

Revenue Management for a Multiclass Single-Server Queue via a Fluid Model Analysis

Constantinos Maglaras

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

Motivated by the recent adoption of tactical pricing strategies in manufacturing settings, this paper studies a problem of dynamic pricing for a multiproduct make-to-order system. Specifically, for a multiclass $M_n/M/1$ queue with controllable arrival rates, general demand curves, and linear holding costs, we study the problem of maximizing the expected revenues minus holding costs by selecting a pair of dynamic pricing and sequencing policies. Using a deterministic and continuous (fluid model) relaxation of this problem, which can be justified asymptotically as the capacity and the potential demand grow large, we show the following: (i) greedy sequencing (i.e., the $c\mu$ -rule) is optimal, (ii) the optimal pricing and sequencing decisions decouple in finite time, after which (iii) the system evolution and thus the optimal prices depend only on the total workload. Building on (i)–(iii), we propose a one-dimensional workload relaxation to the fluid pricing problem that is simpler to analyze, and leads to intuitive and implementable pricing heuristics. Numerical results illustrate the near-optimal performance of the fluid heuristics and the benefits from dynamic pricing.

Subject classifications: revenue management; yield management; dynamic pricing; queuing; sequencing; fluid models.

Area of review: Manufacturing, Service, and Supply Chain Operations.

History: Received November 2003; revisions received December 2004, August 2005; accepted September 2005.

1. Introduction

The last decade has been marked by a growing interest in the adoption of dynamic pricing strategies in such diverse areas as the airline, hotel, and retail industries. In most of these cases, the firm controls a fixed capacity of resources (e.g., the number of seats in a flight) that have to be sold up to a deadline (e.g., the flight departure time). By dynamic pricing, we refer to the tactical optimization of the price of a product or service (e.g., of an airline ticket) as a function of the remaining capacity and time-to-go to maximize the expected revenues extracted from these fixed resources. This practice, often referred to as *revenue management*, is supported by sophisticated information systems processing large amounts of demand data, and relies on an implicit assumption that the firm can apply such price changes in a relatively efficient manner.

More recently, manufacturing firms have also started evaluating the use of such tactical economic optimization tools, with one notable example coming from the automotive industry in the context of its effort to market and produce custom cars in a make-to-order fashion. Broadly speaking, automobile manufacturers try to dynamically adjust the price, target lead time, rebate, etc., for a new order as a function of the existing outstanding orders, and simultaneously select the appropriate production schedule to optimize their profitability. Joint use of economic and operational controls allows the manufacturer to be more responsive to changes in the market conditions and fluctuations in the operating environment due to randomness of the demand

and production functions. Operationally, this raises several interesting questions. For example, in what ways are the economic and operational decisions coupled? What are the benefits of dynamic pricing in a production setting? And what are practical and efficient pricing and sequencing heuristics for such problems? With this motivation, this paper studies the problem of jointly optimizing over the dynamic pricing and sequencing policies for the multiproduct, single-server queuing system. Broadly speaking, this problem lies in the interface of stochastic network theory and revenue management, and the approach taken in this paper combines modelling and analysis techniques from these two areas. Its results illustrate the potential benefits of jointly optimizing pricing and production decisions, and offer some insight on how to practically integrate these two functions that tend to operate separately in many organizations.

We consider a make-to-order firm that produces multiple products, that is modelled as a single-server multiclass $M_n/M/1$ queue. The firm is assumed to operate in a market with imperfect competition, and has power to influence the demand for the various products by varying its price menu. Assuming a general demand curve, linear holding costs incurred by the firm,¹ and convex capacity costs, we study the problem of finding the optimal state-dependent pricing and sequencing strategy as well as a static vector of production rates to optimize the system's long-run expected profit rate. A natural starting point would be to formulate an appropriate control problem for a multiclass

queue within the framework of Markov decision processes (MDPs). While MDPs provide detailed descriptions of the system dynamics and the optimal control problem, they are—with the exception of very restricted examples—not amenable to exact analysis. This paper studies an approximate formulation of the profit maximization problem of interest, posed in the context of the associated deterministic and continuous “fluid model” approximation to the underlying stochastic production system. This can be rigorously justified through a strong-law-of-large-numbers type of scaling in settings where the production rate and potential demand grow proportionally large. Such models have been used successfully in the literature both in revenue management settings that lack production dynamics, and in production systems that do not include the tactical pricing decisions.

The main findings of this paper are the following.

(1) *Capacity choice.* We show that the long-run average profit maximization problem for the fluid model reduces to a static problem of choosing a vector of target demand rates and the service rate vector that maximize profits in the absence of holding costs (Theorem 1). The optimal service rate vector makes the capacity constraint binding (Corollary 1). This problem determines the optimal capacity, but is too coarse to specify good pricing and sequencing policies.

(2) *Structural analysis of the joint pricing and sequencing problem.* Fixing the capacity at the value prescribed above, we then focus on the infinite horizon total profit criterion to determine the optimal sequencing and pricing policies. We show that “reasonable” policies eventually drain the queues (Proposition 1), characterize the properties of the associated value function (Proposition 2), and construct an optimality verification result for this problem and through the associated Bellman equation (Proposition 3). We finally show that sequencing decisions are made according to the $c\mu$ -rule (Proposition 4) and that the demand rates are nonincreasing functions of the queue length (Proposition 5).

(3) *State-space collapse and workload relaxation.* A consequence of the fluid model optimal sequencing and demand controls is that after a finite time the queue length is always in a (efficient) configuration where all of the workload is held at the “cheapest” product class (Theorem 2). This state-space collapse result simplifies the solution of the control problem from then onwards. It also suggests formulating a workload relaxation for our optimal pricing problem, which focuses on the evolution of the workload process, and uses an aggregated demand model and an appropriate holding cost rate. This one-dimensional formulation is simpler to analyze than the multiproduct one (Theorem 3), leads to intuitive and implementable heuristics, and is often solvable in closed form (§4.3 studies the linear demand case).

(4) *Managerial insights.* The key insights gleaned from our analysis are the following: (a) the policy that maximizes

the time-average profits in the deterministic fluid model invests in scarce capacity and operates the system at almost full utilization, increasing prices and reducing demands if backlogs grow large; (b) orders are sequenced according to the greedy ($c\mu$) rule to minimize instantaneous holding costs irrespective of the pricing decisions; (c) the sequencing and pricing decisions are decoupled after an initial transient period whose length is characterized, in the sense that thereafter pricing decisions are made as a function of the aggregate system workload, which does not depend on the sequencing rule. The last two insights are characteristics of the optimal policies in the fluid model formulation of the joint pricing and sequencing problem. Together they suggest a heuristic that sequences jobs according to the $c\mu$ -rule and prices according to the solution of a fluid control problem formulated in terms of the aggregate system workload. The latter is simpler to solve and leads to intuitive policy recommendations. It also minimizes the amount of information sharing between the production and pricing functions of an organization, which is appealing from the viewpoint of operationalizing these joint decisions. The numerical results of §5.2 illustrate the effectiveness of these heuristics when compared to the solution to the original MDP formulation.

The remainder of this paper is structured as follows: This section concludes with a brief literature review. Section 2 describes the model, §3 studies the associated fluid control problem, and §4 studies its workload relaxation. Section 5 summarizes the key insights of our analysis and reports some numerical results.

Literature Review. Our work is related to two bodies of literature focusing on stochastic processing network theory and revenue management, respectively. Standard textbooks on queuing networks provide some background on single-server queues, and standard dynamic programming textbooks, such as Bertsekas (1995), provide the necessary background for the solution of the underlying MDP problem formulations. An early paper from Low (1974) studied a single-product, multiserver problem, and showed the monotonicity of the optimal price policy and proposed an iterative algorithm for computing it. This paper also considered linear holding costs as a surrogate for waiting-time penalties incurred by the firm. The analysis in this paper uses fluid model approximations for queues with state-dependent parameters that were developed by Mandelbaum and Pats (1995) for the case of exponential service times and Poisson arrival streams. In part, this paper extends their results by adding a control dimension to some of their models. The use of fluid models for dynamic pricing in manufacturing systems has been discussed in Kleywegt (2001), while the literature on fluid models for purposes of sequencing and routing control is large; see, e.g., Chen and Yao (1993), Avram et al. (1995), Maglaras (2000), and the references therein. Workload formulations arise in stochastic network control problems that are considered in

the context of their approximate Brownian model formulations. This idea and its consequences in policy design have been pioneered by the work of Harrison (1988, 2000) and Harrison and Van Mieghem (1996). Workload fluid models were first introduced in Harrison (1995), while the use of workload relaxations of fluid model control problems involving sequencing, routing, and admission control were proposed in Meyn (2001).

The work by Gallego and van Ryzin (1994, 1997), the review papers by McGill and van Ryzin (1999), Bitran and Caldentey (2003), and Elmaghraby and Keskinocak (2003), and the book by Talluri and van Ryzin (2004) provide background on pricing and revenue management. Biller et al. (2002) discuss the dynamic pricing problem in the context of the automotive industry, and provide a deterministic, finite horizon analysis of the single-product case. Some of the results in Maglaras and Meissner (2006) that studied multiproduct revenue management problems are used in §4. There are many papers that are tangentially related to ours insofar as they too combine some form of pricing with the analysis of production or inventory systems. Mendelson and Whang (1990) and a stream of related papers have focused on static pricing and sequencing control for social welfare optimization in a multiproduct $M/M/1$ queue facing a market of heterogeneous price- and delay-sensitive users. Afeche (2004) looks at a simplified form of this model under a revenue-maximizing objective, and provides a thorough review of static pricing papers in queues, mostly under an atomistic customer demand model (cf. comment at the end of §2). Chen and Frank (2000) look at dynamic pricing for a single-product $M/M/1$ queue using dynamic programming arguments, and Kalish (1983), Kachani and Perakis (2002), and Kleywegt (2001) are examples of papers that study dynamic pricing issues using some form of fluid or diffusion model. Examples of papers that include inventory control with some element of pricing control decisions are Federgruen and Heching (1999) and Chen and Simchi-Levi (2004a, b).

2. Model Formulation

Consider a single-server production facility (the firm) that offers multiple products, indexed by $i = 1, \dots, I$, to a market of price-sensitive users. It operates in a market with imperfect competition, and has power to influence its vector of demand rates by varying its price menu p . The demand process is assumed to be an I -dimensional nonhomogeneous Poisson process with rate vector $\lambda(p)$ determined through a *demand function* that maps the price vector into a vector of instantaneous demand rates, $\lambda: \mathcal{P} \rightarrow \mathcal{L}$, where $\mathcal{P} \subseteq \mathbb{R}^I$ is the set of feasible price vectors, and $\mathcal{L} = \{x \geq 0: x = \lambda(p), p \in \mathcal{P}\} \subseteq \mathbb{R}_+^I$ is the set of achievable demand rate vectors. Note that $\lambda(\cdot)$ only depends on the time t through the price posted at that instance. We assume that \mathcal{L} is a convex set, the demand function $\lambda(\cdot)$ is continuously differentiable and bounded, and (a) for each

product i , $\lambda_i(p)$ is strictly decreasing in p_i , (b) for each $p_{-i} = (p_1, \dots, p_{i-1}, p_{i+1}, \dots, p_I)$, there exists a *null price* $p_i^\infty(p_{-i}) \in \mathcal{P}$ such that $\lim_{p_i \rightarrow p_i^\infty(p_{-i})} \lambda_i(p_i, p_{-i}) = 0$, and (c) the revenue rate $p \cdot \lambda(p) = \sum_i p_i \lambda_i(p)$ is bounded for all $p \in \mathcal{P}$ and has a finite maximizer. (For any two n -vectors, $x \cdot y$ will denote their inner product.)

Under these assumptions, there exists an 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 $p(\lambda)$. Although, in general, this inverse mapping need not be unique, it turns out that it is for common examples of demand relations; see Talluri and van Ryzin (2004, §7.3.2). Following a standard practice from revenue management, we may then view the demand rate vector as the firm's control, and once this is determined derive the corresponding prices using the inverse demand function. In this case, the expected revenue rate will be denoted by $r(\lambda)$, where $r(\lambda) = \lambda \cdot p(\lambda)$. We will assume that $r(\lambda)$ is continuous, bounded, and strictly concave, and denote its maximizer by $\lambda^\dagger = \arg \max\{r(\lambda): \lambda \in \mathcal{L}\}$.

Infinite capacity buffers are associated with each product, and $Q_i(t)$ will denote the number of product i jobs in the system (i.e., in queue or in service) at time t . Their service times are i.i.d. exponentially distributed with mean m_i (or rate $\mu_i := 1/m_i$). The *load* or *traffic intensity* of the system when the demand vector is λ is defined as $\rho := m \cdot \lambda$. For future use, we define the *aggregate* revenue function as the maximum achievable revenue rate when all products jointly consume capacity at rate ρ ,

$$R(\rho; \mu) = \max \left\{ r(\lambda): \lambda \in \mathcal{L}, \sum_i \lambda_i / \mu_i = \rho \right\}, \quad (1)$$

and denote by $\lambda^r(\rho)$ the corresponding maximizer. We will assume that λ_i^r is nondecreasing in ρ for all products i . This appears to be a mild assumption that is satisfied by many commonly used demand models such as the linear, exponential, pareto, and multinomial logit, to list but a few examples. From the properties of $r(\cdot)$ it follows that $R(\cdot; \mu)$ is concave, bounded, and has a finite maximizer that we denote by $\rho^\dagger := \arg \max\{R(\rho; \mu): \rho = m \cdot \lambda, \lambda \in \mathcal{L}\} = m \cdot \lambda^\dagger$. When $\rho^\dagger \geq 1$, capacity is scarce in the sense that the revenue-maximizing demand rates would make the system unstable, whereas if $\rho^\dagger < 1$, the capacity is ample.

