Influence of dependency between demands and returns in a reverse logistics system

Hichem Zerhouni, Jean-Philippe Gayon* and Yannick Frein

Laboratoire G-SCOP
Grenoble Institute of Technology
46, Avenue Félix Vialet
38031 Grenoble Cedex, France

* Corresponding author, jean-philippe.gayon@grenoble-inp.fr

Abstract

We consider two reverse logistics systems where returned products are as good as new. For the first system, the product return flow is independent of the demand flow. We prove that the optimal policy is of base-stock type and we establish monotonicity results for the optimal base-stock levels, with respect to the system parameters (arrival rate, production rate, return rate, production cost, lost-sale cost, return cost and holding cost). We also provide an efficient algorithm to compute the optimal base-stock level. For the second system, demands and returns are strongly correlated: a satisfied demand induces a product return after a stochastic return lead-time, with a certain probability. When the return lead-time is null, we extend the results obtained for the first system. When the return lead-time is positive, the optimal control problem is more complex and we do not prove that the optimal policy is of base-stock type. However we provide a framework to analyse base-stock policies. Finally, we carry out a numerical study on many scenarios to investigate the impact of ignoring dependency between demands and returns. We observe that ignoring this dependency yields to non-negligible cost increase.

Keywords:
Reverse logistics, Production/Inventory control, Queueing system, Markov decision process.

1 Introduction

Recycling and recovery of used products has drawn attention of companies for several years, not only for ecological reasons, but also for legal and economical ones. At the same time, customers return more and more items to the producers for numerous reasons (DeCroix and Zipkin, 2005). The returned items constitute return flows that must be taken into account. The management of this material flow, opposite to the conventional supply chain flow, is addressed in the rapidly expanding field of reverse logistics (Fleischmann et al., 1997). From a logistic point of view, and regardless of why they occur, product returns complicate the management of an inventory system (DeCroix et al., 2005). First, returns represents an exogenous inbound material flow causing an increase of the inventory between replenishments. Second, returned products - when recovered - give another alternative supply source for replenishing the serviceable inventory (Fleischmann and Kuik, 2003). Several researches investigated the influence of product returns on inventory control. For an overview, we refer the reader to Fleischmann et al. (1997).

Most of the models do not take into account the relation between returns and demand (see for instance Fleischmann et al., 2002). de Brito and Dekker (2001) have explored the assumptions generally made in stochastic models with product returns such as the assumption of independence between returns an demand. They conclude that it is necessary to break with this traditional assumption. Most of the models with product returns that are investigated assume a total, or partial, independence between demanded items and returned ones. This is owed to the great complexity which could be led by the relaxation of this hypothesis.

Among the authors that consider the dependency relation of returns with demand, Simpson (1978) considers a repairable inventory problem where the dependency between the demand process and return process is allowed only in the same period. Kiesmueller and van der Laan (2001) develop a periodic review model with constant return and procurement lead-times. They compare the case of dependent returns with the case of independent returns and obtain numerically that the average cost is smaller in the dependent case. Cheung and Yuan (2003) consider a continuous review model with Poisson demand, exponential return lead-time and instantaneous procurement lead-time. They adopt an (s, S) inventory policy and develop an algorithm to compute the optimal replenishment parameters. However none of these models investigate the impact of neglecting correlation between demand and returns.

In this paper, we relax the instantaneous procurement lead-time assumption of Cheung and Yuan (2003). We use the framework of make-to-stock queues (Veatch and Wein, 1996; Ha, 1997) to model a stochastic and capacitated production process by a single exponential server. This framework allows us to thoroughly characterize the optimal control policy. We consider two make-to-stock systems. In the first one, demands and returns are independent Poisson processes. We prove that the optimal policy is of base-stock type. We establish monotonicity results for the optimal base-stock levels, with respect to the system parameters (arrival rate, production rate, return rate, production cost, lost-sale cost, return cost and holding cost). We then compute analytically the average cost for a given base-stock level and provide properties of the average cost with respect to the base-stock level. In the second model, demands and returns are correlated: a satisfied demand induces a product return with a certain probability after a stochastic return lead-time. We extend the results of the first model when the return lead-time is null. This special case is interesting for several reasons. It is a tractable case where the impact of ignoring dependence between returns and demands is maximum. It also provides a good approximation for short-term returns. When the return lead-time is positive, the structure of the optimal policy is more complex to establish and depends on whether or not we can observe which sold products will be returned. The assumption of observability is not realistic in most of situations and we will not consider this case. When there is no observability, the production decisions can be based only on the inventory level. For this case, we restrict the analysis to base-stock policies and we suggest a numerical procedure to compute the optimal base-stock policy.

Finally, we carry out a numerical study to investigate the impact of ignoring dependency between demand and returns. We first compare the system with independent returns to the system with dependent returns. Then we suggest a heuristic for the system with dependent returns, based on the system with independent returns. We begin by investigating thoroughly the zero return lead-time before looking at the influence of return lead-time.

The remainder of this paper is organized as follows. Section 2 (resp., Section 3) presents the formulation and results for the system with independent (resp. dependent) demand and return processes. These results allow us, in Section 4, to study the impact of ignoring correlation of demand and returns. Finally, in Section 5, we conclude and suggest future research.

2 Model with returns independent of demands

We first consider a simple model where product returns and demand are independent stochastic processes. We will refer to this case as Model 1.

2.1 Formulation

We consider a make-to-stock system producing a single item. The supplier can decide at any time to produce or not this item. The unit production cost is c_p. The processing time is exponentially distributed with mean $1/\mu$ and completed items are stored in a serviceable product inventory, where they incur an holding cost c_h per unit per unit time. Demands for those items arrive according to a Poisson process with rate λ. A demand that cannot be fulfilled immediately, when the inventory is empty, is lost and incurs a lost-sale cost c_l including image cost, penalty cost, etc. We assume that the production cost c_p is smaller than the lost-sale cost c_l, otherwise it is optimal to idle production all the time.

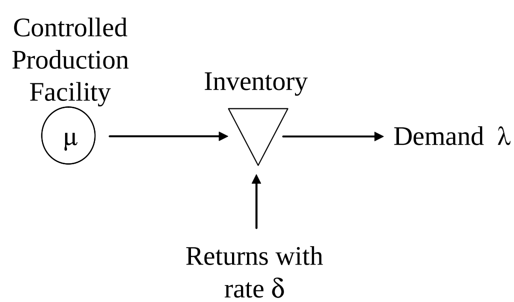

Figure 1: Returns independent of demand

We also suppose that there are random returns of items (Figure 1) that are immediately available to serve customer demand. The inventory is common to new and returned products which are considered as good as new. In this first model, returns occur according to a Poisson process, independent of the demand process, with rate δ. Let $p = \delta/\lambda$ be the proportion of returned products if all demands were satisfied, and $q = 1-p$. We emphasize that p is larger than the proportion of returned products since some demands are not satisfied. We assume that the return rate is smaller than the demand rate, i.e. $\delta < \lambda$ (or equivalently $0 \leq p < 1$). In an infinite planning horizon, this assumption clearly holds if returns are linked to previously satisfied demands. It also guarantees the stability of the stock level. A return incurs a return cost c_r including logistics return costs (repackaging, handling) and possibly the reimbursement of the customer. The state of the system can be summarized by $X(t)$, the stock level at time t (including new and returned products).

A policy π specifies, at any time, when to produce or not. The objective of the supplier is to find the optimal policy minimizing expected discounted costs over an infinite time horizon. We denote by $\beta > 0$ the discount rate.

2.2 Characterization of the optimal policy

We prove in this section that the optimal policy is a base-stock policy.

Definition 1 *A base-stock policy, with base-stock level S, states to produce whenever the stock level is strictly below S and not to produce otherwise.*

The problem of finding the optimal policy can be modelled as a continuous time Markov Decision Process (MDP). We restrict our analysis to stationary markovian policies since there exists an optimal stationary markovian policy (Puterman, 1994).

We define $v^{\pi}(x)$ as the expected total discounted cost associated to policy π and initial state x. We seek to find the optimal policy π^* minimizing $v^{\pi}(x)$ and we let $v^*(x) = v^{\pi^*}(x)$ denote the optimal value function:

$$v^*(x) = \min_{\pi} v^{\pi}(x)$$

We denote by β the discount factor. Then, we can uniformize (Lippman, 1975) the MDP with rate $\gamma > \beta + \lambda + \mu + \delta$ and the optimal value function can be shown to satisfy the following optimality equations:

$$v^*(x) = Tv^*(x), \forall x \in \mathbb{N}$$

where $\mathbb{N}$ is the set of natural numbers and operator T is a contraction mapping defined as

$$Tv(x) = \frac{1}{\gamma}\left[c_h x + \mu T_1 v(x) + \lambda T_2 v(x) + \delta T_3 v(x) + (\gamma - \beta - \lambda - \mu - \delta) v(x)\right] \tag{1}$$

and

$$T_1v(x) = \min[v(x), v(x+1) + c_p]$$
$$T_2v(x) = \begin{cases} v(x-1) & \text{if } x > 0 \\ v(x) + c_l & \text{if } x = 0 \end{cases}$$
$$T_3v(x) = v(x+1) + c_r$$

Operator T_1 corresponds to production decisions while operator T_2 corresponds to demand arrivals. Operator T_3 is associated to product return events.

