

Dynamic Pricing Strategies for Multiproduct Revenue Management Problems

Constantinos Maglaras

Columbia Business School, Columbia University, 409 Uris Hall, 3022 Broadway,
New York, New York 10027, c.maglaras@gsb.columbia.edu

Joern Meissner

Lancaster University Management School, Lancaster University, Room A48,
Lancaster LA1 4XY, United Kingdom, joe@meiss.com

Consider a firm that owns a fixed capacity of a resource that is consumed in the production or delivery of multiple products. The firm strives to maximize its total expected revenues over a finite horizon, either by choosing a dynamic pricing strategy for each product or, if prices are fixed, by selecting a dynamic rule that controls the allocation of capacity to requests for the different products. This paper shows how these well-studied revenue management problems can be reduced to a common formulation in which the firm controls the aggregate rate at which all products jointly consume resource capacity, highlighting their common structure, and in some cases leading to algorithmic simplifications through the reduction in the control dimension of the associated optimization problems. In the context of their associated deterministic (fluid) formulations, this reduction leads to a closed-form characterization of the optimal controls, and suggests several natural static and dynamic pricing heuristics. These are analyzed asymptotically and through an extensive numerical study. In the context of the former, we show that “resolving” the fluid heuristic achieves asymptotically optimal performance under fluid scaling.

Key words: revenue management; dynamic pricing; capacity controls; fluid approximations; efficient frontier

History: Received: July 21, 2003; accepted: November 11, 2005. This paper was with the authors 9 months for 3 revisions.

1. Introduction

Consider a firm that owns a fixed capacity of a certain resource that is consumed in the process of producing or offering multiple products or services, and which must be consumed over a finite time horizon. The firm’s problem is to maximize its total expected revenues by selecting the appropriate dynamic controls. We consider two well-studied variants of this problem. In the first, the firm is assumed to be a monopolist or to operate in a market with imperfect competition, and thus to have power to influence the demand for each product by varying its price. In this setting, the firm’s problem is to choose a dynamic pricing strategy for each of its products to optimize expected revenues. In the second variant, prices are assumed to be fixed either by the competition or through a higher-order optimization problem, and the firm’s problem is now to choose a dynamic capacity allocation rule that controls when to accept

new requests for each of these products. In the sequel, these two problems are referred to as the “dynamic pricing” and “capacity control” formulations, respectively. Revenue management problems of that sort gained interest in the late 1970s in the context of the airline industry, and have since been successfully introduced in a variety of other areas such as hotels, cruise lines, rental cars, retail, etc.

This paper illustrates how these two problems can be reduced to a common formulation, thus connecting prior results that have appeared in the literature under a unified framework, and explores some of the consequences of this formulation. Specifically, we show that the multiproduct dynamic pricing problem introduced by Gallego and van Ryzin (1997) and the capacity control problem of Lee and Hersh (1993) can be recast within this common framework, and be treated as different instances of a single-product pricing problem for appropriate concave revenue

functions (Propositions 1 and 2). Broadly speaking, this is done by decoupling the revenue maximization problems in two parts: First, at each point in time the firm selects an aggregate capacity consumption rate from all products, and second, it computes the vector of demand rates to maximize instantaneous revenues subject to the constraint that all products jointly consume capacity at the aforementioned rate. The latter is akin to the basic microeconomics problem of resource allocation subject to a budget constraint, and gives rise to an appropriate *aggregate revenue rate function* in each case. This common formulation recovers well-known structural results regarding the monotonicity properties of the value function and the associated controls (see Proposition 3 and Corollary 2), which were previously derived in the literature while studying each of these problems in isolation; see, e.g., Gallego and van Ryzin (1994, 1997), Lee and Hersh (1993), Lautenbacher and Stidham (1999), Zhao and Zheng (2000), and the recent book by Talluri and van Ryzin (2004b). It also facilitates the derivation of new results, such as Corollary 1, that extend these properties to the setting of multiproduct pricing policies.

Extending this idea of demand aggregation in considering deterministic and continuous (fluid) approximations of the underlying problems suggests a set of simple pricing and capacity control heuristics for the underlying problems. These stem from the fact that this aggregated formulation leads to closed-form solutions to the fluid model revenue maximization problems (see Proposition 4). This extends the analysis of Gallego and van Ryzin (1997), which offered only an implicit characterization of these policies in the multiproduct setting; see also Bitran and Caldentey (2003) for a discussion of deterministic multiproduct pricing problems. Based on the solution of the fluid formulation, we propose three heuristics: (i) a static pricing heuristic, (ii) a static pricing heuristic applied in conjunction with an appropriate capacity allocation policy, and (iii) a “resolving” heuristic that reevaluates the fluid policy as a function of the current state and time-to-go (which is derived by expressing the fluid solution in feedback form). The first of these heuristics was proposed by Gallego and van Ryzin (1997), while policies that combine static prices with capacity controls as in (ii) have been suggested in other papers such as McGill and van Ryzin

(1999), Feng and Xiao (2004), and Lin et al. (2003). Finally, the “resolving” heuristic (iii) is widely applied in practice, but to the best of our knowledge has not been analyzed theoretically thus far. The only exception was the negative result of Cooper (2002), which illustrated through an example that resolving may in fact do worse than applying the static fluid policy. Propositions 5–7 in §4.2 establishes that all three heuristics achieve asymptotically optimal performance under fluid scaling, i.e., in the spirit of Gallego and van Ryzin (1997) and Cooper (2002). These results show that the phenomenon demonstrated in Cooper’s example does not persist in problems with large capacity and demand, where, in fact, resolving achieves the asymptotically optimal performance. Moreover, the numerical results of §5 illustrate that the dynamic heuristics (ii) and (iii) tend to outperform the static one.

In terms of qualitative insights, we find that the translation of the capacity consumption rate to a set of product-level controls that jointly maximize the instantaneous revenue rate defines an *efficient frontier* for the firm’s optimal pricing and capacity control strategies. This captures in a tractable way the interactions between products due to cross-elasticity effects and the joint capacity constraint. The idea of an efficient frontier has appeared in Talluri and van Ryzin (2004a) in the context of a capacity control problem for a model with customer choice among products, and in Feng and Xiao (2000, 2004) while studying pricing problems with a predetermined set of price points.

The remainder of the paper is structured as follows. This section concludes with some additional bibliographic references. Section 2 describes the two problem formulations. Section 3 demonstrates the reduction of the dynamic programming formulations to that of a single-product pricing problem, and derives some of its structural properties. Section 4 studies the fluid formulations of these problems and analyzes the asymptotic performance of the associated heuristics mentioned earlier. Section 4.3 outlines how the same approach can be extended to a network setting. Section 5 provides some numerical illustration of our results and offers some concluding remarks.

The idea of demand aggregation appeared in Talluri and van Ryzin (2004a) while analyzing a capacity control problem for a system where customer

behavior is captured through a discrete choice model, while similar techniques have been exploited in the past in the numerical solution of the dynamic programs associated with revenue management problems. Similar ideas also arise in problems of revenue management of multiproduct queueing systems; see Maglaras (2006) for a recent example. The papers by Elmaghraby and Keskinocak (2003), Bitran and Caldentey (2003), and McGill and van Ryzin (1999), and the book by Talluri and van Ryzin (2004b), provide comprehensive overviews of the areas of dynamic pricing and revenue management. The modelling framework adopted in this paper closely matches that of Gallego and van Ryzin (1994, 1997). Additional references on the capacity control formulation are Brumelle and McGill (1993) and Lautenbacher and Stidham (1999). Finally, a dynamic pricing heuristic that is related to the one studied in this paper was derived by Reiman (2002) as the solution to a “second-order” control problem that seeks to minimize an appropriate measure of the deviation from the inventory trajectory derived through the deterministic model. Indeed, the single-product policy proposed in Reiman (2002) is the same as the one we derive here. The multiproduct formulation in Reiman (2002) differs from ours, and therefore the resulting pricing policies also differ.

