IIE Transactions (2009) **41**, 1096–1109
Copyright © "IIE"
ISSN: 0740-817X print / 1545-8830 online
DOI: 10.1080/07408170902800279

Stock rationing in an M/E_r/1 multi-class make-to-stock queue with backorders

JEAN-PHILIPPE GAYON[1], FRANCIS DE VÉRICOURT[2] and FIKRI KARAESMEN[3,*]

[1]*Laboratoire G-SCOP, Grenoble INP, 46 avenue Félix Viallet, 38 031 Grenoble Cedex, France*
E-mail: jean-philippe.gayon@g-scop.fr
[2]*European School of Management and Technology, Schlossplatz 1, 10178 Berlin, Germany*
E-mail: francis.devericourt@esmt.org
[3]*Department of Industrial Engineering, Koç University, 34450, Sariyer, Istanbul, Turkey*
E-mail: fkaraesmen@ku.edu.tr

Received August 2006 and accepted December 2008

A model of a single-item make-to-stock production system is presented. 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. It is shown that certain structural properties of optimal stock and capacity allocation policies exist for the case where production may be interrupted and restarted. Also, a complete characterization of the optimal policy in the case of uninterrupted production when excess production can be diverted to a salvage market is presented. A heuristic policy is developed and assessed based on the results obtained in the analysis. Finally the value of production status information and the effects of processing time variability are investigated.

[Supplemental materials are available for this article. Go to the publisher's online edition of *IIE Transactions* for the following free supplemental resource: Appendix]

Keywords: Inventory rationing, optimal control, queuing theory

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 arise in a number of supply chain settings. For instance, delayed product differentiation often results in maintaining a stock of generic components for multiple end-products (De Véricourt *et al.*, 2002). Spare parts inventory management, as in Deshpande *et al.* (2003) for instance, is another situation where inventory allocation is critical. The design of supply contracts in the settings with different retailers can also entail a stock allocation problem at the supplier (Cachon and Lariviere, 1999).

Stock and capacity allocation problems are very challenging and are sometimes 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 typically hard to implement. Indeed, the supplier needs to take many dimensions into account when deciding to allocate stock: The inventory level, the number of waiting demands in the system, the current status of the production process, etc. The complexity of such problems depends on the number of customers sharing the common stock (Ha, 1997a), and on the nature of the production cycle time (Ha, 2000).

In this paper, we consider the model of 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. Different customer classes generate different backorder penalties for the supplier. The objective is to find the stock and capacity allocation policies to minimize the expected discounted (or average) 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.

We provide a partial characterization of the optimal stock and production policy for the above described system which is an $M/E_r/1$ make-to-stock queue with backorders. While this characterization yields some basic properties of the optimal policy, the model turns out to be too

*Corresponding author

challenging for a finer characterization. In order to further enhance the understanding of this model, we then focus on a related problem where production cannot be interrupted but excess inventory can be diverted to an ample market at no cost. For this auxiliary problem, it is shown that the optimal stock and capacity allocation policy can be completely characterized: there exist work-storage thresholds for each class that determine how production and inventory should be allocated in a simple way. In addition, these threshold parameters are easily computable. To our knowledge, this is the first such characterization for a multi-dimensional make-to-stock queue problem with non-exponential production times. Finally, it turns out that similar threshold policies lead to extremely effective heuristics for the standard problem.

The ample salvage market that allows absorption of production excess can be considered an approximation for the original model where production is stopped whenever required. The approximate model is more amenable to analysis than the original model. Furthermore, the approximate model may be of interest in itself if the salvage market assumption is justified. One example of this may be the situation where the supplier can divert inventory to a speculative (spot) market. In recent years, speculative markets for non-commodity items have developed rapidly. For instance, Milner and Kouvelis (2007) mention that 80% of electronic component parts (e.g., memory chips) are sold through contract purchasing while the rest are diverted to a spot market. In particular, suppliers may still conduct their main business through long-term contracts with established customers but can also easily get rid of excess inventories in the speculative market (for which no backorder cost exists). The assumption that the system never stops working is also relevant when the production setup cost is very high. Gupta and Wang (2007) present a model motivated by an integrated steel mill where primary processes remain continuously operational but production has to be allocated between contract and spot (i.e., transactional) customers. Of course, such spot markets may manifest other complications such as different lead time requirements or fluctuating prices that are not taken into account in our approximate model.

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. Since then, there has been considerable research on similar systems under the assumption of exogenous lead times (uncapacitated replenishment) problems. Deshpande *et al.* (2003) present a brief review of this literature.

In the case of endogenous replenishment lead times or limited production capacity, queuing-based models provide a powerful framework which allows explicit modeling of 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 demand classes as introduced by Ha (1997a, 1997b) in the stock rationing context.

Rationing strategies also appear in inventory transshipment problems, which have attracted a lot of attention from researchers and practitioners recently. Zhao *et al.* (2008) characterize the structure of the optimal stock allocation and production policies for a problem with two make-to-stock queues. In this system, each processor primarily serves its own class of customers but is also allowed to serve the other class at an additional cost. Hu *et al.* (2008) study a similar problem in discrete time where production capacity in each period is uncertain. They characterize optimal transhipment and rationing policies.

Ha (1997b) 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. Huang and Iravani (2007) investigate rationing decisions for a two-echelon supply chain with batch ordering and characterize the optimal policy. Gayon *et al.* (2009) investigate a rationing problem with imperfect advance demand information. Cil *et al.* (2009) present some structural results for batch demand arrivals in the lost-sales case with exponential processing times.

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 (1997a) 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, 1997b) and De Véricourt *et al.* (2002) assume exponential processing times. Because of the memoryless property of the exponential distribution, the supplier does not need the current production status (i.e., elapsed processing time) in that case. Information technologies in real production systems, however, can provide constant access to information on the status of the production process which would be needed with non-exponential processing times. In this paper, we consider a multi-class $M/E_r/1$ make-to-stock queue (with an Erlang-r processing time). We assume the supplier knows the current stage (phase) of the Erlang distribution exactly. This allows us to model the information on the production status. In addition, Erlang distributions provide some flexibility in modeling the production process variability. De Véricourt *et al.* (2001) provide insights about 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 under the exponential distribution assumption.

To our knowledge Ha (2000) is the only paper that addresses optimality issues in a 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. This reduces the problem to a single-dimensional Markov Decision Process model. The optimal stock allocation policy is then fully characterized: for each customer class there exists a work–storage threshold level at which it is optimal to reject a demand of this class. More recently, Abouee-Mehrizi *et al.* (2008) investigate the rationing problem for an $M/G/1$ make-to-stock queue from a performance evaluation perspective and compare different policies.

Our model differs from that of Ha (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 the presence of waiting demands from different classes. This issue does not exist when demands are lost.

