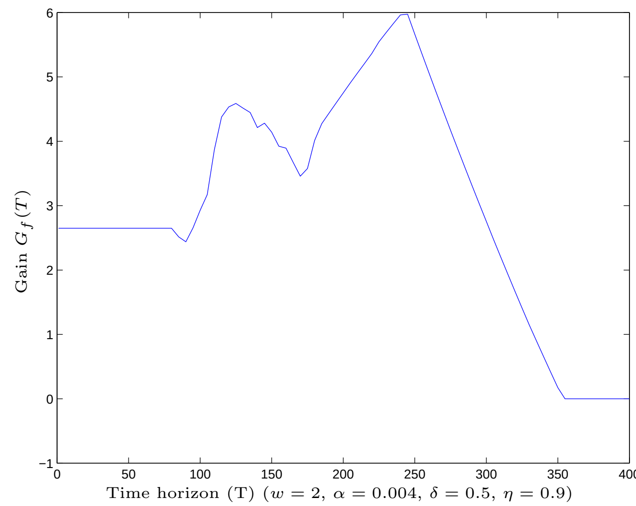

Figure 11: Gain $G_f(T)$ as a function of T

4.1 Gain as a function of the time horizon T

The first set of numerical experiments investigate the relationship between the gain $G_f(T)$ and the time horizon T, keeping all other variables, including the capacity c, constant. A typical plot for the gain $G_f(T)$ as a function of the time horizon T is shown in Figure 11. The gain curve starts out flat; next, it enters a phase where the gain could both increase or decrease; and, finally the gain monotonically decreases to zero. In the instance displayed in Figure 11 the slope of the gain curve is negative immediately after the flat section but this need not always be the case.

In order to gain more insight into the structure of the gain curve, the total revenue $R_f(T)$ and $R_{nf}(T)$ as a function in T is displayed in Figure 12. From LP duality it follows that $R_f(T) = u_f'c + \beta_f T$ (resp. $R_{nf}(T) = u_{nf}'c + \beta_{nf}T$), where (u_f, β_f) (resp. (u_{nf}, β_{nf})) are the optimal dual variables corresponding to the capacity constraints and time constraint in Flex-LP (resp. NoFlex-LP). Since the optimal dual variables (u_f, β_f) and (u_{nf}, β_{nf}) are piecewise constant functions of T, it follows that $R_f(T)$ and $R_{nf}(T)$ are piecewise linear functions of T. For all sufficiently small T, the capacity constraints are slack in both Flex-LP and NoFlex-LP, i.e. the realized demand is insufficient to fill any of the flights in the system. Thus, $R_f(T) = \beta_f(T)T$ and $R_{nf}(T) = \beta_{nf}(T)T$, and the gain

$$G_f(T) = \Big(\frac{R_f(T) - R_{nf}(T)}{R_f(T)}\Big) \times 100 = \Big(\frac{\beta_f - \beta_{nf}}{\beta_{nf}}\Big) \times 100,$$

is constant. In this regime, the gain is primarily due to demand induction as opposed to higher

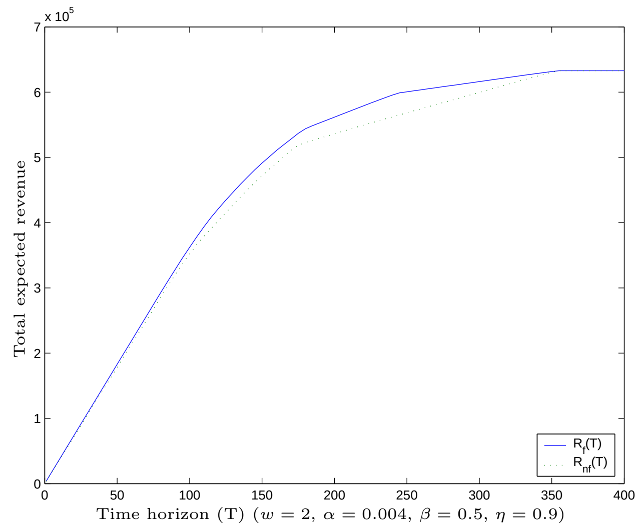

Figure 12: Revenue as a function of T

capacity utilization. We will return to this issue in Section 4.5.