2. Single Resource, Multiproduct Model

This section formulates the multiproduct dynamic pricing and capacity control problems studied by Gallego and van Ryzin (1997) and Lee and Hersh (1993), respectively, among others.

2.1.1. Dynamic Pricing Problem. Consider a firm endowed with C units of capacity of a single resource used in producing or offering multiple products or services, indexed by $i = 1, \dots, n$. Each product i request requires one unit of capacity. There is a finite horizon T over which the resources must be used, and capacity cannot be replenished up to that time. The salvage value of remaining capacity at time T is assumed to be zero. (A constant per-unit salvage value would also result in formulations similar to those developed below.) The firm is either a monopolist or is assumed to operate in a market with imperfect competition, and, therefore, has

power to influence the demand for each product by varying its menu of prices. Let $p(t) = [p_1(t), \dots, p_n(t)]$ denote the vector of prices at time t . The demand process is assumed to be n -dimensional nonhomogeneous Poisson process with rate vector λ determined through a demand function $\lambda(p(t))$, where $\lambda: \mathcal{P} \rightarrow \mathcal{L}$, $\mathcal{P} \subseteq \mathbb{R}^n$ is the set of feasible price vectors, and $\mathcal{L} = \{x \geq 0: x = \lambda(p), p \in \mathcal{P}\} \subseteq \mathbb{R}_+^n$ is the set of achievable demand rate vectors. We assume that \mathcal{L} is a convex set. For ease of exposition, the demand function $\lambda(\cdot)$ is assumed to be stationary; an extension to allow for nonstationarities could follow Gallego and van Ryzin (1994, 1997) and Zhao and Zheng (2000). This class of demand functions incorporates product complementarity and substitution effects. Following Gallego and van Ryzin (1994, 1997), we consider *regular* demand functions that satisfy some additional conditions. In the sequel, x' denotes the transpose of any matrix x , for any real number y , $y^+ := \max(0, y)$, e is the vector of ones of appropriate dimension, and “a.s.” stands for almost surely.

DEFINITION 1. A demand function is said to be regular if it is a continuously differentiable, bounded function, and: (a) for each product i , $\lambda_i(p)$ is strictly decreasing in p_i , (b) $\lim_{p_i \rightarrow \infty} \lambda_i(p) = 0$ (i.e., consumers have bounded wealth), and (c) the revenue rate $p' \lambda(p) = \sum_{i=1}^n p_i \lambda_i(p)$ is bounded for all $p \in \mathcal{P}$ and has a finite maximizer \bar{p} .

We assume that there exists a continuous inverse demand function $p(\lambda)$, $p: \mathcal{L} \rightarrow \mathcal{P}$, that maps an achievable vector of demand rates λ into a corresponding vector of prices $p(\lambda)$. This allows us to view the demand rate vector as the firm’s control, and infer the appropriate prices using the inverse demand function. The expected revenue rate can be expressed as a function of the vector of demand rates λ as $R(\lambda) := \lambda' p(\lambda)$, and is assumed to be continuous, bounded, and strictly concave.

EXAMPLE. Under a linear demand model, the demand for product i is given by

$$\lambda_i(p) = \Lambda_i - b_{ii}p_i - \sum_{j \neq i} b_{ij}p_j \quad \text{or}$$

$$\text{(in vector form)} \quad \lambda(p) = \Lambda - Bp,$$

where Λ_i is the market potential for product i and b_{ii} , b_{ij} are the price and cross-price sensitivity parameters. The inverse demand and revenue functions are

$p(\lambda) = B^{-1}(\Lambda - \lambda)$ and $R(\lambda) = \lambda' B^{-1}(\Lambda - \lambda)$, respectively. Assumption 1 requires that $b_{ii} > 0$ for all i . To ensure that the inverse demand function is well defined and the revenue function is concave, we require that either $b_{ii} > \sum_{j \neq i} |b_{ji}|$ or $b_{ii} > \sum_{j \neq i} |b_{ij}|$ for all i ; both conditions guarantee that B is invertible and that its eigenvalues have positive real parts (Horn and Johnson 1994, Thm. 6.1.10).

The problem that we address is roughly described as follows: Given an initial capacity C , a selling horizon T , and a demand function that maps a vector of prices to a vector of demand rates, the firm's goal is to choose a nonanticipating dynamic pricing strategy for each product in order to maximize its total expected revenues.

We adopt a discrete-time formulation, i.e., one where time has been discretized in small intervals of length δt , indexed by $t = 1, \dots, T$, such that $\mathbb{P}(\text{product } i \text{ arrival in } [0, \delta t]) = \lambda_i \delta t + o(\delta t)$ for all i , and $\mathbb{P}(\text{product } i \text{ and } j \text{ arrivals in } [0, \delta t]) = \lambda_i \lambda_j (\delta t)^2 + o((\delta t)^2)$, where $o(x)$ implies that $o(x)/x \rightarrow 0$ as $x \rightarrow 0$. With slight abuse of notation, we write λ_i in place of $\lambda_i \delta t$, and refer to λ_i either as the demand or the buying probability for product i . The random demand vector in period t , denoted by $\xi(t; \lambda)$, is Bernoulli with probabilities $\lambda(t) = \lambda(p(t))$, and $\mathbb{P}(\xi_i(t) = 1) = \lambda_i(p(t))$ and $\mathbb{P}(\xi_i(t) = 0) = 1 - \lambda_i(p(t))$ for all i . Treating the demand rates λ_i as the control variables (prices are inferred via the inverse demand relationship), the discrete-time formulation of the dynamic pricing problem of Gallego and van Ryzin (1997) is

$$\max_{\{\lambda(t), t=1, \dots, T\}} \left\{ \mathbb{E} \left[\sum_{t=1}^T p(\lambda(t))' \xi(t; \lambda) \right] : \sum_{t=1}^T e^t \xi(t; \lambda) \leq C \text{ a.s. and } \lambda(t) \in \mathcal{L} \forall t \right\}. \quad (1)$$

2.1.2. Capacity Control Problem. The second problem that we consider is the one studied by Lee and Hersh (1993), where the price vector p and the demand rate vector $\lambda = \lambda(p)$ are fixed, and the firm optimizes over capacity allocation decisions. For this problem and without any loss of generality, we assume that products are labelled such that $p_1 \geq p_2 \geq \dots \geq p_n$. The firm has discretion as to which product requests to accept at any given time. This is modelled through the control $u_i(t)$ that is equal to the

probability of accepting a product i request at time t . It is customary to assume that the firm is "opening" or "closing" products, thus considering controls $u_i(\cdot)$ that are zero or one, but this need not be imposed as a restriction. The dynamic capacity control problem is the following:

$$\max_{\{u(t), t=1, \dots, T\}} \left\{ \mathbb{E} \left[\sum_{t=1}^T p' \xi(t; u\lambda) \right] : \sum_{t=1}^T e^t \xi(t; u\lambda) \leq C \text{ a.s. } 1 \text{ and } u_i(t) \in [0, 1] \forall t \right\}, \quad (2)$$

where $u\lambda$ above denotes the vector with coordinates $u_i \lambda_i$.

3. Analysis of the Pricing and Capacity Control Problems

This section describes how to reduce (1) and (2) into dynamic optimization problems where the control is the (one-dimensional) aggregate capacity consumption rate. Subsequently, we derive some structural properties for these two problems through a unified analysis.