To prove that the optimal policy is of base-stock type, it is sufficient to show that the optimal value function $v^*(x)$ is convex in the stock level x. A function v in $\mathbb{N}$ is said to be convex if and only if $\Delta v(x) = v(x+1) - v(x)$ is non-decreasing in x. We will also use the notation $\Delta^2 v(x) = \Delta v(x+1) - \Delta v(x)$. With this notation, v is convex if and only if $\Delta^2 v(x) \geq 0$, for all x.

Let us explain briefly why convexity of the optimal value function implies the base-stock policy structure of the optimal policy. Convexity of v^* ensures the existence of a threshold $S^* = \min[x|\Delta v^*(x) + c_p > 0]$, possibly infinite, such that $\Delta v^* + c_p \leq 0$ if and only if x is below this threshold. Optimality equations state to produce when $\Delta v^* + c_p < 0$ and to idle production when $\Delta v^* + c_p > 0$. If $\Delta v^*(x) + c_p = 0$ then it is equal to produce or not in state x. We decide arbitrarily to produce in this case since it does not affect the optimal cost but increases the percentage of satisfied demand.

To prove convexity of v^*, we define $\mathcal{U}$ a set of real-valued functions in $\mathbb{N}$, with the following properties.

Definition 2 *$v \in \mathcal{U}$ if and only if, for all $x \in \mathbb{N}$, v satisfies the following conditions:*

- *Condition C.1:* $\Delta v(x+1) \geq \Delta v(x)$ $(\Leftrightarrow \Delta^2 v(x) \geq 0)$
- *Condition C.2:* $\Delta v(x) \geq -c_l$

The first condition states convexity of v. The second condition can be rewritten as $(v(x-1) \leq v(x)+c_l)$ and means that it is preferable to satisfy a demand rather than to reject it with cost c_l. We know (Puterman, 1994) that a sequence of real-valued functions $v^{n+1} = Tv^n$ converges to the optimal value function, v^*, for all v^0. In order to prove that $v^* \in \mathcal{U}$, it is therefore sufficient to prove the following lemma.

Lemma 1 *If $v \in \mathcal{U}$ then $Tv \in \mathcal{U}$.*

All the proofs can be found in Appendix. As a direct consequence of Lemma 1, we obtain the following theorem.

Theorem 1 *The optimal value function v^* belongs to $\mathcal{U}$ and the optimal policy is a base-stock policy.*

2.3 Influence of system parameters on optimal base-stock levels

Now, we aim to study the influence of system parameters on the optimal policy. For instance how is influenced the optimal base-stock level by a demand rate increase ? The methodology of this section is inspired by Cil et al. (2009).

The optimal base-stock level and value function corresponding to a given system parameter α will be denoted respectively S_α^* and v_α^*, where α belongs to the set of parameters $\{\lambda, \mu, \delta, c_h, c_r, c_l, c_p\}$ (demand rate, production rate, return rate, holding cost, return cost, lost-sale cost, production cost). The optimal value functions v_α^* satisfy the following optimality equations:

$$v_\alpha^*(x) = T_\alpha v_\alpha^*(x)$$

where T_α corresponds to the operator T defined in Equation (1), indexed by α, the parameter under consideration.

To study the influence of system parameters on the optimal base-stock level, we need to define the submodularity and supermodularity of a value function v_α with respect to the state x and the parameter under consideration α as follows. The value function v_α is submodular in α and x (denoted $SubM(\alpha, x)$) if and only if

$$\Delta v_\alpha(x) \geq \Delta v_{\alpha+\epsilon}(x), \forall x \in \mathbb{N}, \forall \alpha, \forall \epsilon \geq 0$$

The supermodularity in α and x (denoted $SuperM(\alpha, x)$) corresponds to the opposite inequality:

$$\Delta v_\alpha(x) \leq \Delta v_{\alpha+\epsilon}(x), \forall x \in \mathbb{N}, \forall \alpha, \forall \epsilon \geq 0$$

When $v_\alpha^*(x)$ is $SubM(\alpha, x)$ (resp., $SuperM(\alpha, x)$), the optimal base-stock level $S_\alpha^* = \min[x|\Delta v_\alpha^*(x) > 0]$ is non-decreasing (resp., non-increasing) in α.

The sequence of real-valued functions $v_\alpha^{n+1} = T_\alpha v_\alpha^n$ converges to the optimal value function v_α^*. In order to prove that v_α^* has the desired modular properties, it is therefore sufficient to prove the following lemma.

Lemma 2
- *$\forall \alpha \in \{\mu, \delta, c_h, c_r, c_p\}$, if v_α is $SuperM(\alpha, x)$ and belongs to $\mathcal{U}$ then $T_\alpha v_\alpha(x)$ is $SuperM(\alpha, x)$ and belongs to $\mathcal{U}$.*
- *$\forall \alpha \in \{\lambda, c_l, c_r\}$, if v_α is $SubM(\alpha, x)$ and belongs to $\mathcal{U}$ then $T_\alpha v_\alpha(x)$ is $SubM(\alpha, x)$ and belongs to $\mathcal{U}$.*

We deduce directly from Lemma 2 the following theorem.

Theorem 2 *The optimal value function v_α^* is $SuperM(\alpha, x)$ for $\alpha \in \{\mu, \delta, c_h, c_r\}$ and $SubM(\alpha, x)$ for $\alpha \in \{c_l, \lambda, c_r\}$.*

As a result, the optimal base-stock level is independent of the return cost, c_r, non-increasing with the service rate, μ, the return rate, δ, the holding cost, c_h, the production cost, c_p, and non-decreasing with the arrival rate, λ, the lost-sale cost, c_l.

There exists a simple alternative to prove that the optimal base-stock level is independent of c_r: It suffices to notice that the expected return costs are independent of the production policy. The expected discounted return costs can then be considered as a constant in our minimization problem and does not influence the production policy. Therefore, the optimal base-stock level is not sensitive to the return cost.

2.4 Performance evaluation for the average cost problem

The average cost optimal policy can be obtained as the limit of the discounted cost optimal policy when the discounted rate β goes to zero (Weber and Stidham, 1987). As a result, for the average cost problem, the optimal policy is of base-stock type and optimal base-stock levels have the monotonicity properties presented in Theorem 2.

The average cost optimal policy can be computed by various algorithms, for instance by value iteration, as for the discounted cost problem (Puterman, 1994). We provide now an alternative to compute more efficiently the optimal policy.

If we assume a base-stock policy with base-stock level S, the stock level evolves according to a continuous-time Markov chain with transition rates given on Figure 2.

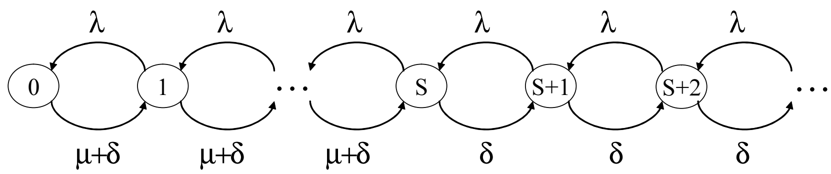

Figure 2: Graph of the Markov chain

Denote by $\pi_x(S)$ the stationary probability to be in state x when the base-stock level is S. We obtain:

$$\pi_x(S) = \begin{cases} \rho_1^{-x}\pi_0(S) & \text{if } 1 \leq x \leq S \\ \rho_1^{-S} p^{x-S} \pi_0(S) & \text{if } x \geq S \end{cases}$$

$$\pi_0(S) = \begin{cases} \rho_1^S \left(\frac{1-\rho_1^{S+1}}{1-\rho_1} + \frac{p}{q} \right)^{-1} & \text{if } \rho_1 \neq 1 \\ \frac{q}{1+Sq} & \text{if } \rho_1 = 1 \end{cases}$$

where $\rho_1 = \frac{\lambda}{\mu+\delta}$.

Now, we can compute the average holding cost, $C_h(S)$, lost-sale cost, $C_l(S)$, production cost, $C_p(S)$, and return cost, $C_r(S)$:

$$C_h(S) = c_h \sum_{x=0}^{\infty} x \ \pi_x(S) = \begin{cases} c_h \pi_0(S) \left(\frac{\rho_1^{S+1} - \rho_1 - \rho_1 S + S}{\rho_1^S \ (1-\rho_1)^2} + \frac{1}{\rho_1^S} \ \frac{p(1+qS)}{q^2} \right) & \text{if } \rho_1 \neq 1 \\ c_h \pi_0(S) \left[\frac{S(S+1)}{2} + \frac{p(1+qS)}{q^2} \right] & \text{if } \rho_1 = 1 \end{cases} \tag{2}$$

$$C_l(S) = \lambda c_l \pi_0(S)$$

$$C_p(S) = \mu \ c_p \sum_{x=0}^{S-1} \pi_x(S) = \mu c_p \left(\frac{1-\rho_1^{-S}}{1-\rho_1^{-1}} \right) \pi_0(S)$$

$$C_r(S) = \delta \ c_r$$

The total average cost, $C(S)$, is then the sum of these three costs. In order to compute efficiently the optimal average cost, C^*, and the optimal base-stock level, S^*, we use the following property.

Property 1 *The average cost $C(S)$ is convex in S if $\rho_1 \leq 1$.*

When $\rho_1 > 1$, convexity does not hold systematically. For example, the average cost is not convex for the following instance: $c_h = 1, c_l = 16, c_r = 4, c_p = 0, \lambda = 1.4, \mu = 1, p = 0.1$.