Let T_0 denote the time horizon at which some capacity constraint first becomes active in one of the two LPs. (In Figure 11 this occurs at $T_0 = 80$). Since $R_f(T)$ and $R_{nf}(T)$ are concave, it follows that the right derivative of the gain curve at $T = T_0$ is positive if the capacity constraint becomes active in the NoFlex-LP and negative if it becomes active in Flex-LP. The gain $G_f(T)$ is typically high in this regime because one observes both demand induction and higher capacity utilization. However, in this regime the gain curve can either increase or decrease as one or the other effect becomes dominant.

Let T_1 denote the time horizon when all capacity constraints are active in FlexLP (In Figure 11, $T_1 = 245$). Then for all $T \geq T_1$ the gain decreases monotonically because the NoFlex revenue $R_{nf}(T)$ increases at a faster rate. When all constraints are active in NoFlex-LP (in Figure 11 this occurs at $T = 360$), $R_f(T) = R_{nf}(T)$ and the gain drops to zero. This is consistent with the intuition that for sufficiently large T, the airline can sell all of its capacity using only specific products and gains no advantage from offering flexible products.

In summary, an airline should consider offering flexibles when it anticipates that the probability of selling out its capacity is low. It is usually even more advantageous to offer flexibles in the intermediate region $T_0 < T < T_1$ because the airline stands to benefit from both demand induction and improved capacity utilization. Extrapolating from this study, we would expect that when

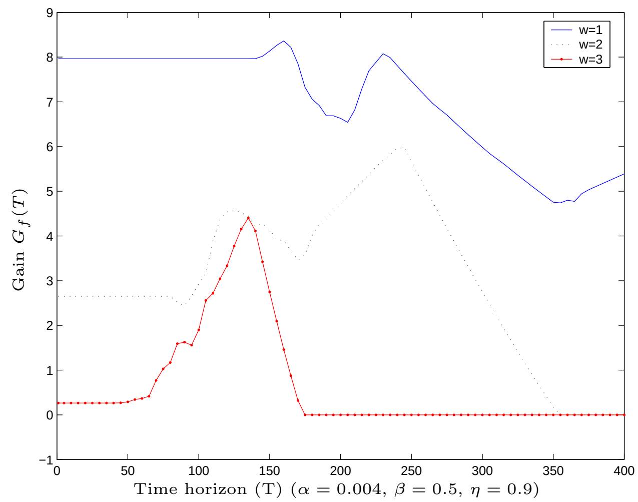

Figure 13: Gain $G_f(T)$ as a function of w

T	$w = 1$	$w = 2$	$w = 3$
51	4	3	1
101	4	2	1
201	3	1	0
301	2	1	0
401	1	0	0

Table 2: Number of Flexibles in Optimal Set

arrival rates are low, as is the case when the utility w is small, flexible products should be offered to induce demand and the gains will remain positive over a larger time horizon T. We should also expect similar behavior when the outside alternative (i.e. a competing product) is more attractive.

4.2 Gain as a function of the O-D utility w

From (31) and (32) it follows that the arrival rate is an increasing function of the utility w. Since the total capacity c is fixed we would anticipate that for a sufficiently high value of w, there would be no gain from offering flexible products. On the other hand, when trip utility is low, flexible products should be attractive because of the demand induction effect, and consequently, the gain $G_f(T)$ due to flexibles should be high. Moreover, since reducing rates is equivalent to scaling up capacity, we expect that for low values of w flexible products will remain attractive over a larger

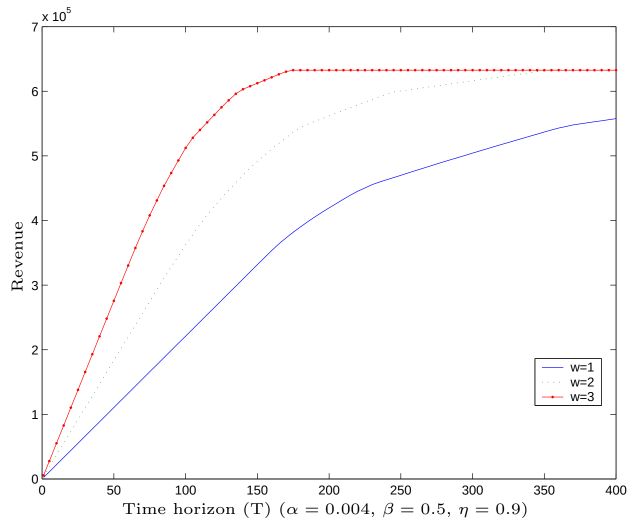

Figure 14: Revenue as a function of w

time horizon T.