3.1. A Common Formulation in Terms of the Aggregate Capacity Consumption

3.1.1. Dynamic Pricing Problem. Let x denote the number of remaining units of capacity at the beginning of period t , and $V(x, t)$ be the expected revenue-to-go starting at time t with x units of capacity left. Then, the Bellman equation associated with (1) is

$$V(x, t) = \max_{\lambda \in \mathcal{L}} \left\{ \sum_{i=1}^n \lambda_i [p_i(\lambda) + V(x-1, t+1)] + (1 - e'\lambda)V(x, t+1) \right\}, \quad (3)$$

with the boundary conditions

$$V(x, T+1) = 0 \quad \forall x \quad \text{and} \quad V(0, t) = 0 \quad \forall t. \quad (4)$$

Letting $\Delta V(x, t) = V(x, t+1) - V(x-1, t+1)$ denote the marginal value of one unit of capacity as a function of the state (x, t) , (3) can be rewritten as

$$\begin{aligned} V(x, t) &= \max_{\lambda \in \mathcal{L}} \left\{ R(\lambda) - \sum_{i=1}^n \lambda_i \Delta V(x, t) \right\} + V(x, t+1) \\ &= \max_{\rho \in \mathcal{R}} \{ R^t(\rho) - \rho \Delta V(x, t) \} + V(x, t+1), \end{aligned} \quad (5)$$

where $\rho := \sum_{i=1}^n \lambda_i$ is the aggregate rate of capacity consumption, $\mathcal{R} := \{\rho: \sum_{i=1}^n \lambda_i = \rho, \lambda \in \mathcal{L}\}$ is the set of achievable capacity consumption rates, and

$$R^r(\rho) := \max_{\lambda} \left\{ R(\lambda): \sum_{i=1}^n \lambda_i = \rho, \lambda \in \mathcal{L} \right\} \quad (6)$$

is the maximum achievable revenue rate subject to the constraint that all products jointly consume capacity at a rate ρ . Note that (6) is a concave maximization problem over a convex set, and its solution is readily computable, often in closed form (examples are given in §5). Moreover, $R^r(\cdot)$ is a concave function and satisfies the conditions of Assumption 1. The optimal vector of demand rates, denoted by $\lambda^r(\rho)$, is unique and continuous in ρ .

PROPOSITION 1. *The dynamic pricing problem (1) can be reduced to the dynamic program (5) and (4) expressed in terms of the aggregate consumption rate. In particular, if $\rho^*(x, t)$ denotes the associated optimal control and $\lambda^*(x, t)$ and $p^*(x, t)$ denote the optimal demand rate and price vectors associated with (1), then $\lambda^*(x, t) = \lambda^r(\rho^*(x, t))$ and $p^*(x, t) = p(\lambda^r(\rho^*(x, t)))$.*

3.1.2. The Capacity Control Problem. Similarly, the Bellman equation associated with (2) is

$$V(x, t) = \max_{u_i \in [0, 1]} \left\{ \sum_{i=1}^n \lambda_i u_i [p_i + V(x-1, t+1)] + (1 - u' \lambda) V(x, t+1) \right\}, \quad (7)$$

with the boundary condition (4), which, using the marginal value of capacity ΔV becomes

$$\begin{aligned} V(x, t) &= \max_{u_i \in [0, 1]} \left\{ \sum_{i=1}^n \lambda_i u_i p_i - u' \lambda \Delta V(x, t) \right\} \\ &\quad + V(x, t+1) \quad (8) \\ &= \max_{0 \leq \rho \leq \sum_{i=1}^n \lambda_i} \{ R^a(\rho) - \rho \Delta V(x, t) \} + V(x, t+1), \quad (9) \end{aligned}$$

where $\rho = u' \lambda$ and

$$R^a(\rho) = \max_u \left\{ \sum_{i=1}^n u_i \lambda_i p_i: u' \lambda = \rho, u_i \in [0, 1] \right\}$$

is the maximum revenue rate when the capacity is consumed at a rate equal to ρ , and $u^a(\rho)$ is the corresponding control.

PROPOSITION 2. *The capacity control problem (2) can be reduced to the dynamic program (9) and (4) expressed in terms of the aggregate consumption rate ρ . In particular, if $\rho^*(x, t)$ denotes the optimal solution of (9) and (4) and $u^*(x, t)$ denote the optimal policy for (2), then $u^*(x, t) = u^a(\rho^*(x, t))$.*

A similar result was derived by Talluri and van Ryzin (2004a) for a capacity control problem for a model with customer choice.

3.2. A Unified Analysis of the Pricing and Capacity Control Problems

The preceding analysis illustrates that both problems can be reduced to “appropriate” single-product pricing problems, highlighting their common structure and enabling a unified treatment. As a starting observation, for both (5) and (9) the optimal control $\rho^*(x, t)$ is computed from

$$\rho^*(x, t) = \arg \max_{\rho \in \mathcal{R}} \{ R(\rho) - \rho \Delta V(x, t) \},$$

where $R(\cdot)$ is a concave increasing revenue function. Using the properties of $R(\cdot)$, one gets that $\rho^*(x, t)$ is decreasing in $\Delta V(x, t)$, which using a backward induction argument in t gives that $\Delta V(x, t)$ is decreasing in x and t . These standard results for single-product dynamic pricing problems are summarized below; a proof can be found in Talluri and van Ryzin (2004b, Prop. 5.2, Ch. 4).

PROPOSITION 3 (TALLURI AND VAN RYZIN 2004B, PROP. 5.2, CH. 4). *For both problems defined in (1) and (2), we have that*

1. $\rho^*(x, t)$ is decreasing in the marginal value of capacity $\Delta V(x, t)$, and
2. $\Delta V(x, t)$ is decreasing in x and t .

Structural results for the pricing and capacity allocation policies follow from the properties of R^r , λ^r and R^a , u^a , respectively. For example, consider the pricing problem for the case where the products are nonsubstitutes, i.e., the demand for product i is only a function of the price for that product p_i . In that case, the Lagrangian associated with (6) is $L(\lambda, x, y) = R(\lambda) + x(\rho - \sum_{i=1}^n \lambda_i) - y' \lambda$, with first-order conditions given by $\partial R(\lambda) / \partial \lambda_i = x + y_i$, for some $x \geq 0$ and $y_i \leq 0$ with $y_i = 0$ if $\lambda_i > 0$. It is easy to show that x is decreasing in ρ (i.e., the shadow price for

the capacity consumption constraint decreases as the rate ρ increases), and that $\lambda_i^r(\rho)$ is decreasing in x .

COROLLARY 1. *Consider the problem specified in (3) and (4), and further assume that the products are non-substitutes, i.e., $\lambda_i(p) = \lambda_i(p_i)$ for all i . Then, $\lambda_i^*(x, t)$ is nondecreasing in $\rho^*(x, t)$ (and nonincreasing in $\Delta V(x, t)$).*

A similar result can be obtained when products are substitutable provided that the demand model satisfies certain conditions analogous to those of the sensitivity matrix B of the linear model in §2.

