STOCK RATIONING IN A MULTI-CLASS MAKE-TO-STOCK QUEUE WITH INFORMATION ON THE PRODUCTION STATUS

August 2006

Jean-Philippe Gayon*
Francis de Véricourt†
Fikri Karaesmen‡

* INPG, Grenoble, FRANCE
gayon@gilco.inpg.fr

† Fuqua School of Business
Duke University
Durham, NC, USA
fdv1@mail.duke.edu

‡ Department of Industrial Engineering
Koç University
34450, Sariyer, Istanbul, TURKEY
fkaraesmen@ku.edu.tr

Stock Rationing in a Multi-class Make-to-Stock Queue with Information on the Production Status

August 17, 2006

Abstract

We consider a single-item make-to-stock production system. The item is demanded by several classes of customers arriving according to Poisson processes with different backorder costs. Item processing times have an Erlang distribution. This allows us model the information on the production status in a tractable way and to investigate the impact of processing time variability. Using a Markov Decision Process formulation, we show some structural properties of the optimal stock allocation policy. We also present a heuristic and discuss how to optimize its parameters. We show by numerical experiments the effectiveness of this heuristic policy. In addition, we study the influence of processing time variability on the performance of the system.

1 Introduction

A stock and capacity allocation problem occurs when a common stock and the production capacity of a supplier must be shared among different markets/customers. Such problems remain at the heart of many supply chain management issues. For instance, delayed product differentiation often results in maintaining a stock of generic components for multiple end-products (de Véricourt, 2002). The design of supply contracts in presence of different retailers can also entail a stock allocation problem at the supplier (Cachon and Lariviere 1999). More recently, Desphande et al. (2003) provide an example of inventory rationing for the U.S. military.

Stock and capacity allocation problems are however very challenging and generally considered intractable as explained by Tsay et al. (1999), especially when customer demands can be backordered. Even when optimal allocation strategies can be characterized, they are usually hard to implement. Indeed, the supplier needs to take many dimensions into account (the inventory level, the number of waiting demands in the system, but also the current status of the production process, etc.) when deciding to allocate stock to some customers while backordering demands from others. The complexity of such problems greatly depends on the number of customers sharing the common stock (Ha 1997b), and on the nature of the production cycle time (Ha 2000).

In this paper, we consider a supplier that produces a standard item in a make-to-stock environment for several classes of customers. Demands for each class are Poisson processes and item processing times have an Erlang distribution. The supplier has a finite production capacity and has some information on the status of the current production. The customer classes have different values and generate different backorder penalties for the supplier. The objective is to minimize the expected discounted holding and backorder costs over an infinite horizon. At each time instant the optimal decision depends on the inventory level, the number of waiting demands of each class and the current production stage.

Under the above assumptions, we provide a partial characterization of the optimal stock allocation and production policy. We also derive a heuristic that is very efficient, with an error typically less than 1%. This heuristic is easy to compute and to implement. It is based on a related problem where the manager can serve at any time an ample market with zero backorder cost. We also analyze the impact of the production time variability on the system performance which is numerically shown, maybe surprisingly, to be quasi-linear for both problems.

Stock and capacity allocation problems were first introduced in the context of inventory control. Topkis (1968) provides one of the earliest formulations of an optimal stock rationing problem for an uncapacitated system in discrete time. He analyzes a system with two classes of customers and shortage costs. Nahmias and Demmy (1981) also consider a rationing problem in an uncapacitated setting. They analyze the cost improvement under (r, Q) policies with rationing. Frank et al. (2004) propose effective heuristics for a system with two customer classes where the demands of the first class must be fully satisfied while demands of the second class can be partially satisfied. Melchiors et al. (2000) propose a performance evaluation method for critical level policies for continuous review systems under (r, Q)-type policies. Melchiors (2003) proposes an alternative rationing policy and assesses the performance of this policy for a similar system. Deshpande et. al (2003) optimize the parameters of (r, Q) policies with rationing, and analyze the benefit of applying this policy for a military logistics system.

These previous works assume uncapacitated replenishment systems (with exogenous replenishment lead times). For limited production capacity, on the other hand, queuing-based models provide a powerful framework which allows modeling explicitly the production capacity and the randomness of the supply process (see Buzacott and Shanthikumar, 1993). We follow this approach and model our system as a single server, single-product, make-to-stock queue with multiple demands as introduced by Ha (1997a, 1997b, 2000) in the stock rationing context.

Rationing strategies also appear in inventory transshipment problems, which has attract a lot of attention from researchers and practitioners recently. Zhao et al. (2004) characterize the structure of the optimal stock allocation and production strategies for a problem with two make-to-stock queues each serving a class of customers, and where inventory transshipment is allowed. Hu et al. (2004) study a similar problem where production capacity is uncertain. They also identify and explain counter intuitive behaviors that can appear in this context.

Ha (1997a) characterizes the optimal rationing and production policy of a multi-class $M/M/1$ make-to-stock queue with lost sales. He shows that there are thresholds for each customer class such that it is optimal to reject an arriving demand from a customer if the on-hand inventory is below the threshold for that customer (and to satisfy the demand with the stock otherwise). Carr and Duenyas (2000) analyze the structure of the optimal admission/sequencing policy for a related problem where demands from one class can be rejected. Lee and Hong (2003) numerically study the performance of a lost-sales system with Coxian processing times operating under critical level rationing policies.

When backorders are allowed, the problem of characterizing the optimal policy becomes significantly more difficult because the number of waiting demands has to be tracked for each customer class. For the backorder case, Ha (1997b) shows that the optimal stock and capacity allocation for two customer classes has a monotone structure. de Véricourt et al. (2002) generalize this result and provide a full characterization of the optimal stock and capacity allocation for n customer classes. The optimal policy specifies threshold levels such that it is optimal to satisfy an arriving demand from a customer if the on-hand inventory is above the threshold for that customer and to backorder the demand otherwise. These threshold levels also determine production priority for waiting demands in a simple way.

The models in Ha (1997a,b) and de Véricourt et al. (2002) assume exponential processing times. Because of the memoryless property of the exponential distribution, the supplier does not need to make decisions based on the current status of the production process. Information Technologies in real production systems can provide however constant access to information on the status of the production process. The manager can then exploit this knowledge and make more accurate inventory allocation decisions. We consider in this paper a multi-class M/E_r/1 make-to-stock queue (with an Erlang-r processing time). We assume the supplier exactly knows the current stage (phase) of the Erlang distribution (which can also correspond to an actual stage of the production process) and therefore, the remaining number of stages to go before completion. This approach allows us to model the information on the production status in a tractable way.

In addition, Erlang distributions provide some flexibility in modeling the production process variability. de Véricourt et al. (2001) provide insights onto the benefit of stock allocation policies when the utilization rate and the relative importance of the customer classes vary. Because of the exponential assumption therein, the impact of production time variability in this comparison is not addressed. In this paper, we evaluate the performance of optimal stock rationing policies when the production time variability increases and the mean stays constant. These two features of the Erlang distribution (information on the production status and production time variability) yield insights that cannot be obtained with the exponential distribution assumption.

To our knowledge Ha (2000) is the only paper that has addressed dynamic optimality issues in stock allocation problem for the make-to-stock queue where the processing time has an Erlang distribution. He assumes lost sales and shows that a single state variable, the work storage level, can fully capture the inventory level and the status of the current production of the system. The problem reduces then to a single dimensional MDP. He then fully characterizes the optimal stock allocation policy: for each customer class there exists a work-storage threshold level at which it is optimal to reject a demand of this class.

Our model differs from Ha's (2000) in the assumption that demands are backordered. The backordering assumption is fundamental from an inventory management perspective and merits attention but it makes the analysis much more challenging for two reasons. First, as mentioned earlier, we deal with an $n+1$ dimensional state space since we need to keep track of the waiting demands of each class. Second, backorders require addressing a new type of decision which corresponds to the production allocation in presence of waiting demands from different classes. This problem does not exist when demands are lost.

Using a Markov Decision Process framework, we present a partial characterization for the optimal stock allocation policy which suggests a work-storage rationing policy where an arriving demand is backordered when the current inventory level is below or at the corresponding threshold. Focusing on this policy leads to the construction of a heuristic exploiting the nested structure of this policy and using a geometric tail approximation for the underlying queueing system. The resulting heuristic policy turns out to be very effective and enables insights into the impact of production time variability on the performance of the system.