Our experimental results are consistent with these expectations. From the plot in Figure 14 it is clear that revenue is an increasing function of w. On the other hand, the gain $G_f(T)$ from selling flexibles plotted in Figure 13 generally (but not always) decreases in w. The gain curves corresponding to different value of w tend to cross in the intermediate range of T where capacity utilization effects begin to dominate demand induction. The gains in this region can be high even for low values of w. Table 2 lists the number of flexible products in the optimal collection. For fixed trip utility w, the number of flexibles used decreases as the time horizon T increases (this is consistent with results in the previous section); and, for a fixed time horizon T, the number of flexibles used decreases with increase in w.

In summary, flexibles are attractive for low values of utility w and small time horizon T because of demand induction. At intermediate values of T, improved capacity utilization may increase the attractiveness of flexibles; however, this effect is not robust. Finally, for all values of w, there is a maximum value of T beyond which flexibles do not add any additional revenue. (In Figure 14, this maximum value is beyond the range of the graph for $w = 1$.)

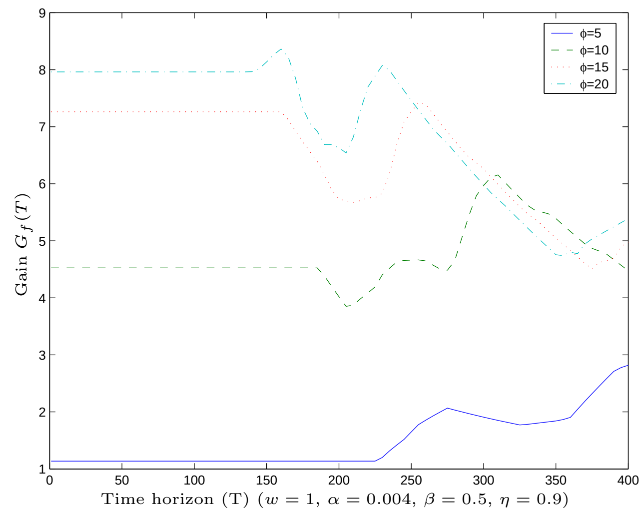

Figure 15: Gain $G_f(T)$ as a function of ϕ

4.3 Effect of outside alternatives on gain

Recall that (see (31) and (32)) we had initially set $\frac{\tilde{g}_j}{g_j} = \frac{\tilde{h}_k}{h_k} = 20$. In this section, we investigate the effect of setting $\frac{\tilde{g}_j}{g_j} = \frac{\tilde{h}_k}{h_k} = \phi$, where $\phi \in \{5, 10, 15, 20\}$. For the choice model in (18) and (19), increasing ϕ increases the rate at which customers leave the network without purchasing any product. This can be interpreted as the presence of an increasingly attractive outside alternative – i.e., increased competition.

Figure 15 plots the gain $G_f(T)$ vs ϕ for trip utility $w = 1$. Figure 16 is the same plot for trip utility $w = 2$ (Note that the gain curve corresponding to $\phi = 20$ is the same as that in Figure 11). From these plots it appears that increasing competition generally increases the gain but also reduces the time T_0 at which capacity utilization effects begin to become important. For small ϕ, i.e. when the outside alternative is not very attractive, the customer pool is essentially captive and there is less incentive for offering flexible products; thus, the gain $G_f(T)$ is small. All the gain in this scenario is due to capacity utilization effects rather than demand induction. On the other hand, for large ϕ, i.e. in highly competitive environments, offering flexible products at lower prices can capture customers that would otherwise have gone to the competition. For all values of ϕ, the gain eventually drops to zero for large enough T (For many of the cases in Figures 15 and 16 this value is higher than 400.)