When the production surplus cannot be sold in a salvage market, we obtain some partial structural results for the optimal stock allocation policy. Although these results uncover certain useful properties of the optimal policy, they are not sharp enough to completely define this policy. Unfortunately, a full characterization seems intractable. The approaches that have been successful so far in analyzing optimal policies for make-to-stock queues are all based on the propagation of convexity properties by iterating on the value function. When the state space has more than one dimension (typically two), this approach always requires the introduction of modularity properties (see for instance Ha (1997a), De Véricourt *et al.* (2000, 2002) or Zhao *et al.* (2004)). It turns out that the optimal value function of the problem without a salvage market does not satisfy some of these modularity properties.

On the other hand, when the production surplus can be sold in an ample salvage market, we show that these modularity properties hold for the optimal control policy. In this case, the production decision is replaced by the simpler decision of diverting inventory to the salvage market. Our analysis of this problem follows a decomposition technique introduced by De Véricourt *et al.* (2002), which consists of relating an n-dimensional control problem to an $(n-1)$-dimensional subproblem and then iterating on the number of demand classes n. The adaptation of this double induction (on time and on the dimension of the problem) to our case necessitates many subtleties and adjustments, and the introduction of more complex modularity conditions. As a result, the analysis of an n-dimensional problem makes use of the optimal structures of the k-dimensional subproblems, $k < n$. In De Véricourt *et al.* (2002), on the other hand, the iteration is mainly based on a single $(n-1)$-dimensional subproblem.

More precisely, we show that the optimal allocation policy of the multi-class problem with a salvage market is characterized by n work-storage rationing thresholds corresponding to the n demand classes. The work-storage level is the total number of completed production stages that are required to produce the current on-hand inventory plus the work in progress. The optimal policy backorders an arriving demand when the current work-storage level is below or at the corresponding threshold. This characterization leads to the construction of a heuristic using a queueing-based approximation. It is observed in a numerical study that this heuristic is extremely effective for the problem without a salvage market. In fact, our numerical results lead us to conjecture that the original problem has the identical optimal policy structure as the salvage market problem. The solution of the salvage market problem therefore not only leads to insights on the optimal policy structure of the original problem but also yields accurate approximations. Finally, we investigate the effects of production status information and processing time variability.

In the next section, we introduce the models and formulate the stock rationing problems with or without a salvage market. Some properties of the optimal policy for the problem without a salvage market are presented in Section 3. The structure of the optimal policy for the system with a salvage market is then characterized in Section 4. Based on this result, we suggest a heuristic for the problem without a salvage market in Section 5. In Section 6, we evaluate the performance of this heuristic, compare the optimal policies of both models and investigate the effects of processing time variability. We conclude the paper in Section 7.

2. Model formulation

2.1. *The problem without a salvage market*

Consider a supplier who produces a single item at a production 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 be backordered. Items held in stock induce holding costs at rate h (per item per unit of time). It is assumed that there are no holding costs for the jobs in process. Demands of Class i, $1 \le i \le n$, 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 $\mathbf{b} = (b_1, \ldots, b_n)$ the vector of backorder costs and by $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_n)$ the vector of demand rates. We will also use the notations $\boldsymbol{\lambda}^k = (\lambda_1, \ldots, \lambda_k)$ and $\mathbf{b}^k = (b_1, \ldots, b_k)$ for $1 \le k \le n$.

The production process consists of r identical stages in series. The processing time of each stage is 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 modeled by a single server whose processing time has an r-stage Erlang distribution 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. This condition is crucial when an average cost criterion is considered.

When the processor becomes idle, the manager of the system must decide whether or not to continue production. When a production stage is completed and if there is at least one part on hand, he/she can choose to satisfy a backordered demand, if any. Finally, when the demand of a customer arrives to the system, the manager has to choose between satisfying it with the on-hand inventory or backordering 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:

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

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 completed ($i(t) = 3$). 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 $-x_i(t)$, $1 \le i \le n$, be the number of backorders of Class i, $1 \le i \le n$, at time t. Hence, we can exhaustively describe the system state with $\mathbf{x}(t) = (x_0(t), x_1(t), \ldots, x_n(t))$ and the state space is $S_n = \mathbb{N}_r \times (\mathbf{Z}^-)^n$, where $\mathbf{Z}^-$ represents the set of non-positive integers. We will also use the notation $\mathbf{x}^k = (x_0, \ldots, x_k)$ when $0 \le k \le n$.

A control policy describes the action to take at any time given the current state $\mathbf{x}(t)$. We restrict the analysis to stationary Markovian policies since the optimal policy is known to belong to this class (Puterman, 1994). Let $\mathbf{a}^\pi(\mathbf{x}) = (a_0^\pi(\mathbf{x}), \ldots, a_n^\pi(\mathbf{x}))$ be the control (action) corresponding to a policy π where $a_0^\pi(\mathbf{x})$ is the action taken when an event occurs (arrival or end of production):

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

Notice that, when $x_0 = 1 - 1/r$, there is no inventory ($s(t) = 0$) and $r - 1$ stages of production are completed ($i(t) = r - 1$). Thus, there remains one more stage of production to be performed before one item is available. This item can then be used to satisfy a backordered demand or to increase the inventory by one unit.

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

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

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

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

The objective is to find a control policy that minimizes the expected discounted costs over an infinite horizon. This problem will be denoted $\mathbf{P}_n(\mu, \boldsymbol{\lambda}, h, \mathbf{b}, r, \alpha)$. We will also be interested in the closely related average cost problem.

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

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

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

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

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

and $\mathbf{e}_i$, $0 \leq i \leq n$, is the ith unit vector. For example, $\mathbf{e}_1$ denotes the $(n+1)$-dimensional vector $(0, 1, 0, \ldots, 0)$. Operator T_0 corresponds to the 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:

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

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}, $0 \leq i, j \leq n+1$, 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 unless stated otherwise.

In what follows, we will frequently refer to the class with the highest backorder cost among all classes that have backordered demands. This class is given by the following function m:

$$\forall \mathbf{x} \in S_n, m(\mathbf{x}) = \begin{cases} \min_{i\in\{1,\ldots,n\}: x_i<0} (i) & \text{if } \exists i \in \{1,\ldots,n\}, x_i < 0, \\ n+1 & \text{otherwise.} \end{cases}$$

2.2. *The problem with a salvage market*

There are a number of situations where shutting down production may be costly and the excess inventory can be sold relatively easily. A typical example occurs when the supplier can sell the item through a spot market in addition to its main business with long-term customers who have specific contracts. This induces a slightly different stock rationing problem where production never stops but has to be allocated between inventory for regular customers and a lower-priority salvage market with ample demand. Motivated by such a setting, for this model, we assume that there exists a customer class with zero backorder cost and ample demand. We also assume that the system produces all the time.