For the capacity control problem, it is easy to recover some well-known structural properties of the optimal policy, see, e.g., Lee and Hersh (1993). Our derivation based on the capacity consumption rate offers new intuition as to why they hold. Specifically, $R^a(\cdot)$ is a knapsack solution for which

$$R^a(\rho) = \min_i c_i + p_i \rho \quad \text{and} \quad (10)$$

$$u_k^a(\rho) = \min\left(\frac{(\rho - \sum_{i < k} \lambda_i)^+}{\lambda_k}, 1\right),$$

where $c_1 = 0$ and $c_i = \sum_{k < i} \lambda_k (p_k - p_i)$, and for any $x \in \mathbb{R}$, $x^+ := \max(x, 0)$, and the optimal control $\rho^*(x, t)$ reduces to the solution to $\max\{\min_i c_i + (p_i - \Delta V(x, t))\rho : 0 \leq \rho \leq \sum_{i=1}^n \lambda_i\}$. Let $i^*(x, t) = \max\{i \geq 1 : p_i \geq \Delta V(x, t)\}$. Then, by inspecting the form of the piecewise-linear objective function involved in the calculation of $\rho^*(x, t)$, we get that $\rho^*(x, t) = \sum_{i \leq i^*(x, t)} \lambda_i$. That is, the solution is “bang-bang” in the sense that the form of the optimal control is such that $u_i^*(x, t)$ is zero if $i > i^*(x, t)$ and one if $i \leq i^*(x, t)$. In addition, from Proposition 3, Part 1, we see that $i^*(x, t)$ is decreasing in the marginal value of capacity $\Delta V(x, t)$.

COROLLARY 2. *For the capacity control problem (2) or, equivalently, (9) and (4), the optimal allocation policy is nested in that $u_i^*(x, t) = 1$ if $i \leq i^*(x, t)$, and $u_i^*(x, t) = 0$ otherwise, and $i^*(x, t)$ is decreasing in the marginal value of capacity $\Delta V(x, t)$.*

REMARK. The subproblem of computing the optimal revenue subject to a constraint on the aggregate capacity consumption rate specified in (6) and (10) defines an *efficient frontier* $(\rho, R^r(\rho))$ and $(\rho, R^a(\rho))$ for the dynamic pricing and capacity allocation problems, respectively. As in the context of portfolio

optimization, the efficient frontier provides a systematic framework for comparing different policies and highlights the structure of the respective optimal controls. It may also lead to computational improvements if this subproblem can be solved efficiently, preferably in closed form. This is possible in some common demand models such as the linear and the multinomial logit. Such an efficient frontier has also appeared and been discussed in more detail in Feng and Xiao (2000, 2004) and Talluri and van Ryzin (2004a). Finally, we note that the structure of the dynamic programs studied in this section has been observed in other papers, such as Lin et al. (2003) and their study of single-resource capacity control problems where each arrival may request multiple units of capacity; and Vulcano et al. (2002) and their analysis of optimal dynamic auctions. The latter involves an analysis of a discrete-time, batch-demand analog to the dynamic program studied here.

4. Analysis of the Pricing Problem Using Its Fluid Approximation

This section studies deterministic (fluid model) formulations of multiproduct revenue management problems to provide some structural results (§4.1) and suggest simple and implementable heuristics for the pricing and capacity control problems (§4.2). The latter have desirable theoretical performance guarantees and are shown to perform well in the numerical experiments of the next section. Finally, §4.3 sketches how to extend these ideas to the network setting. Following the results of §3.2 that relate the pricing and capacity control formulations, hereafter we will focus exclusively on the former.

4.1. Solution to the Deterministic Multiproduct Pricing Problem

The “fluid” model has deterministic and continuous dynamics, and is obtained by replacing the discrete stochastic demand process by its rate, which now evolves as a continuous process. It is rigorously justified as a limit under a law-of-large-numbers type of scaling as the potential demand and the capacity grow proportionally large; see Gallego and van Ryzin (1994, 1997) and §4.2. It is simplest to describe the fluid model in continuous time (this is consistent with Gallego and van Ryzin 1994, 1997). In more detail, the

realized instantaneous demand for product i at time t in the fluid model is deterministic and given by $\lambda_i(t)$. We allow product i requests to consume capacity at a rate of $a_i > 0$ units per unit of demand, and denote by a the vector $[a_1, \dots, a_n]$. This is a generalization of the model considered thus far, which assumed uniform capacity requirements (all equal to 1). With a general capacity requirement vector a , the capacity consumption rate is defined by $\rho = a'\lambda$, and the definitions of R^r and λ^r can be appropriately adjusted to reflect that change. The system dynamics are given by $dx(t)/dt = -\sum_{i=1}^n a_i \lambda_i(t)$, $x(0) = C$, together with the boundary condition that $x(T) \geq 0$. The firm selects a demand rate $\lambda_i(t)$ (or a price) at each time t . The fluid formulation of the multiproduct pricing problem is the following:

$$\max_{\{\lambda(t), t \in [0, T]\}} \left\{ \int_0^T R(\lambda(t)) dt : \int_0^T a' \lambda(t) dt \leq C \text{ and } \lambda(t) \in \mathcal{L} \forall t \right\}. \quad (11)$$

Single-product problem. In this case, Gallego and van Ryzin (1994) showed that a constant price (and thus a constant demand rate) is optimal for (11). Specifically, let $\hat{\lambda} = \arg \max\{R(\lambda) : \lambda \in \mathcal{L}\}$ and $\hat{p} = p(\hat{\lambda})$ be the demand rate and price, respectively, that maximize the revenue rate disregarding any capacity considerations. Also, let $\lambda^0 = C/T$ be the *run-out* rate that depletes capacity at time T , and $p^0 = p(\lambda^0)$. Then, Gallego and van Ryzin (1994) showed that the optimal controls are constant over time and given by $\bar{\lambda} = \min(\hat{\lambda}, \lambda^0)$ and $\bar{p} = \max(\hat{p}, p^0)$. (The overbar notation denotes the optimal fluid solution.) Intuitively, the firm uses the revenue-maximizing price \hat{p} unless this would deplete the capacity too soon, in which case it increases its unit price to p^0 and sells its capacity by time T , while accruing higher total revenues. Gallego and van Ryzin (1997, §4.5) extended these results to multiple products, but in that case without providing such a succinct solution.

Multiproduct problem. Following the approach of §3, we can reduce the multiproduct problem to an appropriate single-product one, and thus solve it in closed form. Specifically, recalling the definitions of the aggregate revenue function $R^r(\rho)$ and optimal demand rate vector $\lambda^r(\rho)$ in (6), adjusted for the fact

that $\rho = a'\lambda$, (11) can be rewritten as

$$\max_{\{\rho(t), t \in [0, T]\}} \left\{ \int_0^T R^r(\rho(t)) dt : \int_0^T \rho(t) dt \leq C, \rho(t) \in \mathcal{R} \forall t \right\}. \quad (12)$$

Note that (12) is identical to a single-product problem with revenue function R^r , and thus is solvable using the approach described above. Let $\rho^0 := C/T$ and $\hat{\rho} = \arg \max_{\rho} R^r(\rho)$. Then, the optimal solution to (12) is to consume capacity at a constant rate $\bar{\rho}$ given by

$$\bar{\rho}(t) := \min(\hat{\rho}, \rho^0) \quad \forall t, \quad (13)$$

the corresponding vector of demand rates is $\lambda^r(\bar{\rho})$, while the price vector is $p(\lambda^r(\bar{\rho}))$. A direct verification that this solution satisfies the optimality conditions for (11) establishes the following:

PROPOSITION 4. *Let $\bar{\lambda}(\cdot)$ and $\bar{p}(\cdot)$ denote the optimal vectors of demand rates and prices for (11). Then, $\bar{\lambda}$, \bar{p} are constant over time and are given by $\bar{\lambda}(t) = \lambda^r(\bar{\rho})$ and $\bar{p}(t) = p(\lambda^r(\bar{\rho}))$.*

4.2. Asymptotically Optimal Heuristics Extracted from the Deterministic Model

Based on the preceding analysis, we discuss three heuristics for the underlying revenue management problems, which we analyze in the asymptotic setting introduced in Gallego and van Ryzin (1997) and Cooper (2002). Among other things, it is shown that the dynamic heuristic that “resolves” the fluid policy as t progresses is asymptotically optimal in an appropriate sense.

a. *The Static Pricing Heuristic of Gallego and van Ryzin (1997).* This policy implements the static prices \bar{p} specified in Proposition 4. (This is the “make-to-order” heuristic in Gallego and van Ryzin 1997.)