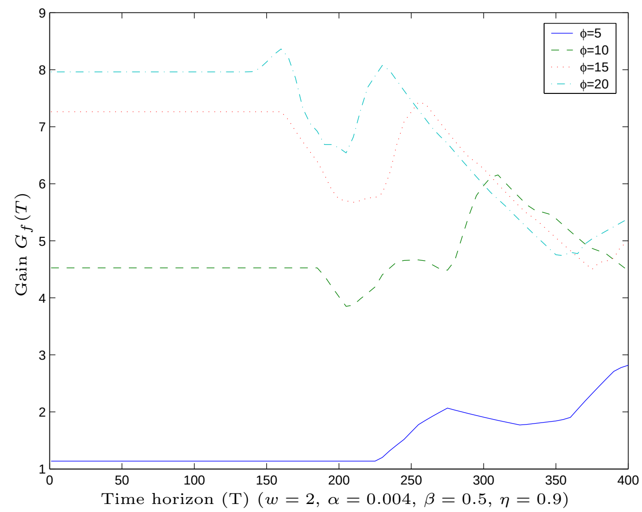

Figure 16: Gain $G_f(T)$ as a function of ϕ

4.4 Effect of discount on gain

Figure 17 and Figure 18 plot the gain $G_f(T)$ when the revenue $r_k = (1-\epsilon)\min_{j\in F_k}\{p_j\}$, $\epsilon \in \{0, 0.25, 0.5, 0.75\}$. Not surprisingly, increasing the discount on flexibles decreases the gain. Interestingly though, even offering 75% discount on produces a positive gain over certain time horizons T. In particular, when $w = 0.5$ (see Figure 17) a 75% discount produces positive gains over the same range as a 0% discount.

4.5 Demand induction vs capacity utilization

As noted earlier, the gain from flexibles can be attributed to two distinct effects: demand induction and higher capacity utilization. In order to disentangle the relative contribution from each of these sources, we conducted the following experiment. We computed the revenue enhancement $G_d(T)$ when no flexibles are offered but the vector $g \in \Re^n_+$ ($\tilde{g}$ was held constant) was modified as follows.

$$g_j = \begin{cases} g_j + h_k, & j = \text{argmax}_{\{j\in F_k\}}\{p_j\} \text{ for some } k \in F, \\ g_j & \text{otherwise,} \end{cases} \tag{33}$$

i.e. all the demand induced by a flexible is routed to the highest-fare specific product. Since only specific products are offered, the gain $G_d(T)$ is only due to demand induction. Note that since it is possible the optimal solution of Flex-LP never offers a particular flexible product, the new g defined in (33) overestimates the demand that can be induced by flexible products. Furthermore,

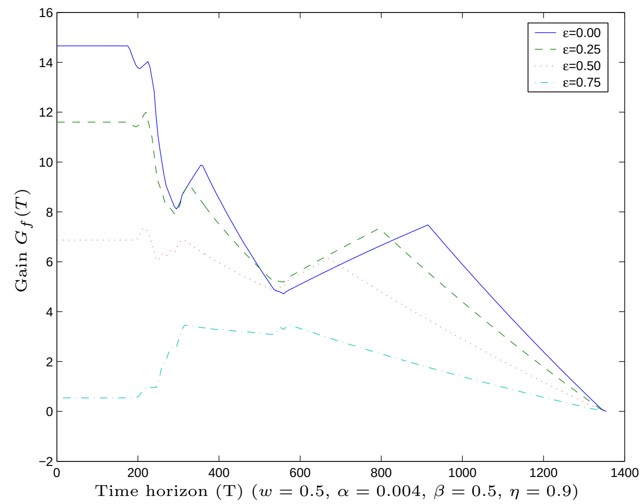

Figure 17: Gain $G_f(T)$ as a function of discount ϵ

since the induced demand is valued at the highest fare specific product rather than at a discount, the gain from the induced demand will be considerably overvalued. Thus, if the gain $G_d(T) \approx 0$, the contribution of demand induction in $G_f(T)$ can be expected to be minimal. Also, $G_f(T) \geq G_d(T)$, at a given horizon T, implies that capacity utilization effects dominate at T.