However we can re-use the upper bound, on the optimal base-stock level, developed by Ha (1997) for a system similar to ours but without product returns ($\delta = 0$). This upper bound also holds for our problem since the optimal base-stock level is smaller when product returns are taken into account (see Theorem 2).

Property 2 *(Directly adapted from Ha, 1997) Let S_u be the smallest non-negative integer larger than $(\rho - 1)/(h'\rho) + 1/\ln\rho - 1$ with $\rho = \lambda/\mu$ and $h' = c_h/(\lambda c_l)$. For $\rho > 1$, $S^* \leq S_u$.*

If we consider the extreme case ($\lambda = 10, \mu = 1, c_h = 1, c_l = 1000, c_p = 0$), then $S_u = 9000$. The time to compute the optimal base-stock level for this instance requires to evaluate $C(S)$ for $S \in \{0, 1, \cdots, 9000\}$, which takes less than one second on a standard personnel computer.

3 Model with returns dependent on demands

3.1 Formulation

We consider a second model which takes into account the correlation between returns and demand (Figure 3). We will refer to this model as Model 2.

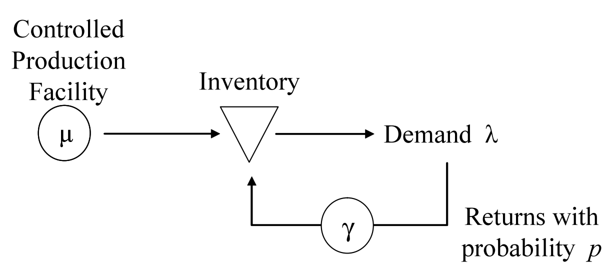

Figure 3: Returns dependent on demand

The main difference with the previous model is that a satisfied demand leads, with probability p, to a product return after a stochastic return lead-time. This return lead-time is assumed to be exponentially distributed with rate γ. de Brito and Dekker (2001) have tested, on real data, the assumption that the time to return is exponential. They conclude, with a statistical test, that this assumption can be rejected for some products and can not be rejected for other products.

The state of the system can now be summarized by two variables $(X(t), Y(t))$ where $X(t)$ is the stock level (including new and returned products) and $Y(t)$ is the number of demands that have been satisfied and will be returned. We can distinguish two cases:

- $Y(t)$ is observable. The decision maker knows exactly which customers will return a product and which customers won't.

- $Y(t)$ is not observable. The decision maker has for only information that a proportion p of customers return a product.

It is generally not realistic to assume that we can observe how many products in the market will be returned. In this case, we have observed numerically that the optimal policy is a state-dependent base-stock policy, i.e. the base-stock level depends on $Y(t)$. We focus on the case where $Y(t)$ is not observable and the decisions are based only on the inventory level.

3.2 Zero return lead-time

Though we were not able to prove structural results for the optimal policy in the general case, we have established the structure of the optimal policy for the limit case with zero return lead-time (equivalent to $1/\gamma = 0$). This limit case simplifies the analysis and also provides a good approximation to systems with short-term returns that typically occur when a customer has the option to return a product, but in a very short period after purchasing. The problem simplifies to a single variable problem where we can get rid of variable $Y(t)$. The analysis of such a problem is similar to the one of Section 2.

Denote by $\tilde{v}^*(x)$ the optimal value function when the initial inventory is x. The optimal value function satisfies the following optimality equations.

$$\tilde{v}^*(x) = \tilde{T}\tilde{v}^*(x), \forall x \in \mathbb{N}$$

where

$$\tilde{T}v(x) = \frac{1}{\gamma}\left[c_h x + \mu T_1 v(x) + \lambda \tilde{T}_2 v(x) + (\gamma - \beta - \lambda - \mu - \delta)v(x)\right] \quad (3)$$

$$\tilde{T}_2 v(x) = \begin{cases} (1-p)v(x-1) + p[v(x) + c_r] & \text{if } x > 0 \\ v(x) + c_l & \text{if } x = 0 \end{cases} \quad (4)$$

and T_1 is defined as in Section 2.2.

To prove convexity of $\tilde{v}^*$, we define $\tilde{\mathcal{U}}$ a set of real-valued functions in $\mathbb{N}$, with the following properties.

Definition 3 *$v \in \tilde{\mathcal{U}}$ if and only if, for all $x \in \mathbb{N}$, v satisfies the following conditions:*

- *Condition C.1:* $\Delta^2 v(x) \geq 0$
- *Condition C.2:* $\Delta v(x) \geq -\eta$ *with* $\eta = \frac{c_l - pc_r}{q}$

The second condition states that it is preferable to satisfy an arriving demand.

Lemma 3 *Assume that $\eta \geq 0$. If $v \in \tilde{\mathcal{U}}$ then $Tv \in \tilde{\mathcal{U}}$.*

As a direct consequence of Lemma 1, we obtain the following theorem.

Theorem 3 *Assume that $\eta \geq 0$. The optimal value function v^* belongs to $\tilde{\mathcal{U}}$ and the optimal policy is a base-stock policy.*

To complete the proof, we had to make the additional assumption that $\eta \geq 0$ which is equivalent to $c_l \geq pc_r$. Otherwise when $c_l < pc_r$, it is optimal to never produce since satisfying a demand incurs (in expectation) a higher return cost pc_r than not satisfying a demand which incurs a lost-sale cost c_l.

The methodology to study the influence of system parameters is similar to the one adopted for Model 1.

Theorem 4 *Assume that $\eta \geq 0$. The optimal value function $\tilde{v}^*_\alpha$ is $SuperM(\alpha, x)$ for $\alpha \in \{\mu, p, c_h, c_r\}$ and $SubM(\alpha, x)$ for $\alpha \in \{c_l, \lambda\}$.*

As a result, the optimal base-stock level is non-increasing with the service rate, μ, the return probability, p, the holding cost, c_h, the return cost, c_r, and non-decreasing with the arrival rate, λ, the lost-sale cost, c_l.

Contrary to the problem with independent returns, the optimal base-stock level is non-increasing with the return cost c_r. When c_r is increasing, producing new items leads to higher return costs and it is therefore better to reduce production. An extreme case would be to consider an infinite return cost for which it is clearly optimal to set the optimal base-stock level to 0, to avoid any return. This phenomenon is illustrated in the numerical study (Figure 6).

If we assume a base-stock policy with base-stock level S, the stock level evolves as in a basic make-to-stock queue without product returns and demand rate λq where $q = 1 - p$ (Figure 4). The transition rates are given on Figure 4.

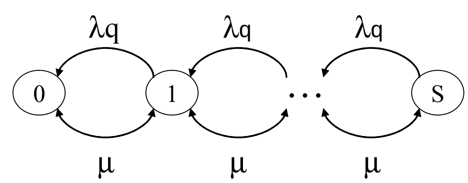

Figure 4: Graph of the Markov chain with zero return lead-time

For convenience, we define the ratio $\rho_2 = \frac{q\lambda}{\mu}$ and the stationary probabilities are given by:

$$\tilde{\pi}_x(S) = \rho_2^{-x}\tilde{\pi}_0(S) \text{ when } 1 \leq x \leq S$$
$$\tilde{\pi}_0(S) = \begin{cases} \frac{\rho_2^S(1-\rho_2)}{1-\rho_2^{S+1}} & \text{if } \rho_2 \neq 1 \\ \frac{1}{S+1} & \text{if } \rho_2 = 1 \end{cases}$$

As for the independent case, we can compute the average holding, lost-sale, production and return costs:

$$\tilde{C}_h(S) = \begin{cases} c_h \frac{\rho_2^{S+1}-\rho_2-\rho_2 S+S}{(1-\rho_2)(1-\rho_2^{S+1})} & \text{if } \rho_2 \neq 1 \\ c_h \frac{S}{2} & \text{if } \rho_2 = 1 \end{cases}$$
$$\tilde{C}_l(S) = \lambda\ c_l\ \tilde{\pi}_0(S) \tag{5}$$
$$\tilde{C}_p(S) = \mu c_p(1 - \tilde{\pi}_S)$$
$$\tilde{C}_r(S) = \lambda\ p\ c_r\ (1 - \tilde{\pi}_0(S)) \tag{6}$$

The average cost, $\tilde{C}(S)$, is again the sum of these three costs.

Unfortunately, the average cost is not systematically convex in S. The average costs $\tilde{C}_r(S)$ and $\tilde{C}_l(S)$ are never simultaneously convex since their second derivative are of opposite sign, from Equations (5) and (6). However we can use again the upper bound S_u of Section 2.4.

3.3 Positive return lead-time

When the return lead-time is positive ($1/\gamma > 0$), we restrict the analysis to base-stock policies such that the system produces if and only if the stock level is smaller than S. However we don't claim that the optimal policy is base-stock.

The continuous-time Markov chain $(X(t), Y(t))$ has a more complex structure than for the zero return lead-time. There does not exist simple analytical formulas for the stationary probabilities. We denote again by $\tilde{C}(S)$ the average cost when base-stock level S is used. In order to compute $\tilde{C}(S)$, we use a dynamic program (detailed in Appendix) where we truncate the state space to $\{0, 1, \cdots, S_u\} \times \{0, 1, \cdots, M\}$. One can then search for the optimal base-stock level $\tilde{S}$ minimizing $\tilde{C}(S)$. Larger and larger values of M are tested until the results become insensitive to increasing the state space.

