Influence of product return lead-time on inventory control

Mohamed Hichem Zerhouni, Jean-Philippe Gayon, Yannick Frein

International Conference
on Industrial Engineering and Systems Management
IESM 2009
May 13-15
MONTRÉAL - CANADA

Influence of product return lead-time on inventory control

Hichem Zerhouni, Jean-Philippe Gayon, Yannick Frein

Laboratoire G-SCOP, Grenoble INP-UJF-CNRS, 46 avenue Félix Viallet - 38031 Grenoble Cedex 1 - France

Abstract

We consider a production-inventory system with capacity constraint and item returns correlated to demand. The system is modelled by an $M/M/1$ make-to-stock queue with lost sales. A satisfied demand incurs an item return with a certain probability, after an exponentially distributed return lead-time.

We distinguish two cases: When the number of items to be returned is observable and when it is not observable. For the first case, we partially characterize the optimal policy. For the second case, we consider a simple base-stock policy as a heuristic. Finally, we carry out a numerical study in order to investigate the impact of return lead-time on these policies. In particular, we exhibit interesting limit behavior, when the return lead-time is either small or large.

Key words: Inventory control, Reverse logistics, Queueing system, Markov decision process.

1 Introduction

We develop, in this article, a model which can be included in the growing field of reverse logistics [3]. Reverse logistics is the management of products that can be returned by the customers for different reasons (ecological, legal, economical). From a logistic point of view, and regardless of why they occur, product returns complicate the management of an inventory system [2]. First, returns represent 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 [4]. Several researches have investigated the influence of item returns on inventory control. For a complete overview, we refer the reader to Fleischmann et al [3].

Problems encountered in reverse logistics can be classified into three main categories [3]:

(1) Problems related to transportation and distribution of the returned items.
(2) Problems related to inventory control.
(3) Problems related to treatment of the returned items.

We focus on the second category which represents more than 50% of published articles these ten last years in the field of reverse logistics [9].

Email addresses: `hichem.zerhouni@g-scop.inpg.fr` (Hichem Zerhouni), `jean-philippe.gayon@g-scop.inpg.fr` (Jean-Philippe Gayon), `yannick.frein@g-scop.inpg.fr` (Yannick Frein).

We are interested by a system where the producer must face its own item returns. One of the specificity of our model is that item returns are directly correlated with demand. An item can be returned, with a certain probability, after a random return lead-time.

Most of the literature assume that item returns are uncorrelated with demand. DeBrito & Dekker [1] have explored this assumption and conclude that it is not realistic in many real situations. However, they do not explore the influence, of ignoring the correlation between demand and returns, on inventory control policies.

A few authors have investigated inventory control problems with returns correlated to demand. Kiesmüller & van der Laan [6] develop a periodic review model with fixed 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. Yuan & Cheung [10] consider a continuous review model that interests us particularly because they assume dependent returns, exponential return lead-time, Poisson demand and instantaneous procurement lead-time. They adopt an (s, S) inventory policy and develop an algorithm to compute the optimal replenishment parameters.

Gayon [5] considers an $M/M/1$ make-to-stock queue with returns. Demand and returns are assumed to be independent Poisson processes. He proves that the optimal policy is a base-stock policy and computes analytically the optimal base-stock level for both average and discounted cost problem. Zerhouni et al. [11] consider the lost-sale version of [5]. They assume that returns are either independent of demand or can be returned instantaneously. For both models, they establish the optimality of base-stock policies and derive monotonicity results of the optimal base-stock levels, with respect to system parameters.

In this paper, we generalize [11] by considering non-zero return lead-times, assumed to be exponentially distributed. We distinguish the cases (1) when the decision maker knows, when a demand is satisfied, if the item will be returned or not, and (2) when he does not have this information. In the *observable* case, we characterize the optimal policy. In the *non-observable* case, we propose a simple base-stock policy heuristic which reveals to be very performant. Moreover, we exhibit numerically interesting behaviors when the return lead-times are either small or large.

The remainder of this paper is organized as follows. Section 2 gives a formal definition of the models we are interested in. In Section 3, we characterize the optimal policy for the observable case. In Section 4, we compare numerically the optimal policy of the observable case with the heuristic for the non-observable case. Finally, Section 5 concludes and suggests future research.