In the next section, we introduce the models and formulate the stock rationing problems as a Markov Decision Process. Some structural properties of the optimal policy for the problem are presented in Section 3. Based on this result, we develop an efficient heuristic for the problem in Section 4. The performance of the heuristic is assessed in Section 5.

We investigate the impact of production time variability in Section 6 and we conclude the paper in Section 7.

2 Model Formulation

Consider a supplier who produces a single item at a single facility for n different classes of customers. The finished items are placed in a common stock. When the inventory is empty, demands are backordered. When it is not, an arriving demand can be either satisfied by the on-hand inventory or can be backordered. Items held in stock induce holding costs at rate h (per item per unit of time). Demands of Class i, $1 \leq i \leq n$, arrive according to a Poisson process with rate λ_i and have a unit backorder cost of b_i (per item per unit of time). Suppose without loss of generality that the backorder costs are ordered such that $b_1 > \ldots > b_n$, that is customer classes are ordered from the most valuable to the least valuable one. We denote by $\mathbf{b} = (b_1, \cdots, b_n)$ the n-dimensional vector of backorder costs.

The production process consists of r identical stages in series, each exponentially distributed with mean $1/r\mu$, and the manager of the system can observe the current stage of the production process. The supplier's facility is thus modelled by a single server whose processing time is r-Erlang distributed with mean $1/\mu$. In order to ensure stability of the system, we assume that $\rho = \sum_{i=1}^{n} \lambda_i/\mu < 1$ where ρ is the utilization rate of the system.

At any time instant, the manager of the system must decide whether to produce or not. When a part is completed, he also must decide between satisfying the waiting demand of a customer or increasing the on-hand inventory level. On the other hand, when the demand of a customer arrives to the system, the manager can either satisfy it with the on-hand inventory or backorder it in order to reserve the stock for future (more valuable) customers.

Let $i(t)$ be the number of stages completed by the part under current production at time t and $s(t)$ be the on-hand inventory at time t. We can aggregate $s(t)$ and $i(t)$ in a single variable $x_0(t) = s(t) + i(t)/r$. In the following, $x_0(t)$ will be referred to as the work-storage level. Furthermore, $i(t)$ and $s(t)$ can be inferred from $x_0(t)$ in the following way:

$$s(t) = \lfloor x_0(t) \rfloor \text{ and } i(t) = r\left(x_0(t) - \lfloor x_0(t) \rfloor\right)$$

where $\lfloor y \rfloor$ denotes the largest integer that is less than or equal to y. For example, if $r = 5$ and $x_0 = 2.6$, the inventory consists of two parts ($s(t) = 2$) and the third stage of production is accomplished ($i(t) = 3$). The work-storage level $x_0(t)$ takes its values in the set $\mathbb{N}_r = [x_0 | r x_0 \in \mathbb{N}]$, where $\mathbb{N}$ represents the set of non-negative integers. Let $-x_i(t)$, $1 \leq i \leq n$, be the number of backorders of Class i, $1 \leq i \leq n$, at time t. Hence we can describe exhaustively the system state with $\mathbf{x}(t) = (x_0(t), x_1(t), \cdots, x_n(t))$ and the state space is $S_n = \mathbb{N}_r \times (\mathbf{Z}^-)^n$, where $\mathbf{Z}^-$ represents the set of non-positive integers. Let $\mathbf{X}$ represent the random variable corresponding to $\mathbf{x}$.

A control policy states the action to take at any time given the current state $\mathbf{x}(t)$. We restrict the analysis to Markovian policies since the optimal policy belongs to this class (Puterman 1994). Let $\mathbf{a}^\pi(\mathbf{x}) = (a_0^\pi(\mathbf{x}), \cdots, a_n^\pi(\mathbf{x}))$ be the control associated with a policy π where $a_0^\pi(\mathbf{x})$ is the action to be followed each time a stage of production is completed

$$a_0^\pi(\mathbf{x}) = \begin{cases} 0 & \text{to allocate the produced item to the on-hand inventory} \\ & \text{(possible only when } (x_0 + 1/r) \in \mathbb{N}) \\ k & 1 \leq k \leq n, \text{ to satisfy a backordered demand of Class } k \\ & \text{(possible only when } x_k < 0 \text{ and } x_0 \geq 1 - 1/r) \\ n+1 & \text{not to produce} \\ & \text{(possible only when } x_0 \in \mathbb{N}) \end{cases} \tag{1}$$

Notice that, when $x_0 = 1-1/r$, there is no inventory ($s(t) = 0$) and $r-1$ stages of production are accomplished ($i(t) = r-1$). Thus there is only one more stage of production to be done before one item is available either to satisfy one demand or to increase the inventory by one unit.

$a_k^\pi(\mathbf{x})$, $1 \leq k \leq n$, is a rationing action to be taken each time a demand of Class k arrives

$$a_k^\pi(\mathbf{x}) = \begin{cases} 0 & \text{to satisfy an arriving demand of Class } k \\ & \text{(possible only when } x_0 \geq 1) \\ k & \text{to backorder an arriving demand of Class } k \end{cases} \tag{2}$$

In state $\mathbf{x}$, the system incurs a cost rate

$$c(\mathbf{x}) = h\lfloor x_0 \rfloor - \sum_{i=1}^{n} b_i x_i \tag{3}$$

The objective is to find a control policy, π, which minimizes the expected discounted costs over an infinite horizon:

$$\min_\pi \lim_{T\to\infty} E^\pi_{\mathbf{x}(0)} \left[\int_0^T e^{-\alpha t} c(\mathbf{X}(t)) dt \right] \tag{4}$$

where α is the interest rate. Without loss of generality, we can rescale time by taking $r\mu + \sum_{i=1}^n \lambda_i + \alpha = 1$ and using uniformization (see Lippman 1975), the optimal value function v^* can be shown to satisfy the following optimality equations:

$$v^*(\mathbf{x}) = c(\mathbf{x}) + r\mu T_0 v^*(\mathbf{x}) + \sum_{k=1}^{n} \lambda_k T_k v^*(\mathbf{x}) \tag{5}$$

where the operators T_0 and T_k, $1 \leq k \leq n$, are

$$T_0 v(\mathbf{x}) = \begin{cases} \min\limits_{1 \leq i \leq n : x_i < 0} [v(\mathbf{x} + \mathbf{e}_0/r), v(\mathbf{x} + \mathbf{e}_0/r - \mathbf{e}_0 + \mathbf{e}_i)] & \text{if } x_0 \notin \mathbb{N}\ x_0 \geq 1 - 1/r \\ \min\limits_{1 \leq i \leq n : x_i < 0} [v(\mathbf{x}), v(\mathbf{x} + \mathbf{e}_0/r), v(\mathbf{x} + \mathbf{e}_0/r - \mathbf{e}_0 + \mathbf{e}_i)] & \text{if } x_0 \in \mathbb{N} \text{ and } x_0 > 0 \\ \min [v(\mathbf{x}), v(\mathbf{x} + \mathbf{e}_0/r)] & \text{if } x_0 = 0 \\ v(\mathbf{x} + \mathbf{e}_0/r) & \text{if } 0 < x_0 < 1 - 1/r \end{cases} \tag{6}$$

$$T_k v(\mathbf{x}) = \begin{cases} \min [v(\mathbf{x} - \mathbf{e}_0), v(\mathbf{x} - \mathbf{e}_k)] & \text{if } x_0 \geq 1 \\ v(\mathbf{x} - \mathbf{e}_k) & \text{if } x_0 < 1 \end{cases} \tag{7}$$

where $\mathbf{e}_i$, $0 \leq i \leq n$, is the i-th unit vector. For example, $\mathbf{e}_1$ denotes the $(n+1)$-dimensional vector $(0, 1, 0, \cdots, 0)$. Operator T_0 is associated with production action a_0^π and T_k, $1 \leq k \leq n$, is associated with the rationing action a_k^π. We also define the operator T such that $Tv = c + r\mu T_0 v + \sum_{k=1}^n \lambda_k T_k v$. Notice that $\mathbf{x} + \mathbf{e}_0$ corresponds to $\mathbf{x}$ increased by one unit of stock whereas $\mathbf{x} + \mathbf{e}_0/r$ corresponds to $\mathbf{x}$ increased by one stage of production.