EXAMPLE. The linear demand model is given by

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

or in vector form $\lambda(p) = \Lambda - Bp$, 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 $r(\lambda) = \lambda \cdot B^{-1}(\Lambda - \lambda)$, respectively. To ensure that these expressions are well defined and satisfy our assumptions, we will require that $b_{ii} > 0$, and either $b_{ii} > \sum_{j \neq i} |b_{ji}|$ or $b_{ii} > \sum_{j \neq i} |b_{ij}|$ for all i . Both conditions relate to the marginal

effect of price changes to individual and total demand, and guarantee that B^{-1} exists and has eigenvalues with positive real parts (Horn and Johnson 1994, Theorem 6.1.10). Finally, the aggregate revenue function defined through (1) is

$$R(\rho) = -\alpha_i \rho^2 + \beta_i \rho + \gamma_i \quad \text{for } \rho \in [r_{i-1}, r_i),$$

with $0 = r_0 \leq r_1 \leq r_2 \leq \dots \leq r_I$, and the constants $(\alpha_i, \beta_i, \gamma_i)$ and r_i depend on the model parameters Λ, B, μ , and are such that $R(\rho)$ is continuous, almost everywhere differentiable, and increasing for all $\rho \leq \rho^\dagger$. (The calculation of these constants is given in the appendix.)

The firm has discretion with respect to the sequencing of jobs at the server, and pricing decisions for each product. Within each product, orders are processed in first-in-first-out (FIFO), the server can only work on one job at any given time, and preemptive-resume type of service is allowed. Under these assumptions, a sequencing policy takes the form of the I -dimensional cumulative allocation process $(T(t): t \geq 0)$ with $T(0) = 0$, where $T_i(t)$ denotes the cumulative time that the server has allocated to class i jobs up to time t . In addition, $T(t)$ is continuous and non-decreasing, and satisfies the capacity constraint

$$\sum_i T_i(t) - \sum_i T_i(s) \leq t - s \quad \text{for } 0 \leq s \leq t < \infty. \quad (2)$$

Let $p(t)$ be the vector of prices posted at time t and $\lambda(t)$ be the corresponding vector of demand rates. As mentioned above, we will treat the demand rate vector as the control, and infer the corresponding price vector via the inverse demand function. The demand policy is the I -dimensional process $(\lambda(t): t \geq 0)$. Both T and λ are restricted to be nonanticipating controls; i.e., decisions at time t can only depend on information that is available up to that time.

Finally, the firm incurs two types of cost. The first is a linear congestion cost given by $\sum_i c_i Q_i(t)$, where $c_i > 0$ for all products. It either captures the weighted delay costs incurred by all outstanding orders or—if appropriate—some notion of cost associated with work-in-progress inventory involved in production. The second is the cost of operating a facility with processing capabilities equal to μ , which is $h(\mu)$ per unit time, where $h: \mathbb{R}_+^I \rightarrow \mathbb{R}_+$ is a convex, strictly increasing in each of its arguments, differentiable function. The processing capability vector μ is static, i.e., selected at time $t = 0$ and fixed thereafter.

The firm's problem is to select the (static) capacity vector μ , the sequencing policy $T(\cdot)$, and the demand rates $\lambda(\cdot)$ to maximize the long-run average profit rate, viz

$$\text{maximize } \lim_{t \rightarrow \infty} \frac{1}{t} \mathbb{E} \left[\int_0^t (r(\lambda(s)) - c \cdot Q(s)) ds \right] - h(\mu). \quad (3)$$

REMARK 1. Let $A_i(t)$ be the number of product i orders that have arrived up to time t . The firm's problem is to choose $\mu, p(\cdot), T(\cdot)$ to maximize $\lim_{t \rightarrow \infty} (1/t) \mathbb{E} \left[\int_0^t p(s) \cdot dA(s) - \int_0^t c \cdot Q(s) ds \right] - h(\mu)$. Using a standard result for

intensity control problems (see Brémaud 1980, §II.2), this can be rewritten as (3). For stable, stationary Markov policies, i.e., demand rates that can be expressed in the form $\lambda(\cdot) = \lambda(Q(\cdot))$, (3) reduces to the steady-state expected profit criterion. We will not rigorously justify these points because our subsequent analysis will not address (3) directly but instead rely on the use of deterministic and continuous fluid model approximations.

Discussion of Modelling Assumptions. In terms of probabilistic assumptions, the one regarding the Poisson nature of the demand processes is important to be able to justify the deterministic fluid models used in this paper as rigorous limits under dynamic pricing policies; this uses the results from Mandelbaum and Pats (1995). The assumption on exponential service times could be extended to allow for general distributions at little additional cost, by easily adjusting the asymptotic results in Mandelbaum and Pats (1995). (The exponential service time assumption in Mandelbaum and Pats was imposed because the service rate was allowed to be state dependent; with constant service rates—as in this paper—one can get the strong-law type of limit for the service processes using renewal theory.) Because we make use of the exponential assumption in the numerical experiments, where we compare the performance of our derived heuristics against the solution of the dynamic program associated with (3), we will proceed under that simplifying assumption. In any case, allowing for general service time distributions would have no effect on the fluid model analysis that disregards the second moment information. As in most papers on pricing in queues and revenue management, our model assumes that self-interested customers decide whether to place an order based solely on the price vector at the time of their arrival; i.e., they are strategic in making purchase selections by explicitly or implicitly optimizing some form of a personal utility function, but they are not strategic in selecting the timing of their arrival in response to the firm's pricing strategy. This allows one to address the firm's pricing problem as an optimal intensity control problem not involving a game-theoretic analysis; see Lariviere and Van Mieghem (2004) for a discussion of this point and a justification of the Poisson arrival process assumption as the solution of such a game-theoretic analysis for a related model. Also, in our model the demand relationship λ captures the aggregate behavior of all potential customers. This is in contrast to more detailed models, such as the one used by Mendelson (1985), Mendelson and Whang (1990), and Van Mieghem (2000), that obtain the demand rates through a customer-by-customer analysis based on more primitive model elements such as personal utility functions and price and delay sensitivity parameters.

Finally, because fluid models best capture the transient behavior of the underlying system, it seems more natural to consider an undiscounted criterion in the fluid model formulation, which motivated the objective given above. An alternate formulation could consider an infinite horizon discounted profit criterion. Numerical tests showed that the

fluid heuristics obtained for the undiscounted criterion in §3–§4 performed well even when compared to the stochastic dynamic programming solution for a discounted objective provided that the discount factors were moderate, i.e., where the time scale for discounting is long compared to job service and interarrival times.

3. Analysis of the Associated Fluid Model Control Problem

The control problem posed above could be addressed using the theory of MDPs (if one restricts attention to exponentially distributed service times), but this is both analytically and computationally hard due to the multidimensional nature of the state space. The approach taken in this paper relies on the associated fluid control problem that will be developed below and studied over the next two sections.

3.1. Formulation of the Associated Fluid Control Problem

The fluid model is derived by replacing the discrete and stochastic demand and production processes by continuous flows with the corresponding deterministic rates. It is rigorously derived as a limit under a strong-law-of-large-numbers type of scaling when we let the production rate and demand grow proportionally large. This amounts to embedding our problem in a sequence of systems with model primitives that scale according to

$$\lambda^n(\cdot) = n\lambda(\cdot), \quad c^n = c, \quad \text{and} \quad h^n(n\cdot) = nh(\cdot). \quad (4)$$

These scaling relations imply that it is reasonable to also grow the capacity proportionally to n according to $\mu^n = n\mu$, and would result in a problem where the various revenue and cost contributions are all of the same order of magnitude that is itself proportional to n . Scaling the initial condition according to $Q^n(0) = nz$, and studying the limit of $q^n(t) := Q^n(t)/n$ as the scaling parameter n grows large, we get that $q^n(t)$ converges to a continuous fluid limit process denoted by $q(t)$.²

In the sequel, let $M = \text{diag}(\mu_1, \dots, \mu_I)$, and $u_i(t)$ denote the fraction of the server capacity dedicated to processing an order of product i at time t ; in the fluid model the server is allowed to split its effort across different product classes. The fluid model equations (see Mandelbaum and Pats 1995) are

$$q(t) = z + \int_0^t \lambda(s) ds - M \int_0^t u(s) ds, \quad q(0) = z, \quad (5)$$

$$q(t) \geq 0, \quad 0 \leq u_i(t) \leq 1, \quad \sum_i u_i(t) \leq 1, \quad \lambda(t) \in \mathcal{L}. \quad (6)$$

Two other quantities of interest are the traffic intensity at time t defined as $\rho(t) := m \cdot \lambda(t)$, and the server workload $w(t) := m \cdot q(t)$, which is the amount of time needed for the

server to clear the current backlog disregarding any future arrivals. For future reference, we note that

$$w(t) = w(0) + \int_0^t \rho(s) ds - t + I(t), \quad (7)$$

where $w(0) = m \cdot z$ and $I(t) := t - \sum_i \int_0^t u_i(s) ds$ is the cumulative server idleness up to time t .

In the spirit of (3), it is natural to consider maximizing the time-average profit criterion

$$\lim_{t \rightarrow \infty} \frac{1}{t} \left[\int_0^t (r(\lambda(s)) - c \cdot q(s)) ds \right] - h(\mu). \quad (8)$$

In considering this problem, we will also restrict attention to controls λ, u that are right-continuous functions with left limits (RCLL) and are nonanticipating, i.e., decisions at time t can only use information that has been made available up to that time, which we summarize below:

$$\lambda(t), u(t) \text{ are nonanticipating and RCLL for } t \geq 0. \quad (9)$$

The next theorem shows that this criterion is useful in selecting the optimal capacity vector μ , but, as its proof highlights, is too coarse to identify “good” pricing and sequencing policies. The latter is addressed in a more “refined” formulation given in (13).

THEOREM 1. *Consider the problem of maximizing (8) over $\mu, \lambda(\cdot)$ and $u(\cdot)$ subject to the fluid model equations (5)–(6) and (9). Then, the optimal time-average profit rate is $\hat{\pi}$ defined by*

$$\hat{\pi} := \max_{\mu, \lambda} \left\{ r(\lambda) - h(\mu) : \lambda \in \mathcal{L}, \sum_i \lambda_i / \mu_i \leq 1, \mu > 0 \right\}, \quad (10)$$

while the optimal capacity vector $\hat{\mu}$ is the associated optimizer, which is unique.

PROOF. *Step 1.* We demonstrate a feasible control for this problem with finite average profit rate. Pick any vector μ and consider the policy: $\lambda(t) = 0$ and $u(t)$ is any nonidling rule (i.e., $\sum_i u_i(t) = 1$) for $t \leq w(0)$; and,

$$\lambda(t) = \hat{\lambda}(\mu) := \arg \max_{\lambda} \left\{ r(\lambda) : \lambda \in \mathcal{L}, \sum_i \lambda_i / \mu_i \leq 1 \right\},$$

and $u(t) = M^{-1}\lambda(t)$ for all $t > w(0)$. From (5), it follows that under that policy $q(t) = 0$ for all $t \geq w(0)$, and that $\lim_{t \rightarrow \infty} (1/t) \int_0^t (r(\lambda(s)) - c \cdot q(s)) ds - h(\mu) = r(\hat{\lambda}(\mu)) - h(\mu)$, which is finite.

Step 2. We show that any optimal control must satisfy the stability condition

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \int_0^t \rho(s) ds \leq 1. \quad (11)$$

Suppose for some converging subsequence,

$$\limsup_{t \rightarrow \infty} (1/t) \int_0^t \rho(s) ds > 1.$$

Then, from (7) we get that $w(t) \rightarrow \infty$ and $q(t) \rightarrow \infty$ as $t \rightarrow \infty$. Because $r(\lambda)$ is bounded, it follows that $\limsup_{t \rightarrow \infty} (1/t) \int_0^t (r(\lambda(s)) - c \cdot q(s)) ds = -\infty$, which is suboptimal. By contradiction, we establish (11).