4 Numerical study

In this numerical study, we focus on the average cost problems. To compute the optimal policies, we use the analytical results of previous sections. First, we compare the influence of return cost, c_r, return probability, p, and expected return rate, $1/\gamma$ on the two systems. Second, we look at the impact of using the optimal policy of System 1 (independent returns) as a heuristic for System 2 (dependent returns).

4.1 Comparison of models

Figure 5 presents the influence of the return probability, p, on the optimal base-stock levels, S^* and $\tilde{S}^*$, and the optimal average costs, C^* and $\tilde{C}^*$.

When $p = 0$, systems 1 and 2 are equivalent and thus $C^* = \tilde{C}^*$, $S^* = \tilde{S}^*$. When p increases, we observe that S^* and $\tilde{S}^*$ are non-increasing in p. On the other hand, C^* and $\tilde{C}^*$ are not monotonic in p. When p increases, the average return cost and holding cost tend to increase while the average lost-sale cost tends to decrease. Depending on the relative importance of these average costs, the total average cost might increase or decrease. When p goes to 1, the holding costs go to infinity for Model 1 since the inventory queue is unstable (demand rate equal to return rate). It is not the case for Model 2 for which the inventory level remains bounded since returns are correlated to demands.

In the literature, when there is no production capacity constraint, the average cost is very sensitive to p when p is close to 1 (Fleischmann et al., 2002) but not when p is close to 0. In our model, the average cost might also be very sensitive to p when p is small. This is due to the production capacity constraint assumption (the production rate is bounded by μ). In this case, increasing slightly the return rate might decrease drastically the lost-sale costs and hence the total cost.

Figure 6 presents the optimal base-stock levels and average costs versus the return cost, c_r. When $c_r = 0$ and $1/\gamma = 0$, systems 1 and 2 behave as a basic make-to-stock queue with modified demand rate $q\lambda$, production rate μ, lost-sale cost c_l and holding cost c_h. Therefore, we have $C^* = \tilde{C}^*$ and $S^* = \tilde{S}^*$ when $c_r = 0$. When c_r increases, we observe that S^* is constant and $\tilde{S}^*$ is non-increasing in c_r, which is consistent with theorems 2 and 4. We also observe that $\tilde{S}^*$ equals 0 when $pc_r \geq c_l$. In this case, the expected return cost for each satisfied demand becomes higher than the potential lost-sale cost and it is no more interesting to produce any item.

Figure 7 presents the optimal base-stock levels and average costs versus the expected return lead-time, $1/\gamma$. System 1 is independent of γ and thus S^* and C^* are also independent of $1/\gamma$. The difference of behavior between the two systems is maximum when the expected return lead-time equals 0. When the expected return lead-time increases, the differences between the two systems diminish since the return process and the demand process become almost uncorrelated.

In conclusion of this section, we can say that Systems 1 and 2 may have very different behaviors, for example for large c_r, large p or small $1/\gamma$.

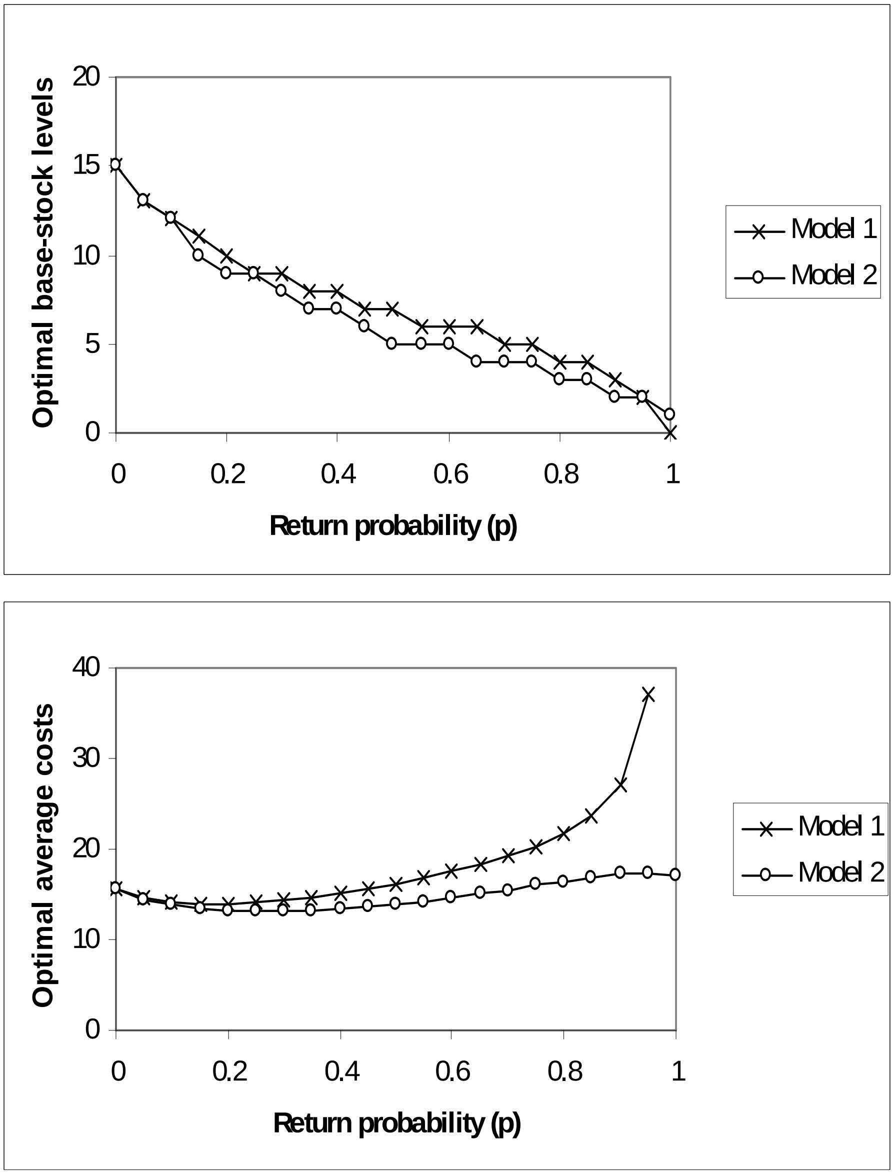

Figure 5: Influence of return probability on optimal policies ($\lambda = \mu = 1, 1/\gamma = 0, c_h = 1, c_r = 16, c_l = 128, c_p = 0$)

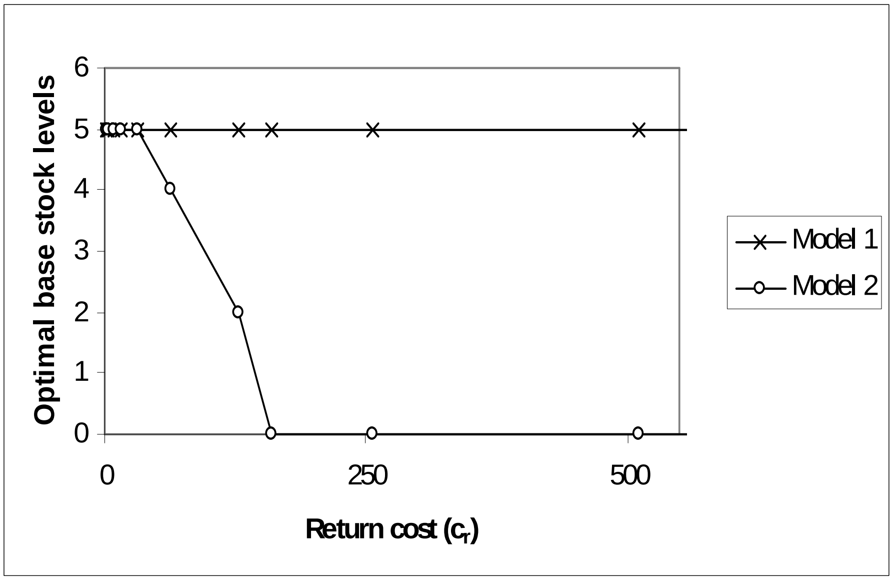

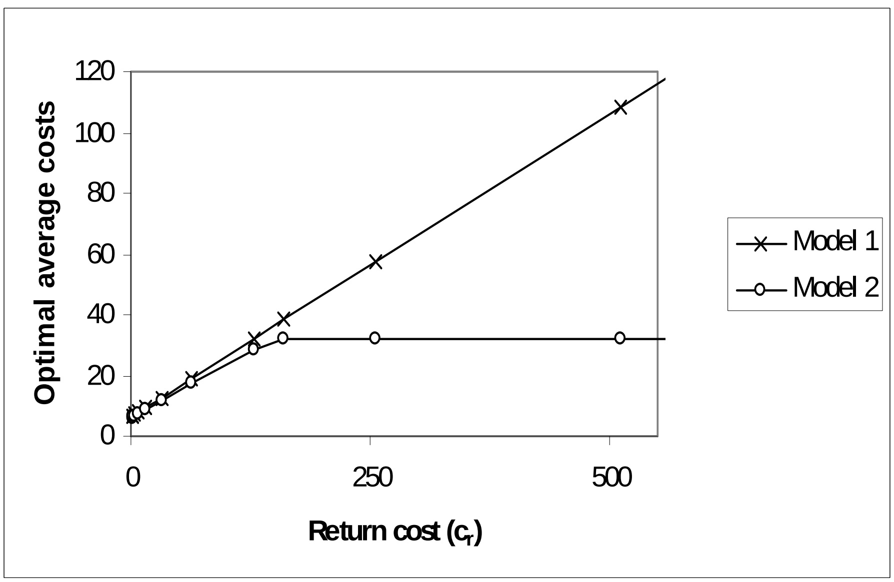

Figure 6: Influence of return cost on optimal policies ($p = 0.2, \lambda = \mu = 1, 1/\gamma = 0, c_h = 1, c_l = 32, c_p = 0$)