In addition, by introducing the change of variable $\mathbf{w} = \mathbf{x} + \mathbf{e}_0/r - \mathbf{e}_0$, operator T_0 can be simplified as follows:

$$T_0 v(\mathbf{x}) = \begin{cases} \min\limits_{1 \leq i \leq n : x_i < 0} [v(\mathbf{w}), v(\mathbf{w} + \mathbf{e}_i)] & \text{if } x_0 \notin \mathbb{N} \text{ and } x_0 \geq 1 - 1/r \\ \min\limits_{1 \leq i \leq n : x_i < 0} [v(\mathbf{x}), v(\mathbf{w} + \mathbf{e}_0), v(\mathbf{w} + \mathbf{e}_i)] & \text{if } x_0 \in \mathbb{N} \text{ and } x_0 > 0 \\ \min [v(\mathbf{x}), v(\mathbf{x} + \mathbf{e}_0/r)] & \text{if } x_0 = 0 \\ \min [v(\mathbf{x} + \mathbf{e}_0/r)] & \text{if } 0 < x_0 < 1 - 1/r \end{cases} \tag{8}$$

It is also convenient to define the operators Δ_i, $0 \leq i \leq n+1$, for the real-valued function v such that $\Delta_i v(\mathbf{x}) = v(\mathbf{x} + \mathbf{e}_i) - v(\mathbf{x})$. We also define the operators Δ_{ij}, $1 \leq i, j \leq n+1$, such that $\Delta_{ij} v(\mathbf{x}) = \Delta_i v(\mathbf{x}) - \Delta_j v(\mathbf{x}) = v(\mathbf{x} + \mathbf{e}_i) - v(\mathbf{x} + \mathbf{e}_j)$. When $j > n$, we take $\Delta_{ij} v(\mathbf{x}) = \Delta_i v(\mathbf{x})$ (for instance $\Delta_{i(n+1)} v = \Delta_i v$). To simplify the notation, we will implicitly assume that $x_i < 0$ for $1 \leq i \leq n$ and $x_j < 0$ for $1 \leq j \leq n$ when we consider $\Delta_{ij} v(\mathbf{x})$ or $\Delta_i v(\mathbf{x})$ (otherwise these quantities are not defined). The number of customer classes of the underlying problem will also be implicit, when no confusion is possible.

Finally, in the rest of the paper, we will frequently refer to the class with the highest backorder cost which has backordered demands. This class is given by the following function m:

$$\forall \mathbf{x} \in S_n, m(\mathbf{x}) = \begin{cases} \min\limits_{i \in \{1, \cdots, n\} : x_i < 0} (i) & \text{if } \exists i \in \{1, \cdots, n\}, x_i < 0 \\ n+1 & \text{otherwise} \end{cases} \tag{9}$$

3 A Partial Characterization of the Optimal Policy

3.1 The Single-class Problem

We start by studying the problem with one class of demands. When there is only one class of customers, the problem is to decide when to satisfy demands of Class 1 and when to produce. First we show by a sample path argument that it is always optimal to satisfy a Class 1 demand. Second we prove by a value iteration argument that the optimal policy is base-stock. The single-class problem also sheds some light into the difficulties to analyze the multi-class case.

Consider two processes defined on the same probability space. Process 1 uses a policy Π_1 that does not always satisfy Class 1 demands. Process 2 uses a policy Π_2 that mimics production decisions of Π_1 and, whenever it is possible, satisfies Class 1 demands. Since the processes are defined on the same probability space, they see the same service times and arrivals of demands. We can decompose the cost function as follows:

$$c(\mathbf{x}) = h\lfloor x_0 \rfloor - b_1 x_1 = (h + b_1)\lfloor x_0 \rfloor - b_1(\lfloor x_0 \rfloor + x_1) \tag{10}$$

Then we notice that the quantity $(\lfloor x_0 \rfloor + x_1)$ is the same for any process that have the same service times and arrivals of demands. Thus Processes 1 and 2 can be compared through the quantity $\lfloor x_0 \rfloor$ which is always smaller (or equal) for Process 2. As a result, it is optimal to always satisfy demands of Class 1. Therefore, we can not have both inventory and backorders of Class 1 and the state variable of the system can be described by a single variable x_0 with $\lfloor x_0 \rfloor^+ = \max(0, \lfloor x_0 \rfloor)$ the inventory level and $\lfloor x_0 \rfloor^- = -\min(0, \lfloor x_0 \rfloor)$ the number of backorders of Class 1. Whatever the sign of $\lfloor x_0 \rfloor$, the number of stages completed by the part under current production is $r(x_0 - \lfloor x_0 \rfloor)$.

To identify the optimal policy, we introduce the set of functions, $\mathcal{V}_0$, defined with the following property:

$$v(\mathbf{x} + \mathbf{e}_0 + \mathbf{e}_0/r) - v(\mathbf{x} + \mathbf{e}_0/r) \geq v(\mathbf{x} + \mathbf{e}_0) - v(\mathbf{x}) \tag{11}$$

The following lemma states that operator T preserves $\mathcal{V}_0$ for the single-class problem.

Lemma 1 *If $v \in \mathcal{V}_0$, then $Tv \in \mathcal{V}_0$*

Proof: See Appendix A

Using value iteration and Lemma 1, we obtain that the optimal value function belongs to $\mathcal{V}_0$. As a result, the optimal policy is of base-stock type: there exists a base-stock level S^* such that it is optimal to produce if the work-storage level x is smaller than S^* and to idle production otherwise.

3.2 The Multi-Class Problem

As expected, the multi-class problem turns out to be much more challenging than the single class problem. Nevertheless, we are able to establish a number of basic results on the structure of the optimal policy for this case. Lemma 2 establishes three basic properties described in Definition 1 for the optimal policy (where the last two are consequences of the first one).

Definition 1 *Let $\mathcal{U}_n$ be a set of functions such that $v \in \mathcal{U}_n$ if and only if*

1. $\Delta_{ij} v(\mathbf{x}) < 0$ *when* $1 \leq i < j \leq n$
2. $\Delta_{0j} v(\mathbf{x}) < \Delta_{0i} v(\mathbf{x})$ *when* $1 \leq i < j \leq n$
3. $\Delta_{0j} v(\mathbf{x} - \mathbf{e_j}) < \Delta_{0i} v(\mathbf{x} - \mathbf{e}_i)$ *when* $1 \leq i < j \leq n$

The following lemma states that operator T preserves $\mathcal{U}_n$.

Lemma 2 *If $v \in \mathcal{U}_n$, then $Tv \in \mathcal{U}_n$*

Proof: See Appendix B.

The structural properties suggested by Lemma 2 are fairly intuitive. Assume that there are backorders of classes i and j with $1 \leq i < j \leq n$ ($b_i > b_j$), the first property states that it is better to satisfy Class i, the more expensive one. The second property states that if increasing the inventory when there are Class i backorders in the system decreases costs, then increasing the inventory when there are Class j backorders in the system also decreases costs. The third property is symmetrical to the second one: if the policy states to satisfy an arriving demand of Class j with the on-hand inventory, it also states to satisfy the arriving demands of more expensive classes.

Even though Lemma 2 establish basic properties on how to prioritize items in a multi-class systems, a complete characterization of the optimal policy requires several additional properties which turn out to be impossible to establish by our approach. In particular, for the single-class problem Equation (11) implies that v is supermodular in the production status ($\mathbf{e}_0/r$) and the inventory level ($\mathbf{e}_0$). In order to generalize Lemma 1 to the multi-dimensional problem, more modularity properties are required to ensure that Equation (11) can be propagated. For instance, with 2 demand classes, we would need v to be also supermodular in the production status ($\mathbf{e}_0/r$) and the number of waiting demands of Class 2 ($\mathbf{e}_2$). Unfortunately, the optimal value function does not necessarily satisfy these additional modularity properties (For example, a numerical study shows that the optimal value function for $\mu = 1$, $\lambda_1 = 0.3$, $\lambda_2 = 0.3$, $h = 0.01$, $b_1 = 10$, $b_2 = 1$, $\alpha = 0.01$, is not supermodular in $\mathbf{e}_0/r$ and $\mathbf{e}_2$). This means that the marginal cost of continuing production can increase in the number of waiting demands, especially when the inventory level is already high. As a result, a full characterization of the optimal policy seems extremely difficult in the multi-class case.

4 A Plausible Heuristic Policy

The results of the previous section partially uncover the priority properties of production and stock allocation policies but do not suggest a precise policy. On the other hand, Ha (2000) presents the characterization of the optimal policy in the lost sales case as a work-storage threshold type policy. The same structure is shown to be optimal by Gayon et al. (2006a) in a similar problem to the one considered but where production cannot be interrupted and excess inventory is diverted to a salvage market. It is therefore plausible that this type of policy is near-optimal and potentially very effective if its parameters can be computed efficiently. Hereon we focus on the investigation of this particular class of policies whose precise definition is given below.