Step 3. Assume $\mu > 0$ (because otherwise $\lambda(t) = 0$ for all $t \geq 0$) such that $\rho(\cdot)$ is well defined. Recall the definition of $R(\rho)$ in (1) and note that

$$\begin{aligned} & \frac{1}{t} \int_0^t (r(\lambda(s)) - c \cdot q(s)) ds - h(\mu) \\ & \leq \frac{1}{t} \int_0^t r(\lambda(s)) ds - h(\mu) \\ & \leq \frac{1}{t} \int_0^t R(\rho(s)) ds - h(\mu) \\ & \leq R\left(\frac{1}{t} \int_0^t \rho(s) ds\right) - h(\mu) \\ & \leq \hat{\pi}. \end{aligned}$$

The first inequality follows the fact that $q(t) \geq 0$, the second from the definition of $R(\cdot)$, the third from Jensen's inequality, and the last one by noting that (10) can be rewritten as $\max_{\mu, \lambda, \rho} \{R(\rho) - h(\mu) : \lambda \in \mathcal{L}, \sum_i \lambda_i / \mu_i = \rho, \rho \leq 1, \mu \geq 0\}$, where the constraint $\rho \leq 1$ is needed due to (11), and $\sum_i \lambda_i / \mu_i \leq 1$ is convex in (λ, μ) for $\lambda, \mu \geq 0$. The control specified in Step 1 for $\mu = \hat{\mu}$ achieves the upper bound $\hat{\pi}$. Finally, the concavity of $r(\lambda) - h(\mu)$ implies the uniqueness of $\hat{\mu}$, which completes the proof. \square

Hereafter we will fix the capacity vector to $\hat{\mu}$ defined via (10), but to simplify notation we will denote it by μ and use $\hat{\lambda}$ in place of $\hat{\lambda}(\hat{\mu})$. We also relabel the products such that

$$c_1 \mu_1 \leq c_2 \mu_2 \leq \dots \leq c_I \mu_I. \quad (12)$$

The structure of (10) and the assumptions on $r(\cdot)$ and $h(\cdot)$ imply the following:

COROLLARY 1. *The capacity is scarce in the sense that (i) $m \cdot \hat{\lambda} = 1$ and (ii) $\hat{\lambda} \leq \lambda^*$.*

It is worth remarking that this fluid model analysis will invest in zero processing capacity if the marginal revenue rate at $\rho = 0$ is smaller than the marginal cost of capacity at $\mu = 0$. This was also observed in the numerical results of §5.

While this criterion is useful in optimizing over the optimal level of production capacity, it is too coarse to use in the fluid model to construct good pricing and sequencing policies for the underlying multiproduct revenue management problem. Indeed, any control (λ, u) that drains the queue and then uses $(\hat{\lambda}, \hat{u})$ thereafter is optimal for the time-average criterion because the transient phase until the queue is emptied is washed out of the objective. Instead, we will consider the total profit criterion

$$\int_0^\infty (\tilde{r}(\lambda(t)) - c \cdot q(t)) dt, \quad (13)$$

where $\tilde{r}(\lambda) = r(\lambda) - r(\hat{\lambda})$. (The capacity has been optimized via the time-average problem, and its associated cost

need not be included in this formulation. Note that including it in the total profit criterion could result in a different optimum for μ , but this would no longer be optimal for the higher-order criterion given in (8).) Total profit criteria are common in fluid model control problems in the literature because they emphasize the transient system behavior, which is well captured through the fluid equations. The remainder of this section will study the problem of maximizing (13) over (λ, u) subject to the fluid model equations (5)–(6) and (9).

3.2. Characterization of the Fluid Optimal Controls

I. Preliminary Structural Results. The first result establishes that good control policies must be “stable” in that they eventually empty the queue-length vector, and that there exists an optimal solution. Restricting attention to such policies, we can then characterize the optimal one as the solution to the associated Bellman equation.

PROPOSITION 1. *Consider the problem of maximizing (13) over $\lambda(\cdot)$ and $u(\cdot)$ subject to (5)–(6) and (9). Then: (i) it suffices to restrict attention to controls under which $q(t) \rightarrow 0$ as $t \rightarrow \infty$, and (ii) there exists an optimal pair (λ^*, u^*) and an associated trajectory q^* in the sense that*

$$\int_0^\infty (\tilde{r}(\lambda^*(t)) - c \cdot q^*(t)) dt \geq \int_0^\infty (\tilde{r}(\lambda(t)) - c \cdot q(t)) dt,$$

where (λ, u) is any other feasible control and q is the associated trajectory.

The proof is relegated to the appendix. To facilitate analysis and the numerical computations that will follow, we will restrict attention to queue lengths that lie in a large but bounded domain. In detail, assuming that each queue is no greater than a large constant N , we will express the set of possible queue-length vectors by $\mathcal{Q} := \{q \geq 0 : m \cdot q \leq N_w\}$, where $N_w = m \cdot N$. That is, we express the domain as a function of the aggregate workload $m \cdot q$ as opposed to the individual queue lengths, because as it will turn out the workload is nonincreasing, whereas the queue lengths are not. For concreteness, we will enforce this condition by imposing the constraint that $\lambda(t) \leq \hat{\lambda}$ if $m \cdot q(t) = N_w$; we will show at the end of this subsection that this condition is never invoked, and therefore does not change the structure of the optimal policy; see the remark following Proposition 5.

Next, we proceed with an informal derivation of the Bellman equation associated with (13) that characterizes the optimal policy. Let $V(q)$ denote the optimal profit extracted under (13) starting from $q \in \mathcal{Q}$. The existence of this function follows from Proposition 1. The next proposition summarizes the main structural properties of the value function that are used in subsequent analysis. The proof is given in the appendix.

PROPOSITION 2. Let $V(q)$ denote the optimal profit extracted under (13) starting from $q \in \mathcal{Q}$. Then: (i) for all $q \in \mathcal{Q}$, $-\infty < V(q) \leq 0$, and $V(0) = 0$; (ii) $V(\cdot)$ is Lipschitz continuous, and therefore it is almost everywhere (a.e.) differentiable; (iii) $V(\cdot)$ is concave; and (iv) $\nabla V(z) \leq \nabla V(0)$ for all $z \geq 0$ (componentwise).

Using the definition of $V(\cdot)$ and its a.e. differentiability, a standard dynamic programming argument gives that

$$V(q(t)) = \max_{\lambda \in \mathcal{L}, u \in \mathcal{U}(q)} [(\tilde{r}(\lambda) - c \cdot q(t))\delta t + V(q(t)) + \nabla V(q(t)) \cdot (\lambda - Mu)\delta t] + o(\delta t),$$

which leads to the Bellman equation

$$c \cdot q = \max_{\lambda \in \mathcal{L}, u \in \mathcal{U}(q)} [\tilde{r}(\lambda) + \nabla V(q) \cdot (\lambda - Mu)] \quad \forall q \in \mathcal{Q} \text{ and } V(0) = 0, \quad (14)$$

where $\mathcal{U}(q) = \{u: 0 \leq u \leq 1, \sum_i u_i \leq 1, \lambda_j - \mu_j u_j \geq 0 \forall j \text{ s.t. } q_j = 0\}$. (The last set of constraints ensures that $q(t) \geq 0$.) Let λ^*, u^* denote the maximizers in (14).

PROPOSITION 3. Consider the problem of maximizing (13) subject to (5)–(6) and (9). Let $q^*(t)$ denote the queue-length trajectory under (λ^*, u^*) defined through (14). Then, (λ^*, u^*) is optimal in the sense that for any $q(0) = z \in \mathcal{Q}$ and any feasible (λ, u) under which $q(t) \rightarrow 0$ as $t \rightarrow \infty$,

$$V(z) = \int_0^\infty (\tilde{r}(\lambda^*(t)) - c \cdot q^*(t)) dt \geq \int_0^\infty (\tilde{r}(\lambda(t)) - c \cdot q(t)) dt.$$

PROOF. From (5), we have that $dq(t)/dt = \lambda(t) - Mu(t)$ a.e., and the definition of the Bellman equation gives that

$$0 \geq \tilde{r}(\lambda(t)) - c \cdot q(t) + \nabla V(q(t)) \cdot (\lambda(t) - Mu(t)) \quad \forall t \geq 0. \quad (15)$$

Integrating (15) over t , we get that

$$\int_0^t (\tilde{r}(\lambda(s)) - c \cdot q(s)) ds \leq - \int_0^t dV(q(s)) = V(z) - V(q(t)). \quad (16)$$

Letting $t \rightarrow \infty$ and using the facts that $q(t) \rightarrow 0$, $V(0) = 0$, and $V(\cdot)$ is continuous, we conclude that for all $z \geq 0$, $V(z) \geq \int_0^\infty (\tilde{r}(\lambda(t)) - c \cdot q(t)) dt$. To complete the proof we need to show that $V(z)$ is indeed the total profit under (λ^*, u^*) . Because (16) holds with equality under λ^*, u^* , it suffices to show that $q^*(t) \rightarrow 0$ as $t \rightarrow \infty$. We argue by contradiction. Suppose that $\limsup_t q^*(t) = q^\infty \neq 0$. Then, this implies that $\limsup_t V(q^*(t)) = V(q^\infty) < 0$, and therefore that $\int_0^\infty (\tilde{r}(\lambda^*(t)) - c \cdot q^*(t)) dt > V(z) > -\infty$. On the other hand, from Proposition 1, if $q^*(t) \not\rightarrow 0$ as $t \rightarrow \infty$, then $\int_0^\infty (\tilde{r}(\lambda(t)) - c \cdot q(t)) dt = -\infty$, which leads to a contradiction. Hence, $q^*(t) \rightarrow 0$ and consequently $V(q^*(t)) \rightarrow 0$ as $t \rightarrow \infty$, which completes the proof. \square

II. Characterization of Optimal Sequencing and Pricing Policies. Using the characterization of the optimal controls in terms of the Bellman equation given in (14), we will next show the following properties for the optimal policies: (a) sequencing decisions are made according to the $c\mu$ -rule, and (b) the demand vector is bounded above by $\hat{\lambda}$.

PROPOSITION 4. Fix any demand rate trajectory $\lambda(t)$ for $t \geq 0$. Then, the $c\mu$ -rule defined by

$$u^*(t) = \arg \max \left\{ \sum_i (c_i \mu_i) u_i: u \geq 0, \sum_i u_i \leq 1, \lambda_j(t) - \mu_j u_j \geq 0 \forall j \text{ s.t. } q_j(t) = 0 \right\}, \quad (17)$$

is pathwise optimal in that $c \cdot q^*(t) \leq c \cdot q(t)$ for all $t \geq 0$, where $q^*(t)$ and $q(t)$ denote the queue-length trajectories under $u^*(t)$ and any other feasible allocation $u(t)$, respectively.

SKETCH OF PROOF. The proof follows along the lines of Avram et al. (1995, §4.1) using the associated Hamiltonian formulation; the extension to an infinite horizon is done along the lines of Seierstad and Sydsaeter (1987, §§6.5–6.6). See also Chen and Yao (1993, §3). \square

A consequence of Proposition 4 is that $\nabla V(0)_i \mu_i$ is non-increasing in i . This is established as follows. For $z^i = e_i \cdot (w \mu_i)$, where e_i is the i th unit vector and $w > 0$ is small, let $\lambda^i(\cdot)$ be the optimal demand vector trajectory starting from that initial condition and $q^i(t)$ be the corresponding queue-length trajectory under λ^i and the $c\mu$ -rule. Then, for $j < i$,

$$V(z^j) \geq \int_0^\infty (\tilde{r}(\lambda^i(t)) - c \cdot q^j(t)) dt \geq \int_0^\infty (\tilde{r}(\lambda^i(t)) - c \cdot q^i(t)) dt = V(z^i),$$

where the first inequality follows because $\lambda^i(\cdot)$ need not be optimal starting from z^j , and the second can be established using the properties of the $c\mu$ -rule (e.g., using an induction argument on the number of classes). For small w , we also have that $V(z^i) = V(0) + \nabla V(0)_i (w \mu_i) + O(w^2)$, which together with the above inequality can be used to obtain the desired result.

The next result establishes that the optimal demand rate is upper bounded by $\lambda^*(0) = \hat{\lambda}$. The proof is given in the appendix.

PROPOSITION 5. The optimal demand rates $\lambda^*(q)$ defined via (14) satisfy $\lambda^*(q) \leq \hat{\lambda}$ for all $q \geq 0$.

A consequence of this last result is that the aggregate load into the system at any time t is $\rho^*(t) = m \cdot \lambda^*(t) \leq 1$. Using (7), this implies that the workload in the system is nonincreasing in time, and, in turn, that if $q(0) \in \mathcal{Q}$, then $q(t) \in \mathcal{Q}$ for all $t \geq 0$. As a result, the solution of the fluid optimal control problem enforces the exogenous constraint

imposed on the behavior of the system when it reaches the boundary of \mathcal{Q} , and moreover the optimal control does not depend on the size of the set \mathcal{Q} , as measured by the parameter N_w .

Both of the properties derived in the last two propositions are exploited in the next section to characterize the evolution of an optimally controlled system, and to propose an appropriately constructed one-dimensional relaxation to the fluid control problem. They are also helpful in numerically computing the optimal demand rates $\lambda^*(q)$.

III. Comments on Computing the Optimal Pricing Policy. Despite the simple structure of this fluid control problem, it is still not solvable in closed form, and one has to resort to numerical optimization techniques to compute the optimal pricing strategy. The simplest way to do so would be to discretize over time and solve a tractable concave maximization problem. Specifically, the objective comprises of the concave revenue term minus the linear holding cost for each time period. The fluid model dynamics are captured through a set of linear equality constraints of the form $q(t+1) = q(t) + \lambda(t) - Mu(t)$, where t is now a discrete index. Additional constraints for the allocation rule are that $u(t) \geq 0$ and $\sum_i u(t) \leq 1$, for the demand rates that $\lambda(t) \in \mathcal{L}$, and for the state that $q(t) \geq 0$. The complexity of this problem grows with the number of products and the number of time periods. Instances of modest size with tens of products and a few hundred time periods that result in a few thousand variables and constraints can be computationally tractable for commonly used demand models such as the linear, exponential, and isoelastic ones.