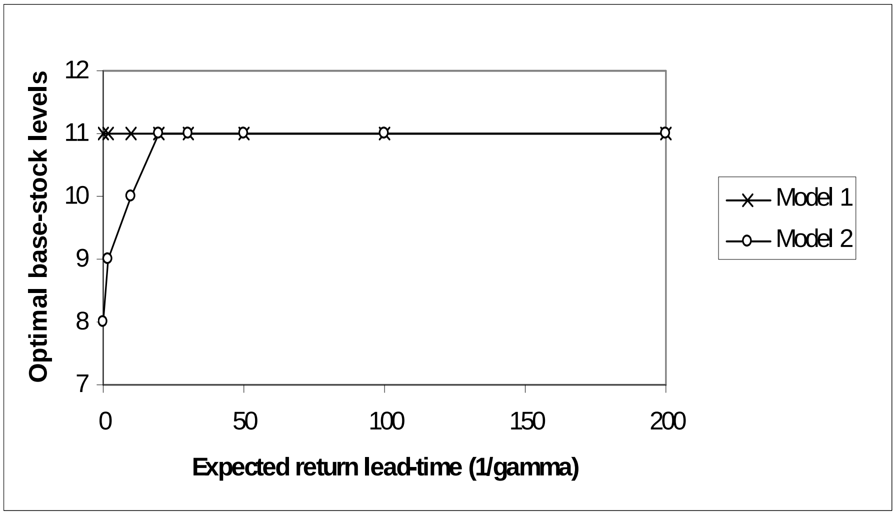

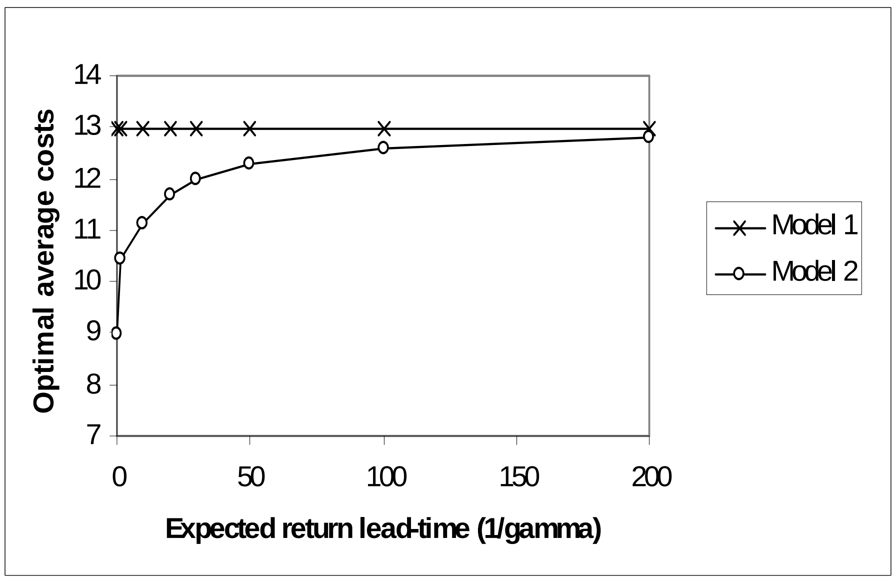

Figure 7: Influence of return cost on optimal policies ($p = 0.5, \lambda = \mu = 1, c_h = 1, c_l = 1000, c_r = 0, c_p = 0$)

4.2 Heuristic policy

The question which naturally arises is: can we use Model 1 as a good approximation to design an efficient policy for Model 2 ? To answer this question, we consider the two following policies for Model 2 :

- The optimal policy of Model 2, with base-stock level $\tilde{S}^*$ and average cost $\tilde{C}^*$
- A heuristic policy for Model 2, with base-stock level S^* (the optimal base-stock level of Model 1). The average cost associated to this heuristic is then $\tilde{C}(S^*)$.

In order to compare these two policies, we consider the relative cost increase for using the heuristic, defined by:

$$\Delta C = \frac{\tilde{C}(S^*) \ - \tilde{C}^*}{\tilde{C}^*}$$

The lower ΔC is, the better the independent model approximates the dependent model.

In the following, we set without loss of generality $\mu = 1$ and $c_h = 1$, which is equivalent to set time and monetary units. We also set the return lead-time to 0 ($1/\gamma = 0$). This case give an upper-bound to ΔC with respect to other values of $1/\gamma$. We also set $c_p = 0$. The optimal policy of a problem instance with parameters c_l, c_r and c_p is the same as the optimal policy associated to a problem instance with lost sale cost $(c_l + c_p)$, return cost $(c_r - c_p)$ and null production cost (other things being equal).

We have evaluated the performance of the heuristic by varying the other parameters for 22990 instances corresponding to all the combinations of the following sets:

- $\lambda \in \{0.2, 0.4, 0.6, 0.8, 1, 1.2, 1.4, 1.6, 1.8, 2\}$
- $p \in \{0.05, 0.1, 0.15, 0.20, 0.25, 0.30, 0.35, 0.40, 0.45, 0.50, 0.55, 0.60, 0.65, 0.70, 0.75, 0.80, 0.85, 0.90, 0.95\}$
- $c_l \in \{1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024\}$
- $c_r \in \{1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024\}$

With this range of parameters, we investigate systems in overloaded regime (when $\mu > \lambda$) and in underloaded regime (when $\mu < \lambda$). We also consider systems with a small return rate as well as systems with a return rate close to the demand rate. Finally, we consider different scenarios for the cost parameters ($c_r < c_l$, $c_r = c_l$ and $c_r > c_l$). The relation between the return cost, c_r, and the lost-sale cost, c_l, is not obvious and depends on the relation between logistics return costs, remanufacturing costs, image/penalty costs and reimbursement costs.

We restrict our analysis to the 12951 instances where there is an interest to produce for both systems ($S^*, \tilde{S}^* > 0$). In Table 1, we give the distribution of ΔC. For 62.1% of the instances, the heuristic performs very well ($\Delta C \leq 1\%$). However for 10.7% of the instances, the heuristic performs poorly with $\Delta C \geq 10\%$.

ΔC	0-1 %	1-5%	5-10%	10-20%	20-50%	> 50%
Number of instances	8041	2707	823	699	552	129
% of instances	62.1%	20.9%	6.4%	5.4%	4.3%	1.0%

Table 1: Distribution of ΔC for instances with $S^*, \tilde{S}^* > 0$

Some other important conclusions, based on the tested instances, are the following:

- The maximum error is of 97%.
- Even with a small return probability of $p = 5\%$ (925 instances), the percentage cost increase ΔC can go up to 10 %.
- For a return probability $p \leq 0.35$ (5555 instances), the percentage cost increase ΔC is bounded by 11%.

For the moment, we have assumed a zero return lead-time. When the return lead-time increases, we systematically observe that the performance of the heuristic increases. Figure 8 plots the influence of the expected return lead-time on the relative cost increase, ΔC.

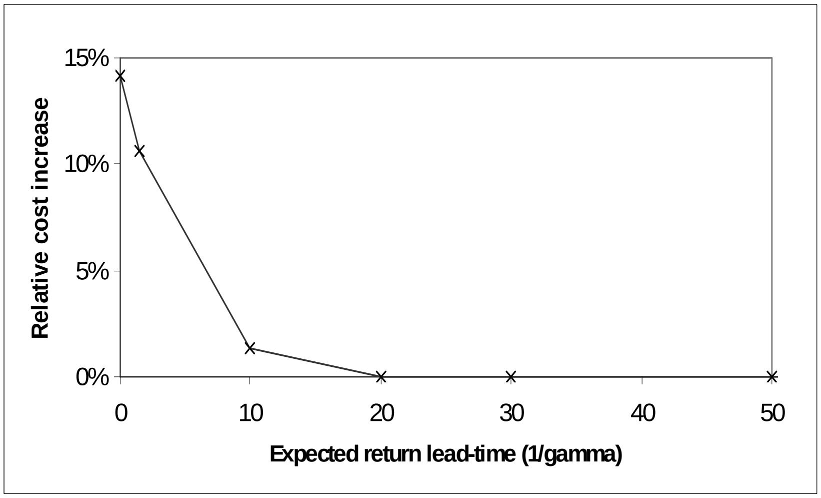

Figure 8: Influence of expected return lead-time on relative cost increase ($p = 0.5, \lambda = \mu = 1, c_h = 1, c_l = 1000, c_r = 0, c_p = 0$)

In conclusion, ignoring dependency between returns and demands may yield to very bad performances and should not be a systematic assumption.