A Work Storage Rationing (WR) policy is shown to be optimal in Gayon et al. (2006a). A WR-policy is defined by $n+1$ parameters, one corresponding to each type of demand. Let $z_k \in \mathbb{N}_r$ be the work-storage rationing level of Class k, $1 \leq k \leq n+1$, that is, all arriving demands of this type are backordered when the work-storage level is below z_k. Note that Class $(n+1)$ corresponds to the salvage market. Moreover, when a part is produced it is allocated to a backordered demand of Class k, only if the work storage level x_0 is larger than or equal to z_k. It is allocated to the stock otherwise. If some of these parameters are equal, the resource is allocated to the most expensive customer class (that is to the class $m(\mathbf{x})$ in state $\mathbf{x}$). Note that, in a WR-policy, demands of Class 1 are always satisfied, whenever inventory is available. A formal definition can be found in Gayon et al. (2006a).

We propose, for the problem without a salvage market, a modified WR-policy as a heuristic policy. A modified WR (MWR) policy is a WR-policy except that the salvage market rationing level is replaced by an (integer) base-stock level. When there are no backordered demands ($m(\mathbf{x}) = n+1$), this modified policy states to produce if the inventory level is strictly smaller than the base-stock level, and not to produce otherwise. All the other controls are the same as those in the original WR-policy. This heuristic is justified by the fact that both systems are governed by very similar equations.

To compute the parameters of the heuristic MWR-policy, we will first compute the optimal policy parameters of the WR-policy for the problem with a salvage market. We consider the average cost minimization criterion which has a simpler interpretation: the optimal average cost does not depend on the initial conditions and the optimal policy parameters do not depend on the discount factor selected which facilitates various comparisons (see Ha 1997b or Ha 2000 for a similar approach).

Restricting ourselves to the class of WR-policies, let us now turn to the problem of computing optimal values of the policy parameters. However, due to the curse of dimensionality, it rapidly becomes impossible to obtain these values using numerical techniques when the number of classes increases. To get around this difficulty, we use a result of Gayon et al. (2006a) who establish a strong relationship between a k-class problem and a $(k-1)$- class subproblem. The essence of the approximation leading to the heuristic policy is then a successive computation of the rationing levels $z_1, \cdots, z_{k+1}$. When the rationing level vector $z_1, \cdots, z_k$ of the $(k-1)$-class subproblem and the corresponding average cost g_{k-1} have been evaluated, the next rationing level z_{k+1} and optimal average cost g_k for the k-class problem can be computed by solving a single dimensional problem. Indeed, when the work-storage level is larger than z_k, all demands are satisfied with the stock and there are no backorders in recurrent states. When the work-storage level is lower than z_k, the average cost is given by g_{k-1}. By iterating this step for each subproblem, we obtain the following algorithm, whose full justification is given in Appendix C.

Heuristic 1 *Consider an n-class problem. Construct the sequences ρ_k, η_k and $\tilde{z}_k$ as follows:*

Initialize $\tilde{z}_1 = 1 - 1/r$, $\rho_0 = 1$.

For $k = 1, \ldots, n$ *do,*

$$\rho_k = \frac{1}{\mu} \sum_{i=1}^{k} \lambda_i$$

η_k *is the solution in the interval (0,1) of the equation* $\left(\frac{r}{r + \rho_k(1 - 1/\eta_k)} \right)^r = 1/\eta_k$

$$\tilde{z}_{k+1} = \tilde{z}_k + \log_{\eta_k} \frac{\eta_k(h + b_{k+1})}{\rho_k(h + b_k) \left[\eta_k + (1 - \eta_k) \frac{1-\rho_{k-1}}{1-\eta_{k-1}} \right]}$$

The heuristic rationing levels z_k*,* $k \geq 1$*, are then given by*

$$z_1 = \tilde{z}_1$$

$$z_k = \frac{\lfloor r\tilde{z}_k + 1 \rfloor}{r} \quad \textit{for } k = 2, 3..., n$$

$$z_{n+1} = \frac{\lfloor r\tilde{z}_{n+1} + 1 \rfloor}{r}$$

The MWR heuristic for our problem, without a salvage market, is obtained by rounding-off z_{n+1} in order to obtain the base-stock level.

Let us note that the above algorithm can easily be adapted to any M/G/1 make-to-stock queue. We do not pursue this adaptation here since testing the performance of the algorithm in other settings than M/E$_r$/1 would require the understanding and the computation of the optimal policy. Section 5 presents a summary of the results on the performance of the MWR policy for the M/E$_r$/1 case.

5 The Performance of the Heuristic Policy

In order to evaluate the performance of the heuristic policy described in Section 5, we compared the average cost g^* of the optimal policy of the problem (evaluated numerically by a value iteration algorithm) with the average cost g, where g is obtained by numerically computing the performance of the modified WR policy with the rationing levels $\tilde{z}_i$'s of the previous heuristic. We denote then by $\Delta g = (g - g^*)/g^*$, the relative cost increase when using the heuristic policy instead of the optimal policy.

In order to systematically test the performance of the heuristic, we investigated several cases with 2 or 3 classes of customers by varying the different parameters of the problems. In all our numerical studies we set for any n-class problem $\mu = 1$ and $b_n = 1$ without loss of generality. We can then define exhaustively the problem with the parameters $(r, \rho, \lambda_1/\lambda_2, \cdots, \lambda_{n-1}/\lambda_n, b_1/b_2, \cdots, b_{n-1}/b_n, h')$ where h' is the relative holding cost defined by:

$$h' = h\frac{\sum_{i=1}^{n} \lambda_i}{\sum_{i=1}^{n} \lambda_i b_i}.$$

This quantity expresses the relative importance of the holding cost compared to the backlog costs.

Tables 1 and 2 report typical results for the rationing levels z_i and Δg where the number of stages r are varied. Table 1 considers a 2-class problem while Table 2 presents the results for a problem with 3 classes. The heuristic procedure appears to perform very well in general, providing an error less than 1%.

r	z_2	z_3	Δg
1	2.00	15	0.00%
2	2.00	12	0.00%
3	2.00	11	0.00%
5	1.80	10	0.17%
10	1.80	9	0.35%
20	1.85	9	0.21%

Table 1: Performance of the heuristic for a 2-class problem with $\rho = 0.8, \lambda_1 = \lambda_2, b_1/b_2 = 10, h' = 0.01$

r	z_2	z_3	z_4	Δg
1	1.00	4.00	8	0.00%
2	1.50	3.50	7	0.00%
3	1.33	3.67	6	0.40%
5	1.40	3.40	6	0.37%
10	1.35	3.35	6	1.01%

Table 2: Performance of the heuristic for a 3-class problem with $\rho = 0.75, \lambda_1 = \lambda_2 = \lambda_3, b_1/b_2 = 10, b_2/b_3 = 10, h' = 0.01$

Stock rationing is significantly influenced by the asymmetries in the arrival rates of different customer classes and in their backorder costs. Tables 3 and 4 present results on the performance of the heuristic for a two-class systems where the ratios b_1/b_2 and λ_1/λ_2 are varied. As can be observed in Tables 3 and 4, the performance of the heuristic is excellent regardless of these in the investigated cases.

b_1/b_2	z_2	z_3	Δg
1	0.80	13	0.00%
2	0.80	12	0.00%
10	1.80	10	0.17%
20	2.40	9	0.06%
50	2.80	7	0.49%
100	3.20	6	0.54%

Table 3: Performance of the heuristic for a 2-class problem with $r = 5, \rho = 0.8, \lambda_1 = \lambda_2, h' = 0.01$

λ_1/λ_2	z_2	z_3	Δg
10	4.80	11	0.16%
5	4.00	10	0.03%
2	2.60	10	0.00%
1	2.00	10	0.15%
0.5	1.60	10	0.64%
0.2	1.20	11	0.06%
0.1	0.80	11	0.67%

Table 4: Performance of the heuristic for a 2-class problem with $r = 5,\ \rho = 0.8, b_1/b_2 = 10, h' = 0.01$

The heuristic policy seems to perform in all the previous examples. The last investigation considers the other two parameters h' and ρ which also influence rationing phenomena considerably. In particular whenever h' or ρ is extremely low, there is little need to carry a large inventory, and a very small base-stock level is sufficient. This, in turn, leaves little room for rationing. In Table 5 we investigate utilization rates less than or equal to 0.6, and vary the holding cost. It can be observed in this table that the approximation can sometimes lead to large percentage errors in certain cases. In particular, the most problematic cases are when ρ is extremely low (0.1 or 0.2) and the relative holding cost h' is very high (1 or 2). The rationing levels are very low in these examples and a small approximation error in the rationing levels may lead to a magnified percentage error in terms of the average cost. While this is a limitation of the heuristic, it may be argued that these situations are not the most appropriate for stock rationing, and that the absolute error in terms of cost will be relatively small since the average inventory level will be very low in such cases.