4.2.1. Dynamic Heuristics. The static nature of policy (a) is desirable for implementation purposes (see Gallego and van Ryzin 1997), but also lacks the capability of corrective action against stochastic fluctuations. This does not arise in the fluid formulation, where the capacity is drained along the optimal deterministic trajectory, but it is relevant for the stochastic problems of original interest. The next heuristics provide two possible adjustments to this static policy

that add such control capability and seem of practical interest. We start by observing that the solution of the fluid pricing problem of §4.1 can also be described in feedback form as

$$\bar{\rho}(x, t) = \min\left(\hat{\rho}, \frac{x}{T-t}\right), \quad (14)$$

where x is the remaining capacity at time t . The deterministic trajectory of the fluid model is, of course, such that $x/(T-t) = C/T$ for all t if $\hat{\rho} \geq C/T$, and $x/(T-t) = (C - \hat{\rho}t)/(T-t) \geq C/T$ if $\hat{\rho} < C/T$. In both cases, $\bar{\rho}(x, t) = \min(\hat{\rho}, C/T)$ for all x, t along the optimal fluid trajectory of the remaining capacity process, and thus (14) is identical to the static control derived in (13).

b. *A List-Price Capacity Control (LPCC) Heuristic.* One way to implement (14) is by introducing capacity control capability on top of the static prices given in (a). Specifically, our second heuristic is defined as follows:

1. price according to \bar{p} and label products such that $\bar{p}_1/a_1 \geq \bar{p}_2/a_2 \geq \dots \geq \bar{p}_n/a_n$, and
2. compute $\bar{\rho}(x, t)$ and use the capacity controls $u_i(x, t) = 1$ if $x > 0$, $u_i(0, t) = 0$, and

$$u_i(x, t) = \begin{cases} 1 & \text{if } \bar{\rho}(x, t) - \sum_{j<i} a_j \bar{\lambda}_j \geq 0 \\ 0 & \text{otherwise} \end{cases} \quad \text{for } i \geq 2. \quad (15)$$

Note that this policy can only reduce the aggregate capacity consumption rate from its nominal value of $\sum_{i=1}^n a_i \bar{\lambda}_i$, but can never increase it. A product is made available only if the fluid solution starting from that state would choose to sell this product in all future time periods, and “closes” the product if the fluid solution would dictate only partial acceptance of the associated demand.

This policy is a refinement of the static pricing policy in (a) and the make-to-order heuristic of Gallego and van Ryzin (1997). Other examples of joint pricing and capacity controls can be found in the recent papers by Vulcano et al. (2002), Lin et al. (2003), and Feng and Xiao (2004).

c. *A Dynamic Pricing Heuristic.* The third policy translates the aggregate control $\bar{\rho}(x, t)$ into product-level rates (and prices) through

$$\lambda(x, t) = \lambda^r(\bar{\rho}(x, t)) \quad \text{and} \quad p(x, t) = p(\lambda(x, t)), \quad (16)$$

where the mapping $\lambda^r(\cdot)$ was the maximizer in (6) and it is continuous in ρ . This corresponds to the idea of “resolving” the fluid problem as we step through time, which is widely applied in practice, where, however, the resolving occurs at discrete points in time, e.g., daily or weekly depending on the application setting. Despite its practical appeal and use, to the best of our knowledge policies that use this resolving idea have not been analyzed theoretically, other than the isolated example provided by Cooper (2002). Our preceding discussion illustrates that “resolving” is nothing but implementing the fluid policy in feedback form. The analysis that follows provides a characterization of its asymptotic behavior, while the numerical results of the next section demonstrate that it tends to outperform the other two candidate policies. The single-product version of this policy has also been studied by Reiman (2002).

4.2.2. Asymptotic Performance Analysis of the Pricing Heuristics. The remainder of this subsection offers a brief asymptotic characterization of the performance under these three heuristics that shows that all three are (fluid-scale) asymptotically optimal in a regime where the potential demand and capacity grow proportionally large; this is the first-order asymptotic optimality criterion of Gallego and van Ryzin (1997) and Cooper (2002). Specifically, using k as an index, we consider a sequence of problems with demand model $\lambda^k(\cdot) = k\lambda(\cdot)$ and capacity $C^k = kC$, and let k increase to infinity; hereafter, a superscript k will denote quantities that scale with k . Let N_i for $i = 1, \dots, n$ denote independent unit rate Poisson processes, and recall that the functional strong-law-of-large numbers for the Poisson process asserts that as $k \rightarrow \infty$ and for all $t \geq 0$,

$$\frac{N_i(kt)}{k} \rightarrow t \quad \text{a.s.} \quad (17)$$

For all of the candidate policies, the capacity dynamics can be expressed as follows. The cumulative demand for that product up to time t is equal to $N_i(A_i^k(t))$, where $A_i^k(t) = \int_0^t \lambda_i^k(s) ds$, and the remaining capacity at time t is

$$X^k(t) = C^k - \sum_{i=1}^n a_i N_i(A_i^k(t)). \quad (18)$$

Our goal here is to analyze the “fluid-scale” behavior of the capacity process defined as $\bar{X}^k(t) := X^k(t)/k$ under the three candidate policies. The asymptotic optimality of the static policy (a) was established by Gallego and van Ryzin (1997). Nevertheless, we include our analysis, which is different than theirs and serves to introduce the ideas used in studying the dynamic policies (b) and (c).

Analysis of the static heuristic (a). The firm uses the constant price vector \bar{p} , which induces the demand rates $\lambda^k(\bar{p}) = k\bar{\lambda}$. Under this policy, $A_i^k(t) = \bar{\lambda}_i kt$, and the capacity dynamics are $X_a^k(t) = (C^k - \sum_{i=1}^n a_i N_i \cdot (A_i^k(t)))^+$. (The subscript is used to identify the policy.) As $k \rightarrow \infty$

$$\frac{N_i(A_i^k(t))}{k} \rightarrow \bar{\lambda}_i t \quad \text{a.s., uniformly in } t \in [0, T], \quad (19)$$

from which it follows that as $k \rightarrow \infty$ and for all $t \in [0, T]$

$$\bar{X}_a^k(t) \rightarrow \left(C - \sum_{i=1}^n a_i \bar{\lambda}_i t \right)^+ = C - \bar{\rho} t \quad \text{a.s.}$$

(The $(\cdot)^+$ was removed because from (13) $\bar{\rho} t \leq C$ for all $t \in [0, T]$.) Let R_a^k denote the revenues extracted under policy (a), and $\tau^k := \inf\{s \geq 0: \sum_{i=1}^n a_i N_i(\bar{\lambda}_i ks) \geq C^k\}$ be the random time where the aggregate capacity requested reaches or exceeds the available capacity C^k . Then, $R_a^k := \sum_{i=1}^n \bar{p}_i N_i(k\bar{\lambda}_i \min(T, \tau^k)) - \delta$, where δ is a random variable that corrects revenues for the case where $\tau^k < T$, which is bounded above by $\max_i \bar{p}_i$. (We will not delve into an accurate description of δ , because it is asymptotically negligible.) Using (19) and arguing by contradiction, one can easily conclude that $(T - \tau^k)^+ \rightarrow 0$ a.s., as $k \rightarrow \infty$. Using the expression for R_a^k we get the following:

PROPOSITION 5. *Suppose that demand and capacity are scaled according to $\lambda^k(\cdot) = k\lambda(\cdot)$ and $C^k = kC$, and consider the static pricing policy $p^k(x, t) = \bar{p}$ for all x, t and all k . Then, as $k \rightarrow \infty$, $\bar{X}_a^k(t) \rightarrow C - \bar{\rho} t$ a.s., uniformly in t , and $(1/k)R_a^k \rightarrow \sum_{i=1}^n \bar{p}_i \bar{\lambda}_i T$ a.s.*

Recall the fact established in Gallego and van Ryzin (1997) that the solution of the deterministic pricing problem serves as an upper bound for the revenues extracted in the stochastic system, i.e., $R_a^k \leq k \sum_{i=1}^n \bar{p}_i \bar{\lambda}_i T$. Applying the bounded convergence theorem gives that $\mathbb{E}R_a^k \rightarrow \sum_{i=1}^n \bar{p}_i \bar{\lambda}_i T$, and establishes the asymptotic optimality of the static pricing heuristic (a).