5 Conclusion and future research

In this paper, we have investigated the impact of dependency between returns and demands in a reverse logistics context. We have considered two models: One with returns independent of demand and the other with returns dependent on demand. For the first model, we show that the discounted cost optimal policy is base-stock and we establish monotonicity results of the optimal base-stock level with respect to system parameters. These results pertain to the average cost problem, for which we derive additional analytical results allowing to compute efficiently the optimal policy. For the second model, we obtain similar results when the return lead-time is equal to zero. When the return lead-time is positive, we restrict the analysis to base-stock policies and we provide a numerical method to compute the optimal base-stock level.

In a numerical study, we show that System 2 can not be well approximated by System 1, especially with high probabilities of return and short return lead-times. However, through a systematic factorial analysis, we show that, when the probability of return is less than 35 % and the return lead-time is null, the optimal base-stock level of Model 1 provides a good heuristic to Model 2 and yields to errors smaller than 11 %. When the return lead-time increases, these errors are becoming smaller and smaller. We also show a non-monotonic behavior of the average cost for Model 1, due to the production capacity constraint.

There are several possible avenues for research. It would be interesting to study if these results pertain to other probability distributions of return lead-times, production lead-times, and demand processes. Another option would be to add a control on returns and to consider the joint problem of controlling manufacturing and remanufacturing.

References

K. L. Cheung and X. M. Yuan. An infinite horizon inventory model with periodic order commitment. *European Journal of Operational Research*, 146:52–66, 2003.

E.B. Cil, E.L. Ormeci, and F. Karaesmen. Effects of system parameters on the optimal policy structure in a class of queueing control problems. *Queueing Systems*, 61:273–304, 2009. URL `http://home.ku.edu.tr/~fkaraesmen/`.

M. P. de Brito and R. Dekker. Modelling product returns in inventory control exploring the validity of general assumptions. *Economic Institute Report*, 27, 2001.

G. DeCroix and P. H. Zipkin. Inventory management for an assembly system with product or component returns. *Management Science*, 51:1250–1265, 2005.

G. DeCroix, J-S. Song, and P. Zipkin. A series system with returns: Stationary analysis. *Operations Research*, 52(2):350–362, 2005.

M. Fleischmann and R. Kuik. On optimal inventory control with independent stochastic item returns. *European Journal of Operational Research*, 151: 25–37, 2003.

M. Fleischmann, J.M. Bloemhof-Ruwaard, R. Dekker, E. van der Laan, J.A.E.E. van Nunen, and L.N. Van Wassenhove. Quantitative models for reverse logistics: A review. *European Journal of Operational Research*, 103:1–17, 1997.

M. Fleischmann, R. Kuik, and R. Dekker. Controlling inventories with stochastic item returns: A basic model. *European Journal of Operational Research*, 138:63–75, 2002.

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, 1997.

G. P. Kiesmueller and E. A. van der Laan. An inventory model with dependent product demands and returns. *International journal of production economics*, 72:73–87, 2001.

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

M.L. Puterman. *Markov Decision Processes*. Wiley, 1994.

V.P. Simpson. Optimum solution structure for a repairable inventory problem. *Operations Research*, 26(2):270–281, 1978.

M.H. Veatch and L.M. Wein. Scheduling a make-to-stock queue: index policies and hedging points. *Operations Research*, 44(4):634–647, 1996.

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

Appendix

Proof of Lemma 1

Assume that $v \in \mathcal{U}$. We first prove that Tv satisfies Condition 1 of $\mathcal{U}$ (convexity). To that end, we separately show that $T_1 v$ and $T_2 v$ are convex.

Let $S = \min[x : \Delta v(x) + c_p > 0]$ (possibly infinite). We can rewrite operator T_1 as

$$T_1 v(x) = \begin{cases} v(x+1) + c_p & \text{if } x < S \\ v(x) & \text{if } x \geq S \end{cases}$$

and deduce

$$\Delta T_1 v(x) = \begin{cases} \Delta v(x+1) \geq -c_l & \text{if } x < S-1 \\ -c_p \geq -c_l & \text{if } x = S-1 \\ \Delta v(x) \geq -c_l & \text{if } x \geq S. \end{cases} \tag{7}$$

Inequalities in Equation 7 come directly from Condition 2 of $\mathcal{U}$ satisfied by v and from the assumption that the production cost c_o is smaller than the lost-sale cost c_l. Then we have

$$\Delta^2 T_1 v(x) = \begin{cases} \Delta^2 v(x+1) \geq 0 & \text{if } x < S-2 \\ -\Delta v(x+1) \geq 0 & \text{if } x = S-2 \\ \Delta v(x+1) \geq 0 & \text{if } x = S-1 \\ \Delta^2 v(x) \geq 0 & \text{if } x \geq S. \end{cases} \tag{8}$$

Inequalities in Equation 8 come directly from convexity of v and definition of S. For all $x \in \mathbb{N}$, we have therefore $\Delta^2 T_1 v(x) \geq 0$ and we conclude that $T_1 v$ is convex. Furthermore

$$\Delta T_2 v(x) = \begin{cases} \Delta v(x-1) \geq -c_l & \text{if } x \geq 1 \\ -c_l & \text{if } x = 0 \end{cases} \tag{9}$$

and

$$\Delta^2 T_2 v(x) = \begin{cases} \Delta^2 v(x-1) \geq 0 & \text{if } x \geq 1 \\ \Delta v(x) + c_l \geq 0 & \text{if } x = 0. \end{cases}$$

The quantity $\Delta^2 T_2 v(x)$ is non-negative from Condition C.1 when $x \geq 1$ and from Condition C.2 when $x = 0$, thus $T_2 v$ satisfies C.1. Otherwise $T_3 v$ and

$c_h x$ are trivially convex. Finally Tv, as a non-negative linear combination of convex functions, is also convex. Let us prove now that Tv satisfies Condition C.2. $T_1 v$ and $T_2 v$ satisfy C.2 from Equations (7) and (9). On the other hand $T_3 v$ clearly satisfies Condition 2 and $\Delta c_h x = c_h \geq 0$ since $c_h x$ is non-decreasing in x. Therefore, we have

$$\begin{aligned}\Delta Tv(x) &= \frac{1}{\gamma}\left[\Delta c_h x + \mu \Delta T_1 v(x) + \lambda \Delta T_2 v(x) + \delta T_3 v(x) + (\gamma - \beta - \lambda - \mu - \delta)\Delta v(x)\right] \\ &\geq -\frac{\gamma - \beta}{\gamma} c_l \geq -c_l\end{aligned}$$

and Tv satisfies Condition C.2.

Proof of Lemma 2

Let $v_\alpha \in \mathcal{U}$. Then $T_\alpha v_\alpha \in \mathcal{U}$ from Lemma 1. We need now to show that T_α propagate submodular and supermodular properties.

We first prove that operators $c_h x$, T_1, T_2, T_3 preserve the modular properties summarized in Table 2. Notice that these operators might depend on the parameter α under consideration.

	$SuperM(\alpha, x)$	$SubM(\alpha, x)$
$c_h x$	$\forall \alpha$	$\forall \alpha \neq c_h$
$T_1 v_\alpha(x)$	$\forall \alpha$	$\forall \alpha \neq c_p$
$T_2 v_\alpha(x)$	$\forall \alpha \neq c_l$	$\forall \alpha$
$T_3 v_\alpha(x)$	$\forall \alpha$	$\forall \alpha$

Table 2: Preservation of submodularity and supermodularity by the operators

$\Delta c_h x = c_h$ is non-decreasing in c_h and independent of the other parameters of the system. Therefore $c_h x$ is $SuperM(\alpha, x)$ for all α and $SubM(\alpha, x)$ for all $\alpha \neq c_h$.

From Cil et al. (2009) (we adapt their maximization problem into a minimization problem), we know that T_1 preserve $SubM(\alpha, x)$ and $SuperM(\alpha, x)$ for all $\alpha \neq c_p$.

When $\alpha = c_p$, we show now that T_1 preserves $SuperM(c_p, x)$. Define $S = \min[x : \Delta v_{c_p}(x) + c_p > 0]$ and $S_\epsilon = \min[x : \Delta v(x) + c_p + \epsilon > 0]$. From

Equation (7), we have then

$$\Delta T_1 v_{c_p+\epsilon}(x) - \Delta T_1 v c_p(x) = \begin{cases} \Delta v_{c_p+\epsilon}(x+1) - \Delta v_{c_p}(x+1) \geq 0 & \text{if } x < S_\epsilon - 1 \leq S - 1 \\ -c_p - \Delta v_{c_p}(x+1) \geq 0 & \text{if } x = S_\epsilon - 1 < S - 1 \\ \Delta v_{c_p+\epsilon}(x) - \Delta v_{c_p}(x+1) \geq 0 & \text{if } x \geq S_\epsilon, x < S - 1 \\ 0 & \text{if } x = S_\epsilon - 1 = S - 1 \\ \Delta v_{c_p+\epsilon}(x) + c_p \geq 0 & \text{if } x \geq S_\epsilon, x \leq S - 1 \end{cases} \tag{10}$$

The inequalities in Equation (10) come from the definitions of the thresholds S and S_ϵ.