ρ	$h = 0.01$	$h = 0.05$	$h = 0.20$	$h = 0.50$	$h = 1.00$	$h = 2.00$
0.1	0.00%	18.52%	0.00%	0.00%	66.38%	0.00%
0.2	0.00%	0.00%	17.87%	0.00%	0.00%	51.84%
0.3	0.11%	0.14%	0.01%	0.00%	0.00%	0.00%
0.4	0.57%	0.47%	1.57%	0.33%	0.00%	0.00%
0.5	0.54%	1.13%	0.84%	0.38%	4.11%	0.00%
0.6	0.60%	0.69%	0.59%	1.69%	0.30%	0.00%

Table 5: Performance of the heuristic for a 2-class problem with $r = 4, \lambda_1 = \lambda_2, b_1/b_2 = 10$

A simpler alternative to the previous heuristic would be to consider the policy specified by Véricourt et al. (2002) with the mean of the exponential processing time distribution matched to the mean of the Erlang processing time distribution. This "exponential" heuristic facilitates the computation a little but does not perform well since it does not take into account neither the processing time variability nor the information on the production stage. Table 6 reports the performance of such a policy for parameters similar to Table 1. In this table, Δg^{ML} denotes the relative cost increase when using this exponential heuristic policy instead of the optimal policy. This comparison shows that ignoring the information on the production status leads to significant costs.

6 Effect of the Processing Time Variability

In this section, we study the impact of the processing time variability on the optimal cost. We measure the processing time variability with its squared coefficient of variation which is equal in our settings to $c_v^2 = 1/r$. When $c_v^2 = 1$, the processing time is exponential, while when c_v^2 approaches 0, the processing time approaches a deterministic value. Note that when c_v^2 increases, the information accuracy on the production status deteriorates, which

r	z_2	z_3	Δg^{ML}
1	2.00	15	0.00%
2	2.00	15	23.2%
3	2.00	15	30.9%
5	2.00	15	37.3%
10	2.00	15	42.1 %
20	2.00	15	44.4 %

Table 6: Performance of an alternative heuristic for a 2-class problem with $\rho = 0.8, \lambda_1 = \lambda_2, b_1/b_2 = 10, h' = 0.01$

reinforces the impact of the variability in the system.

We start by investigating the effect of c_v^2 on the optimal average cost by varying various parameters of the problem. Figure 1, represents for a 2-class problem the optimal cost g^* as a function of c_v^2 for different values of ρ, where $\lambda_1/\lambda_2 = 1, b_1/b_2 = 10$ and $h' = 0.01$.

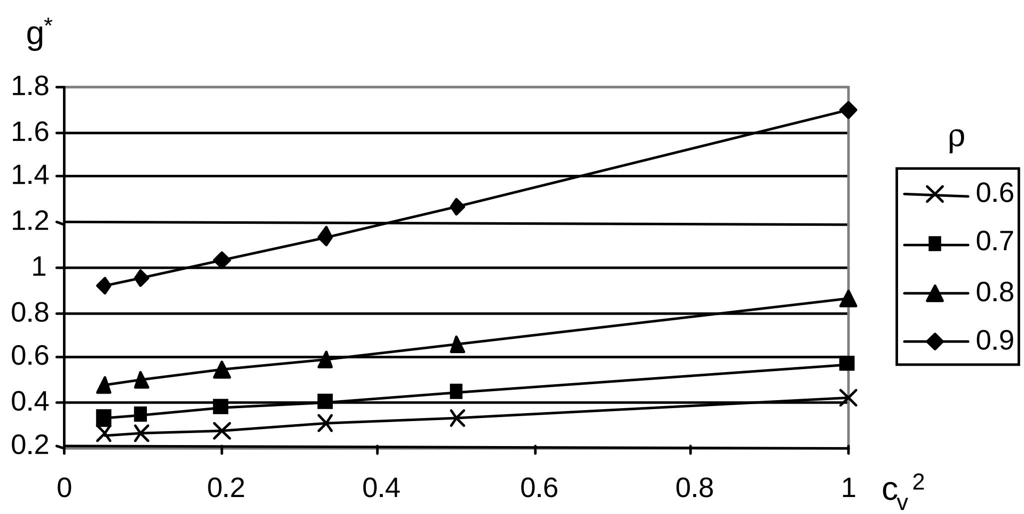

Figure 1: Impact of c_v^2 on g^* for different values of ρ

The interesting observation in Figure 1 is that, the impact of the squared coefficient of variation on the average cost appears to be quasi-linear. This is reminiscent of the relationship between the average cost and the mean size of an $M/G/1$ queue which is linear in c_v^2 according to the well-known Pollaczek-Khinchin formula. The piecewise linear cost structure of the system makes however this analogy non-trivial. The linear effect has also been observed on examples where other parameters are varied (i.e. $\lambda_1/\lambda_2, b_1/b_2$ and h') and also in examples with three customer classes.

This linear relationship can be exploited to evaluate the average costs of problems with large r. These cases are indeed difficult to analyze due to the size of the state space. For instance, if we want to evaluate the average cost when processing time is deterministic, we can compute the costs for $r = 1$ and $r = 2$ only, and then deduce the result with a linear approximation based on these two points.

We finally analyze the average cost increase due to processing time variability. To that end we denote by $g^*(c_v)$ the optimal average cost of a problem with a coefficient of variation c_v. Hence, $g^*(0)$ represents the average cost of problems with a deterministic processing time and is obtained using the previous linear approximation. Figure 2 depicts the relative cost increase $\Delta g^*(c_v) = (g^*(c_v) - g^*(0))/g^*(0)$ when the coefficient of variation varies and for different values of ρ. Δg^* appears to significantly increase as c_v^2 increases. For instance, when $\rho = 0.9$, the relative cost increase for an exponential processing time ($c_v^2 = 1$) reaches 94%. Other similar examples, not reported here, show that ρ has a much higher impact on Δg^* than other parameters such as $b_1/b_2, \lambda_1/\lambda_2$ and h'.

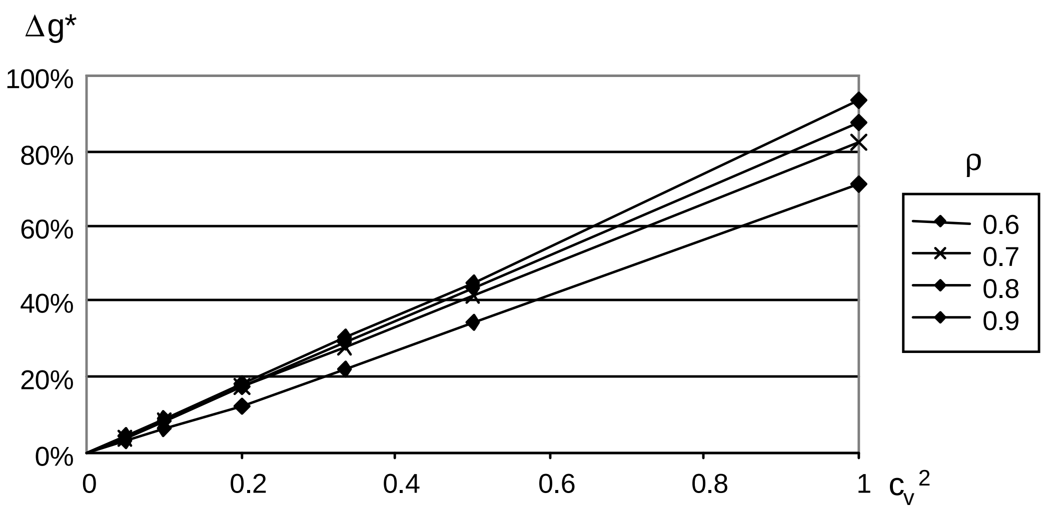

Figure 2: Cost penalty for variability

7 Conclusion and future research

In this paper, we have analyzed a stock rationing problem with several customer classes where the processing times have an Erlang distribution. The Erlang distribution assumption allows us to model the information on the production status in a tractable manner and enables modeling production time variability.