Given a problem without a salvage market $\mathbf{P}_n(\mu, \boldsymbol{\lambda}, h, \mathbf{b}, r, \alpha)$, we denote the corresponding problem with a salvage market by $\tilde{\mathbf{P}}_n(\mu, \boldsymbol{\lambda}, h, \mathbf{b}, r, \alpha)$. For problem $\tilde{\mathbf{P}}_n$, the actions corresponding to a policy π are denoted by $\tilde{\mathbf{a}}^\pi(\mathbf{x})$ and are defined as in Equations (1) and (2). The only exception is when $\tilde{a}_0^\pi = n+1$ which corresponds to satisfying a demand from the salvage market (whereas $a_0^\pi = n+1$ corresponds to not producing in the problem without a salvage market). The salvage market will be referred to as the $(n+1)$th class of customers with zero backorder cost $b_{n+1} = 0$. The objective is to characterize the optimal policy which minimizes the expected discounted cost and the optimal value functions of problem $\tilde{\mathbf{P}}_n$ can similarly be shown to satisfy the following optimality equations:

$$\tilde{v}^*(\mathbf{x}) = c(\mathbf{x}) + r\mu \tilde{T}_0 \tilde{v}^*(\mathbf{x}) + \sum_{k=1}^n \lambda_k \tilde{T}_k \tilde{v}^*(\mathbf{x}),$$

where $\tilde{T}_k = T_k$, $1 \leq k \leq n$, and operator $\tilde{T}_0$ is

$$\tilde{T}_0 v(\mathbf{x}) = \begin{cases} \min[v(\mathbf{x}+\mathbf{e}_0/r), \min_{1\leq i\leq n+1: x_i<0} v(\mathbf{x}+\mathbf{e}_0/r-\mathbf{e}_0+\mathbf{e}_i)] \\ \quad \text{if } x_0 \geq 1-1/r, \\ v(\mathbf{x}+\mathbf{e}_0/r) \quad \text{if } x_0 < 1-1/r, \end{cases}$$

and $\mathbf{e}_{n+1} = \mathbf{0}$. Since the system is assumed to always produce, there is a term $v(\mathbf{x})$ in $T_0 v(\mathbf{x})$ that corresponds to the option of not producing in the problem without a salvage market. This term is not present in $\tilde{T}_0 v(\mathbf{x})$. On the other hand, the term $v(\mathbf{x}+\mathbf{e}_0/r-\mathbf{e}_0+\mathbf{e}_{n+1})$ in $\tilde{T}_0 v(\mathbf{x})$ does not appear in $T_0 v(\mathbf{x})$ and corresponds to the decision of selling the produced part on the salvage market.

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

$$\tilde{T}_0 v(\mathbf{x}) = \begin{cases} \min[v(\mathbf{w}+\mathbf{e}_0), \min_{1\leq i\leq n+1: x_i<0} v(\mathbf{w}+\mathbf{e}_i)] \\ \quad \text{if } x_0 \geq 1-1/r, \\ v(\mathbf{w}+\mathbf{e}_0) \quad \text{if } x_0 < 1-1/r. \end{cases}$$

We also define the operator $\tilde{T}$ such that $\tilde{T}v = c + r\mu \tilde{T}_0 v + \sum_{k=1}^n \lambda_k \tilde{T}_k v$.

Finally, the operators Δ_i and Δ_{ij} as well as the function $m(\mathbf{x})$ are still well defined. It should be noted that for problem $\tilde{\mathbf{P}}_n$, $m(\mathbf{x}) = n + 1$ designates the salvage market, whose corresponding backorder cost is zero.

3. A partial characterization of the optimal policy for the problem without a salvage market

3.1. *The single-class problem*

We start by studying the problem with a single demand class. When there is a single class of customers, the problem is to decide when to satisfy demands of Class 1 and when to produce. A simple sample path argument (not detailed) shows that it is always optimal to satisfy a Class 1 demand. Therefore, we cannot 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. Regardless of 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:

$$\Delta_0 v(\mathbf{x} + \mathbf{e}_0/r) \geq \Delta_0 v(\mathbf{x}). \tag{3}$$

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

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

The proof of this result and all subsequent proofs, unless stated otherwise, are included in the online Appendix.

Using value iteration and Proposition 1, we obtain that the optimal value function belongs to $\mathcal{V}_0$. As a result, the optimal policy is of threshold type: there exists a threshold 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. Proposition 2 establishes three 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. $\Delta_{ij} v(\mathbf{x}) \leq 0$ when $1 \leq i < j \leq n$.
2. $\Delta_{0j} v(\mathbf{x}) \leq \Delta_{0i} v(\mathbf{x})$ when $1 \leq i < j \leq n$.
3. $\Delta_{0j} v(\mathbf{x} - \mathbf{e}_j) \leq \Delta_{0i} v(\mathbf{x} - \mathbf{e}_i)$ when $1 \leq i < j \leq n$.

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

We obtain by value iteration that the optimal value function belongs to $\mathcal{U}_n$. The structural properties suggested by Proposition 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 establishes that it is better to satisfy Class i, the more expensive one. The second property implies that if increasing the inventory level 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 it is optimal to satisfy an arriving demand of Class j with a given on-hand inventory, it is also optimal to satisfy the arriving demands of more expensive classes at the same inventory level.

Even though Proposition 2 establishes basic properties on how to prioritize items, a sharper characterization of the optimal policy requires several additional properties which turn out to be difficult to establish by this approach. For instance, Proposition 2 implies that it is optimal to employ a strict priority rule for production allocation when all classes are backordered but it is not clear what rule to follow in deciding when to increase inventory instead of satisfying a Class i backorder. Since such rules may in general depend on the entire state vector, additional effort is needed to completely define an optimal policy.

In order to outline some of the technical challenges in the multi-class problem, let us describe some analogies with the single-class problem. Before that, we need to define the notions of submodularity and supermodularity. The value function v is supermodular in vectors $(\mathbf{u}, \mathbf{v})$, with $\mathbf{u} \neq \mathbf{v}$, if for all $\mathbf{x} \in S_n$ such that $\mathbf{x} + \mathbf{v}$, $\mathbf{x} + \mathbf{u}$ and $\mathbf{x} + \mathbf{u} + \mathbf{v}$ are in S_n, we have:

$$v(\mathbf{x} + \mathbf{u} + \mathbf{v}) + v(\mathbf{x}) \geq v(\mathbf{x} + \mathbf{u}) + v(\mathbf{x} + \mathbf{v}).$$

The definition of submodularity is the same but with opposite inequality (see Veatch and Wein (1992) for more on these notions). For the single-class problem, Equation (3) implies that v is supermodular in $(\mathbf{e}_0, \mathbf{e}_0/r)$ (i.e., in the inventory level, s, and in the production status, i). In order to generalize Proposition 1 to the multi-dimensional problem, more modularity properties are required to ensure that Equation (3) can be propagated.

For instance, with two demand classes, a first step to this generalization would be to show that v is supermodular in $(\mathbf{e}_i, \mathbf{e}_0/r)$, i.e., the marginal benefit of continuing production increases in the number of waiting demands. Unfortunately, the optimal value function does not systematically satisfy this property. For example, a numerical study shows that the optimal value function for $x_0 = 9.5$, $x_2 = -1$, $r = 2, \mu = 1, \lambda_1 = 0.3, \lambda_2 = 0.3, h = 0.01, b_1 = 10, b_2 = 1$, $\alpha = 0.01$, is not supermodular in $(\mathbf{e}_i, \mathbf{e}_0/r)$. As a result, a better characterization of the optimal policy seems difficult in the multi-class case using this approach. In the next section, we will prove that the above modularity property holds when a salvage market is considered (see 3(c) in Definition 4).