2 Model formulation

We consider an $M/M/1$ make-to-stock queue producing a single item. The supplier can decide at any time to produce this item. The processing time is exponentially distributed with mean $1/\mu$ and completed items are stored in a serviceable products inventory. Demands for those items follow a Poisson distribution with rate λ. A demand that cannot be fulfilled immediately, when the inventory is empty, is lost. We also suppose that each satisfied demand incurs, with a certain probability p, a product return with a lead-time exponentially distributed with rate γ. The returned items may subsequently be used to serve the demand also (Figure 1). The inventory is common to new and returned items. The state of the system can 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 have to distinguished two cases:

- The observable case: The decision maker knows $Y(t)$ at each time and makes his decisions based on this information.
- The non observable case: The decision maker does not know the number of $Y(t)$.

The cost structure is detailed in the following table which also summarizes the main notations used in this paper.

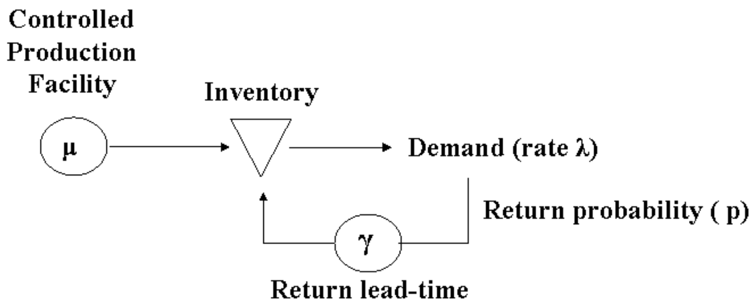

Fig. 1. Dependent returns with non-zero lead-time

x	= stock level (including new and returned products)
y	= number of products to be returned
μ	= production rate
λ	= demand rate
γ	= return lead-time rate
p (q)	= probability that a satisfied demand is (not) returned
c_h	= inventory holding cost per unit per time unit
c_p	= unit production cost
c_l	= unit lost sales cost
c_r	= unit return cost

A policy π specifies when to produce or not, for each state of the system. The objective of the decision maker is to find the policy minimizing expected discounted cost over an infinite time horizon. This problem can be modelled as a continuous-time Markov Decision Process (MDP) in the observable case. In the non-observable case, the problem could be modelled as a Partially Observed MDP (POMDP).

3 Algorithm for the optimal discounted cost in the observable case

The structure of the optimal policy in the observable case is discussed in Section 4. In the current section, we present an algorithm which allows to compute the optimal discounted cost associated to the optimal policy in the observable case. We restrict our analysis to stationary markovian policies since there exists an optimal stationary markovian policy [8]. We define the following stochastic processes when policy π is applied, initial stock level is x and number of products to be returned is y:

- $U^{\pi}_{x,y}(t)$: number of customer demand that have not been satisfied up to time t
- $V^{\pi}_{x,y}(t)$: number of products that have been returned up to time t
- $W^{\pi}_{x,y}(t)$: number of products that have been produced up to time t
- $X^{\pi}_{x,y}(t)$: stock level at time t

We define also $v^{\pi}(x,y)$ as the expected total discounted cost associated to policy π, initial state (x,y). If α is the discounted factor ($\alpha > 0$) then $v^{\pi}(x,y)$ is the sum of production costs, holding costs, lost-sale costs and return costs:

$$v^{\pi}(x,y) = E\left[\int_0^{\infty} e^{-\alpha t}\left(c_l\ dU^{\pi}_{x,y}(t) + c_r\ dV^{\pi}_{x,y}(t) + c_p\ dW^{\pi}_{x,y}(t) + c_h\ dX^{\pi}_{x,y}(t)\right)\right]$$

We seek to find the optimal policy π^* minimizing $v^{\pi}(x,y)$ and we let $v^*(x,y) := v^{\pi^*}(x,y)$ denote the optimal value function:

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

This problem can be modelled as a continuous-time Markov Decision Process (MDP). 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 [7] the MDP with rate $c = \lambda + \mu + \gamma M$ and the optimal value function can be shown to satisfy the following optimality equations:

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

where the operator T is a contraction mapping defined as

$$Tv(x,y) := \frac{1}{c+\alpha} \; [c_H \; x + \mu \; T_0 \; v(x,y) + \lambda \; p \; T_1 \; v(x,y) + \lambda \; (1-p) \; T_2 \; v(x,y) + \; \gamma \; T_3 \; v(x,y)]$$

and

$$\begin{aligned}
T_0 \; v(x,y) \; &:= \; \min\left[v(x,y); c_p + v(x+1,y)\right] \\
T_1 \; v(x,y) \; &:= \; \begin{cases} v(x-1,y+1) & \text{if } x > 0 \;\text{ and } y < M \\ v(x,y) + c_l & \text{if } x = 0 \\ v(x-1,y) & \text{if } x > 0 \;\text{ and } y = M \end{cases} \\
T_2 \; v(x,y) \; &:= \; \begin{cases} v(x-1,y) & \text{if } x > 0 \\ v(x,y) + c_l & \text{if } x = 0 \end{cases} \\
T_3 \; v(x,y) \; &:= \; \begin{cases} y \; [v(x+1,y-1) + c_r] + (M-y) \; v(x,y) & \text{if } y > 0 \\ M \; v(x,y) & \text{if } y = 0 \end{cases}
\end{aligned}$$

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, the optimal value function v^* and the optimal policy can be computed by value iteration [8].

4 Structure of the optimal policy in the obervable case

On the basis of all the numerical studies that we carried out according to the approach presented in section 3, we conjecture that the optimal policy in the observable case is of the following type:

Conjecture 1 *There exists a switching curve $S(y)$ such that it is optimal to produce if and only if $x < S(y)$. Moreover, the switching curve has the following property:*

$$S(y) - 1 \leq S(y+1) \leq S(y)$$

The term ”switching” in Conjecture 1 refers to the border between producing and not producing.
Figure 2 presents the optimal policy for a numerical exemple.

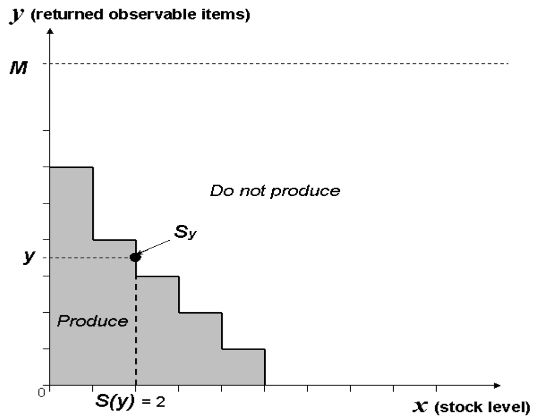

Fig. 2. Optimal policy : dicounted problem & observable returns

$$(\lambda = 0.9, \mu = 1, p = 0.10, \alpha = 0.1, \gamma = 10, c_h = 1, c_p = 0, c_l = 45, c_r = 15)$$

The analytical proof of the Conjecture 1 is hard to establish due the the high correlation between demand and returns. Hence, we consider the particular case of instantaneous returns ($1/\gamma \to 0$) and we prove that the structure of the optimal policy in the observable case with immediate returns is a base-stock policy. When items are instantaneously returned, the optimal discounted cost does not depend any more of y and is given by $v^*(x) = Tv^*(x),\ \forall\ (x)\ \in\ \mathbb{N}$. Where the operator T is defined as:

$$Tv(x) := \frac{1}{\lambda + \mu + \lambda\ p + \alpha}\ \left[c_H\ x + \mu\ T_0\ v(x) + \lambda\ T_1\ v(x)\right]$$

and

$$T_0\ v(x)\ :=\ \min\left[v(x); c_p + v(x+1)\right]$$

$$T_1\ v(x)\ :=\ \begin{cases} q\ v(x-1) + p\ v(x) + p\ c_r & \text{if } x > 0 \\ v(x) + c_l & \text{if } x = 0 \end{cases}$$

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.*

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 $S^* = \min[x \mid \Delta v^*(x) > 0]$, possibly infinite, such that $v^* \leq 0$ if and only if x is below this threshold. Optimality equations state to produce when $\Delta v^* < 0$ and to idle production when $\Delta v^* > 0$. If $\Delta v^*(x) = 0$ 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:* $\Delta v(x+1) \geq \Delta v(x)$ *(*$\Delta^2 v(x) \geq 0$*)*
- *Condition C.2:* $q\Delta v(x) + c_l - p\ cr \geq 0$

The first condition states convexity of v while the second one means that it is preferable to satisfy an arriving demand. 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.