4. State-Space Collapse and a One-Dimensional Workload Relaxation of the Fluid Model Profit Maximization Problem

The structural properties of the fluid optimal sequencing and pricing policies have one important implication about the optimal system behavior that we show in Theorem 2: The optimal queue-length trajectory couples in finite time with a trajectory that holds all of its workload in the “cheapest” product class, and subsequently evolves as the optimal solution to an appropriately defined single-product problem. Motivated by this state-space collapse property, we subsequently propose and analyze a relaxation to the fluid pricing problem that is based on its one-dimensional workload rather than the I -dimensional queue-length vector.

4.1. Fluid-Scale State-Space Collapse

We first introduce some useful notation. For any workload position $w \geq 0$, we define

$$\begin{aligned} \Delta(w) &= \arg \min \{c \cdot q : q \geq 0, m \cdot q = w\} \\ &= [w\mu_1, 0, \dots, 0] \end{aligned} \quad (18)$$

to be the queue-length vector that holds workload w and has minimum cost. For the linear holding cost structure of our model, this corresponds to keeping all the workload into the “cheapest” and lowest-priority class, which by our labelling convention is Class 1; i.e., $c \cdot \Delta(w) = c_1\mu_1 w$.

THEOREM 2. Let $q^*(t)$ denote the optimal trajectory for the problem of maximizing (13) subject to (5)–(6). Then, for any $z \geq 0$, $q^*(t) = \Delta(m \cdot q^*(t))$ for all $t \geq T(z) := \mu_1(\sum_{i>1} m_i z_i) / \hat{\lambda}_1$.

PROOF. From Propositions 4 and 5, we know that $q^*(t)$ is the trajectory under the $c\mu$ -rule with $\lambda^*(t) \leq \hat{\lambda}$ for all t . Consider any $z \geq 0$ with initial workload $w(0) = m \cdot z$, and note that the lowest priority class is not served until all high-priority classes have been drained, i.e., $u_i^*(t) = 0$ for all $t \leq \tau$, where $\tau = \inf\{t \geq 0 : q_i(t) = 0 \forall i > 1\}$. Now, the high-priority class queue lengths $q_i^*(t)$ for $i = 2, \dots, I$ are upper bounded by the queue length of an $(I-1)$ -class queue with initial condition $\bar{q}(0) = [z_2, z_3, \dots, z_I]$ and $\lambda(t) = \hat{\lambda}$. The latter system has traffic intensity $1 - \hat{\lambda}_1/\mu_1$, and drains its initial workload in $T(z) = \mu_1(\sum_{i>1} m_i z_i) / \hat{\lambda}_1$ time under any nonidling sequencing rule (i.e., $\sum_i u_i(t) = 1$ if $\bar{q}(t) \neq 0$). Moreover, $\bar{q}(t) = 0$ for all $t \geq T(z)$. It follows that $\tau \leq T(z)$ and $q_i^*(t) = 0$ for all $i > 1$ and $t \geq T(z)$. This completes the proof. \square

The dynamics of the optimally controlled multiclass fluid queue for $t \geq \tau$ are given by

$$\dot{w}^*(t) = -1 + \rho^*(t) \quad \text{and} \quad q^*(t) = \Delta(w^*(t)) \quad (19)$$

with $\rho^*(t) = m \cdot \lambda^*(t)$. That is, the queue length evolves on the hyperplane defined through $q = \Delta(m \cdot q)$, or in other words, the state space has effectively “collapsed” to that of the one-dimensional workload process. Given Theorem 2 and the definition of $R(\rho)$, it is easy to see that for all $t \geq \tau$, the profit maximization problem posed through (13) reduces to that of maximizing

$$\int_{\tau}^{\infty} [\tilde{R}(\rho(t)) - c_1\mu_1 w(t)] dt,$$

subject to (19) and where $\tilde{R}(\rho) := R(\rho) - R(1)$ is analogous to the definition of $\tilde{r}(\cdot)$. The state of this problem is the one-dimensional workload $w(t)$, and the control is the one-dimensional aggregate capacity consumption rate $\rho(t)$. The aggregate control is mapped to product level demand rates (and thus prices) using the mapping λ^r defined through (1), which selects the demand rate vector that maximizes the instantaneous revenue rate subject to the constraint $m \cdot \lambda = \rho(t)$. The workload is mapped to a queue-length vector through the mapping $\Delta(w)$. Optimizing the total system profits for $t \geq \tau$ is a one-dimensional control problem, which is much simpler analytically and computationally than the multidimensional (queue-length) formulation that one started with.

4.2. The Workload Relaxation of the Fluid Pricing Problem

Workload formulations of network control problems have been used extensively in the context of heavy-traffic theory, where they emerge naturally as equivalent formulations to the approximating Brownian control problem formulations; see Harrison (1988, 2000) and Harrison and Van Mieghem (1996). Roughly speaking, Brownian approximations are derived in a way that involves a compression of time that makes the initial transient period of length $T(z)$ negligible, and where the queue length can be shown to always evolve on that minimum cost manifold, $q = \Delta(m \cdot q)$. Motivated by these results, Meyn (2001) recently proposed to use such workload relaxations to fluid model network control problems involving sequencing and routing control, implicitly disregarding this short initial transient period where the queue length is not on that efficient configuration. In this paper, we extend Meyn’s idea to incorporate the dynamic pricing element, and propose a workload relaxation to our profit maximization problem. The sequencing decisions are greedy (as shown in Proposition 4), and are suppressed in this model posed directly in terms of the aggregate workload process for the appropriate cost parameter. The latter is partially justified by Theorem 2.

Specifically, we propose the following relaxation to the fluid control problem posed through (13) and (5)–(6): Choose the load $\rho(t)$ for $t \geq 0$ to

$$\text{maximize } \int_0^\infty (\tilde{R}(\rho(t)) - \zeta w(t)) dt \tag{20}$$

subject to (7), which describes the workload dynamics, the control constraint

$$0 \leq \rho(t) \leq \rho_{\max}, \quad \text{where} \tag{21}$$

$$\rho_{\max} := \max\{\rho: \rho = m \cdot \lambda, \lambda \in \mathcal{L}\},$$

and $\zeta := \min_i c_i \mu_i$. This is equivalent to the queue-length formulation of (13) and (5)–(6) if the queue starts at a minimum cost configuration $z = \Delta(m \cdot z)$, because in that case $q^*(\cdot) = \Delta(m \cdot q^*(\cdot))$. On the contrary, if $z \neq \Delta(m \cdot z)$, (20) is only an approximation to (13) for some initial transient period of length no more than $T(z)$, where the queue length reaches a minimum cost state. Maximizing (20) subject to (7) and (21) is a single-product problem for the revenue function $R(\rho)$ and the holding cost rate ζ . It is much easier to analyze, often in closed form, and the numerical results of the next section will illustrate that it generates heuristics with very good performance. Moreover, its solution translates into easily implementable policies for the underlying multiproduct profit maximization problem by: (a) orders are sequenced according to the $c\mu$ -rule, and (b) products are priced to induce the demand rates $\lambda(t) = \lambda^r(\rho(t))$.

The remainder of this section studies the single-product workload relaxation. We use the notation \bar{V} for the associated value function and $\bar{\rho}(w)$ for its optimal control. Our first result specializes the findings of Propositions 2 and 5 to the single-product case.

THEOREM 3. Consider the problem of maximizing (20) subject to (7) and (21), and let $\bar{V}(w)$ denote the associated value function starting from an initial workload position w . Then: (i) $\bar{V}(\cdot)$ is concave and nonincreasing in w , and (ii) $\bar{\rho}(w)$ is nonincreasing in w .

PROOF. The concavity follows from Proposition 2. Consider two initial conditions $w^1 > w^2$, and denote by $\bar{w}^1(t)$, $\bar{w}^2(t)$ the associated optimal trajectories. Let $\tau = \inf\{t \geq 0: \bar{w}^1(t) = \bar{w}^2(0)\}$, which is well defined and finite because from Proposition 1, $\bar{w}^1(t) \rightarrow 0$ as $t \rightarrow \infty$. Then,

$$\bar{V}(w^1) = \int_0^\tau (\tilde{R}(\rho(t)) - \zeta w(t)) dt + \bar{V}(w^2).$$

Because $w^1 > w^2$, it follows that $\int_0^\tau \rho^*(s) ds < \tau$, which implies from the concavity of \tilde{R} that

$$\int_0^\tau \tilde{R}(\rho(s)) ds \leq \tau \tilde{R}\left(\frac{1}{\tau} \int_0^\tau \rho(s) ds\right) \leq 0,$$

and therefore that $\bar{V}(w^1) \leq \bar{V}(w^2)$. This proves property (i).

To prove part (ii), we note that from Proposition 2 we get that \bar{V} is Lipschitz continuous, and therefore a.e. differentiable. This implies that part (i) can be restated as $\bar{V}'(w) \leq 0$ and is nonincreasing in w . Adapting (14), we get that the Bellman equation for the workload control problem is

$$\zeta w = \max_{\rho, y} [\tilde{R}(\rho) + \bar{V}'(w)(\rho - y)], \tag{22}$$

where the maximization is over $\rho \in [0, \rho_{\max}]$ and $y \in [0, 1]$; $y(t)$ is interpreted as the total effort exerted by the server on all product classes, i.e., $y(t) = \sum_i u_i(t)$. As in Proposition 2, it is easy to show that $\bar{\rho}(0) = \hat{\rho} = m \cdot \lambda$, and that $y(w) = 1$ for all $w \geq 0$. Using the properties of \tilde{R} and the nonincreasing nature of $\bar{V}'(w)$, we readily conclude that $\bar{\rho}(w)$ is nonincreasing in w and complete the proof. \square

An interesting issue that arises is to compare the time it takes until state-space collapse is achieved, with the time it takes for the entire system to drain, which itself follows from Proposition 1. We have been unable to get a crisp characterization or bound for the relation between the two, however, a back-of-the-envelope analysis shows that the difference can be significant. Specifically, the rate at which the system is drifting towards the state-space collapse position is $1 - \rho_{-1}(t) \geq \hat{\lambda}_1 / \mu_1$, where $\rho_{-1}(t) = \rho(t) - \lambda'_1(\rho(t)) / \mu_1$ is the aggregate load due to products $i = 2, \dots, I$. Similarly, the rate at which the system is drained is $1 - \rho(t)$. The time taken to reach each of these two states is inversely proportional to these loads. Suppose, for example, that the average values along the optimal workload trajectory for $\rho(t)$ and $\rho_{-1}(t)$ are 0.95 and 0.75, respectively (i.e., Product 1 demand consumed 20% of the system’s processing capacity). Then, the time required to drain the system would be five times longer than the time needed to achieve state-space collapse. This simple example illustrates that the relative difference between these two

quantities is likely to be significant when it is optimal to price in a way that Product 1 consumes a significant portion of the system's processing capacity, and where the holding cost parameters are such that the average load over the optimal trajectory is high.

The monotonicity of $\bar{\rho}(w)$ implies that the product-level demands have a nested structure.

COROLLARY 2. *Let $\mathcal{F}(w) = \{i: \lambda_i^r(\bar{\rho}(w)) > 0\}$ be the set of products offered when the workload is $w \geq 0$. Then, $\mathcal{F}(w) \supseteq \mathcal{F}(w+x)$ for all $w, x \geq 0$.*

PROOF. This is a direct consequence of the fact that $\bar{\rho}(w)$ is nonincreasing in w and (by assumption) $\lambda_i^r(\rho)$ is nondecreasing in ρ . \square

The last result has important practical implications in terms of the optimal set of product offerings at increasing levels of congestion as measured by the aggregate system workload. Specifically, according to the solution to the workload fluid model pricing problem, the firm will offer products in a nested structure, by tactically removing the ones that become not profitable as congestion increases and the target aggregate capacity consumption $\bar{\rho}(w)$ decreases; once a product is removed, it is never reintroduced at higher levels of congestion.

Finally, $R'(\bar{\rho}(w))$ is interpreted as the marginal—or opportunity—cost of additional work arriving when the workload is equal to w . Often the simple structure of the workload formulation can be exploited to solve this pricing problem in closed form. Even if that is not possible depending on the form of the demand model, this one-dimensional problem is much simpler to tackle numerically using the discretization approach mentioned earlier in comparison to the multiproduct formulation.

4.3. Closed-Form Solution of Workload Formulation for the Linear Demand Model

The linear demand model is often used in practice due to its simple and intuitive structure, its tractability when embedded in mathematical optimization formulations of revenue management problems, and the fact that its parametric form is suitable for statistical estimation. This subsection provides a closed-form solution to the workload relaxation for the case of the linear demand model.