4. Characterization of the optimal policy for the problem with a salvage market

4.1. *Preliminary results*

We now extend Proposition 2 to the case with an ample salvage market.

Definition 2. Let $\tilde{\mathcal{U}}_n$ be a set of functions such that $v \in \tilde{\mathcal{U}}_n$ if and only if:

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

Property 1 of $\tilde{\mathcal{U}}_n$ applied when $j = n+1$ implies that it is better to satisfy backorders of Class i, $1 \leq i \leq n$ rather than demands from the salvage market. It should be noted that we implicitly take $\Delta_{i(n+1)} v(\mathbf{x}) = \Delta_i v(\mathbf{x})$ for $0 \leq i \leq n$.

Proposition 3. *If* $v \in \tilde{\mathcal{U}}_n$*, then* $\tilde{T} v \in \tilde{\mathcal{U}}_n$

The proof of Proposition 3 is similar to the proof of Proposition 2. The only difference is in showing that $\Delta_{ij} \tilde{T}_0 v(\mathbf{x}) \leq 0$ but the arguments are the same.

A direct application of value iteration implies that the optimal value function also belongs to $\tilde{\mathcal{U}}_n$. A useful property is that for $v \in \tilde{\mathcal{U}}_n$, the operators satisfy:

$$\tilde{T}_0 v(\mathbf{x}) = v(\mathbf{w} + \mathbf{e}_i) + \min\left[0, \Delta_{0i} v(\mathbf{w})\right] \text{ with } i = m(\mathbf{x}),$$
$$\tilde{T}_k v(\mathbf{x}) = v(\mathbf{x} - \mathbf{e}_0) + \min\left[0, \Delta_{0k} v(\mathbf{x} - \mathbf{e}_k - \mathbf{e}_0)\right] \quad \text{for } 1 \leq k \leq n.$$

This implies that the optimal actions are entirely determined by the sign of $\Delta_{0i}\tilde{v}^*$, for $1 \leq i \leq n+1$.

4.2. *Work-storage rationing policies*

Consider a particular class of policies entirely described 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 (or equal) z_k. 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 $\mathbf{x}$). This class of policies will be referred to as Work-Storage Rationing (WR) policies. In a WR policy, the decisions depend on the current production status (i.e., stage of production) in addition to the current inventory level. Definition 3 gives a formal description of WR policies.

Definition 3. A WR policy π is characterized by a $(n+1)$-dimensional rationing vector $\mathbf{z} = (z_1, \ldots, z_{n+1})$ where $z_1 = 1 - 1/r \leq z_2 \leq \ldots \leq z_{n+1}$ such that:

$$\tilde{a}_0^{\pi}(\mathbf{x}) = \begin{cases} 0 & \text{if } x_0 < z_i \quad \text{and } i = m(\mathbf{x}), \\ i & \text{if } x_0 \geq z_i \quad \text{and } i = m(\mathbf{x}). \end{cases}$$
$$\tilde{a}_k^{\pi}(\mathbf{x}) = \begin{cases} k & \text{if } x_0 \leq z_k, \\ 0 & \text{if } x_0 > z_k \text{ and } m(\mathbf{x}) \geq k. \end{cases}$$

In a WR policy, demands of Class 1 are always satisfied when inventory is available, since $z_1 = 1 - 1/r$. According to such a policy and assuming that $\mathbf{x}(t=0) = \mathbf{0}$, the recurrent region of the space is $[\mathbf{x} \in S_n | x_0 \leq z_{m(\mathbf{x})}]$. The definition leaves the policy unspecified for $x_0 > z_k^*$ and $m(\mathbf{x}) < k$. A precise definition for these states is not necessary if the initial state is $\mathbf{0}$.

We claim that the optimal policy is a WR policy. To prove the claim, we will argue inductively on the number of customer classes. The construction of the proof is based on the following key property: The optimal value function of an n-class problem is closely related to the optimal value function of a k-class problem, in the region of the state space where $x_0 \leq z_k^*$. In particular, it will be shown that for this region, the optimal actions do not depend on the demands of classes strictly greater than k. The 0-class subproblem corresponds to a problem with the salvage market only and no other customer class. The transformation, which relates an n-class problem $\tilde{\mathbf{P}}_n(\mu, \boldsymbol{\lambda}, h, \mathbf{b}, r, \alpha)$ to a $(n-1)$-class subproblem, is based on the decomposition of the cost function $c(\mathbf{x})$:

$$c(\mathbf{x}) = c_{n-1}(\mathbf{x}^{n-1}) - b_n \left(\lfloor x_0 \rfloor + \sum_{i=1}^{n} x_i \right),$$

where c_{n-1} is the cost function of the $(n-1)$-class subproblem $\tilde{\mathbf{P}}_{n-1}(\mu, \boldsymbol{\lambda}^{n-1}, h + b_n, \mathbf{b}^{n-1} - b_n \mathbf{1}_{n-1}, r, \alpha)$ and $\mathbf{1}_{n-1} = \sum_{i=1}^{n-1} \mathbf{e}_i$. We can iterate this decomposition for $k < n$:

$$c_k(\mathbf{x}^k) = c_{k-1}(\mathbf{x}^{k-1}) - (b_k - b_{k+1}) \left(\lfloor x_0 \rfloor + \sum_{i=1}^{k} x_i \right).$$

It follows that:

$$\begin{cases} c_0(\mathbf{x}^0) = (h + b_1)\lfloor x_0 \rfloor \\ c_k(\mathbf{x}^k) = (h + b_{k+1})\lfloor x_0 \rfloor - \sum_{i=1}^{k} (b_i - b_{k+1}) x_i \\ \qquad \text{for } 1 \leq k \leq n-1. \end{cases} \tag{4}$$

Therefore, c_k is the cost of a k-class problem $\tilde{\mathbf{P}}_k(\mu, \boldsymbol{\lambda}^k, h + b_{k+1}, \mathbf{b}^k - b_{k+1}\mathbf{1}_k, r)$.

Hence, for any n-class problem we have defined n subproblems with the number of customer classes equal to $0, 1, \ldots, n-1$ respectively. We denote by v_k^* (resp. π_k^*)

the optimal value function (resp. optimal policy) of the $k - class$ subproblem.

We will show that the optimal policy is a WR policy by iterating on the number of classes. To start the induction, assume that the optimal policy of any $(n-1)$-class problem is a WR policy. In particular the optimal policy π^*_{n-1} of the $(n-1)$-class subproblem defined above is a WR policy, with $(z^*_1, \ldots, z^*_n)$ its rationing vector.

Based on the optimal value function v^*_{n-1} of the $(n-1)$-class subproblem, we introduce $\tilde{\mathcal{V}}_n$, a structured set of value functions. We will again employ value iteration to show that the optimal value function of the n-class problem belongs to $\tilde{\mathcal{V}}_n$. In the following definition, $[\mathbf{x}]_0$ designates the first component of vector $\mathbf{x}$.