Figure 19 plots the gains $G_f(T)$ and $G_d(T)$ as a function of the time horizon T. At $T = 150$, the gain curve $G_d(T)$ drops below $G_f(T)$ and stays below for all $T \geq 150$. Therefore, we conclude that capacity utilization is the primary explanation for increase in revenue for moderate to large values of T. The gain $G_d(T) \approx 0$ for $T \geq 180$, therefore, it follows that demand induction is important for small T (if at all).

4.6 Summary of computational study results.

The computational study described above illustrates how the gain from flexible products depends on the various system parameters. The gain from offering flexible products is moderate for small values of the time horizon T, increases (although not monotonically) for intermediate values of T and drops to 0 for high T (see Figure 11). Small values of T correspond to the case where capacity is very large relative to demand, and in this situation, the gain from flexibles is entirely due to demand induction. In contrast, for high values of T the capacity is small relative to demand, and in this case there is no benefit from offering flexibles because the airline can fill all of its flights with higher-paying specific product customers. Flexibles provide highest value in the intermediate

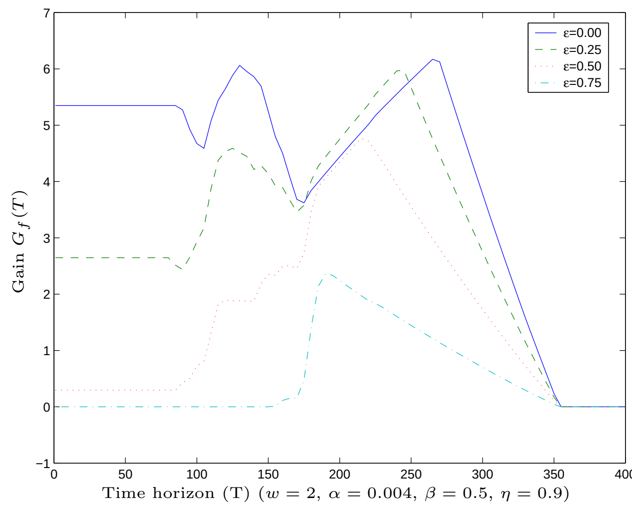

Figure 18: Gain $G_f(T)$ as a function of discount ϵ

region when both the demand induction and higher capacity utilization effects are important. This interpretation is confirmed by isolating the demand induction gain and comparing it to the total gain as shown in Figure 19. Similar results also hold for changing the trip utility parameter w, as shown in Figure 13.

The gain from flexibles also depends upon the attractiveness of the outside alternative and the flexible discount. When the outside alternative is unattractive, the customer pool is essentially captive and offering flexibles has little or no benefit. As the outside alternative becomes increasingly attractive, the benefits from offering flexibles increases, as shown in Figures 15 and 16. Finally, increasing the discount offered for flexibles tends to decrease the overall gain as shown in Figures 17 and 18, although the gain is still positive even when flexibles are offered at a 75% discount.

Note that the numerical results discussed here refer to the deterministic value function $V^D(c)$. However, the asymptotic optimality result (see Lemma 3 and Corollary 2) implies that, at least for long time horizons T, similar conclusions will hold for the stochastic value function $V(c)$.

5 Conclusion and extensions

The main goal of this paper is to show how to formulate and solve a network revenue management model with flexible products. We have shown that the optimal value of the stochastic problem is closely approximated (in the limit of long time horizon and large capacity) by the optimal value of