Recall from the description in §2 that the linear demand model is of the form $\lambda(p) = \Lambda - Bp$, and its revenue function is $r(\lambda) = \lambda \cdot B^{-1}(\Lambda - \lambda)$. Its associated aggregate revenue function $R(\rho)$ is defined through (1) and is shown in the appendix that it can be expressed as

$$R(\rho) = -\alpha_i \rho^2 + \beta_i \rho + \gamma_i \quad \text{for } \rho \in [r_{i-1}, r_i]$$

for $0 = r_0 \leq r_1 \leq r_2 \leq \dots \leq r_I$, and constants $(\alpha_i, \beta_i, \gamma_i)$ and r_i that depend on the model parameters Λ, B, μ , and are such that $R(\rho)$ is continuous, almost everywhere differentiable, and increasing for all $\rho \leq \rho^\dagger$. The value of r_{i-1}

is that of the smallest aggregate capacity consumption rate above which it is optimal to start offering the i most profitable products.

Recall that $\zeta := \min_i c_i \mu_i$. Starting with the Bellman equation in (22) and expanding $\bar{R}(\rho)$ into $R(\rho) - R(1)$, we get that for all $w \geq 0$,

$$\bar{\rho}(w) = \arg \max_{\rho} \{R(\rho) + \bar{V}'(w)\rho: 0 \leq \rho \leq \rho_{\max}\} \quad (23)$$

and

$$R(1) + \zeta w = R(\bar{\rho}(w)) + \bar{V}'(w)(\bar{\rho}(w) - 1). \quad (24)$$

We will use the first expression to express $\bar{V}'(w)$ in terms of $\bar{\rho}(w)$, and the second one to pointwise solve for $\bar{\rho}(w)$ for all $w \geq 0$. Specifically, the first-order optimality condition for $\bar{\rho}(w)$ is that $R'(\rho) = -\bar{V}'(w)$, which gives that at the optimum and for some $i \in [1, 2, \dots, I]$,

$$\bar{V}'(w) = 2\alpha_i \bar{\rho}(w) - \beta_i. \quad (25)$$

Using (24) and (25), we get that the optimal drift $\bar{\rho}(w)$ satisfies a quadratic equation of the form

$$\alpha_i \rho^2 - 2\alpha_i \rho + \beta_i + \gamma_i - R(1) - \zeta w = 0,$$

the solution of which is that

$$\rho(w) = 1 - \sqrt{\frac{\zeta w}{\alpha_i} + \delta_i},$$

$$\text{where } \delta_i = (R(1) + \alpha_i - \beta_i - \gamma_i)/\alpha_i.$$

The value of the index i above is implicitly determined such that the solution $\rho(w) \in [r_{i-1}, r_i]$. It is straightforward but tedious to show that there is a unique value of i for which this expression is consistent; i.e., if one computes $\rho(w)$ through the above expression, taking the value of i as given, then $\rho(w)$ is indeed in the interval $[r_{i-1}, r_i]$. This is captured in the following definition. First, for $i = 1, \dots, I-1$, define $w_i = \alpha_i \alpha_{i+1} (\delta_i - \delta_{i+1}) / ((\alpha_i - \alpha_{i+1}) \zeta)$ and $w_0 = \inf\{w \geq 0: \bar{\rho}(w) = 0\}$. Second, set

$$\bar{\rho}(w) = \left[1 - \sqrt{\frac{\zeta w}{\alpha_i} + \delta_i} \right]^+, \quad \text{whenever } w \in [w_i, w_{i-1}], \quad (26)$$

where this last expression has incorporated the implicit constraint that $\rho \geq 0$. Note that the w_i s are such that $\bar{\rho}(w)$ is continuous and decreasing in w , $0 = w_I \leq w_{I-1} \leq \dots \leq w_0$, and that $1 - \sqrt{\zeta w_{i-1} / \alpha_i} + \delta_i = r_i$. The optimality of this control is established using the verification result of Proposition 3 for the function

$$\bar{V}'(w) = \begin{cases} 2\alpha_i \bar{\rho}(w) - \beta_i, & w \in [w_i, w_{i-1}), \\ -R(1) - \zeta w, & w \geq w_0. \end{cases}$$

Finally, $\bar{\rho}(w)$ is disaggregated into product demands via the mapping λ^r as follows:

$$\lambda_i^r(\bar{\rho}(w)) = \frac{\Lambda_i}{2} - \kappa \frac{m_i b_{ii}}{2} \quad \text{for } i \leq \hat{i}(\bar{\rho}(w)) \quad \text{and} \\ \lambda_i^r(\bar{\rho}(w)) = 0 \quad \text{otherwise,}$$

where $\hat{i}(\bar{\rho}(w)) = \max\{i: \bar{\rho}(w) \geq r_{i-1}\}$ defines the set of products the firm will offer, and

$$\kappa = \left(\sum_{j \leq \hat{i}(\bar{\rho}(w))} m_j \Lambda_j - 2\bar{\rho}(w) \right) \left(\sum_{j \leq \hat{i}(\bar{\rho}(w))} m_j^2 b_{jj} \right)^{-1};$$

a derivation of these expressions is given at the end of the appendix. Finally, prices are inferred through the inverse demand function.

5. Discussion and Numerical Results

The first part of this section offers a short summary of the key insights gleaned from our analysis regarding the structure of the optimal pricing and sequencing policies. The second part reports on a set of numerical results that compare the pricing and sequencing heuristics extracted from the fluid model analysis to the solution of the underlying dynamic program, as well as the best static pricing policy.

5.1. Main Insights

The analysis of the two preceding sections leads to several insights briefly discussed below.

(1) *Invest in scarce capacity.* If one assumes that the demand model is known and it is stationary (as considered in this paper), then it is optimal in the fluid model to invest in production capacity that is scarce. That is, the demand vector that would maximize the instantaneous revenues in the absence of the capacity constraint would make the system unstable. This is intuitive because the revenue function is concave and the capacity cost is convex, making marginal revenue contribution close to the revenue-maximizing demand vector small compared to the marginal cost of the extra capacity needed to cope with that demand.

(2) *High resource utilization and congestion pricing.* Operationally, this choice of capacity vector induces the firm to operate its processing resources at close to full utilization, moderating excess backlogs through a (dynamic) increase in one or more prices. This behavior was established under the optimal fluid control policy, but was numerically observed to also hold under the optimal control of the MDP formulation associated with the original problem of §2.

(3) *Nested pricing policy.* As the system backlog grows large, under the fluid optimal pricing policy the firm increases prices in a way that effectively removes products from the market in a nested fashion, according to their marginal revenue contribution.

(4) *Sequencing and pricing decisions decouple.* These two elements of control are essentially decoupled in the fluid control formulation insofar as sequencing is done according to the greedy $c\mu$ -rule independent of the pricing decisions, and pricing eventually depends only on the system workload rather than the individual queue lengths, which are themselves insensitive to the choice of the sequencing rule.

(5) *Pricing as function of workload.* A reasonable relaxation to the original problem is to sequence jobs using the greedy policy and price according to the solution of a fluid control problem formulated in terms of the aggregate system workload, which is simpler to solve and leads to intuitively appealing policies. It is also practical to implement because it only requires modest information sharing between the pricing and production functions of an organization. The numerical results that follow illustrate the effectiveness of this heuristic when compared to the solution to the original MDP formulation, as well as to the optimal static pricing policies.

(6) *Workload relaxations of revenue management problems.* A generalization of the last few remarks suggests the use of workload relaxations for revenue management of make-to-order systems. While workload fluid models are essentially heuristically derived because fluid-scale state-space collapse results similar to that of Theorem 2 offer only partial support for their validity, they seem to capture some of the essential elements of the underlying pricing and operational control problems while maintaining a fair amount of tractability. (Further justification can be obtained through a diffusion analysis, by adopting the arguments of Harrison and Van Mieghem 1996 to systems with dynamic pricing (drift) control capability.) Finally, workload fluid model relaxations lead to practically implementable solutions with modest coordination requirements between the pricing and production functions based on the aggregated workload information.

5.2. Implementable Heuristics and Numerical Results

We conclude with a numerical study of the performance of the capacity choice, pricing, and sequencing policies that are extracted via the fluid model analysis. Conceptually, this can be separated into two issues: (a) How good are the heuristics that are derived from analysis of fluid model profit maximization formulations, and (b) what is the impact of the workload relaxation on the performance of these heuristics. This subsection is split into two parts, focusing on single-product and two-product problems, respectively, that effectively address these issues in sequence.

Implementation of the Fluid Heuristics. The preceding analysis has culminated in two heuristics that are based on the solution of the (multidimensional) fluid control profit maximization problem of §3 and its workload relaxation of §4, respectively. This section studies the performance of the

pricing rule extracted from the workload relaxation, which is summarized through the aggregate control $\bar{\rho}(w)$ and is mapped to a vector of demand rates through $\lambda^r(\bar{\rho}(w))$. Instead of selecting the price vector that maps into this vector of demands, we propose to introduce a small tunable parameter $\theta \in \mathbb{R}$ and price to induce the demand vector

$$\lambda(w) = \lambda^r(\bar{\rho}(w) + \theta); \tag{27}$$

i.e., we propose to use the state-dependent control as extracted from the workload fluid model formulation, but perturb its aggregate decision $\bar{\rho}(w)$ by a constant shift θ that tries to correct for some of the idealizations of the deterministic fluid model. This leaves the shape and structure of the proposed control unchanged, as computed through the fluid model analysis. Its use is fairly economical in the sense that it uses a single parameter to perturb the aggregated control decision, and use the mapping from ρ to λ s to disaggregate its effect to the multiple products in a way that exploits the structure of the proposed policy. Finally, the magnitude of this adjustment is “small” in comparison to the dynamic component $\bar{\rho}(w)$, and one would expect that in the asymptotic regime sketched in §3.1 it is of order $o(1)$ —i.e., $\theta^n \rightarrow 0$ as $n \rightarrow \infty$ in the context of (4), making the resulting shifts to the demand rates to be of order $o(n)$, while the demand rates themselves are of order n . Similar adjustments to static fluid prices have appeared in other papers, and were either heuristically proposed as in Gallego and van Ryzin (1994) for the perishable, single-product revenue maximization problem, or analytically derived as in Maglaras and Zeevi (2003) in the context of optimal static pricing for a single-product stochastic service system.

This tunable parameter is selected either numerically or through simulation. For example, for the single-product problem with Markovian assumptions, it is straightforward to compute the steady-state profit rate under the state-dependent demand policy $\lambda(Q)$ defined through (27) and numerically optimize over θ . A similar calculation is possible but tedious for the multiproduct case, where a simulation-based optimization approach may be more suitable. The latter can also be used if the service times follow general distributions.

Single-Product Problems. This first part studies the capacity investment decision under different policies, as well

as the relative performance of the fluid pricing heuristics against the solution of the “MDP” formulation associated with the problem of maximizing (3) described in §2. “FM(0)” refers to a direct implementation of the fluid control law $\lambda_f(Q)$, i.e., with $\theta = 0$, while “FM(θ^*)” is the fluid policy with the optimally tuned parameter θ^* . These three dynamic pricing policies will also be compared to “Static,” the optimal static pricing policy without admission control, and “Static + Adm.,” which refers to the optimal static price with admission control. The two static-price systems behave like $M/M/1$ and $M/M/1/K$ queues, respectively.

(1) *Performance with optimized capacity.* Table 1 compares the capacity choices and profit rates under these five candidate policies in a single-product model with a linear demand function and linear cost of capacity.³ We make a few observations:

(a) The fluid pricing heuristics outperform the static pricing policies by 0.5%–5%, while admission control adds up to 1% to the profit rate achieved under a static pricing policy. Such performance gains have significant impact to the firm’s profitability.

(b) The capacity under the FM(0) policy is computed through (10), and seems to systematically underestimate the ones under the MDP and FM(θ^*) policies that were both computed through a numerical search.

(c) The suboptimality gaps increased as the holding cost parameter c grew larger, and the optimal capacity level μ decreased as a function of its cost parameter h . Note that for large h , it may become unprofitable to operate the firm, i.e., the optimal capacity investment is $\mu = 0$. For example, according to (10), the firm would not invest in processing capacity if $r'(0) \leq h$.

(d) The expected traffic intensity in all these test cases ranged from 0.70 to 0.95 (see also Table 2).

(2) *Performance comparison with common capacity.* Table 2 compares the profit rate under these policies operating under a common value of processing capacity that was computed based on our fluid model analysis using (10). This isolates the performance effect of the pricing decision under each candidate policy. Our results show that the relative advantage of the fluid heuristics over the static pricing policies increases when considered under a common capacity choice. In addition, as the price sensitivity parameter gets large (in the lower half of this table), the performance gain due to dynamic pricing increases and the gaps between the fluid and static heuristics widen, and also

Table 1. Profit rate and capacity investment under a linear demand model: $\lambda(p) = 20 - 4 \cdot p$.

c, h	MDP		FM		FM(θ^*)		Static		Static + Adm.	
	$\mathbb{E}\pi$	μ	Gap (%)	μ	Gap (%)	μ	Gap (%)	μ	Gap (%)	μ
0.1, 0.5	19.1	9.7	1.6	9.0	0.04	9.7	0.8	10.1	0.8	10.1
0.5, 0.5	17.5	11.0	3.7	9.0	0.50	10.6	1.3	11.7	1.2	11.5
0.1, 1	14.6	8.2	3.8	8.0	0.10	8.2	2.6	8.6	2.4	8.6
0.5, 1	12.6	8.6	5.2	8.0	0.10	8.8	4.5	9.4	3.9	9.2

Table 2. Performance comparison with common capacity selected using (10).