For the problem without a salvage market, we have fully characterized the optimal policy for a single-class problem and have provided a partial characterization of the optimal policy for the multi-class problem. A full characterization, in a general setting, seems to be extremely difficult since the approach, used to propagate modularity properties for the problem with a salvage market, does not work. However, based on the findings of the problem with a salvage market, we have presented a heuristic which is observed to perform very well in a numerical study. We also have identified a linear effect of the processing time variability on the system performance for both problems. Moreover, we have proposed an efficient heuristic evaluation of the corresponding optimal parameters. This heuristic procedure allows addressing problems with a large number of customer classes that would not be tractable otherwise.

Finally, our results constitute a useful benchmark for systems with more general processing times than Erlang distributions. These problems can be non-Markovian and are extremely difficult to analyze since the optimal decisions should take the actual processing time into account. Even if they could be characterized, these policies would most likely be hard to implement. For the deterministic case, our heuristic procedure should already perform well as Erlang distributions approach deterministic times for large numbers of stages. For the more general case, multi-stage distributions with different exponential processing times provide a promising alternative to approximate the processing time. Our heuristic can in fact be directly extended to this case. In general, the nested approach using an $M/G/1$ approximation presented in this paper offers a tractable framework to evaluate the optimal rationing levels in multi-class make-to-stock queues with generally distributed processing times.

Acknowledgments: The authors thank Yves Dallery and Paul Zipkin for their insightful comments on this paper.

References

J.A. Buzacott and J.G. Shanthikumar. *Stochastic Models of Manufacturing Systems.* Prentice Hall, 1993.

G.P. Cachon and M.A. Lariviere. Capacity choice and allocation: Strategic behavior and supply chain performance. *Management Science*, 45:1091–1108, 1999.

S. Carr and I. Duenyas. Optimal admission control and sequencing in a make-to-stock / make-to-order production system. *Operations Research*, 48(5):709–720, 2000.

F. de Véricourt, F. Karaesmen, and Y. Dallery. Assessing the benefits of different stock-allocation policies for a make-to-stock production system. *Manufacturing and Service Operations Management*, 3(2):105–121, 2001.

F. de Véricourt, F. Karaesmen, and Y. Dallery. Optimal stock allocation for a capacitated supply system. *Management Science*, 48(11):1486–1501, 2002.

V. Deshpande, M.A. Cohen, and K. Donohue. A threshold inventory rationing policy for service-differentiated demand classes. *Management Science*, 49:683–703, 2003.

K.C. Frank, R.C. Zhang, and I. Duenyas. Optimal policies for inventory systems with priority demand classes. *Operations Research*, 51:993–1002, 2003.

J.P. Gayon, F. de Véricourt, and F. Karaesmen. On the structure of the optimal stock rationing policy in an $\mathrm{M/E}_r/1$ make-to-stock queue with a salvage market. *Working Paper*, 2006.

A.Y. Ha. Inventory rationing in a make-to-stock production system with several demand classes and lost sales. *Management Science*, 43(8):1093–1103, 1997a.

A.Y. Ha. Stock rationing policy for a make-to-stock production system with two priority classes and backordering. *Naval Research Logistics*, 44, 1997b.

A.Y. Ha. Stock rationing in an $M/E_k/1$ make-to-stock queue. *Management Science*, 46(1): 77–87, 2000.

X. Hu, I. Duenyas, and R. Kapuscinski. Optimal joint inventory and transshipment control under uncertain capacity. Technical report, University of Michigan Business School, Ann Arbor, 2004.

F. Karaesmen, G. Liberopoulos, and Y. Dallery. Production/inventory control with advance demand information. In *Stochastic Modelling and Optimization of Manufacturing Systems and Supply Chains*, pages 243–270. J.G. Shanthikumar, D.D. Yao and W.H.M. Zijm (eds.), Kluwer Academic Publishers, Boston, MA, 2003.

J.E. Lee and Y. Hong. A stock rationing policy in a (s,s)-controlled stochastic production system with 2-phase coxian processing times and lost sales. *International Journal of Production Economics*, 83:299–307, 2003.

S. Lippman. Applying a new device in the optimization of exponential queueing systems. *Operations Research*, 23(4):687–710, 1975.

P. Melchiors. Restricted time-remembering policies for the inventory rationing problem. *International Journal of Production Economics*, 81:461–468, 2003.

P. Melchiors, R. Dekker, and M.J. Kleijn. Inventory rationing in an (s,q) inventory model with lost sales and two demand classes. *Journal of the Operational Research Society*, 51: 111–122, 2000.

S. Nahmias and W.S. Demmy. Operating characteristics of an inventory system with rationing. *Management Science*, 27:1236–1245, 1981.

M.L. Puterman. *Markov Decision Processes*. John Wiley and Sons, 1994.

H.C. Tijms. *Stochastic models: an algorithmic approach*. John Wiley and Sons, 1994.

D.M. Topkis. Optimal ordering and rationing policies in a nonstationary dynamic inventory model with n demand classes. *Management Science*, 15:160–176, 1968.

A. Tsay, S. Nahmias, and N. Agrawal. Modeling supply chain contracts: A review. In *Quantitative Models for Supply Chain Management*. S. Tayur, R. Ganeshan and M. Magazine, Kluwer's International Series, Boston, 1999.

H. Zhao, V. Deshpande, and J.K. Ryan. Ptimal dynamic production and inventory transshipment policies for a two-location make-to-stock system. *Working Paper, Krannert School of Management, Purdue University W. Lafayette*, 2004.

A Proof of Lemma 1

For the single-class problem, the optimality equations simplifies to: $v^*(\mathbf{x}) = c(\mathbf{x}) + \mu T_0 v^*(\mathbf{x}) + \lambda_1 v^*(\mathbf{x} - \mathbf{e}_0)$. Assume that $v \in \mathcal{V}_0$. Let $\delta_0 v$ be the operator such that $\delta_0 v(\mathbf{x}) = v(\mathbf{x} + \mathbf{e}_0/r) - v(\mathbf{x})$. The quantity $S = \min[x_0 \in \mathbb{N} : \delta_0 v(\mathbf{x}) \geq 0]$ is well defined and policy π, associated to v, states to produce if the work-storage level x_0 is smaller than S and to idle production otherwise. Following the definition of S, we can rewrite the operator T_0:

$$T_0 v(\mathbf{x}) = \begin{cases} v(\mathbf{x} + \mathbf{e}_0/r) & \text{if } x_0 \notin \mathbb{N} \text{ or } x_0 < S \\ v(\mathbf{x}) & \text{if } x_0 \in \mathbb{N} \text{ and } x_0 > S \end{cases} \tag{12}$$

Then we have

$$\Delta_0 T_0 v(\mathbf{x}) = \begin{cases} \Delta_0 v(\mathbf{x}+\mathbf{e}_0/r) & \text{if } x \notin \mathbb{N} \text{ or } x_0+1<S \\ v(\mathbf{x}+\mathbf{e}_0) - v(\mathbf{x}+\mathbf{e}_0/r) & \text{if } x_0+1=S \\ \Delta_0 v(\mathbf{x}) & \text{if } x \in \mathbb{N} \text{ and } x_0+1>S \end{cases} \tag{13}$$

and

$$\Delta_0 T_0 v(\mathbf{x}+\mathbf{e}_0/r) - \Delta_0 T_0 v(\mathbf{x})$$

$$= \begin{cases} \Delta_0 v(\mathbf{x}+2\mathbf{e}_0/r) - \Delta_0 v(\mathbf{x}+\mathbf{e}_0/r) \geq 0 & \text{if } x+1+1/r<S \\ & \text{or if } (x \notin \mathbb{N} \text{ and } x_0+1/r \notin \mathbb{N}) \\ -\delta_0 v(\mathbf{x}+\mathbf{e}_0/r) \geq 0 & \text{if } x_0+1+1/r=S \\ \Delta_0 v(\mathbf{x}+2\mathbf{e}_0/r) - \Delta_0 v(\mathbf{x}+\mathbf{e}_0/r) & \\ \qquad\qquad +\delta_0 v(\mathbf{x}+\mathbf{e}_0/r) \geq 0 & \text{if } x_0+1=S \\ 0 & \text{if } x_0+1>S \text{ and } x_0+1/r \in \mathbb{N} \\ \Delta_0 v(\mathbf{x}+2\mathbf{e}_0/r) - \Delta_0 v(\mathbf{x}) \geq 0 & \text{if } x_0+1>S \text{ and } x_0 \in \mathbb{N} \end{cases} \tag{14}$$

The inequalities in (14) come from the definition of S and the assumption $v \in \mathcal{V}_0$, so that $T_0 v \in \mathcal{V}_0$. Since the cost function $c(.)$ also belongs to $\mathcal{V}_0$, we obtain the result.