Definition 4. Let $\tilde{\mathcal{V}}_n \subset \tilde{\mathcal{U}}_n$ such that $v \in \tilde{\mathcal{V}}_n$ if and only if:

1. $\Delta_{ij} v(\mathbf{x}) = \Delta_{ij} v^*_{n-1}(\mathbf{x})$ for $0 \leq i < j \leq n$ and $[\mathbf{x} + \mathbf{e}_i]_0 \leq z^*_n$.
2. $\Delta_{0i} v(\mathbf{x}) \geq 0$ for $i = m(\mathbf{x}) < n+1$ and $x_0 > z_n{}^* - 1$.
3. For $x_0 > z^*_n - 1$ and $m(\mathbf{x}) \geq n$
   (a) $\Delta_{0n} v(\mathbf{x} + \mathbf{e}_0/r) \geq \Delta_{0n} v(\mathbf{x})$;
   (b) $\Delta_{0n} v(\mathbf{x} + \mathbf{e}_n) \leq \Delta_{0n} v(\mathbf{x})$;
   (c) $\Delta_n v(\mathbf{x} + \mathbf{e}_0/r) \geq \Delta_n v(\mathbf{x})$;
   (d) $\Delta_0 v(\mathbf{x} + \mathbf{e}_0/r) \geq \Delta_0 v(\mathbf{x})$;
   (e) $\Delta_n v(\mathbf{x} + \mathbf{e}_n) \geq \Delta_n v(\mathbf{x})$.

Condition 1 of the previous definition links any function $v \in \tilde{\mathcal{V}}_n$ to the optimal value function of the $(n-1)$-class subproblem, when the work-storage level is below the last optimal rationing level. When the work-storage level is above the optimal rationing level, Condition 2 states that it is always better to satisfy a waiting demand of the class with the highest backorder cost rather than to increase the inventory level.

Condition 3 describes monotonicity properties that the value functions must satisfy in the directions of $\mathbf{e}_0$ and $\mathbf{e}_n$, when $x_k = 0$ for $0 < k < n$. These conditions guarantee in turn that the optimal rationing decisions for the demand class with the lowest backorder cost can be described with a monotone switching curve, when no demands of other classes are waiting. More formally, Condition 3 states that $\Delta_{0n} v$ is increasing in x_0/r and decreasing in x_n, $\Delta_0 v$ is increasing in x_0/r and $\Delta_n v$ is increasing in both x_0/r and x_n. Condition 3 may be also interpreted in terms of submodularity and supermodularity. Conditions 3(a), 3(c) and 3(d) establish that v is supermodular in $(\mathbf{e}_0 - \mathbf{e}_n, \mathbf{e}_0/r)$, $(\mathbf{e}_n, \mathbf{e}_0/r)$ and $(\mathbf{e}_0, \mathbf{e}_0/r)$ respectively. Condition 3(b) establishes that v is submodular in $(\mathbf{e}_0 - \mathbf{e}_n, \mathbf{e}_n)$ and Condition 3(e) establishes that v is convex in x_n. It can be shown that Conditions 3(a) and 3(c) imply Condition 3(d) and that Conditions 3(b) and 3(c) imply Condition 3(e). Finally, Condition 4 implies that demands from the salvage market should be satisfied if and only if there are no other backordered demands and if $x_0 \geq z_{n+1}$.

Note that for the lost-sales case studied by Ha (2000), Condition 3(d) is the only modularity property that v needs to verify, since that model has a single-dimensional state space. In our case, the multi-dimensional aspect of the problem requires the value function to satisfy more conditions that are also less typical.

We will next establish that, for any number of customer classes, the optimal value function belongs to $\tilde{\mathcal{V}}_k$ and therefore the optimal policy is a WR policy. We denote by $P(n)$ this property.

Definition 5. We say that property $P(n)$ is true if for all k-class subproblems, $k \leq n$:

1. the optimal value function v^*_k belongs to $\tilde{\mathcal{V}}_k$,
2. the optimal policy is a WR policy.

We prove in the online Appendix that $P(0)$ is true. For the 0-class problem, the WR policy with $z^*_1 = 1 - 1/r$ is optimal. It is hence optimal to allocate all the produced items to the salvage market. If we assume that $P(n-1)$ is true, then $\tilde{\mathcal{V}}_n$ is well defined and not empty since v^*_{n-1}, the optimal policy of the $(n-1)$-class problem, belongs to $\tilde{\mathcal{V}}_n$.

Lemma 1. *If $P(n-1)$ is true then*:

1. $v \in \tilde{\mathcal{V}}_n$ *implies that* $\tilde{T}v \in \tilde{\mathcal{V}}_n$.
2. $v^*_n \in \tilde{\mathcal{V}}_n$.

Lemma 2. *Let* $z^*_{n+1} = \min[x_0 \in I\!N_r \,|\, \Delta_0 v^*_n(\mathbf{w}) > 0$ *and* $m(\mathbf{x}) = n+1]$. *If $P(n-1)$ is true then the WR policy with rationing vector $(z^*_1, \ldots, z^*_{n+1})$ is optimal for the n-class problem.*

Lemmas 1 and 2 imply that if $P(n-1)$ is true then $P(n)$ is true. Finally $P(n)$ is true for all n. We can now state our main result.

Theorem 1. *For all n-class problems $\tilde{\mathbf{P}}_n(\mu, \boldsymbol{\lambda}^n, h, \mathbf{b}^n, r, \alpha)$, the optimal policy is a WR policy with rationing vector $(z^*_1, \ldots, z^*_{n+1})$. In addition, for $k < n$, the projection $(z^*_1, \ldots, z^*_{k+1})$ is the optimal rationing vector of the k-class subproblem $\tilde{\mathbf{P}}_k(\mu, \boldsymbol{\lambda}^k, h + b_{k+1}, \mathbf{b}^k - b_{k+1}\mathbf{1}_k, r, \alpha)$.*

From a technical point of view, this result can also be interpreted in terms of switching surfaces dividing the state space into different regions for which the optimal action is constant. Under this interpretation, Theorem 1 indicates that all switching surfaces are vertical planes defined by the equations $x_0 = z^*_k$. In particular, our result is consistent with the corresponding characterization of De Vricourt *et al.* (2002). This simplifies the policy structure in a significant manner because the precise description of a generic switching surface may require infinitely many parameters whereas the vertical line is described by a single parameter. The entire optimal policy is then completely characterized by n parameters in the sense that given only the values of these parameters the optimal policy can be implemented.

Finally, it should be noted that so far we concentrated on the discounted cost problem. Fortunately, there are existing results for controlled queueing systems that ensure that the average cost problem retains the same optimal policy structure as the discounted problem (Weber and Stidham, 1987). Our model satisfies these conditions and the optimal policy structure is retained for the average cost problem. The average cost case is attractive since the optimal average cost does not depend on the initial conditions and the optimal policy parameters do not depend on the discount factor selected. This facilitates comparisons and interpretations (see Ha (1997a, 2000) and De Véricourt *et al.* (2002) for similar comparisons.)