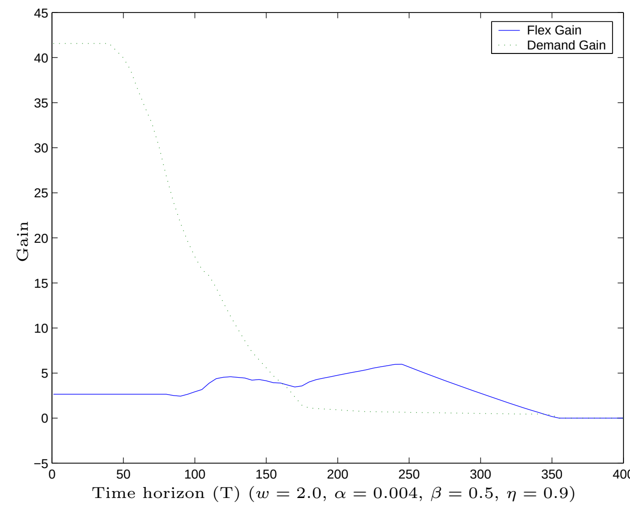

Figure 19: Demand induction vs Flexible products

an appropriately defined deterministic control problem. The optimal controls for the deterministic problem can be calculated by solving an LP. When the demands for the specific and flexible products do not depend on which other products are available (see Section 2), the LP is relatively small and therefore can be solved to optimality by standard methods. When the demands for the various products are given by a customer choice model (see Section 3) the demand for each product will depend upon which other products are available. In this case, LP has one variable for each possible subset of products that can be offered. Consequently, the number of variables grows exponentially with the number of products offered. However, we show that the LP can be solved efficiently by a column-generation algorithm for a broad class of consumer choice models. Section 4 reports the results of a detailed computational study of a small subnetwork of American Airlines.

A possible shortcoming of our model is the assumption that relative arrival rates are constant over the entire time horizon. Typically, customers who book early are more price-sensitive when compared with those who book later and, consequently, airlines make discounted fares available earlier rather than later [2, 5, 11, 15]. This situation can be modeled by dividing the time horizon $[0, T]$ into K subintervals of length T_l, $l = 1, \ldots, L$ and assuming that the arrival rates $(\lambda_l(S), \gamma_l(S))$ are a function of *both* the set of products S offered and the interval l. Thus, the corresponding LP is given by

$$
\begin{array}{rrcll}
\max & \sum_l \sum_S \big(p'\lambda_l(S) + r'\gamma_k(S)\big) t_l(S) & & & \\
\text{subject to} & A\Big(\sum_l \sum_S \lambda_l(S) t_l(S)\Big) \;+\; Bz & \leq & c, & \\
& \sum_l \sum_S \gamma_l(S) t_l(S) \;-\; Uz & = & 0, & \\
& \sum_S t_l(S) & \leq & T_l, & l = 1, 2, \ldots, L \\
& t_l(S) & \geq & 0, & \forall S \subseteq N \cup F,\ l = 1, 2, \ldots, L \\
& z & \geq & 0. &
\end{array}
\tag{34}
$$

The choice models determining the rates $(\lambda_l(S), \gamma_l(S))$ could be independent across the intervals T_l, or they could explicitly incorporate more complex behavior, e.g. observing which products are available during this epoch and deciding whether or not to wait until a later epoch to purchase. Note that, although the number of variables in (34) has increased K *times* compared to (15), the number of constraints has only increased by $K - 1$. This would indicate that column-generation would also be an efficient way to solve the LP (34).

Further extensions of the flexible product concept might provide additional benefits. The flexible products introduced in this paper are available to all customers. However, in Section 2.5 we saw that offering the flexible product only to a subset of customers can significantly improve benefits relative to offering it to all customers. However, such a modification is likely to result in a reduction in demand induction. Another modification would be to offer different flexible products, each of which contains the same set of specific products but that differ in the time that the buyer learns to which specific product he has been assigned. For the "standard" flexible product, the buyer would learn which specific product he has been assigned at time T. For a premium price, he could find out earlier. This would require a choice model that included a utility on "assignment time". One possible approach would be to adopt a model similar to that of Gale and Holmes [7, 8] who posit a two-flight model in which, at the time of booking, consumers know that they will have a higher preference for one flight than the other, but do not know which flight they will actually prefer until closer to the time of departure. While these approaches might be effective, in practice the airline would need to be sure that they were not overwhelming or confusing the customer with too many product alternatives.

Designing appropriate flexible products, i.e. choosing the specific products that constitute a given flexible product, is a hard and challenging problem. Intuitively, one would expect that flexible products could designed by clustering together specific products that are considered substitutes. However, assessing the degree to which a pair of products are considered substitutes could be quite hard. Moreover, one must also guard against creating too many flexible products.