B Proof of Lemma 2

Assume that $v \in \mathcal{U}_n$ and $1 \leq i < j \leq n$. Let us show that Tv verifies the first condition of $\mathcal{U}_n$.

First of all, we have $\Delta_{ij} c(\mathbf{x}) = b_j - b_i < 0$. Let us show now that $\Delta_{ij} T_0 v(\mathbf{x}) < 0$. To that end, we have to distinguish four cases.

1. $x_0 = 0$

   $$\Delta_{ij} T_0 v(\mathbf{x}) = \min\left[v(\mathbf{x}+\mathbf{e_i}), v(\mathbf{x}+\mathbf{e_i}+\mathbf{e}_0/r)\right] - \min\left[v(\mathbf{x}+\mathbf{e_j}), v(\mathbf{x}+\mathbf{e_j}+\mathbf{e}_0/r)\right]$$

   If $\min\left[v(\mathbf{x}+\mathbf{e_j}), v(\mathbf{x}+\mathbf{e_j}+\mathbf{e}_0/r)\right] = v(\mathbf{x}+\mathbf{e_j})$, then

   $$\Delta_{ij} T_0 v(\mathbf{x}) \leq \Delta_{ij} v(\mathbf{x}) \leq 0$$

   If $\min\left[v(\mathbf{x}+\mathbf{e_j}), v(\mathbf{x}+\mathbf{e_j}+\mathbf{e}_0/r)\right] = v(\mathbf{x}+\mathbf{e_j}+\mathbf{e}_0/r)$, then

   $$\Delta_{ij} T_0 v(\mathbf{x}) \leq \Delta_{ij} v(\mathbf{x}+\mathbf{e}_0/r) \leq 0$$

   Therefore $\Delta_{ij} T_0 v(\mathbf{x}) \leq 0$.

2. $0 < x_0 < 1 - 1/r$

$$\Delta_{ij} T_0 v(\mathbf{x}) = \Delta_{ij} v(\mathbf{x} + \mathbf{e}_0/r) \le 0$$

3. $x_0 \in \mathbb{N}$, $x_0 > 0$

   Let $m(\mathbf{x} + \mathbf{e}_i) = p$ and $m(\mathbf{x} + \mathbf{e_j}) = q$. Notice that $q \le p \le j$. We have

$$\begin{aligned} \Delta_{ij} T_0 v(\mathbf{x}) &= \min[v(\mathbf{x} + \mathbf{e}_i), v(\mathbf{w} + \mathbf{e}_i + \mathbf{e}_0), v(\mathbf{w} + \mathbf{e}_i + \mathbf{e_p})] \\ &\quad - \min[v(\mathbf{x} + \mathbf{e}_j), v(\mathbf{w} + \mathbf{e_j} + \mathbf{e}_0), v(\mathbf{w} + \mathbf{e_j} + \mathbf{e_q})] \end{aligned}$$

   If $\min[v(\mathbf{x} + \mathbf{e}_j), v(\mathbf{w} + \mathbf{e_j} + \mathbf{e}_0), v(\mathbf{w} + \mathbf{e_j} + \mathbf{e_q})] = v(\mathbf{x} + \mathbf{e}_j)$, then

$$\Delta_{ij} T_0 v(\mathbf{x}) \le \Delta_{ij} v(\mathbf{x}) \le 0$$

   If $\min[v(\mathbf{x} + \mathbf{e}_j), v(\mathbf{w} + \mathbf{e_j} + \mathbf{e}_0), v(\mathbf{w} + \mathbf{e_j} + \mathbf{e_q})] = v(\mathbf{w} + \mathbf{e_j} + \mathbf{e}_0)$, then

$$\Delta_{ij} T_0 v(\mathbf{x}) \le \Delta_{ij} v(\mathbf{w} + \mathbf{e}_0) \le 0$$

   If $\min[v(\mathbf{x} + \mathbf{e}_j), v(\mathbf{w} + \mathbf{e_j} + \mathbf{e}_0), v(\mathbf{w} + \mathbf{e_j} + \mathbf{e_q})] = v(\mathbf{w} + \mathbf{e_j} + \mathbf{e_q})$, then

$$\Delta_{ij} T_0 v(\mathbf{x}) \le v(\mathbf{w} + \mathbf{e}_i + \mathbf{e_p}) - v(\mathbf{w} + \mathbf{e_j} + \mathbf{e_q})$$

   If $p = q$, then $v(\mathbf{w} + \mathbf{e}_i + \mathbf{e_p}) - v(\mathbf{w} + \mathbf{e_j} + \mathbf{e_q}) = \Delta_{ij} v(\mathbf{w} + \mathbf{e_p}) \le 0$. If $p > q$, then $q = i$ and $v(\mathbf{w} + \mathbf{e}_i + \mathbf{e_p}) - v(\mathbf{w} + \mathbf{e_j} + \mathbf{e_q}) = \Delta_{pj} v(\mathbf{w} + \mathbf{e}_i) \le 0$ (since $p \le j$). Therefore $\Delta_{ij} T_0 v(\mathbf{x}) \le 0$.

4. $x_0 \notin \mathbb{N}$, $x_0 \ge 1 - 1/r$

   With the same notations, we have

$$\begin{aligned} \Delta_{ij} T_0 v(\mathbf{x}) &= \min[v(\mathbf{w} + \mathbf{e}_i + \mathbf{e}_0), v(\mathbf{w} + \mathbf{e}_i + \mathbf{e_p})] \\ &\quad - \min[v(\mathbf{w} + \mathbf{e_j} + \mathbf{e}_0), v(\mathbf{w} + \mathbf{e_j} + \mathbf{e_q})] \end{aligned}$$

   If $\min[v(\mathbf{w} + \mathbf{e_j} + \mathbf{e}_0), v(\mathbf{w} + \mathbf{e_j} + \mathbf{e_q})] = v(\mathbf{w} + \mathbf{e_j} + \mathbf{e}_0)$, then

$$\Delta_{ij} T_0 v(\mathbf{x}) \le \Delta_{ij} v(\mathbf{w} + \mathbf{e}_0) \le 0$$

   Otherwise

$$\Delta_{ij} T_0 v(\mathbf{x}) \le v(\mathbf{w} + \mathbf{e}_i + \mathbf{e_p}) - v(\mathbf{w} + \mathbf{e_j} + \mathbf{e_q}) \le 0$$

   We obtain the previous inequality using the same argument as in case 3.

Let us show now that $\Delta_{ij} T_k v(\mathbf{x}) < 0$ for $1 \le k \le n$.

$$\begin{aligned} \Delta_{ij} T_k v(\mathbf{x}) &= \min[v(\mathbf{w} + \mathbf{e}_i - \mathbf{e}_0), v(\mathbf{w} + \mathbf{e}_i - \mathbf{e}_k)] \\ &\quad - \min[v(\mathbf{w} + \mathbf{e_j} - \mathbf{e}_0), v(\mathbf{w} + \mathbf{e_j} - \mathbf{e_k})] \end{aligned}$$

If $\min[v(\mathbf{w}+\mathbf{e_j}-\mathbf{e}_0), v(\mathbf{w}+\mathbf{e_j}-\mathbf{e_k})] = v(\mathbf{w}+\mathbf{e_j}-\mathbf{e}_0)$, then

$$\Delta_{ij}T_k v(\mathbf{x}) \leq \Delta_{ij} v(\mathbf{w}-\mathbf{e}_0) \leq 0$$

Otherwise

$$\Delta_{ij}T_k v(\mathbf{x}) \leq \Delta_{ij} v(\mathbf{w}-\mathbf{e}_k) \leq 0$$

We conclude that $\Delta_{ij}Tv(\mathbf{x}) = \Delta_{ij}c(\mathbf{x}) + r\mu\Delta_{ij}T_0v(\mathbf{x}) + \sum_{i=1}^{n}\lambda_i T_i v(\mathbf{x}) \leq 0$ and Tv verifies the first condition of $\mathcal{U}_n$. Conditions 2 and 3 are direct consequences of Condition 1, applied respectively in $\mathbf{x}$ and in $\mathbf{x}+\mathbf{e}_0-\mathbf{e}_i-\mathbf{e}_j$. Finally $Tv \in \mathcal{U}_n$.