Theorem 1 *If $v \in U$ then $Tv \in U$.*

The proof is given in Appendix.

5 Heuristic for the non-observable case

When the number of products to be returned is not observable, we consider a heuristic, based only on the inventory level. This heuristic is a base-stock policy such that the system produces if and only if the stock level x is smaller than S. We denote by $\tilde{v}_S(x, y)$ the infinite horizon discounted cost associated to this policy. The value function $\tilde{v}_S(x, y)$ satisfies the same equations as v^* except that operator T_0, associated to the production decision, is replaced by operator $\tilde{T_0}$ defined as:

$$\tilde{T_0}v(x, y) := \begin{cases} v(x+1, y) + c_p & \text{if } x < S \\ v(x, y) & \text{otherwise} \end{cases}$$

One can then search for the optimal base-stock level $\tilde{S}$ minimizing $\tilde{v}_S(0, 0)$. We then denote $\tilde{v}(0, 0) := \min_S \tilde{v}_S(0, 0)$.
We emphasize that the computation of $\tilde{S}$ does not use informations on y, though the computation of the heuristic cost requires y.

A more elaborate policy could take into account $N(t)$, the number of items that are in the market (number of sold items minus number of returned items). Then we know that $Y(t)$, the number of products to be returned, is distributed according to a binomial distribution $\mathcal{B}(N(t), p)$. We don't consider this policy in this paper.

6 Numerical study

In this section, we compare the performance (cost) of the base-stock policy presented in Section 5 with the optimal cost calculated in Section 3. We have computed the optimal policy and the heuristic policy for more than 20000 instances corresponding to the combinations of : $c_p, c_r, c_l \in \{0, 10, \cdots, 100\}$, $\lambda \in \{0.8, 1, 1.2\}$, $p \in \{0, 0.3, 0.6, 0.9\}$, $\gamma \in \{0.0001, 0.001, 0.1, \cdots, 1000\}$. In all our numerical studies we set $\mu = 1$ and $c_h = 1$ without loss of generality, since it is equivalent to set a monetary and a time unit.

In Figure 3, we represent the following discounted costs with respect to the average return lead-time, $1/\gamma$, for an initial state of the system $(0, 0)$:

- $v^*(0, 0)$: Optimal discounted cost in the observable case (see Section 3)
- $\tilde{v}(0, 0)$: Discounted cost of the best heuristic policy in the non-observable case (see Section 5)
- $v^*_{p=0}(0, 0)$: Optimal discounted cost when there is no product return. This cost is computed by setting the return probability to 0 in Section 3
- $v^*_{\gamma=\infty}(0, 0)$: Discounted cost of the optimal policy with zero return lead-time (see [11] for a detailed description of this special case), which can be seen as the limit of $v^*(0, 0)$ when γ goes to infinity

We also represent in Figure 3 the optimal base-stock levels corresponding to the discounted costs described above.