Remark 1. In particular, a critical condition of Weber and Stidham for the passage from the discounted cost to the average cost is the existence of a stationary policy with a finite average cost. This condition is satisfied in our case by taking a strict priority allocation policy with a base-stock level of zero. The resulting system is related to an $M/E_r/1$ queue with pre-emptive priorities. Recall that since allocation takes place at the end of processing completion, a higher-class customer can always preempt a service process that was initiated by a lower-class customer. In addition, a customer that pre-empts, only experiences the remaining part of the on-going service time. The resulting queueing system generates stochastically smaller queue lengths than the standard strict priority system. Since it was assumed that $\rho < 1$, this system generates a finite number of expected backorders and therefore a finite average cost.

5. Heuristic policy for the problem without a salvage market

The results of Section 3 partially uncover the priority properties of production and stock allocation policies but do not suggest a precise policy. In Section 4, we show that a WR policy is optimal for the problem with a salvage market. Based on this fact, we propose, for the problem without a salvage market, a modification of the WR policy as a heuristic. In particular, we define a Modified Work-Storage Rationing (MWR) policy as a WR policy where the salvage market rationing level is replaced by an integer-valued base-stock level. When there are no backordered demands ($m(\mathbf{x}) = n + 1$), this modified policy recommends producing if the inventory level is strictly smaller than the base-stock level, and not producing otherwise. All the other controls are the same as those in the original WR policy.

The MWR policy is clearly plausible since it satisfies the properties established in Section 3. However, more importantly, our numerical results suggest that it may even be the optimal policy. In particular, we are not able to find any counterexamples to the optimality of the MWR policy in a fairly exhaustive numerical study. In addition, if $(z_1^*, \ldots, z_{n+1}^*)$ is the optimal rationing vector of the problem with salvage market $\tilde{\mathbf{P}}_n(\mu, \boldsymbol{\lambda}, h, \mathbf{b}, r, \alpha)$, we systematically obtained numerically that $(z_1^*, \ldots, z_n^*, \lfloor z_{n+1}^* \rfloor)$ is the optimal rationing vector of the problem without a salvage market $\mathbf{P}_n(\mu, \boldsymbol{\lambda}, h, \mathbf{b}, r, \alpha)$. These results can be explained by the fact that both systems are governed by very similar equations.

In principle, the optimal parameters of the MWR policy can be obtained algorithmically by the characterization in Theorem 1. However, as the number of classes increases, it rapidly becomes difficult to obtain the parameters. In the following, we first develop a heuristic to compute the optimal policy parameters of the WR policy for the problem with a salvage market and then suggest another heuristic to compute the optimal policy parameters of the MWR policy for the problem without a salvage market.

To compute the policy parameters of the heuristic MWR policy, we adapt the exact algorithm presented by De Véricourt *et al.* (2002) for an $M/M/1$ make-to-stock queue. The main procedure exploits the relationship between the k-class problem and the $(k-1)$-class problem in the same manner as in De Véricourt *et al.* (2002). However, the Erlang distribution brings several computational complications. In order to obtain a tractable iterative approach, we employ a geometric approximation as in Tijms (1994) and Karaesmen *et al.* (2003) for the $M/E_r/1$ queue. In addition, the optimal thresholds should correspond to work-storage levels, which requires keeping track of the phases of the distribution. The details are provided in the online Appendix.

To outline the main points of the approach, we construct a method to successively compute the rationing levels $z_1, \ldots, z_{k+1}$. When the rationing vector $(z_1, \ldots, z_k)$ 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 to compute the parameters of the WR policy. The full justification for this algorithm is given in the online Appendix.

Heuristic 1. *Consider an n-class problem. Construct the sequences* ρ_k, η_k *and* $\tilde{z}_k$ *as follows:*
Initialize $\tilde{z}_1 = 1 - 1/r$, $\rho_1 = 1$, $\eta_0 = 0$.
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*

$$(r/(r + \rho_k(1 - 1/\eta_k)))^r = 1/\eta_k$$

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

Table 1. Problem instances

	Number of classes (n)		
	1	*2*	*3*
$r \in$	{2, 3, 4, 5, 10, 15, 20}	{2, 3, 5, 10}	{2, 3, 5}
$\rho \in$	{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9}	{0.2, 0.4, 0.6, 0.8}	{0.4, 0.6, 0.8}
$h \in$	{0.0001, 0.0005, 0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 5}	{0.001, 0.01, 0.1, 1}	{0.01, 0.1, 1}
$b_1/b_2 \in$	NA	{1, 5, 25, 100}	{2, 5, 10}
$b_2/b_3 \in$	NA	NA	{2, 5, 10}
$\lambda_1/\lambda_2 \in$	NA	{0.2, 1, 5}	{1}
$\lambda_2/\lambda_3 \in$	NA	NA	{1}
Number of instances	567	768	243

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

$$z_1 = \tilde{z}_1$$
$$z_k = \max\{1 - 1/r, \lfloor r\tilde{z}_k + 1\rfloor/r\} \quad \text{for } k = 2, 3 \ldots, n$$
$$z_{n+1} = \lfloor r\tilde{z}_{n+1} + 1\rfloor/r.$$

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.

6. Numerical results

In this section, we focus on the average cost problems and consider the instances summarized in Table 1. For example, we consider for the single-class problem $567 = 7 \times 9 \times 9$ instances corresponding to the combinations of: $r \in \{2, 3, 4, 5, 10, 15, 20\}$, $\rho \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$, $h \in \{0.0001, 0.0005, 0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 1, 5\}$. Without loss of generality, we set $\mu = 1$ and $b_n = 1$ for all instances where n is the number of customer classes.

6.1. *Performance of the heuristic policy*

In order to evaluate the performance of the MWR heuristic proposed in Section 5, we compare the average cost g^* of the optimal policy of the problem with the average cost g^{H} of the heuristic policy. The average costs are computed numerically by value iteration. We denote by $\Delta g = (g^{\mathrm{H}} - g^*)/g^*$, the relative cost increase when using the heuristic policy instead of the optimal policy. Another useful benchmark for the heuristic performance is the rationing level differences $\Delta z_i = z_i - z_i^*$.

We evaluated the performance of the heuristic with one, two and three classes of customers for the 1578 instances described in Table 1.

In Table 2, we summarize our main results. For one, two and three classes of customers, the heuristic finds rationing levels and base-stock levels with a maximum error of 1 unit for all the 1578 instances tested. The relative cost increase for using the heuristic policy is always less than 2% when the base-stock level is higher than ten. However, when the base-stock level is low, a small approximation error in the base-stock level may lead to a magnified percentage error in terms of the average cost. This situation occurs whenever the holding cost h is very high or the utilization rate ρ is very low. Although this is a limitation of the heuristic, it may be argued that these situations are not the most relevant for stock rationing, and that when the base-stock level is low, the few rationing alternatives can all be individually evaluated and compared.