C Heuristic Evaluation for the optimal rationing levels

For a k-class problem, if we consider the WR policy whose rationing work level vector is $(\mathbf{z}^k, z)$, the corresponding average cost $g(z)$ can be written as (see Gayon et al. 2006a):

$$g(z) = E\left[c_{k-1}(\mathbf{X}^{k-1})\right] - (b_k - b_{k+1})E\left[\sum_{i=1}^{k} X_i + \lfloor X_0 \rfloor\right] \tag{15}$$

We momentarily consider that X_0 and z_k are integers and we can then approximate $g(z)$ by

$$g(z) \simeq P\{X_0 \leq z_k\}g_{k-1} + \sum_{s=z_k+1}^{z}(h+b_k)sP\{X_0 = s\} - (b_k - b_{k+1})E\left[\sum_{i=1}^{k} X_i + X_0\right] \tag{16}$$

where $Y = z - \sum_{i=1}^{k} X_i - X_0$ is an $M/E_r/1$ queue-length process with arrival rate $\sum_{i=1}^{k}\lambda_i$. When $X_0 > z_k$, $Y = z - X_0$ because of the WR policy structure. No simple analytical expressions for the distribution of Y exists though, except for the exponential case ($r = 1$). We use then a geometric tail approximation for the queue length distribution of an M/G/1 queue (see Tijms 1994 , Karaesmen, Liberopoulos and Dallery 2003). Following this approach, it is then possible to approximate the value of z which minimizes $g(z)$.

Let us present briefly the approximation of the queue length distribution $\pi(n)$ of an $M/G/1$ queue described in detail in Tijms (1994). We denote by λ the arrival rate and by ρ the utilization rate. Let $f(t)$ be the probability density function of the processing time and L^* be its Laplace transform, then the approximation is given by

$$\pi(n) = \sigma\eta^n \quad \text{for } n \text{ sufficiently large} \tag{17}$$

where $\tau = 1/\eta$ is the smallest real solution strictly larger than 1 of the below equation

$$L^*(\lambda(1-\tau)) = \tau$$

In the case of an r-Erlang processing time with mean, (17) becomes

$$\left(\frac{r\mu}{r\mu+\lambda(1-\tau)}\right)^r = \tau$$

In general, there is no closed form solution of the last equation but its numerical solution is straightforward.

Tijms also proposes an expressions for the constant σ that is asymptotically exact. In order to simplify the final form, we simply assume that the approximation given by equation (17) is valid for all $n > 1$ as in Karaesmen et al. (2003) and that $\pi(0) = 1 - \rho$ where ρ is the utilization rate. With the normalization condition, we obtain

$$\sigma = \frac{\rho}{\eta}(1-\eta)$$

The random variable Y is approximated by an $M/E_r/1$ queue-length process with utilization rate $\rho_k = \sum_{i=1}^{k} \lambda_i/\mu$. With the queue size distribution heuristic introduced above, we have

$$\begin{cases} P[Y=0] = 1-\rho \\ P[Y=j] = \frac{\rho_k}{\eta_k}(1-\eta_k)\eta_k^j \quad \text{if } j>0 \end{cases} \tag{18}$$

where η_k is the real solution strictly smaller than 1 of the following equation

$$\left(\frac{r}{r+\rho_k(1-1/\eta_k)}\right)^r = 1/\eta_k$$

In general, there is no closed form solution of the last equation but its numerical solution is straightforward. However, it is also possible to obtain an approximate value $\tilde{\eta_k}$ of η_k by taking $\tilde{\eta_k}$ such that the mean size of the approximate queue be equal to the exact mean size of the $M/G/1$ queue given by the Pollaczek-Khintchine formula

$$\frac{\rho_k}{1-\eta_k} = \rho_k + \frac{\rho_k^2(1+c_v^2)}{2(1-\rho_k)} \Rightarrow \tilde{\eta_k} = \rho_k \frac{2-\rho_k(1-c_v^2)}{2-\rho_k^2(1-c_v^2)}$$

where c_v is the coefficient of variation of the processing time. In the case of an r-Erlang processing time, $c_v^2 = 1/r$. We have tested but not reported the performance of the heuristic with an approximated $\tilde{\eta_k}$. It works well when r is small and tends to deteriorate when r is increasing.

Using (18), (16) becomes

$$\begin{aligned} g(z) &= P\{Y > z - z_k\}g_{k-1} + \sum_{s=z_k+1}^{z}(h+b_k)sP\{Y = z-s\} - (b_k - b_{k+1})E\left[z-Y\right] \\ &= g_{k-1}\frac{\rho_k}{\eta_k}\eta_k^{z-z_k} + (h+b_k)\sum_{s=z_k+1}^{z-1} s\frac{\rho_k}{\eta_k}(1-\rho_k)\rho_k^{z-s} + (h+b_k)(1-\rho_k)z \\ &\quad + (b_k - b_{k+1})\left(\frac{\rho_k}{1-\eta_k} - z\right) \end{aligned} \tag{19}$$

We can then evaluate the difference $\Delta g(z) = g(z+1) - g(z)$

$$\Delta g(z) = -\frac{\rho_k}{\eta_k}\eta_k^{z-z_k}\left[(1-\eta_k)(g_{k-1} - (h+b_k)z_k) + \eta_k(h+b_k)\right] + h + b_{k+1}$$

which is nondecreasing in z. The cost $g(z)$ is convex and its minimum is attained at $\min\{z \in \mathbb{R} | \Delta g(z) > 0\}$, that is at z where,

$$z = z_k + \frac{\ln \frac{\eta_k}{\rho_k} \frac{h+b_{k+1}}{\eta_k(h+b_k)+(1-\eta_k)(g_{k-1}-(h+b_k)z_k)}}{\ln \eta_k} = z_{k+1} \tag{20}$$

We do not round off z in order to keep track of the production information. Using the value of z_{k+1} and (19), a direct computation leads to the expression of g_{k+1}

$$\begin{aligned} g_k &= \frac{\rho_k}{\eta_k}\left[(\frac{\eta_k}{\rho_k}z_{k+1} - \frac{\eta_k}{1-\eta_k})(h+b_{k+1}) + \left(g_{k-1} - (z_k - \frac{\eta_k}{1-\eta_k})(h+b_k)\right)\eta_k^{z_{k+1}-z_k}\right] \\ &= (h+b_{k+1})\left(z_{k+1} + \frac{1-\rho_k}{1-\eta_k}\right) \end{aligned}$$

If we replace g_k by its value in equation (20), it gives

$$z_{k+1} = z_k + \log_{\eta_k} \frac{\eta_k(h+b_{k+1})}{\rho_k(h+b_k)\left[\eta_k + (1-\eta_k)\frac{1-\rho_{k-1}}{1-\eta_{k-1}}\right]}$$

When we have obtained all the work rationing levels, we do

$$z_k = \frac{\lfloor r z_k + 1 \rfloor}{r}$$

in order to have $z_k \in \mathbb{N}_r$. We initialize the algorithm with $z_1 = 1 - 1/r$ (satisfy, whenever possible, Class 1 demands) and $\rho_0 = 1$.

Bios

Jean-Philippe Gayon is assistant professor in the Department of Industrial Engineering at INPG (Grenoble, France). He has a Ph.D. degree from ECP (Paris, France). His research interests are in stochastic modeling with applications to make-to-stock queues and dynamic pricing.

Francis de Véricourt is Associate Professor of Operation Management at the Fuqua School of Business at Duke University. Prior to joining the faculty of Duke University, he was a visiting researcher at MIT. His main research interest is in the optimization of supply chains and service delivery systems. He has published articles in leading academic and professional journals such as *Management Science*, *Operations Research*, and *IIE Transactions*. He is an Associate Editor for *Operations Research* and *IIE Transactions*. Francis is a native of Paris, France, where he graduated.

Fikri Karaesmen is Associate Professor in the Department of Industrial Engineering at Koç University (Istanbul, Turkey). He has a B.S. degree from METU (Ankara, Turkey) and a Ph.D. degree from Northeastern University (Boston, MA, USA). His research interests are in stochastic models of production and inventory systems and service systems. His papers have appeared in *IIE Transactions*, *Management Science*, *Manufacturing & Service Operations Management*, *Operations Research* and other journals. He is an Associate Editor for *4OR*, *IEEE Transactions on Automation Sciences and Engineering* and *IIE Transactions*.