Analysis of the dynamic heuristic (c). The dynamic nature of this policy requires a more detailed study. The cumulative demand for product i up to time t is equal to $N_i(A_i^k(t))$, where

$$A_i^k(t) = \int_0^t k\lambda_i(s) ds \quad \text{where} \\ \lambda_i(s) = \lambda_i^r(\min(\hat{\rho}, X_c^k(s)/(k(T-s)))) \quad (20)$$

for $\lambda_i^r(\cdot)$ defined in (6), and $X_c^k(t)$ denotes the remaining capacity at time t under policy (c), defined in (18). Now, note that $A_k(0) = 0$, $A_k(t)$ is nondecreasing and $A_i^k(t) - A_i^k(s) \leq (t-s) \cdot k\lambda_{i,\max}$, where $\lambda_{i,\max} = \arg \max\{\lambda_i: \lambda \in \mathcal{L}\}$. This implies that the family of processes $\{(1/k)A^k(t)\}$ is equicontinuous, and therefore relatively compact. It follows that it has a converging subsequence, say k_j , such that $(1/k_j)A_i^{k_j}(t) \rightarrow \bar{A}_i(t)$ for all i . Along this subsequence we get that $N_i(A_i^{k_j}(t))/k_j$ converges to $\bar{A}_i(t)$, and therefore that $\bar{X}_c^{k_j}(t)$ itself converges to a limit $\bar{x}_c(t)$; the last two results hold a.s., uniformly in t . Expression (22) will show that the limit trajectories do not depend on the selection of the converging subsequence itself. In the sequel we denote the converging subsequence by k to simplify notation. Using the continuity of $\lambda_i^r(\cdot)$ and Lemma 2.4 of Dai and Williams (1995), we get that

$$\frac{1}{k} A_i^k(t) = \int_0^t \lambda_i^r \left(\min \left(\hat{\rho}, \frac{\bar{X}_c^k(s)}{T-s} \right) \right) ds \\ \rightarrow \int_0^t \lambda_i^r \left(\min \left(\hat{\rho}, \frac{\bar{x}_c(s)}{T-s} \right) \right) ds \quad (21)$$

as $k \rightarrow \infty$ a.s., uniformly in t . Using (17), (18), and (21), we get that as $k \rightarrow \infty$

$$\bar{X}_c^k(t) = C - \frac{1}{k} \sum_{i=1}^n a_i N_i(A_i^k(t)) \\ \rightarrow C - \int_0^t \min \left(\hat{\rho}, \frac{\bar{x}_c(s)}{T-s} \right) ds = C - \bar{\rho} t, \quad (22)$$

a.s., uniformly in t . The revenues extracted under policy (c) are $R_c^k = \sum_{i=1}^n \int_0^t p_i^k(t) dN_i(A_i^k(t))$, where $p^k(t)$ is the price vector that corresponds to $\lambda^r(\min(\hat{\rho}, \bar{X}_c^k(t)/(T-t))$, the demand rate vector at time t , and the integrals should be interpreted in the Riemann-Stieltjes sense. From (21) and (22) we have that $\lambda_i^r(\min(\hat{\rho}, \bar{X}_c^k(t)/(T-t))) \rightarrow \bar{\lambda}_i$ a.s., uniformly in t . By the continuity of the inverse demand function we

have that $p^k(t) \rightarrow \bar{p}$ a.s., uniformly in t , and therefore, again using Lemma 2.4 of Dai and Williams (1995), we get the result summarized below.

PROPOSITION 6. *Suppose that demand and capacity are scaled according to $\lambda^k(\cdot) = k\lambda(\cdot)$ and $C^k = kC$, and consider the dynamic pricing heuristic defined through (20). Then, $\bar{X}_c^k(t) \rightarrow C - \bar{\rho}(t)$ a.s., uniformly in $[0, T]$, and $(1/k)R_c^k \rightarrow \sum_{i=1}^n \bar{p}_i \bar{\lambda}_i T$ a.s.*

Analysis of the LPCC heuristic (b). This policy is defined through $A_i^k(t) = k\bar{\lambda}_i \cdot \int_0^t u_i^k(t) dt$, where $u_i^k(t)$ was defined in (15) and can be expressed as follows:

$$u_1^k(t) = \mathbf{1}\{\bar{X}_b^k(t) > 0\} \quad \text{and} \\ u_i^k(t) = \mathbf{1}\left\{\min\left(\frac{\bar{X}_b^k(t)}{T-t}\right) - \sum_{j<i} a_j \bar{\lambda}_j \geq a_i \bar{\lambda}_i\right\} \quad \text{for } i \geq 2, \quad (23)$$

where $\mathbf{1}\{\cdot\}$ is the indicator function. Similarly to the analysis of policy (c), the family $\{\bar{X}_b^k(t), t \in [0, T]\}$ is tight, and thus it has a converging subsequence $\{k_j\}$ on which $\bar{X}_b^{k_j}(t) \rightarrow \bar{x}_b(t)$ a.s., uniformly in t . Writing down $\bar{x}_b(t)$ and evaluating $u_i(t)$ reveals that in the limit model $u_i(t) = 1$ for all products i . Let R_b^k be the corresponding revenue. Similar arguments to the ones used above give the following result.

PROPOSITION 7. *Suppose that demand and capacity are scaled according to $\lambda^k(\cdot) = k\lambda(\cdot)$ and $C^k = kC$, and consider the LPCC heuristic defined through (23). Then, $\bar{X}_b^k(t) \rightarrow C - \bar{\rho}(t)$ a.s., uniformly in $[0, T]$, and $(1/k)R_b^k \rightarrow \sum_{i=1}^n \bar{p}_i \bar{\lambda}_i T$ a.s.*

That is, resolving in the context of the dynamic pricing or the LPCC heuristics is fluid-scale asymptotically optimal, as is the static policy (a). This shows that the suboptimal behavior demonstrated by Cooper (2002) does not persist in systems with large capacity and large demand. The same asymptotic performance can be achieved while resolving at discrete points in time, provided that this is done sufficiently frequently. If l^k is the time between resolving epochs, then the type of analysis used in studying discrete-review policies (see Harrison 1996 and Maglaras 2000) can be applied to show that it suffices that $l^k \downarrow 0$. This allows the number of demand requests between resolving periods to be large in absolute terms, but small compared to the capacity; e.g., a period $l^k = \log(k)/k$ corresponds to order $\log(k)$ arrivals.