The pricing of flexible products is also an important issue. In our model, we assume that the prices of all products – flexible and specific – as exogenous. This is consistent with the revenue management literature. It is also consistent with airline practice in which prices are changed relatively infrequently but allocations are changed often. However, it is clear, that pricing is an important decision in itself. It is possible that a bid-price based approach such as that described in Gallego and Van Ryzin [10] and discussed in Section 2 could be used to provide guidance when it would be profitable to raise or lower prices.

Finally, it is likely that a supplier's policy with flexible products is likely to influence future purchasing customer behavior. Customers are more likely to purchase a flexible product if they feel that the chance of being assigned to an attractive specific product is high. Over time, customers would tend to update their probabilities of being assigned to a particular specific product given that they have purchased a flexible product. Let p_i^k be defined as a customer's *ex ante* probability that she will be assigned to specific product i if she purchases flexible product k. An "consistent expectations equilibrium" would be one in which the customers' *ex post* probability of being assigned to i given that he purchases flexible product k is equal to p_i^k. In general, we would expect that if customers' *ex ante* expectations deviated from their *ex post* experiences, that they would update their expectations accordingly. In this case, the assignment of customers to specific products would have implications for future choices that would need to be considered in order for the airline to maximize long-term profitability.

References

[1] S. P. Anderson, A. de Palma, and J-F. Thisse. *Discrete Choice Theory of Product Differentiation.* MIT Press, Cambridge, MA, 1992.

[2] P. P. Belobaba. Application of a Probabilistic Decision Model to Airline Seat Inventory Control. *Operations Research*, 37:183-197, 1989.

[3] M. Ben-Akiva and S.R. Lerman. *Discrete Choice Analysis*. MIT Press, Cambridge, MA, 1985.

[4] A. Ben-Tal and A. Nemirovski. *Modern Convex Optimization*. SIAM.

[5] D. Bertsimas and I. Popescu. Revenue Management in a Dynamic Network Environment. *Transportation Science*, 37(3):257-277, 2003.

[6] Y. Feng and G. Gallego, Optimal Starting Times for End of Season Sales and Optimal Stopping Times for Promotional Fares. *Management Science*, 41:1371-1391, 1995.

[7] I. L. Gale and T. J. Holmes. The Efficiency of Advance-Purchase Discounts in the Presence of Aggregate Demand Uncertainty. *Int. J. of Ind. Org.*, 10:413-437, 1992.

[8] I. L. Gale and T. J. Holmes. Advance-Purchase Discounts and Monopoly Allocation of Capacity. *The American Economic Review*, 83:135-146, 1993.

[9] G. Gallego and R. Phillips. Revenue Management of Flexible Products. *Manufacturing and Service Operations Management* 6(4):321-337, 2004.

[10] G. Gallego and G. J. van Ryzin. A Multi-Product Dynamic Pricing Problem and Its Applications to Network Yield Management. *Operations Research*, 45:24-41, 1997.

[11] T. Lee and M. Hersh. A Model for Dynamic Airline Seat Inventory Control with Multiple Seat Bookings. *Transportation Science*, 27(3):252-256, 1998.

[12] M. Z. F. Li and T. H. Oum. A Note on the Single-Leg, Multifare Seat Allocation Problem. *Transportation Science*, 36:349-353, 2002.

[13] K. Littlewood. Forecasting and Control of Passenger Bookings. In *AGIFORS Symposium Proceedings*, 12:95-117, 1972.

[14] R. Phillips. A Marginal-Value Approach to Origin and Destination Revenue Management. In *System Modeling and Optimization*, eds. J. Henry and J.-P. Yvon, Springer-Verlag, London, 1993.

[15] R. Phillips. *Pricing and Revenue Optimization*. Stanford University Press, Stanford, CA. (2005).

[16] G. J. van Ryzin and Q. Liu. On the choice-based linear programming model for network revenue management. Working Paper. Decision, Risk and Operations Division. Columbia Business School.

[17] K. Talluri and G. J. van Ryzin. An Analysis of Bid-Price Controls for Network Revenue Management. *Management Science*, 44:1577-1592, 1998.

[18] K. Talluri and G. J. van Ryzin. Revenue Management Under a General Discrete Choice Model of Consumer Behavior. To appear in *Management Science*. 2003.

[19] E. Williamson. Airline Network Seat Control. Ph.D Thesis. MIT, Cambridge, MA, 1992.