Λ, b	c, h	MDP		FM(0)		FM(θ^*)		Static		Static + Adm.	
		$\mathbb{E}\pi$	ρ	Gap (%)	ρ	Gap (%)	ρ	Gap (%)	ρ	Gap (%)	ρ
20, 4	0.1, 0.5	19.0	0.93	1.2	0.87	0.04	0.93	1.5	0.90	1.4	0.90
	0.5, 0.5	17.3	0.86	1.6	0.79	0.20	0.85	3.4	0.80	2.8	0.83
	0.1, 1	14.6	0.96	3.8	0.86	0.10	0.95	3.1	0.90	2.8	0.92
	0.5, 1	12.5	0.88	4.7	0.78	0.30	0.86	6.3	0.81	4.8	0.86
20, 8	0.1, 0.5	7.0	0.93	4.2	0.83	0.20	0.92	4.2	0.87	3.6	0.90
	0.5, 0.5	5.5	0.83	5.0	0.73	0.50	0.82	9.2	0.75	6.3	0.81
	0.1, 1	3.5	0.96	19.7	0.80	5.80	0.88	13.7	0.88	9.6	0.96
	0.5, 1	1.9	0.87	33.1	0.69	11.80	0.77	37.5	0.75	21.9	0.84

the gap between the MDP policy and all other heuristics grows larger. This is due to the fact that as the price sensitivity parameter increases, the revenue rates decrease substantially (c.f., $r(\lambda) = \lambda(\Lambda - \lambda)/b$), indirectly making the holding cost term more significant. Indeed, close inspection of the optimal policy shows that in such cases the firm operates the system with very few jobs in the queue, where the nature of the static heuristics and the idealizations of the fluid approximations become more pronounced. In general, the accuracy of the fluid heuristics improved as the “size” of the system as measured by the potential demand Λ and processing capacity μ grew larger, which is consistent with the scaling given in (4). Intuitively, this says that pricing heuristics extracted from fluid approximations are expected to perform well in settings where the actual processing time of each order is much smaller than the actual time it takes for this order to go through the system.

(3) *General demand model.* Table 3 reports on a small set of results for a model with an exponential demand model. We note that in this case the fluid policy was computed numerically by solving the Bellman equation (14) at each queue-length position Q . For the exponential demand model, the fluid pricing policies outperformed the static ones by about 1%–3%, while similar optimality gaps were observed under the isoelastic demand model.

To recapitulate, the main insights extracted from our numerical results of the single-product model are the following:

(1) The dynamic heuristics FM(0) and FM(θ^*) outperform the static pricing policies by 0.5%–5%, which is significant. Admission control adds about 1% of profits to static pricing.

(2) The effect of dynamic pricing and the overall performance gaps increases as functions of the holding cost

and price sensitivity parameters. In both of these cases, the static policies price conservatively to control congestion costs when queues build up.

(3) The fluid model analysis sets capacity according to (10), which tends to underestimate the optimal choices under both MDP and FM(θ^*), albeit by relatively small margins. Its effect on the maximum achievable profit rate under the MDP policy ranged from 0.1%–1.4% in the experiments that we ran. In settings where capacity is difficult to change while demand models and competitive effects may vary substantially over time, it may be practical to adopt the fluid model solution as a way to set capacity, and use pricing to fine-tune the firm’s performance. Under all candidate policies it is optimal to invest in scarce capacity unless the firm is operating in an environment with very small capacity costs and very large congestion (holding) costs.

Multiproduct Problems. The last results look at a two-product system under a linear demand model. They focus on the performance of the policy extracted via the workload relaxation of §4, which sequences jobs according to the $c\mu$ -rule and prices as a function of the system workload. The examples tested below have two nonsubstitutable products that follow a linear demand relationship, that is, $b_{ij} = 0$ for all $i \neq j$. The fluid policy that we tested was the one specified through (27) extracted from the workload relaxation with an optimized parameter θ . For the linear demand model, the expression for $\rho_f(Q)$ becomes

$$\rho_f(Q) = \left[1 + \theta - \sqrt{\frac{\zeta w}{\alpha_i} + \delta_i} \right]^+,$$

where $w = m \cdot Q$, $\delta_i = (R(1) + \alpha_i - \beta_i - \gamma_i)/\alpha_i$, α, β, γ were as defined in the appendix in (32)–(34), $\zeta = \min(c_1\mu_1, c_2\mu_2)$, and $R(1) = \max\{r(\lambda): m \cdot \lambda = 1, 0 \leq \lambda \leq \Lambda\}$.

Table 3. Single-product, exponential demand $\lambda(p) = 20e^{-b \cdot p}$ with common capacity selected using (10).

b	c, h	MDP		FM(0)		FM(θ^*)		Static		Static + Adm.	
		$\mathbb{E}\pi$	ρ	Gap (%)	ρ	Gap (%)	ρ	Gap (%)	ρ	Gap (%)	ρ
1	0.1, 1	6.2	0.76	0.1	0.78	0.1	0.69	0.6	0.74	0.6	0.74
	0.5, 1	5.6	0.56	1.5	0.63	0.2	0.55	4.2	0.65	4.0	0.65
2	0.1, 1	2.6	0.76	0.4	0.76	0.4	0.66	1.9	0.72	1.8	0.72
	0.5, 1	2.0	0.55	2.0	0.62	1.9	0.55	4.2	0.55	3.6	0.55

Table 4. The two-product model with linear demand.

Λ_1, Λ_2	b_1, b_2	c_1, c_2	MDP			FM-work			Static + Adm.		
			$\mathbb{E}\pi$	$\mathbb{E}\lambda_1$	$\mathbb{E}\lambda_2$	Gap (%)	$\mathbb{E}\lambda_1$	$\mathbb{E}\lambda_2$	Gap (%)	$\mathbb{E}\lambda_1$	$\mathbb{E}\lambda_2$
16, 8	1, 1	0.2, 0.4	44.6	3.8	0.2	1.0	3.8	0.1	2.7	3.7	0.2
12, 8	1, 1	0.2, 0.4	30.4	3.0	0.9	0.6	2.9	0.9	2.3	2.9	0.9
8, 16			43.6	0.3	3.6	3.3	0.1	3.8	5.8	0.3	3.4
8, 8			20.9	2.0	1.9	0.6	1.9	1.9	3.4	1.9	1.9
8, 8	2, 1	0.2, 0.4	15.7	1.3	2.5	0.7	1.2	2.6	6.1	1.2	2.4
	1, 2		15.9	2.6	1.2	0.9	2.6	1.2	5.6	2.5	1.0
8, 8	1, 1	0.1, 0.4	21.5	2.0	1.9	0.2	2.1	1.9	1.7	2.0	1.9
		0.4, 0.8	19.6	2.0	1.8	0.6	1.9	1.9	7.1	1.9	1.6

Notes. Fluid policy was computed using the workload relaxation with an optimized θ parameter (cf. (27)). In all experiments, $\mu_1 = \mu_2 = 4$ and $h_1 = h_2 = 0.1$.

The test cases reported in Table 4 correspond to relatively small problem instances that kept the numerical solution of the two-dimensional dynamic program associated with the stochastic formulation of §2 tractable. The selected parameters tested a range of scenarios for the relative revenue and holding cost contribution of each product. The primary observation from the results of Table 4 is that the policy extracted via the workload relaxation had similar performance gaps to those observed for the single-product models for a wide range of parameters, while the optimality gaps of the static pricing policy with admission control

degraded in this multiproduct setting. This suggests that the restriction of the workload relaxation that forces prices to be functions of the workload had little impact on the performance of the fluid heuristics, and plausibly that the pricing decisions under the MDP policy may also mostly depend on the total workload and not the individual queue lengths. This is explored graphically in Figure 1.

Specifically, the plots in Figure 1 explore the form of the MDP demand policy as a function of the system workload, and to a large extent illustrate that the MDP controls indeed resemble those extracted via the fluid workload

Figure 1. MDP controls as functions of the system workload for a two-product model with a linear demand function with $B = [1, 0; 0, 1]$, $\mu_1 = \mu_2 = 4$, $c_1 = 0.2$, $c_2 = 0.4$, and $h_1 = h_2 = 0.1$.

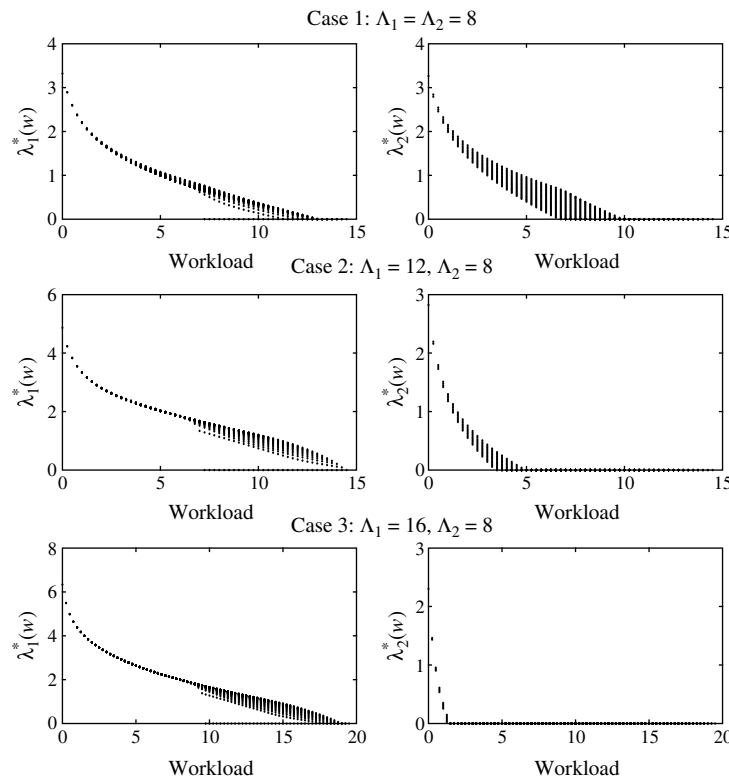



Figure 2. Steady-state probability distribution under MDP controls for a two-product model with a linear demand function with $B = [1, 0; 0, 1]$, $\mu_1 = \mu_2 = 4$, $c_1 = 0.2$, $c_2 = 0.4$, and $h_1 = h_2 = 0.1$.



relaxation of the previous section. The top two panels focus on a problem with a symmetric demand model for two products, while the lower panels are for problems where Product 1 contributes higher revenues than Product 2. In all cases, Product 2 incurred higher holding costs. Figure 1 plots the MDP controls as functions of the total workload. The multiple values for λ_i^* for each w reflect the fact that the MDP solution depends on the two-dimensional queue-length vector (Q_1, Q_2) , and thus the controls at different queue-length configurations that hold the same workload may differ. The relatively narrow spread of λ^* values at each w lends credibility to the proposed approximation of the MDP controls with state-dependent functions of the workload. Moreover, plots of the two-dimensional steady-state distributions shown in Figure 1 reveal that most of the probability is concentrated in states where (a) the system workload is held in the cheaper queue (here Product 1) or, stated differently, where the queue is in an “efficient” configuration; and (b) the aggregate workload is modest ($w \leq 4$ and $w \leq 4$, respectively), allowing the system manager to effectively modulate the traffic intensity into the system by mostly adjusting the demand of Product 2, which contributes less revenue. (We plotted the distributions for the first two parameter sets in Figure 2. The third set with $\Lambda_1 = 16$ and $\Lambda_2 = 8$ is similar, although in that case the system operates almost like a single-product system.)

These observations, of course, depend on the magnitude of the holding cost parameters and the variability of the service time distributions. We tested the former and observed similar behavior for a wide range of holding cost vectors, but did not check for the dependence of these results on the service time distribution, as this would render the underlying problem of §2 intractable. Partial evidence in support of using a pricing policy that depends on the workload rather than the queue-length vector even in the presence of general service times can be obtained using an analysis of a second-order refinement of our formulation based on

a diffusion control problem; see, e.g., Çelik and Maglaras (2005).

The above results demonstrate that the pricing and sequencing policies extracted from the fluid model workload formulation perform well in a variety of parameter settings. An interesting direction for future work would be to extend this analysis to multiproduct stochastic processing networks.

Appendix. Proofs

PROOF OF PROPOSITION 1. Part (i): We start by noting that the control specified in Step 1 of the proof of Theorem 1 can be employed here as well to establish that there exists a pair of feasible control policies λ and u that result in a finite objective function. (Details are omitted.) We will first show that $\int_0^\infty (\rho(s) - 1) ds \leq 0$, and then deduce the desired result.

Step 1. (a) Suppose that $\int_0^\infty (\rho(t) - 1) dt = +\infty$. From (7), it follows that $w(t) \rightarrow \infty$ and $q(t) \rightarrow \infty$ as $t \rightarrow \infty$. Because $\tilde{r}(\lambda) \leq \alpha := r(\lambda^\dagger) - r(\hat{\lambda})$, $\int_0^\infty (\tilde{r}(\lambda(t)) - c \cdot q(t)) dt \leq \int_0^\infty (\alpha - c \cdot q(t)) dt = -\infty$. This is suboptimal, and by contradiction, $\int_0^\infty (\rho(t) - 1) dt < \infty$.