From Equation (9), we have

$$\Delta T_2 v_\alpha(x) = \begin{cases} \Delta v_\alpha(x-1) & \text{if } x \geq 1 \\ -c_l & \text{if } x = 0 \end{cases}$$

When $\alpha \neq c_l$, it is clear that T_2 propagates $SubM(\alpha, x)$ and $SuperM(\alpha, x)$. T_2 propagate also $SubM(c_l, x)$ since

$$\Delta T_2 v_{c_l+\epsilon}(x) - \Delta T_2 v_{c_l}(x) = \begin{cases} \Delta v_{c_l+\epsilon}(x-1) - \Delta v_{c_l}(x-1) \leq 0 & \text{if } x \geq 1 \\ -\epsilon \leq 0 & \text{if } x = 0 \end{cases}$$

At last T_3 preserves $SubM(\alpha, x)$ and $SuperM(\alpha, x)$ for all α since $\Delta T_3 v_\alpha(x) = \Delta v_\alpha(x+1)$ for all α.

When α belongs to $\{c_h, c_r, c_l\}$, $T_\alpha v_\alpha$ is a linear combination of operators $c_h x$, T_1, T_2, T_3 where the coefficients of the combination do not depend on c_h, c_r, c_l. We can conclude that $T_\alpha v_\alpha$ is $SuperM(\alpha, x)$ for $\alpha \in \{c_h, c_r\}$ and $SubM(\alpha, x)$ for $\alpha \in \{c_l, c_r\}$.

When α belongs to $\{\lambda, \mu, \delta\}$, it is more complicated because $T_\alpha v_\alpha$ is a linear combination of operators $c_h x$, T_1, T_2, T_3 where some of the coefficients of the combination depend on λ, μ, δ. Let us consider the propagation of $SuperM(\mu, x)$. We have:

$$T_\mu v_\mu(x) = \frac{1}{\gamma}\left[c_h x + \mu T_1 v_\mu(x) + \lambda T_2 v_\mu(x) + \delta T_3 v_\mu(x) + \epsilon v_\mu(x) + (\gamma - \beta - \lambda - \mu - \delta) v_\mu(x)\right]$$

and

$$
\begin{aligned}
&T_{\mu+\epsilon}v_{\mu+\epsilon}(x)\\
&=\frac{1}{\gamma}[c_h x+(\mu+\epsilon)T_1 v_{\mu+\epsilon}(x)+\lambda T_2 v_{\mu+\epsilon}(x)+\delta T_3 v_{\mu+\epsilon}(x)+(\gamma-\beta-\lambda-\mu-\epsilon-\delta)v_{\mu+\epsilon}(x)]\\
&=\frac{1}{\gamma}[c_h x+\mu T_1 v_{\mu+\epsilon}(x)+\lambda T_2 v_{\mu+\epsilon}(x)+\delta T_3 v_{\mu+\epsilon}(x)+(\gamma-\beta-\lambda-\mu-\delta)v_{\mu+\epsilon}(x)\\
&+\epsilon(T_1 v_{\mu+\epsilon}(x)-v_{\mu+\epsilon}(x))]
\end{aligned}
$$

Then

$$
\begin{aligned}
&\Delta T_{\mu+\epsilon}v_{\mu+\epsilon}(x)\geq\Delta T_\mu v_\mu(x)\\
&\Leftrightarrow\left(\begin{array}{c}\mu\Delta T_1 v_{\mu+\epsilon}(x)\\ +\lambda\Delta T_2 v_{\mu+\epsilon}(x)\\ +\delta\Delta T_3 v_{\mu+\epsilon}(x)\\ +(\gamma-\beta-\lambda-\mu-\delta)\Delta v_{\mu+\epsilon}(x)\\ +\epsilon\Delta[T_1 v_{\mu+\epsilon}(x)-v_{\mu+\epsilon}(x)]\end{array}\right)\geq\left(\begin{array}{c}\mu\Delta T_1 v_\mu(x)\\ +\lambda\Delta T_2 v_\mu(x)\\ +\delta\Delta T_3 v_\mu(x)\\ +(\gamma-\beta-\lambda-\mu-\delta)\Delta v_\mu(x)\end{array}\right)
\end{aligned}
\tag{11}
$$

The four first lines of (11) satisfy the inequality since T_1, T_2, T_3 propagate $SuperM(\mu,x)$. It is therefore sufficient to prove that $\Delta[T_1 v_{\mu+\epsilon}(x)-v_{\mu+\epsilon}(x)]\geq 0$.

The arguments are the same to propagate $SubM(\lambda,x)$ and $SuperM(\delta,x)$. To summarize, we have to prove now the following three additional properties:

1. $\Delta[T_1 v(x)-v(x)]\geq 0$ for $SuperM(\mu,x)$
2. $\Delta[T_2 v(x)-v(x)]\leq 0$ for $SubM(\lambda,x)$
3. $\Delta[T_3 v(x)-v(x)]\geq 0$ for $SuperM(\delta,x)$

We have omitted the subscript α since these properties hold independently of α. From Cil et al. (2009), we know that convexity of $v(x)$ in x (true since $v\in\mathcal{U}$ implies that $T_1 v(x)-v(x)$ is non-decreasing in x. Let's prove now that $T_2 v(x)-v(x)$ is non-increasing in x. From Equation (9), we have:

$$
\Delta(T_2 v(x)-v(x))=\begin{cases}\Delta v(x-1)-\Delta v(x)\leq 0 & \text{if } x\geq 1\\ -c_l-\Delta v(x)\leq 0 & \text{if } x=0\end{cases}
\tag{12}
$$

Inequalities in (12) come from $v\in\mathcal{U}$. Finally, $T_3 v(x)-v(x)=\Delta v(x)+c_r$ is non-decreasing in x since v is convex.

Proof of Property 1

To prove that C is convex, we will separately prove that C_h, C_r and C_l are convex.

Case 1: $\rho_1 < 1$

We have

$$\frac{\partial^2 Ch(S)}{\partial S^2} = A_1 \times A_2 \times A_3 \times A_4$$

with

$$\begin{aligned}
A_1 =& c_h \rho_1 \ln \rho_1 \rho_1^S \\
A_2 =& 1/((q\rho_1^{S+1} + p(\rho_1 - 1) - q)^3) \\
A_3 =& q\rho 1^{S+1} \left[(p(\rho_1 - 1) - q) q \ln \rho_1 S - p \left(q(\ln \rho_1 + 2(\rho_1 - 1))\right.\right. \\
& \left.\left. -(\rho_1 - 1)\ln \rho_1\right) - q^2(\ln \rho_1 - 2)\right] \\
A_4 =& 1 - (p(\rho_1 - 1) - q)
\end{aligned}$$

In order to prove that C_h is convex, it is sufficient to prove that $A_1, A_2 < 0$ and $A_3, A_4 > 0$. It is clear that $A_1, A_2 < 0$ and $A_4 > 0$, due to $\rho_1 < 1$. Now $A_3 > 0$ since we have $A_3 = A_3^1 + A_3^2 + A_3^3$ where

$$\begin{aligned}
A_3^1 =& \underbrace{(p(\rho_1 - 1) - q)}_{<0} \underbrace{q \ln \rho_1 S}_{<0} > 0 \\
A_3^2 =& -p(\underbrace{q(\ln \rho_1 + 2(\rho_1 - 1))}_{<0} \underbrace{-(\rho_1 - 1)\ln \rho_1}_{<0}) > 0 \\
A_3^3 =& -q^2 \underbrace{(\ln \rho_1 - 2)}_{<0} > 0
\end{aligned}$$

We can conclude that C_h is convex.

C_r is clearly convex since it is a constant function of S and C_l is convex since

$$\frac{\partial^2 C_l(S)}{\partial S^2} = \frac{\overbrace{-(p(\rho_1 - 1) - q)}^{>0} \overbrace{q(\rho_1 - 1)}^{<0} \overbrace{(\ln \rho_1)^2}^{>0} \overbrace{\rho_1^S}^{>0} \overbrace{[q\rho_1^{S+1} \overbrace{-p(\rho_1 - 1)}^{>0} + q]}^{>0}}{(\underbrace{q\rho_1^{S+1} - q}_{<0} + \underbrace{p(\rho_1 - 1)}_{<0})^3}$$

$$> 0$$

Finally C is convex for $\rho_1 < 1$.

Case 2: $\rho_1 = 1$

In this case, the second derivatives simplify to:

$$\frac{\partial^2 Ch(S)}{\partial S^2} = c_h \frac{-q(q-1)}{(qS+1)^3} \geq 0$$
$$\frac{\partial^2 C_l(S)}{\partial S^2} = \lambda c_l \frac{2q^3}{(1+qS)^3} \geq 0$$

and we conclude that C is also convex for $\rho_1 = 1$.

Proof of Lemma 3

Assume that $v \in \tilde{\mathcal{U}}$. Then $T_1 v$ is convex (see the proof of Lemma 1). We prove, in the following, that $\tilde{T}_2 v$ is also convex.

We have

$$\Delta \tilde{T}_2 v(x) = \begin{cases} q\Delta v(x-1) + p\Delta v(x) \geq -q\eta - p\eta \geq -\eta & \text{if } x > 0 \\ p\Delta v(x) - q\eta \geq -p\eta - q\eta \geq -\eta & \text{if } x = 0 \end{cases} \tag{13}$$

and

$$\Delta^2 \tilde{T}_2 v(x) = \begin{cases} q\Delta_x^2 v(x-1) + p\Delta^2 v(x) \geq 0 & \text{if } x > 0 \\ q[\Delta v(x) + \eta] + p\Delta^2 v(x) \geq 0 & \text{if } x = 0 \end{cases}$$

The quantity $\Delta^2 \tilde{T}_2 v(x)$ is non-negative from Condition C.1 when $x \geq 1$ and from Condition C.2 when $x = 0$, thus $T_2 v$ satisfies C.1. Otherwise $T_3 v$ and $c_h x$ are trivially convex. Finally $\tilde{T} v$, as a non-negative linear combination of convex functions, is also convex.