4.3. Dynamic Pricing Network Revenue Management Problems

Suppose that the firm is operating a network of resources, indexed by $j = 1, \dots, m$, and that each product i request consumes A_{ij} units of resource j capacity. Let $A := [A_{ij}]$ denote the associated capacity consumption matrix, and assume that the initial capacity for each resource j is C_j . Then, the fluid model formulation of the network dynamic pricing problem is

$$\max_{\{\lambda(t), t \in [0, T]\}} \left\{ \int_0^T R(\lambda(t)) dt : \int_0^T A\lambda(t) dt \leq C \text{ and } \lambda(t) \in \mathcal{L} \forall t \right\}. \quad (24)$$

As before, this problem can be expressed in terms of ρ , which is defined by $\rho := A\lambda$. Specifically, let

$$R^r(\rho) := \max_{\lambda} \{R(\lambda) : A\lambda = \rho, \lambda \in \mathcal{L}\} \quad (25)$$

be the maximum achievable revenue rate when resource capacity is consumed at a rate ρ , and let $\lambda^r(\rho)$ denote the corresponding vector of optimal demand rates. Then, (24) can be reduced to

$$\max_{\{\rho(t), t \in [0, T]\}} \left\{ \int_0^T R^r(\rho(t)) dt : \int_0^T \rho(t) dt \leq C \text{ and } \rho(t) \in \mathcal{R} \forall t \right\}. \quad (26)$$

Let $\bar{\rho}$ denote the solution to (26). Then, $\lambda^r(\bar{\rho})$ is the vector of optimal demand rates for (24). This reduction could prove computationally beneficial because, as is often the case, the number of products (e.g., the number of fare-class and origin-destination pairs) tends to be greater than the number of resources (e.g., number of flights in a hub-and-spoke network). We refer the reader to Gallego and van Ryzin (1997) and Kleywegt (2001) for fluid formulations to multiproduct network revenue management problems.

5. Numerical Examples

This section reports on a set of numerical examples that contrast the performance of the heuristics proposed in the previous section to that of the optimal policy obtained from the dynamic program. Our base model has two products, each consuming one unit of

capacity per request, and a linear demand relationship of the form $\lambda(p) = \Lambda - Bp$ with $\Lambda = [0.3, 0.1]$ and $T = 200$ time periods. The price set is $\mathcal{P} = \{p: \Lambda - Bp \geq 0\}$, the inverse demand is $p(\lambda) = B^{-1}(\Lambda - \lambda)$, and the revenue function is $R(\lambda) = \lambda'B^{-1}(\Lambda - \lambda)$. The policies that we consider are the following:

- “RevMax” corresponds to the monopoly price vector \hat{p} that maximizes the aggregate instantaneous revenue rate disregarding the capacity constraints, given by $\hat{p} = B^{-1}(\Lambda - \hat{\lambda})$, where $\hat{\lambda} = \arg \max\{\lambda'B^{-1} \cdot (\Lambda - \lambda): \lambda \geq 0\} = (1/2)(B^{-1} + B^{-1'})^{-1}\Lambda$.
- “Fluid” implemented the price vector $\bar{p} = p(\lambda^r(\bar{\rho}))$ as defined in Proposition 4.
- “LPCC” is the joint (list) pricing and capacity control heuristic defined through (15).
- “DynPrice” is the dynamic implementation of the fluid policy defined through (16).
- “DP” implemented the solution of the dynamic program outlined in §3.

REMARK (COMMENTS ON COMPLEXITY). The expected revenues under the first two static policies were computed analytically using a binomial model that under DP was simply $V(C, T)$, while for the remaining policies we averaged out revenues over 1,000 simulated sample paths. The effort needed to compute $V(C, T)$ in problems with uniform capacity requirements (Tables 1 to 3) is that of solving a single-product pricing problem, which grows in proportion to $C \times T$. If $R^r(\rho)$ and $\lambda^r(\rho)$ cannot be expressed in closed form and are not precomputed and stored, the state space dimension of the dynamic program stays the same, but its complexity increases by a constant factor that depends on the number of products. Table 4 looks at examples with nonuniform capacity requirements, where the demand aggregation no longer applies and the backward induction step of the dynamic program needs to be changed, increasing the overall complexity again by a constant factor. The complexity of computing LPCC and DynPrice, the performance of which is the emphasis of our study, is negligible.

5.1.1. The Effect of the Joint Capacity Constraint.

Table 1 provides an illustration of the behavior of these heuristics on a particular problem instance as we varied the available capacity. Tables 2 and 3 will provide summary statistics for many randomly generated test cases. The demand model had no cross-price

Table 1 Optimality Gaps Relative to Optimal Policy (DP)

C	RevMax (%)	Fluid (%)	LPCC (%)	DynPrice (%)	DP
25	31.4	3.0	2.6	0.4	417.63
30	22.2	4.4	3.3	1.0	440.53
35	12.6	5.1	1.9	1.3	451.63
40	5.4	5.4	0.7	0.7	457.00
45	1.8	1.8	0.2	0.2	459.50
50	0.7	0.7	0.3	0.3	460.39

Note. $B = \text{diag}(1, 6)$

terms and was given by $\lambda_1(p) = 0.3 - 1p_1$ and $\lambda_2(p) = 0.1 - 6p_2$.

The Fluid pricing problem becomes unconstrained for $C \geq 40$ units, in which cases the RevMax and Fluid prices coincide. We observe the following. First, the relative performance under all heuristics improves as the capacity C increases; this is consistent with the results of Gallego and van Ryzin (1994, 1997) for Fluid and RevMax. Second, while the Fluid heuristic that incorporates the capacity constraint outperforms RevMax when capacity is scarce, its regret over the DP policy can still be substantial (2%–5%). Third, when the capacity is scarce, $C \leq 25$, the fluid prices effectively switch off Product 2 and operate the system as a single-product one. As the capacity increases, it is optimal to offer both products, and the effect of the capacity control of LPCC becomes more evident. Switching from an effectively single-product to a two-product solution causes the optimality gaps to be nonmonotone. Fourth, the dynamic pricing (resolving) heuristic is significantly better than all others when capacity is scarce.

5.1.2. The Cross-Price Elasticity Effects and Multiple Products. Table 2 summarizes results from 100 test cases for systems with fixed capacity $C = 40$, linear demand model with $\Lambda = [0.3, 0.1]$, and randomly

Table 2 Average and Standard Deviations of % Optimality Gaps for Two-Product Examples: $C = 40$, $\lambda(p) = \Lambda - Bp$, $\Lambda = [0.3, 0.1]$, and Random B Matrices

\hat{p}	RevMax	Fluid	LPCC	DynPrice
0.8–0.9	1.5 (0.3)	1.6 (0.3)	1.6 (0.3)	0.8 (0.2)
0.9–1.0	3.6 (0.6)	3.6 (0.5)	3.1 (1.6)	1.1 (0.2)
1.0–1.1	5.2 (1.1)	3.9 (0.3)	3.0 (1.5)	1.0 (0.2)
1.1–1.2	9.8 (1.3)	3.1 (0.4)	2.4 (0.3)	1.0 (0.2)
1.2–1.3	14.0 (1.5)	3.1 (0.4)	2.1 (0.2)	0.9 (0.2)

Table 3 Average and Standard Deviations of % Optimality Gaps for Three-Product Examples: $C = 40$, $\lambda(p) = \Lambda - Bp$, $\Lambda = [0.3, 0.05, 0.05]$, and Random B Matrices

$\hat{\rho}$	RevMax	Fluid	LPCC	DynPrice
0.8–0.9	1.5 (0.3)	1.6 (0.2)	1.5 (0.3)	0.8 (0.1)
0.9–1.0	3.3 (0.7)	3.3 (0.7)	3.3 (1.6)	1.0 (0.2)
1.0–1.1	6.4 (1.4)	3.7 (0.4)	3.6 (2.0)	1.1 (0.2)
1.1–1.2	9.4 (1.3)	3.1 (0.3)	3.0 (2.0)	0.9 (0.2)
1.2–1.3	14.0 (1.3)	3.0 (0.3)	2.1 (1.5)	0.8 (0.2)

generated B matrices; the elements of the main diagonal were drawn from a uniform distribution on $[0, 1]$, while the cross terms were drawn from a uniform distribution on $[-1, 0]$. We then tested that the dominant diagonal condition given in §2 was satisfied, which also assured that an inverse exists. Results have been grouped according to their nominal load factors $\hat{\rho} := (\sum_i a_i \hat{\lambda}_i)/C$, where $\hat{\lambda}$ is the RevMax demand vector that would maximize instantaneous revenue in the absence of the capacity constraint.¹ The table reports the average and standard deviation of the % optimality gap for each candidate policy. The main observation is that DynPrice significantly outperformed all other heuristics not only in terms of the average gap, but also in its standard deviation. In turn, LPCC improved over Fluid in cases where the capacity constraint was binding and the ability to incorporate the extra element of capacity control decisions was beneficial. Table 3 reports results for three-product examples, and illustrates the consistently good performance of the LPCC and DynPrice policies. The resolving structure of the latter is important in problems with multiple products and substitution and/or complementarity effects.