6.2. *Value of production status information*

The optimal WR policy uses information on the production status since the optimal rationing levels depend on work-storage levels in general. This may not be feasible if the production status cannot be observed or may not be desirable due to increased complexity. Ha (2000) has investigated the cost penalty for ignoring information on the current production for a multi-class $M/E_r/1$ make-to-stock queue with lost sales. In this subsection, we perform a corresponding investigation for the backorder case. We define a Critical Stock Level (CSL) policy with parameters $R_2, \ldots, R_n, S$ as follows:

Table 2. Performance of the heuristic policy

	Number of classes		
Performance criteria	*1*	*2*	*3*
Percent of instances for which the heuristic policy is optimal	85.9	52.3	33.7
Percent of instances for which $\lvert\Delta z_i\rvert \leq 1$, for $i \leq n+1$	100	100	100
Maximum Δg (%)	47	139	47
Maximum Δg if $z_{n+1}^* \geq 10$ (%)	1.4	2.0	1.0

Table 3. Cost increase for using CSL policy instead of the optimal policy

	r			
Costs	*2*	*3*	*5*	*10*
Average cost increase for two-class instances (%)	0.15	0.22	0.27	0.35
Maximum cost increase for two-class instances (%)	4.06	3.92	5.5	5.4
Average cost increase for three-class instances (%)	0.23	0.42	0.57	Not evaluated
Maximum cost increase for three-class instances (%)	2.27	2.75	4.00	Not evaluated

1. Produce if and only if the inventory level is strictly smaller than S.
2. Satisfy arriving (or backordered) demands of class i if and only if $x > R_i$ and $m(\mathbf{x}) \geq i$.

A CSL policy with parameters $R_2, \ldots, R_n, S$ is equivalent to an MWR policy with parameters $R_2 - 1/r, \ldots, R_n - 1/r, S$. It is computationally difficult to find the optimal CSL policy since it requires the calculation of the value function of each possible CSL policy. To avoid this difficulty, we only evaluate CSL policies with critical stock levels that are adjacent to the optimal work-storage rationing levels. The best candidate policy might not be optimal among CSL policies but provides a good upper bound. We denote by g^{CSL} the average cost of the best candidate policy and by $\Delta g^{\mathrm{CSL}} = (g^{\mathrm{CSL}} - g^*)/g^*$ the relative cost increase for ignoring production status information.

For the 768 two-class instances (respectively 243 three-class instances) presented in Table 1, we compute Δg^{CSL} and report its average and maximum values in Table 3. We observe that the CSL policy performs very well which confirms the results obtained by Ha (2000) for an $M/E_r/1$ make-to-stock queue with lost sales.

6.3. *Influence of processing time variability on optimal average costs*

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

We start by investigating the effect of c_v^2 on the optimal average cost. Figure 1 represents for a two-class problem the optimal cost g^* as a function of c_v^2 for different values of b_1.

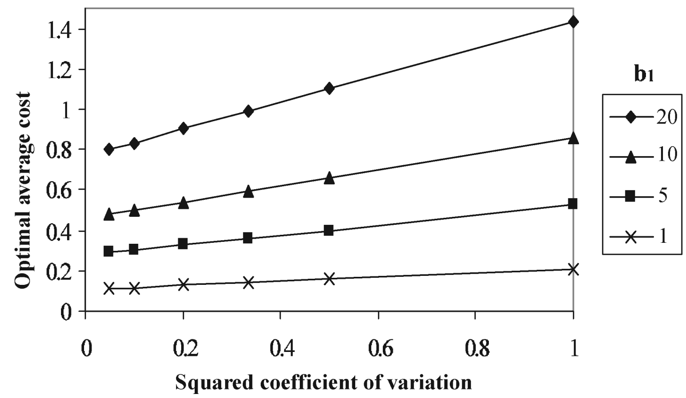

Fig. 1. Impact of c_v^2 on g^* for different values of b_1/b_2 ($\lambda_1 = \lambda_2 = 1$, $\mu = 1$, $h = 0.055$ and $b_2 = 1$).

The interesting observation in Fig. 1 is that the impact of the squared coefficient of variation on the average cost appears to be approximately 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 this analogy non-trivial. The linear effect has also been observed on examples where other parameters are varied 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, in order to evaluate the average cost when the 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 this 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

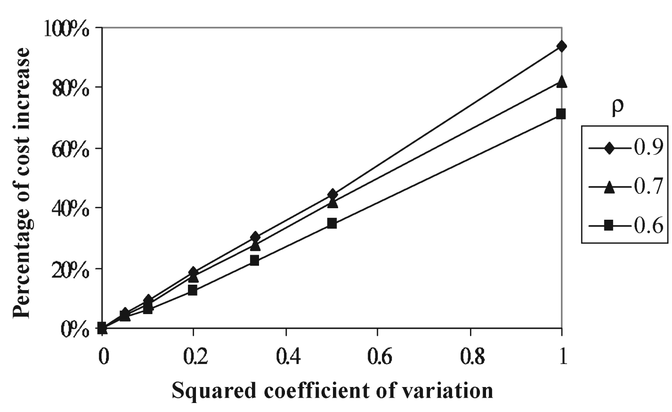

Fig. 2. Cost penalty for variability ($\lambda_1 = \lambda_2$, $\mu = 1$, $h = 0.055$, $b_1 = 10$, $b_2 = 1$).

the relative cost increase $\Delta g^*(c_v) = (g^*(c_v) - g^*(0))/g^*(0)$ as the coefficient of variation changes, 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.

6.4. *Influence of production variability on the benefits of rationing*

Ha (1997b) and De Véricourt *et al.* (2001) have studied the benefits of rationing for multi-class $M/M/1$ make-to-stock queues with lost sales and backorders respectively. To complement these investigations, in this subsection we explore the influence of production variability on these benefits in the context of a multi-class $M/E_r/1$ make-to-stock queue with backorders.

More precisely, we compare the average cost g^* of the optimal policy with the average cost g^{FCFS} of the optimal First-Come First-Served (FCFS) policy. Basically, the FCFS policy serves customers in chronological order of arrival and produces if and only if the stock level x_0 is strictly below the base-stock level z. The following result will help us to evaluate g^{FCFS}.