(b) Similarly, if $\lim_{t \rightarrow \infty} \int_0^t (\rho(s) - 1) ds = N$ for some $N > 0$, then there exists $T > 0$ such that for $t \geq T$, $N/2 \leq \int_0^t (\rho(s) - 1) ds \leq 3N/2$, or equivalently, $1 + N/(2t) \leq (1/t) \int_0^t \rho(s) ds \leq 1 + 3N/(2t)$. Using the definition of $R(\cdot)$ and Jensen’s inequality, we get that for any $t \geq T$,

$$\begin{aligned} \int_0^t r(\lambda(s)) ds &\leq \int_0^t R(\rho(s)) ds \\ &\leq tR\left(\frac{1}{t} \int_0^t \rho(s) ds\right) \\ &\leq t\left(R(1) + R'(1)\frac{3N}{2t} + o(1/t)\right), \end{aligned}$$

where $R'(1) = dR(\rho)/d\rho|_{\rho=1} > 0$. Letting $t \rightarrow \infty$, we get that $\int_0^\infty \tilde{r}(\lambda(s)) ds \leq (3N/2)R'(1)$. In addition, from (7),

we get that $w(t) \geq N/2$ for all $t \geq T$, which, in turn, implies that $c \cdot q(t) \geq c_1 \mu_1 w(t) \geq c_1 \mu_1 N/2$ (follows from the product labelling of (12)). This implies that

$$\int_0^\infty (\tilde{r}(\lambda(t)) - c \cdot q(t)) dt \leq \frac{3N}{2} R'(1) - \int_T^\infty c \cdot q(t) dt = -\infty.$$

Arguing by contradiction, we conclude that $\int_0^\infty (\rho(s) - 1) ds \leq 0$, and as a side result that $\int_0^\infty \tilde{r}(\lambda(t)) dt \leq 0$.

Step 2. Suppose that $\limsup_{t \rightarrow \infty} q_i(t) > \epsilon > 0$ for some product i . Then, $\limsup_{t \rightarrow \infty} \int_0^t c \cdot q(s) ds = +\infty$. Because $\int_0^\infty \tilde{r}(\lambda(t)) dt \leq 0$, it follows that $\liminf_{t \rightarrow \infty} \int_0^t (\tilde{r}(\lambda(s)) - c \cdot q(s)) ds = -\infty$, which is again suboptimal. We conclude that it suffices to restrict attention to controls under which $q(t) \rightarrow 0$ as $t \rightarrow \infty$.

Part (ii) of the proof follows from Seierstad and Sydsæter (1987, Theorem 6.10). \square

PROOF OF PROPOSITION 2. Part (i): From Part 1(b) of Proposition 1, we have that starting from any $z \geq 0$ and under any candidate optimal control, $\int_0^\infty \tilde{r}(\lambda(t)) dt \leq 0$. Because $\int_0^\infty c \cdot q(t) dt \geq 0$, this implies that $\int_0^\infty (\tilde{r}(\lambda(t)) - c \cdot q(t)) dt \leq 0$ and $V(z) \leq 0$. For $z = 0$, the static control $(\hat{\lambda}, M^{-1}\hat{\lambda})$ is feasible, it keeps $q(t) = 0$ for all $t \geq 0$, and $\int_0^\infty (\tilde{r}(\lambda(t)) - c \cdot q(t)) dt = 0$.

Part (ii): Consider two initial conditions z^1, z^2 with $z^1 \neq z^2$, and let $\Delta := V(z^1) - V(z^2)$. We wish to bound $|V(z^1) - V(z^2)| = |\Delta|$. Suppose that $\Delta > 0$ and consider any feasible trajectory from z^2 to z^1 , and let τ_{21} be the first time that this trajectory reaches z^1 . Then, using the optimality of $V(z^2)$, we have that

$$V(z^2) \geq \int_0^{\tau_{21}} (\tilde{r}(\lambda(s)) - c \cdot q(s)) ds + V(z^1),$$

which implies that

$$\Delta \leq - \int_0^{\tau_{21}} (\tilde{r}(\lambda(s)) - c \cdot q(s)) ds \leq \int_0^{\tau_{21}} c \cdot q(s) ds \leq \tau_{21} K'$$

for $N' := \max_i c_i N_w$. To get a bound on τ_{21} , we will construct a control that takes the state from z^2 to z^1 as follows: (a) the system manager will produce $(z_i^2 - z_i^1)^+$ orders of class $i = 1, \dots, I$; and (b) price in a way such that $\lambda_i(t) = \hat{\lambda}_i$ for all $t \leq (z_i^1 - z_i^2)^+ / \hat{\lambda}_i$. Step (a) will be completed in no longer than $\sum_k (z_k^2 - z_k^1)^+ / \mu_k$ time units, while Step (b) will be completed in no longer than $\max_i (z_i^1 - z_i^2)^+ / \hat{\lambda}_i$. Combining the two gives the bound

$$\tau_{21} \leq N'' |z^2 - z^1|,$$

where $N'' = \max_i 1/\hat{\lambda}_i$ and for any vector x , $|x| = \sum_i |x_i|$. A similar bound can be constructed in the case where $\Delta < 0$ to show that $|\Delta| < (N'N'')|z^2 - z^1|$ for all $z^1, z^2 \in \mathcal{Q}$. That is, V is Lipschitz continuous with constant $(N'N'')$.

Part (iii): Suppose that (λ^i, u^i) are the optimal controls starting from initial conditions z^i for $i = 1, 2$, and $q^i(t)$ are the corresponding queue-length trajectories. Let $z = \kappa z^1 + (1 - \kappa)z^2$ and define $\lambda(t) = \kappa \lambda^1(t) + (1 - \kappa)\lambda^2(t)$

and $u(t) = \kappa u^1(t) + (1 - \kappa)u^2(t)$. Then, it is easy to check that because \mathcal{L} is convex, $\lambda(t) \in \mathcal{L}$, and that $0 \leq u(t) \leq 1$ and $\sum_i u_i(t) \leq 1$. Moreover, under λ, u ,

$$\begin{aligned} q(t) &= z + \int_0^t \lambda(s) ds - M \int_0^t u(s) ds \\ &= \kappa q^1(t) + (1 - \kappa)q^2(t) \geq 0 \quad \forall t \geq 0. \end{aligned}$$

Hence, the policy (λ, u) is feasible. The concavity of V is established by noting that

$$\begin{aligned} V(z) &\geq \int_0^\infty (\tilde{r}(\kappa \lambda^1(t) + (1 - \kappa)\lambda^2(t)) \\ &\quad - c \cdot (\kappa q^1(t) + (1 - \kappa)q^2(t))) dt \\ &\geq \kappa \int_0^\infty (\tilde{r}(\lambda^1(t)) - c \cdot q^1(t)) dt \\ &\quad + (1 - \kappa) \int_0^\infty (\tilde{r}(\lambda^2(t)) - c \cdot q^2(t)) dt \\ &= \kappa V(z^1) + (1 - \kappa)V(z^2), \end{aligned}$$

where the first inequality follows from V 's optimality and the second from the concavity of \tilde{r} .

Part (iv): Pick any $\epsilon > 0$ and let $\gamma(\epsilon) = \inf\{|z|: z \geq 0, \nabla V(z)_k \geq \nabla V(0)_k + \epsilon \text{ for some } k\}$, and denote by z^ϵ a limiting vector along any subsequence that achieves the infimum. If $z^\epsilon = 0$, the property we wish to prove holds automatically.

(a) Suppose that $z_k^\epsilon > 0$. Pick $\delta > 0$ small and note that the concavity of V implies that

$$V(z^\epsilon) \leq V(z^\epsilon - \delta e_k) + \delta \nabla V(z^\epsilon - \delta e_k)_k \quad \text{and}$$

$$V(z^\epsilon - \delta e_k) \leq V(z^\epsilon) - \delta \nabla V(z^\epsilon)_k,$$

where e_k is the k th unit vector. Adding these expressions and dividing by δ gives that $\nabla V(z^\epsilon)_k \leq \nabla V(z^\epsilon - \delta e_k)_k$. From the definition of z^ϵ , it follows that $\nabla V(z^\epsilon)_k \geq \nabla V(0)_k + \epsilon$ and $\nabla V(z^\epsilon - \delta e_k)_k < \nabla V(0)_k + \epsilon$, which leads to a contradiction. Therefore, $z_k^\epsilon = 0$.

(b) Suppose that $z_k^\epsilon = 0$ and $z_j^\epsilon \neq 0$ for some $j \neq k$. A first-order Taylor expansion gives that

$$V(z^\epsilon + \delta e_k) = V(z^\epsilon) + \delta \nabla V(z^\epsilon)_k + o(\delta),$$

where we say that $f(x)$ is $o(x)$ if $\lim_{x \rightarrow 0} f(x)/x = 0$. From the concavity of V , we also get that

$$\begin{aligned} V(z^\epsilon + \delta e_k) &\leq V(z^\epsilon - \delta e_j) + \delta \nabla V(z^\epsilon - \delta e_j)_j \\ &\quad + \delta \nabla V(z^\epsilon - \delta e_j)_k. \end{aligned}$$

Combining the last two expressions, we get that

$$\begin{aligned} 0 &\leq [(V(z^\epsilon - \delta e_j) + \delta \nabla V(z^\epsilon - \delta e_j)_j) - V(z^\epsilon)] \\ &\quad + \delta [\nabla V(z^\epsilon - \delta e_j)_k - \nabla V(z^\epsilon)_k] + o(\delta), \end{aligned}$$

where the first term is $o(\delta)$. Divide by δ and let $\delta \downarrow 0$ to get that $\nabla V(z^\epsilon)_k \leq \nabla V(0)_k + \epsilon$, which again leads to a contradiction. Letting $\epsilon \downarrow 0$ gives that $\nabla V(z) \leq \nabla V(0)$ for all $z \geq 0$, which completes the proof of part (iv). \square

PROOF OF PROPOSITION 5. The proof is by induction on $k(q) = \max\{i: q_i > 0\}$. If $k(q) = 0$, $q = 0$ and $\lambda^*(0) = \hat{\lambda}$ (from Propositions 1 and 2). Assume that the property holds if $k(q) = i$ and consider the case $k(q) = i + 1$. In the remainder of this proof, $u(\lambda, q)$ will denote the solution to (17) as a function of the variable λ and the queue-length vector q . The proof of the induction step uses the following result.

LEMMA 1. For any $y \in \mathbb{R}_+^L$, let $\lambda(q; y) = \arg \max\{\tilde{r}(\lambda) + y \cdot (\lambda - Mu(\lambda, q)): \lambda \in \mathcal{L}\}$, and recall that $\nabla V(q) \leq \nabla V(0)$ for all $q \geq 0$. Then, $\lambda(q; \nabla V(q)) \leq \lambda(q; \nabla V(0))$.

From Lemma 1, it suffices to show that $\lambda(q; y) \leq \hat{\lambda}$; where we use the shorthand notation $y = \nabla V(0)$. Using the definition of the $c\mu$ -rule, we first obtain an expression for the allocation control $u(\lambda, q)$:

$$u_j(\lambda, q) = \begin{cases} \frac{\lambda_j}{\mu_j} \wedge \left(1 - \sum_{i>j} \frac{\lambda_i}{\mu_i}\right)^+, & j > k(q), \\ \left(1 - \sum_{i>j} \frac{\lambda_i}{\mu_i}\right)^+, & j = k(q), \\ 0, & j < k(q), \end{cases} \quad (28)$$

which depends on q through $k(q) = i + 1$, the index of the highest-priority nonempty class. Letting $f^{i+1}(\lambda) := \tilde{r}(\lambda) + y \cdot (\lambda - Mu(\lambda, q))$ and $\lambda^{i+1} := \arg \max\{f^{i+1}(\lambda): \lambda \in \mathcal{L}\}$, the induction step reduces to showing that $\lambda^{i+1} \leq \hat{\lambda}$. Using (28), we get that $f^{i+1}(\lambda) = f^i(\lambda) + g^{i+1}(\lambda)$, where

$$\begin{aligned} f^i(\lambda) &= \tilde{r}(\lambda) + y \cdot \lambda - \sum_{j \geq i+1} y_j \left[\mu_j \left(1 - \sum_{l>j} \lambda_l / \mu_l\right)^+ \wedge \lambda_j \right] \\ &\quad - y_i \mu_i \left(1 - \sum_{l>i} \lambda_l / \mu_l\right)^+, \\ g^{i+1}(\lambda) &= y_i \mu_i \left(1 - \sum_{l>i} \lambda_l / \mu_l\right)^+ - y_{i+1} \mu_{i+1} \left(1 - \sum_{l>i+1} \lambda_l / \mu_l\right)^+ \\ &\quad + y_{i+1} \left[\mu_{i+1} \left(1 - \sum_{l>i+1} \lambda_l / \mu_l\right)^+ \wedge \lambda_{i+1} \right]. \end{aligned}$$