Let us prove now that Tv satisfies Condition C.2. We have $\tilde{T}_2 v$ satisfying C.2 from Equations (13). Moreover $T_1 v$ satisfies C.2 since

$$\Delta T_1 v(x) = \begin{cases} \Delta v(x+1) \geq -\eta & \text{if } x < S-1 \\ 0 \geq -c_l & \text{if } x = S-1 \\ \Delta v(x) \geq -\eta & \text{if } x \geq S. \end{cases}$$

Therefore, we have

$$\begin{aligned}\Delta\tilde{T}v(x) &= \frac{1}{\gamma}\left[\Delta c_h x + \mu\Delta T_1 v(x) + \lambda\Delta\tilde{T}_2 v(x) + (\gamma-\beta-\lambda-\mu)\Delta v(x)\right] \\ &\geq \frac{1}{\gamma}\left[-\mu\eta - \lambda\eta - (\gamma-\beta-\lambda-\mu)\eta\right] \\ &\geq -\frac{\gamma-\beta}{\gamma}\eta \geq -\eta\end{aligned}$$

and $\tilde{T}v$ satisfies Condition C.2.

Proof of Theorem 4

Let $v_\alpha \in \tilde{\mathcal{U}}$. Then $\tilde{T}_\alpha v_\alpha \in \tilde{\mathcal{U}}$ from Lemma 3. We need now to show that $\tilde{T}_\alpha$ propagates submodular and supermodular properties.

We first prove that operators $c_h x$, T_1, $\tilde{T}_2$ preserve the modular properties summarized in Table 3. Notice that these operators might depend on the parameter α under consideration.

	$SuperM(\alpha,x)$	$SubM(\alpha,x)$
$c_h x$	$\forall\alpha$	$\forall\alpha \neq c_h$
$T_1 v_\alpha(x)$	$\forall\alpha$	$\forall\alpha$
$\tilde{T}_2 v_\alpha(x)$	$\forall\alpha \neq c_l$	$\forall\alpha \notin \{p, c_r\}$

Table 3: Preservation of submodularity and supermodularity by the operators

$\Delta c_h x = c_h$ is non-decreasing in c_h and independent of the other parameters of the system. Therefore $c_h x$ is $SuperM(\alpha,x)$ for all α and $SubM(\alpha,x)$ for all $\alpha \neq c_h$. From Cil et al. (2009) (we adapt their maximization problem into a minimization problem), we know that T_1 preserve $SubM(\alpha,x)$ and $SuperM(\alpha,x)$ for all α.

From Equation (13), we have

$$\Delta\tilde{T}_2 v_\alpha(x) = \begin{cases} q\Delta v_\alpha(x-1) + p\Delta v_\alpha(x) & \text{if } x>0 \\ p\Delta v_\alpha(x) - q\eta & \text{if } x=0 \end{cases}$$

When $\alpha \notin \{c_l, p, c_r\}$, it is clear that $\tilde{T}_2$ propagate $SubM(\alpha, x)$ and $SuperM(\alpha, x)$. $\tilde{T}_2$ propagates $SubM(c_l, x)$ since

$$\Delta\tilde{T}_2 v_{c_l+\epsilon}(x) - \Delta\tilde{T}_2 v_{c_l}(x) = \begin{cases} q[\Delta v_{c_l+\epsilon}(x-1) - \Delta v_{c_l}(x-1)] \\ \quad +p[\Delta v_{c_l+\epsilon}(x) - \Delta v_{c_l}(x)] \leq 0 & \text{if } x \geq 1 \\ p[\Delta v_{c_l+\epsilon}(x) - \Delta v_{c_l}(x)] - q\epsilon \leq 0 & \text{if } x = 0 \end{cases}$$

$\tilde{T}_2$ propagates $SuperM(c_r, x)$ since

$$\Delta\tilde{T}_2 v_{c_r+\epsilon}(x) - \Delta\tilde{T}_2 v_{c_r}(x) = \begin{cases} q[\Delta v_{c_r+\epsilon}(x-1) - \Delta v_{c_r}(x-1)] \\ \quad +p[\Delta v_{c_r+\epsilon}(x) - \Delta v_{c_r}(x)] \geq 0 & \text{if } x \geq 1 \\ p[\Delta v_{c_r+\epsilon}(x) - \Delta v_{c_r}(x)] + qp\epsilon \geq 0 & \text{if } x = 0 \end{cases}$$

$\tilde{T}_2$ propagates $SuperM(p, x)$ since

$$\Delta\tilde{T}_2 v_{p+\epsilon}(x) - \Delta\tilde{T}_2 v_p(x) = \begin{cases} q[\Delta v_{p+\epsilon}(x-1) - \Delta v_p(x-1)] \\ \quad +p[\Delta v_{p+\epsilon}(x) - \Delta v_p(x)] + \epsilon\Delta^2 v_{p+\epsilon}(x) \geq 0 & \text{if } x \geq 1 \\ p[\Delta v_{p+\epsilon}(x) - \Delta v_p(x)] + \epsilon[\Delta v_{p+\epsilon}(x) + c_r] \\ \quad \geq \epsilon(-\eta + c_r) \geq 0 & \text{if } x = 0 \end{cases}$$

When α belongs to $\{c_h, c_r, c_l\}$, $T_\alpha v_\alpha$ is a linear combination of operators $c_h x$, T_1, $\tilde{T}_2$ where the coefficients of the combination do not depend on c_h, c_r, c_l. We can conclude that $T_\alpha v_\alpha$ is $SuperM(\alpha, x)$ for $\alpha \in \{c_h, c_r\}$ and $SubM(\alpha, x)$ for $\alpha \in \{c_l, c_r\}$.

When α belongs to $\{\lambda, \mu\}$, it is more complicated (see proof of Lemma 2) and we have to prove the following two additional properties:

1. $\Delta[T_1 v(x) - v(x)] \geq 0$ for $SuperM(\mu, x)$
2. $\Delta[\tilde{T}_2 v(x) - v(x)] \leq 0$ for $SubM(\lambda, x)$

We have omitted the subscript α since these properties hold independently of α. From Cil et al. (2009), we know that convexity of $v(x)$ in x (true since $v \in \mathcal{U}$ implies that $T_1 v(x) - v(x)$ is non-decreasing in x. Let's prove now that $\tilde{T}_2 v(x) - v(x)$ is non-increasing in x. From Equation (13), we have:

$$\Delta(\tilde{T}_2 v(x) - v(x)) = \begin{cases} q\Delta v(x-1) + p\Delta v(x) - \Delta v(x) = -\Delta^2 v(x) \leq 0 & \text{if } x \geq 1 \\ -q[\Delta v(x) + \eta] \leq 0 & \text{if } x = 0 \end{cases} \tag{14}$$

Inequalities in (14) come from $v \in \tilde{\mathcal{U}}$.

Value iteration algorithm to compute the average cost of a base-stock policy when the return lead-time is positive

Let $\tilde{C}(S)$ denote the average cost when the base-stock level is S. Let $h(x, y)$ denote the relative value function for initial state (x, y). In order to be able to uniformize this MDP, we assume that y is bounded by M. This is not a crucial assumption since our results will hold for any M. We can now uniformize (Lippman, 1975) the MDP with rate $c = \lambda + \mu + \gamma M$. The optimal value function can be shown to satisfy the following optimality equations:

$$h(x, y) + \tilde{C}(S) = Th(x, y), \; \forall \; (x, y) \; \in \; \mathbb{N} \times \mathbb{N}$$

where the operator T is a contraction mapping defined as

$$Th(x, y) \;=\; \frac{1}{c+\beta}\left[c_H x + \mu T_0 h(x, y) + \lambda p T_1 h(x, y) + \lambda(1-p) T_2 h(x, y) + \gamma T_3 \; h(x, y)\right]$$

and

$$T_0 h(x, y) \;=\; \begin{cases} h(x+1, y) & \text{if } x < S \\ h(x, y) & \text{if } x \geq S \end{cases}$$

$$T_1 h(x, y) \;=\; \begin{cases} h(x-1, y+1) & \text{if } x > 0 \text{ and } y < M \\ h(x, y) + c_l & \text{if } x = 0 \\ h(x-1, y) & \text{if } x > 0 \text{ and } y = M \end{cases}$$

$$T_2 h(x, y) \;=\; \begin{cases} h(x-1, y) & \text{if } x > 0 \\ h(x, y) + c_l & \text{if } x = 0 \end{cases}$$

$$T_3 h(x, y) \;=\; \begin{cases} y\,[h(x+1, y-1) + c_r] + (M-y)h(x, y) & \text{if } y > 0 \\ M h(x, y) & \text{if } y = 0 \end{cases}$$

Operator T_0 is associated to the optimal production decision. Operators T_1 (resp. T_2) is associated to a demand that will (resp. not) lead to a return. Finally, operator T_3 corresponds to the return of a product.

Based on these optimality equations, we used a value iteration algorithm (Puterman, 1994) to compute the average cost.