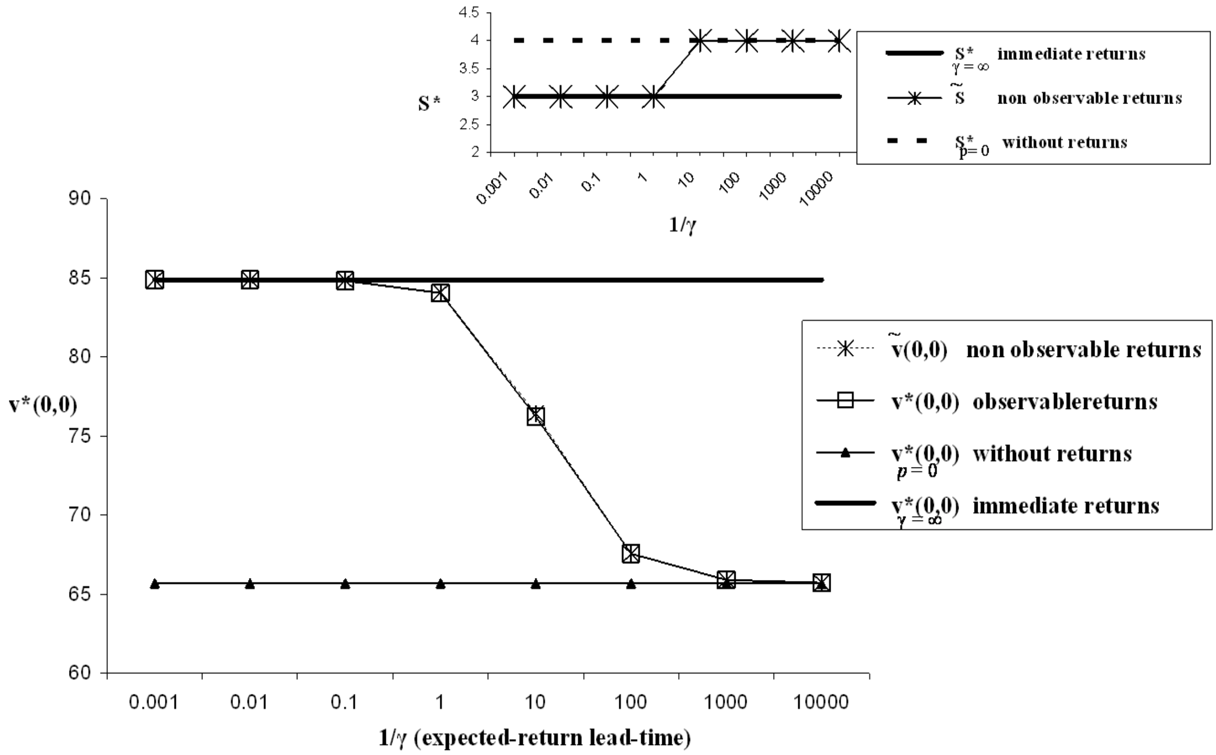

Fig. 3. Comparison oh heuristic and optimal policies

$$(\lambda = 0.9, \mu = 1, p = 0.25, \alpha = 0.1, c_h = 1, c_p = 0, c_l = 20, c_r = 16)$$

We summarize below our main observations related to Figure 3 and to other numerical results:

- The optimal base-stock levels ($S(y)$ and $\tilde{S}$) are non-decreasing in the expected return lead-time. It is clear that shorter return lead-times imply to reduce production anticipation and though the base-stock levels.
- The optimal policy always performs better than the heuristic policy even if the Figure 3 does not show it clearly. This result is obvious since it exploits more informations on the system.
- The performances of the optimal policy and the heuristic policy are very close when the expected return lead-time, $1/\gamma$, is either very small or very large. When expected return lead-time is small, the information on y, the number of products to be returned, can not be exploited in an efficient way. When expected return lead-time is large, the discount factor α decreases the effects of future returns. We have also carried a numerical analysis with the average cost criterion, instead of the discounted criterion. We still observe that the costs of both policies converge when $1/\gamma$ goes to infinity. When expected return lead-time is large, the standard deviation of the return lead-time, $1/\gamma$, is also large. This makes observability of y less useful, since as expected value of return lead-time increases so does the variance.
- We also observe, when the expected return lead-time goes to infinity, that both costs converge to the optimal cost of a system without product returns. This is again due to the effect of the discounted rate, α, and it is not observed with an average cost criterion.
- There are values in the middle range for expected return lead-times for which the relative benefit of observing y is maximum. However, the maximum benefit we have observed was less than 2%. Therefore, a simple base-stock policy is a good heuristic in many cases.

7 Conclusion and future research

In this paper, we have considered production-inventory system with product returns correlated to demand through an exponential return lead-time. We have distinguished two cases: When the number of products to be returned is observable and when it is not observable. We have partially characterized the optimal policy in the observable case. In the numerical study, we have compared this optimal policy with a heuristic for the non-observable case. It appears that the heuristic performs very well when the expected return lead-time is either very small or very large. Even for intermediate values of the expected return lead-time, the relative cost difference does not exceed two percents on the computed instances.

We have assumed in this paper that all returned products are systematically accepted in the inventory and this assumption could be relaxed. Another avenue for research would be to include a remanufacturing stage for returned products and to consider the joint problem of controlling manufacturing and remanufacturing.

References

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

[2] G. DeCroix, J-S. Song, and P. Zipkin. A series system with returns: Stationary analysis. *OPERATIONS RESEARCH*, 52 NO. 2:350–362, 2005.