These expressions will allow us to make use of the induction hypothesis, i.e., that $\lambda^i = \arg \max\{f^i(\lambda): \lambda \in \mathcal{L}\} \leq \hat{\lambda}$. Note that $f^{i+1}(\lambda)$ is the sum of concave functions, and thus it is concave itself. Simple algebraic manipulations give that $g^{i+1}(\lambda) = 0$ if $1 - \sum_{l \geq i+1} \lambda_l / \mu_l \leq 0$, and $g^{i+1}(\lambda) = (y_i \mu_i - y_{i+1} \mu_{i+1})(1 - \sum_{l \geq i+1} \lambda_l / \mu_l) \geq 0$, otherwise. (The last assertion used the fact that $\nabla V(0)_i \mu_i$ is decreasing in i ; cf. the comment after Proposition 4.) It follows that first,

$$\begin{aligned} &\max\{f^{i+1}(\lambda): \lambda \in \mathcal{L}\} \\ &= \max\left\{f^{i+1}(\lambda): \lambda \in \mathcal{L}, 1 - \sum_{l \geq i+1} \lambda_l / \mu_l \geq 0\right\}, \end{aligned}$$

and second, from the properties of λ^i and the functional form of $g^{i+1}(\lambda)$, that $\partial f^{i+1}(\lambda) / \partial \lambda_j|_{\lambda^i} \leq 0$ for all j . This establishes the induction hypothesis $\lambda^{i+1} \leq \lambda^i \leq \hat{\lambda}$, and completes the proof. \square

SKETCH OF PROOF OF LEMMA 1. Similarly to the proof of Proposition 5, this lemma can be proved by induction on $k(q) = \max\{i: q_i > 0\}$. Expression (28) still gives $u(\lambda; q)$, which only depends on q through the index $k(q)$. This implies that $\lambda(q; y)$ is also a function of the queue length through $k(q)$. Denoting $\lambda(q; y)$ by $\lambda^i(y)$, when $k(q) = i$, the induction step that one would wish to show is that $\lambda^i(y) \geq \lambda^i(y - x)$ for any $x \in \mathbb{R}_+^L$. The arguments used above can be adapted in this setting to establish the desired result. Details are omitted. \square

DERIVATION OF AGGREGATE REVENUE FUNCTION ASSOCIATED WITH THE LINEAR DEMAND MODEL. The linear demand model is given by

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

where Λ_i is the market potential for product i and b_{ii} , b_{ij} are the price and cross-price sensitivity parameters. This is expressed in vector form as $\lambda(p) = \Lambda - Bp$ for the obvious choice of Λ , B . Under the assumptions listed in §2, the revenue function $r(\lambda) = \lambda \cdot B^{-1}(\Lambda - \lambda)$, which is a concave quadratic. The aggregate revenue function is defined through (1) (reproduced here):

$$R(\rho) = \max\{r(\lambda): m \cdot \lambda = \rho, \lambda \geq 0\},$$

which can be written as a concave, piecewise quadratic function in ρ of the form

$$R(\rho) = -\alpha_i \rho^2 + \beta_i \rho + \gamma_i \quad \text{for } \rho \in [r_{i-1}, r_i].$$

The derivation given below demonstrates how to compute the constants $(\alpha, \beta, \gamma, r)$ given the model parameters Λ , B , μ . This is done for the special case where there are no product substitution and/or complementarity effects; i.e., $b_{ij} = 0$ for all $i \neq j$, and $B = \text{diag}(b_{11}, \dots, b_{II})$. For convenience, we assume that products are labelled such that $\Lambda_1 / m_1 b_{11} \geq \Lambda_2 / m_2 b_{22} \geq \dots \geq \Lambda_I / m_I b_{II}$. Then, $r_0 = 0$ and the remaining constants α_i , β_i , γ_i , r_i are defined recursively as follows. First,

$$r_1 = \min\left\{\rho \geq 0: \rho = m_1 l_1, \frac{\Lambda_1 - 2l_1}{m_1 b_{11}} = \frac{\Lambda_2}{m_2 b_{22}}\right\}, \quad (29)$$

$$\begin{aligned} r_2 &= \min\left\{\rho \geq 0: \rho = m_1 l_1 + m_2 l_2, \right. \\ &\quad \left. \frac{\Lambda_1 - 2l_1}{m_1 b_{11}} = \frac{\Lambda_2 - 2l_2}{m_2 b_{22}} = \frac{\Lambda_3}{m_3 b_{33}}\right\}, \quad (30) \end{aligned}$$

and so on. Second, the product-level demand rates that correspond to some ρ are given by

$$\begin{aligned} \lambda_i^r(\rho) &= \frac{\Lambda_i}{2} - \kappa \frac{m_i b_{ii}}{2} \quad \text{for } i \leq \hat{i}(\rho) \quad \text{and} \\ \lambda_i^r(\rho) &= 0 \quad \text{otherwise,} \end{aligned}$$

where

$$\kappa = \left(\sum_{j \leq \hat{i}(\rho)} m_j \Lambda_j - 2\rho \right) \left(\sum_{j \leq \hat{i}(\rho)} m_j^2 b_{jj} \right)^{-1}$$

$$\text{for } \hat{i}(\rho) = \max\{i: \rho \geq r_{i-1}\}; \quad (31)$$

the last expressions used our labelling convention. Note that κ is decreasing and $\hat{i}(\rho)$ is increasing in ρ , respectively. Intuitively, the firm starts by offering the most “profitable” product when ρ is very small, and then sequentially introduces more products as the target consumption rate increases. Finally, $R(\rho) = \sum_{j \leq \hat{i}(\rho)} \lambda_j^*(\rho) (\Lambda_j - \lambda_j^*(\rho)) / b_{jj}$, which gives that

$$\alpha_{\hat{i}(\rho)} = \left(\sum_{j \leq \hat{i}(\rho)} m_j^2 b_{jj} \right)^{-1}, \quad (32)$$

$$\beta_{\hat{i}(\rho)} = \left(\sum_{j \leq \hat{i}(\rho)} m_j \Lambda_j \right) \left(\sum_{j \leq \hat{i}(\rho)} m_j^2 b_{jj} \right)^{-1}, \quad (33)$$

$$\gamma_{\hat{i}(\rho)} = \left(\sum_{j \leq \hat{i}(\rho)} \frac{\Lambda_j^2}{4b_{jj}} \right) - \left(\sum_{j \leq \hat{i}(\rho)} \frac{m_j \Lambda_j}{2} \right)^2 \left(\sum_{j \leq \hat{i}(\rho)} m_j^2 b_{jj} \right)^{-1}. \quad (34)$$

For example, when $\rho < r_1$, we have that $\hat{i}(\rho) = 1$, i.e., only Product 1 is offered, and therefore we should get that $R(\rho) = (\rho/m_1)(\Lambda_1 - (\rho/m_1))/b_{11}$, which agrees with the constants $\alpha_1 = 1/(m_1^2 b_{11})$, $\beta_1 = \Lambda_1/(m_1 b_{11})$, and $\gamma_1 = 0$ given by the above expressions. A similar argument can be applied when the cross-price sensitivity parameters are nonzero. \square

Endnotes

1. These capture variable production and work-in-process inventory costs.
2. A detailed derivation of these equations under the assumptions that $\mu^n/n \rightarrow \mu$ and the state-dependent demand rate satisfies $\lambda^n(n)/n \rightarrow \lambda(\cdot)$ as $n \rightarrow \infty$ can be found in Mandelbaum and Pats (1995), which focused on performance analysis of queues with state-dependent parameters in the absence of any economic considerations.
3. The model parameters are not selected to match any particular business application, but rather to be representative of the many test cases that we tried in terms of their suboptimality gaps, traffic intensities, and relative difference of their optimal capacities.

Acknowledgments

The author is grateful to Jan Van Mieghem, the anonymous associate editor, and three referees for their helpful comments during the review of this paper.

References

Afeche, P. 2004. Incentive-compatible revenue management in queueing systems: Optimal strategic idleness and other delaying tactics. Working paper, Kellogg School of Management, Northwestern University, Evanston, IL.

Avram, F., D. Bertsimas, M. Ricard. 1995. Fluid models of sequencing problems in open queueing networks; an optimal control approach. F. Kelly, R. Williams, eds. *Stochastic Networks*, Vol. 71. *Proceedings of the IMA*. Springer-Verlag, New York, 199–234.

Bertsekas, D. P. 1995. *Dynamic Programming and Optimal Control*, Vols. 1 and 2. Athena Scientific, Belmont, MA.

Biller, S., L. M. A. Chan, D. Simchi-Levi, J. Swann. 2002. Dynamic pricing and the direct-to-consumer model in the automotive industry. *E-Commerce J.* (Special edition on dynamic pricing) 5(2) 309–334.

Bitran, G., R. Caldentey. 2003. An overview of pricing models for revenue management. *Manufacturing Service Oper. Management* 5(3) 203–229.

Brémaud, P. 1980. *Point Processes and Queues: Martingale Dynamics*. Springer-Verlag, New York and Heidelberg, Berlin, Germany.

Çelik, S., C. Maglaras. 2005. Dynamic pricing and leadtime quotation for a multi-class make-to-order queue. Working paper, Columbia University Business School, New York.

Chen, H., M. Z. Frank. 2000. State dependent pricing with a queue. *IIE Trans.* 32.

Chen, H., D. Yao. 1993. Dynamic scheduling of a multiclass fluid network. *Oper. Res.* 41(6) 1104–1115.

Chen, X., D. Simchi-Levi. 2004a. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The finite horizon case. *Oper. Res.* 52(6) 887–896.

Chen, X., D. Simchi-Levi. 2004b. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The infinite horizon case. *Math. Oper. Res.* 29(3) 698–723.

Elmaghraby, W., P. Keskinocak. 2003. Dynamic pricing: Research overview, current practices and future directions. *Management Sci.* 49(10) 1287–1309.

Federgruen, A., A. Heching. 1999. Combined pricing and inventory control under uncertainty. *Oper. Res.* 47(3) 454–475.

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. 1988. Brownian models of queueing networks with heterogeneous customer populations. W. Fleming, P. L. Lions, eds. *Stochastic Differential Systems, Stochastic Control Theory and Applications*, Vol. 10. *Proceedings of the IMA*, Springer-Verlag, New York, 147–186.

Harrison, J. M. 1995. Balanced fluid models of multiclass queueing networks: A heavy traffic conjecture. F. Kelly, R. Williams, eds. *Stochastic Networks*, Vol. 71. *Proceedings of the IMA*, Springer-Verlag, New York, 1–20.

Harrison, J. M. 2000. Brownian models of open processing networks: Canonical representation of workload. *Ann. Appl. Probab.* 10 75–103.

Harrison, J. M., J. A. Van Mieghem. 1996. Dynamic control of Brownian networks: State space collapse and equivalent workload formulations. *Ann. Appl. Probab.* 7 747–771.

Horn, R. A., C. R. Johnson. 1994. *Matrix Analysis*. Cambridge University Press, Cambridge, UK.

Kachani, S., G. Perakis. 2002. A fluid dynamics model of dynamic pricing and inventory control for make-to-stock manufacturing systems. Working paper, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA.

Kalish, S. 1983. Monopoly pricing and dynamic demand and production cost. *Marketing Sci.* 2 135–159.

Kleywegt, A. J. 2001. An optimal control problem of dynamic pricing. Working paper, Georgia Institute of Technology, Atlanta, GA.

Lariviere, M., J. A. Van Mieghem. 2004. Strategically seeking service: How competition can generate Poisson arrivals. *Manufacturing Service Oper. Management* 6(1) 23–40.

Low, D. W. 1974. Optimal dynamic pricing policies for an $m/m/s$ queue. *Oper. Res.* 22 545–561.

- 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., J. Meissner. 2006. Dynamic pricing strategies for multi-product revenue management problems. *Manufacturing Service Oper. Management* **8**(2) 136–148.
- Maglaras, C., A. Zeevi. 2003. Pricing and capacity sizing for systems with shared resources: Approximate solutions and scaling relations. *Management Sci.* **49**(8) 1018–1038.
- Mandelbaum, A., G. Pats. 1995. State-dependent queues: Approximations and applications. F. Kelly, R. Williams, eds. *Stochastic Networks*, Vol. 71. *Proceedings of the IMA*, Springer Verlag, New York, 239–282.
- McGill, J., G. van Ryzin. 1999. Revenue management: Research overview and prospects. *Transportation Sci.* **33**(2) 233–256.
- Mendelson, H. 1985. Pricing computer services: Queueing effects. *Comm. ACM* **28**(3) 312–321.
- Mendelson, H., S. Whang. 1990. Optimal incentive-compatible priority pricing for the $M/M/1$ queue. *Oper. Res.* **38**(5) 870–883.
- Meyn, S. P. 2001. Sequencing and routing in multiclass queueing networks: Part I: Feedback regulation. *SIAM J. Control Optim.* **40**(3) 741–776.
- Seierstad, A., K. Sydsaeter. 1987. *Optimal Control Theory with Economic Applications*. Elsevier, Amsterdam, The Netherlands.
- Talluri, K., G. van Ryzin. 2004. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers, Dordrecht, The Netherlands.
- Van Mieghem, J. 2000. Price and service discrimination in queueing systems: Incentive compatibility of $Gc\mu$ policy. *Management Sci.* **46**(9) 1249–1267.