Proposition 4. *The average costs of the two following policies are equal:*

1. *The FCFS policy, with base-stock level z, for an n-class $M/G/1$ make-to-stock queue with arrival rates $\lambda_1, \ldots, \lambda_n$, independent and identically distributed (i.i.d) service times (distribution function F), holding cost h, backorder costs $b_1, \ldots, b_n$.*
2. *The FCFS policy, with base-stock level z, for a single-class $M/G/1$ make-to-stock queue with arrival rate $\hat{\lambda}$, i.i.d service times (distribution function F), holding cost h and backorder cost $\hat{b}$ where:*

$$\hat{\lambda} = \sum_{i=1}^{n} \lambda_i, \quad \hat{b} = \frac{\sum_{i=1}^{n} \lambda_i b_i}{\sum_{i=1}^{n} \lambda_i}.$$

Table 4. Cost increase for using a FCFS policy instead of the optimal policy

	r			
Costs	*2*	*3*	*5*	*10*
Average cost increase for two-class instances (%)	14.7	14.1	13.8	13.5
Maximum cost increase for two-class instances (%)	158.2	154.4	151.6	149.7
Average cost increase for three-class instances (%)	23.7	23.2	22.8	Not evaluated
Maximum cost increase for three-class instances (%)	115.2	113.3	110.1	Not evaluated

Table 5. Cost increase for using a SP policy instead of the optimal policy

	r			
Costs	*2*	*3*	*5*	*10*
Average cost increase for two-class instances (%)	6.92	6.51	6.28	6.14
Maximum cost increase for two-class instances (%)	73.8	72.9	70.4	69.2
Average cost increase for three-class instances (%)	8.27	7.93	7.70	Not evaluated
Maximum cost increase for three-class instances (%)	51.3	49.8	47.6	Not evaluated

Based on Proposition 4, we can compute g^{FCFS} by evaluating the optimal policy of a single-class problem. In Table 4, we report the average and maximum cost increase ($\Delta g^{\text{FCFS}} = (g^{\text{FCFS}} - g^*)/g^*$) for the instances presented in Section 5.

The second policy that is considered is a Strict Priority (SP) policy, which was used as a benchmark by De Véricourt *et al.* (2001). A SP policy does not reserve any inventory for future demands but allocates production to backorders according to a pre-emptive priority rule. Thus, before constituting any inventory, all backorders are cleared employing a priority rule starting from the most expensive class. Table 5 presents the average and maximum cost increase ($\Delta g^{\text{SP}} = (g^{\text{SP}} - g^*)/g^*$) for the instances presented in Section 5.

It can be observed from both Tables 4 and 5 that variability, summarized by the parameter r, does not seem to have a big impact on the percentage suboptimality. There is a slight decrease in the percentage suboptimality as r increases (indicating decreased variability) possibly explained by the corresponding overall decrease in costs. This is an interesting observation since r impacts the overall costs and policy parameters in a significant manner as was seen in Section 6.3. The conclusion is that variability appears to have a parallel impact on all considered policies and the percentage difference stays relatively stable.

7. Conclusion and future research

In this paper, we 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. In addition to the standard problem, we also investigated an auxiliary problem where the manager can sell the production surplus in an ample salvage market.

For the problem with a salvage market, we provided a full characterization of the optimal policy by exploiting the nested structure. The resulting structure of the optimal policy is fairly intuitive and easy to implement. Moreover, we 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.

For the problem without a salvage market, we fully characterized the optimal policy for a single-class problem and provided a partial characterization for the multi-class case. Based on the findings of the problem with a salvage market, we also presented a heuristic that performs very well for the problems without a salvage market. Moreover, based on numerical results, we conjecture that the optimal policy is an MWR policy and that the rationing levels are equal to the optimal rationing levels of the problem with a salvage market.

Our results constitute a useful benchmark for systems with general processing times other 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, because Erlang distributed processing times approach deterministic times when the number of stages is large. For the more general case, multi-stage distributions with different exponential processing times provide a promising alternative to approximate the processing time. Our heuristics 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.

Acknowledgements

The authors thank the two anonymous reviewers and the Associate Editor for their comments leading to an improved exposition of the paper. The authors are also grateful to Yves Dallery and Paul Zipkin for their insightful comments on previous versions of this paper. The research of F. Karaesmen has been partially supported by grants from TÜBA and TÜBİTAK.

References

Abouee-Mehrizi, B., Balcioglu, B. and Baron, O. (2008) Strategies for a single product multi-class M/G/1 make-to-stock queue. Working paper, University of Toronto, Toronto, Canada.

Buzacott, J.A. and Shanthikumar, J.G. (1993) *Stochastic Models of Manufacturing Systems*, Prentice Hall, Englewood Cliffs, NJ.

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

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

Çil, E.B., Örmeci, E.L. and Karaesmen, F. (2009) Effects of system parameters on the optimal policy structure in a class of queueing control problems. *Queueing Systems*, **61**, 273–304.

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

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

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

Gayon, J.P., Benjaafar, S. and de Véricourt, F. (2009) Using imperfect demand information in production-inventory systems with multiple customer classes. *Manufacturing & Service Operations Management*, **11**(1), 128–143.

Gupta, D. and Wang, L. (2007) Capacity management for contract manufacturing. *Operations Research*, **55**, 367–377.

Ha, A.Y. (1997a) Stock rationing policy for a make-to-stock production system with two priority classes and backordering. *Naval Research Logistics*, **44**, 457–472.

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

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

Hu, X., Duenyas, I. and Kapuscinski, R. (2008) Optimal joint inventory and transshipment control under uncertain capacity. *Operations Research*, **56**(4), 881–897.

Huang, B. and Iravani, S.M.R. (2007) Optimal production and rationing decisions in supply chains with information sharing. *Operations Research Letters*, **35**(5), 669–676.

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

Lee, J.E. and Hong, Y. (2003) 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.

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

Milner, J.M. and Kouvelis, P. (2007) Inventory, speculation, and sourcing strategies in the presence of online exchanges. *Manufacturing & Service Operations Management*, **9**(3), 312–331.

Puterman, M.L. (1994) *Markov Decision Processes*, Wiley, New York NY.

Tijms, H.C. (1994) *Stochastic Models: An Algorithmic Approach*. Wiley.

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

Tsay, A., Nahmias, S. and Agrawal, N. (1999) Modeling supply chain contracts: a review, in *Quantitative Models for Supply Chain Management*, Tayur, S., Ganeshan, R. and Magazine, M. (eds.), Kluwer, Boston, MA, pp. 299–336.

Veatch, M.H. and Wein, L.M. (1992) Monotone control of queueing networks. *Queueing Systems*, **12**, 391–408.

Weber, R. and Stidham, S. (1987) Optimal control of service rates in networks of queues. *Advanced applied probabilities*, **19**(1), 202–218.

Zhao, H., Deshpande, V. and Ryan, J.K. (2008) Optimal dynamic production and inventory transshipment policies for a two-location make-to-stock system. *Operations Research*, **56**, 400–410.

Biographies

Jean-Philippe Gayon is an Assistant Professor in the Department of Industrial Engineering at Grenoble INP (France). He has a Ph.D. degree from the École Centrale Paris (France). His research interests lie in optimization of stochastic inventory and production systems. He has had articles published in *Manufacturing & Service Operations Management* and *Probability in the Engineering and Informational Sciences*.

Francis de Véricourt joined the ESMT faculty in September of 2007 as a Professor as well as the First Associate Dean of R&D, where he plays an integral part in guiding the direction of ESMT's research activities. He was previously an Associate Professor at the Fuqua School of Business at Duke University. His research interests broadly lie in developing advanced analytical methods to help make better managerial decisions. His recent work has focused on solving dynamic decision problems in service and supply chain management. He has published articles in leading academic journals such as *Management Science* and *Operations Research*. He also has served in a variety of editorial roles.

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 *IIE Transactions*, *Manufacturing & Service Operations Management*, *4OR* and *IEEE Transactions on Automation Sciences and Engineering*.