[3] M. Fleischmann, J. M. Bloemhof-Ruwaard, R. Dekker, E. A. 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.

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

[5] J-P. Gayon. A make-to-stock queue with product return. *Working paper (http://www.g-scop.inpg.fr/ gayonj/)*, 2007.

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

[7] S.A. Lippman. Applying a new device in optimization of exponential queuing systems. *Operations Research*, 23(4):687–710, 1975.

[8] M.L Puterman. *Markov Decision Processes, discrete stochastic dynamic programming.* Wiley-Interscience, 1994.

[9] Sergio Rubio, Antonio Chamorro, and Francisco J. Miranda. Characteristics of the research on reverse logistics (1995-2005). *International journal of production research*, 46 No. 4:1099–1120, february 2008.

[10] X. M. Yuan and K. L. Cheung. Modeling returns of merchandise in an inventory system. *OR Spektrum*, 20:147–154, 1998.

[11] H. Zerhouni, J.P. Gayon, and Y. Frein. Control of a production-inventory system with product returns dependent on demand. Working paper (http://www.g-scop.inpg.fr/ gayonj/Documents/PublicationsJP/WorkingPapers/ProductReturns.pdf).

[12] H. Zerhouni, J.P. Gayon, and Y. Frein. Control of a production-inventory system with product returns dependent on demand. In *International Conference on Information Systems, Logistics and Supply Chain*, 2008.

Appendix

Optimal policy in the observable case with instantaneous returns.

In order to prove that $v^* \in \mathcal{U}$, let us prove that T preserves conditions C.1 and C.2.
To show that $Tv \in \mathcal{U}$, we have to show that $T_0 v$ and $T_1 v \in \mathcal{U}$ as $Tv = (T_0 + T_1)\ v$. The proof is given for $c_l \geq p\ cr$.

Operator T_0

We define $S_y = \min[x : \Delta v(x, y) > 0]$ (S_y possibly infinite). So we obtain :

$$\Delta_x T_0\ v(x) = \begin{cases} \Delta_x v(x+1) & \text{if } x < S \\ 0 & \text{if } x = S-1 \\ \Delta_x v(x) & \text{if } x \geq S \end{cases}$$

$$\Delta_x^2 T_0\ v(x) = \begin{cases} \Delta_x v(x+1)^2 \geq 0 & \text{if } x < S \\ -\Delta_x v(x+1) \geq 0 & \text{if } x = S-2 \\ \Delta_x v(x+1) \geq 0 & \text{if } x = S-1 \\ \Delta_x v(x)^2 \geq 0 & \text{if } x \geq S \end{cases}$$

$$q\ \Delta_x T_0\ v(x, y) + c_l - p\ c_r = \begin{cases} \Delta_x v(x+1) + c_l - p\ c_r \geq 0 & \text{if } x < S \\ c_l - p\ c_r \geq 0 & \text{if } x = S-1 \\ q\ \Delta_x v(x) + c_l - p\ c_r \geq 0 & \text{if } x \geq S \end{cases}$$

Operator T_1

$$\Delta_x T_1\ v(x) = \begin{cases} q\ \Delta_x v(x-1) + p\ \Delta_x v(x) & \text{if } x > 0 \\ p\ \Delta_x v(x) + p\ cr - c_l & \text{if } x = 0 \end{cases}$$

$$\Delta_x^2 T_1\ v(x) = \begin{cases} q\ \Delta_x v(x-1)^2 + p\ \Delta_x^2 v(x) \geq 0 & \text{if } x > 0 \\ q\ \Delta_x v(x) + c_l - p\ cr + p\ \Delta_x^2 v(x) \geq 0 & \text{if } x = 0 \end{cases}$$

$$q\ \Delta_x T_1\ v(x) + c_l - p\ cr = \begin{cases} q\ [q\ \Delta_x v(x-1) + p\ \Delta_x v(x)] + c_l - p\ c_r \geq q\ \Delta_x v(x-1) + c_l - p\ c_r \geq 0 & \text{if } x > 0 \\ q\ [-c_l + p\ \Delta_x v(x) +\ cr] + c_l - p\ c_r = 0 & \text{if } x = 0 \end{cases}$$

We found that if $v \in \mathcal{U}$ then Tv satisfies confitions C.1, C.2, C.3 and C.4. Hence, $Tv \in \mathcal{U}$. ■