5.1.3. Nonuniform Capacity Requirements. Table 4 summarizes results for two models with two products that have different capacity requirements. This change complicates the associated dynamic programming formulation (see comments in the beginning of this section), but does not affect the fluid analysis of §4.1, and the heuristics extracted therein are still valid. In the notation of §4, product i consumes a_i units of capacity and $a_1 \neq a_2$. Our findings suggest

¹ A file with the test parameters for Tables 2 and 3 is available at <http://www.meiss.com/> or <http://www.gsb.columbia.edu/faculty/cmaglaras/maglaras.htm>.

Table 4 Optimality Gaps Relative to Optimal Policy (DP) (i) $B = [1 \ -0.4; -0.6 \ 6]$ and $a = [1 \ 2]$

C	RevMax (%)	Fluid (%)	LPCC (%)	DynPrice (%)	DP
30	42.7	4.3	3.7	2.7	476.91
40	27.1	5.6	2.1	1.5	496.05
50	12.0	5.7	1.7	1.0	502.77
60	2.9	3.2	1.1	0.3	505.81
70	0.8	0.7	0.4	0.3	506.73
80	0.3	0.2	0.2	0.1	506.86

(ii) $B = [1 \ -0.4; -0.6 \ 6]$ and $a = [2 \ 1]$

C	RevMax (%)	Fluid (%)	LPCC (%)	DynPrice (%)	DP
30	39.5	4.4	4.0	3.9	340.95
40	32.2	3.3	3.0	1.7	406.89
50	23.4	3.0	2.9	0.6	451.85
60	14.2	3.7	3.4	1.2	479.64
70	6.0	4.4	3.6	0.8	495.78
80	1.7	2.0	1.7	0.7	503.60

that the fluid model heuristics perform quite well in cases where the capacity requirements are small compared to the capacity itself. (Eventually, as C and λ grow large, these policies become asymptotically optimal in the sense of Propositions 5 to 7.)

The last three tables illustrate that static pricing is close to optimal in problem instances that are either very under- or overloaded, and that the effect of dynamic pricing, either through the optimal policy (DP) or the DynPrice heuristic, is most pronounced when $\hat{\rho}$ is close to 1. This is intuitive: When $\hat{\rho} \ll 1$, the system's capacity far exceeds the nominal requirement under the prices that maximize revenues per unit time (implemented in RevMax), and thus one would expect that the static pricing heuristics RevMax and Fluid to be close to optimal. In contrast, if $\hat{\rho} \gg 1$, the system's capacity is too low, and the static prices implemented through the Fluid heuristic are close to optimal. Finally, we did not observe any problem instances where the performance under the DynPrice heuristic was worse (in a statistically significant manner) than that under the LPCC or Fluid heuristics. The narrowest gap between LPCC and DynPrice was observed in the first row of Table 4, which, due to the nonuniform capacity requirements, corresponds to the problems with the lowest overall capacity in terms of the number of sales; this is the closest to the

setting in Cooper's example (2002), which had low capacity and a short selling horizon.

References

- Bitran, G., R. Caldentey. 2003. An overview of pricing models for revenue management. *Manufacturing Service Oper. Management* 5(3) 203–229.
- Brumelle, S., J. McGill. 1993. Airline seat allocation with multiple nested fare classes. *Oper. Res.* 41(1) 127–137.
- Cooper, W. 2002. Asymptotic behavior of an allocation policy for revenue management. *Oper. Res.* 50(4) 720–727.
- Dai, J. G., R. J. Williams. 1995. Existence and uniqueness of semi-martingale reflecting Brownian motions in convex polyhedrons. *Theory Probab. Appl.* 40(1) 1–40.
- Elmaghraby, W., P. Keskinocak. 2003. Dynamic pricing: Research overview, current practices and future directions. *Management Sci.* 49(10) 1287–1389.
- Feng, Y., B. Xiao. 2000. A continuous-time yield management model with multiple prices and reversible price changes. *Management Sci.* 46(5) 644–657.
- Feng, Y., B. Xiao. 2004. Integration of pricing and capacity allocation for perishable products. Preprint, Chinese University of Hong Kong.
- Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* 40(8) 999–1020.
- Gallego, G., G. van Ryzin. 1997. A multiproduct dynamic pricing problem and its applications to network yield management. *Oper. Res.* 45(1) 24–41.
- Harrison, J. M. 1996. The BIGSTEP approach to flow management in stochastic processing networks. F. Kelly, S. Zachary, I. Ziedins, eds. *Stochastic Networks: Theory and Applications*. Oxford University Press, Oxford, UK, 57–90.
- Horn, R. A., C. R. Johnson. 1994. *Matrix Analysis*. Cambridge University Press, Cambridge, UK.
- Kleywegt, A. J. 2001. An optimal control problem of dynamic pricing. Working paper, Georgia Institute of Technology, Atlanta, GA.
- Lautenbacher, C., S. Stidham. 1999. The underlying Markov decision process in the single-leg airline yield-management problem. *Transportation Sci.* 33(2) 136–146.
- Lee, T., M. Hersh. 1993. A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Sci.* 27(3) 252–265.
- Lin, G. Y., Y. Lu, D. D. Yao. 2003. The stochastic knapsack revisited: Structure, switch-over policies, and dynamic pricing. Preprint, Columbia University, New York.
- Maglaras, C. 2000. Discrete-review policies for scheduling stochastic networks: Trajectory tracking and fluid-scale asymptotic optimality. *Ann. Appl. Probab.* 10(3) 897–929.
- Maglaras, C. 2006. Revenue management for a multi-class single-server queue via a fluid model analysis. *Oper. Res.* Forthcoming.
- McGill, J., G. van Ryzin. 1999. Revenue management: Research overview and prospects. *Transportation Sci.* 33(2) 233–256.
- Reiman, M. I. 2002. Asymptotically optimal dynamic pricing for a wholesale/retail telecommunications service provider. *2nd INFORMS Conf. Revenue Management*, New York.
- Talluri, K., G. van Ryzin. 2004a. Revenue management under a general discrete choice model of consumer behavior. *Management Sci.* 50(1) 15–33.
- Talluri, K., G. van Ryzin. 2004b. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers, Boston/Dordrecht/London.
- Vulcano, G., G. van Ryzin, C. Maglaras. 2002. Optimal dynamic auctions for revenue management. *Management Sci.* 48(11) 1388–1407.
- Zhao, W., Y.-S. Zheng. 2000. Optimal dynamic pricing for perishable assets with nonhomogeneous demand. *Management Sci.* 46(3) 375–388.