Optimal Control of a Production-Inventory System with Customer Impatience

Saif Benjaafar[†], **Jean-Philippe Gayon**[*] **and Seda Tepex**[†]

[†] Industrial and System Engineering, University of Minnesota, Minneapolis, MN 55455, USA
saif@umn.edu, tepex007@umn.edu

[*] Laboratoire G-scop, Grenoble INP, 46 Avenue Félix Vialet, 38031 Grenoble Cedex, France
jean-philippe.gayon@grenoble-inp.fr

Abstract

We consider the optimal control of a production-inventory system with a single product where items are produced one unit at a time. Upon arrival, customer orders can be fulfilled from existing inventory, if there is any, backordered, or rejected. Although customers are willing to wait if their orders are backlogged, they are not infinitely patient. In particular, customers cancel their orders if their waiting time in backlog exceeds a certain patience time. This patience time is random and varies from one customer to another. At each decision epoch, we must determine whether or not to produce an item and, should an order arise and there is no inventory on-hand, whether to reject the order or backorder it. Rejecting an order incurs a rejection cost. If an order is backordered but the order ends up being cancelled by the customer, we incur a cancellation cost. We can mitigate both costs by holding inventory but we incur an inventory holding cost. We formulate the problem as a Markov decision process. We show that the optimal policy can be described by two thresholds: a production base-stock level that determines when production takes place and an admission threshold that determines when orders should be accepted. We also characterize analytically the sensitivity of these thresholds to operating parameters, including the demand and production rates and the various cost parameters. Using the structure of the optimal policy, we formulate the dynamics of the corresponding production-inventory system as a Markov chain, which allows us to compute efficiently the performance of the system for any choice of base-stock level and admission threshold, In a numerical study, we compare the performance of the optimal policy against several other policies and show that those that ignore customer impatience can perform poorly.

Keywords: Production-inventory systems, customer impatience, optimal control, make-to-stock queues, Markov decision processes

1 Introduction

Inventory problems treated in the literature fall mostly into two categories. One deals with systems where customers are assumed to be infinitely patient, so that a customer whose order is backlogged is willing to wait for that order to be fulfilled no matter how long it takes. The other deals with systems where customers have zero patience, so that a customer whose order cannot be fulfilled immediately is considered lost. However, in practice, it is more common for customers to be willing to wait, but only up to a point. Customers whose orders are backordered eventually cancel their orders and leave if their waiting time in backlog exceeds a certain *patience time*. This patience time usually varies from one customer to another and, for the same customer, may vary from one ordering instance to the next. Despite the prevalence of such behavior in practice, there is limited literature that deals with this issue. Consequently very little is known about optimal control policies, or even effective heuristics, for such systems. Very little is also known about the impact of not accounting for customer impatience in making inventory decisions.

In this paper we address some of these limitations in the context of a production-inventory system with a single product. In particular, we consider a continuous time and continuous review system where demand orders arrive continuously over time one unit at a time with stochastic inter-arrival times. With each order arrival, a decision must be made regarding whether to fulfill the order from on-hand inventory, backorder it, or reject it. If an order is rejected, a rejection cost (e.g., a lost sale cost) is incurred. If an order is backordered but the order ends up being cancelled by the customer due to impatience, a cancellation cost (e.g., the sum of a lost sale cost and a penalty for loss of goodwill) is incurred. We can mitigate both costs by increasing the number of items held in inventory, but there is a linear cost for keeping inventory on-hand. Inventory is replenished from a production facility that produces units one at a time with stochastic production times. At any point in time, the system manager must decide on whether or not to produce and whether or not to accept an incoming order, should one arise.

We formulate the problem as a Markov decision process (MDP) and use it to characterize the structure of the optimal policy. We show that the optimal policy can be described by two thresholds: a production base-stock level and an admission threshold. The production base-stock level determines when production takes place while the admission threshold determines when orders should be accepted. We also characterize analytically the sensitivity of these thresholds to operating parameters, including the demand and production rates and the various cost parameters. Using the structure of the optimal policy, we model the dynamics of the corresponding production-inventory system as a Markov chain which allows us to compute effeciently the performance of the system for any choice of base-stock level and admission threshold. Using numerical results, we compare the performance of the optimal policy against several other policies and show that those that do not account for impatience can perform poorly.

In the existing inventory literature, the issue of customer impatience has been treated mostly in the context of so called inventory systems with *partial backordering*. Under partial backordering (see for example Montgomery et al. (1973), Moinzadeh (1989), Smeitink (1990), Nahmias and Smith (1994), and the references therein), an arriving customer that faces a stockout is backordered with a certain probability and is lost otherwise. In situations where multiple orders are placed at once, this means that a fraction of customers are backordered while the remainder is lost. These models capture the simplest case of customer impatience with a mixture of only two types of customers: some that are infinitely patient and, therefore, can be backordered, and some that have zero patience and, therefore, are lost if they cannot be fulfilled immediately. This obviously ignores the possibility of having customers who are willing to wait but with varying degrees of patience. Posner et al. (1972) and Das (1977) do consider systems where customers are initially willing to wait, but if their demand is not fulfilled within their patience time, they leave the system. However, in their case, they assume a particular inventory control policy, either a (q, r) or a base-stock policy, and do not allow for the possibility of rejecting customers. To our knowledge, our paper is the first to characterize the optimal policy for an inventory system with customer impatience.

Although the modeling of customer impatience is surprisingly limited in the inventory literature, there is significant and growing literature that models impatience in the context of queueing systems; see for example Gans et al. (2003), Garnet et al. (2002), Mandelbaum and Zeltyn (2005), Jouini et al. (2007a, 2007b), Armony et al. (2007), Ward and Kumar (2008), and the references therein. A queueing system can be viewed as a make-to-order version of the system we consider in this paper, where inventory is not allowed to be held in anticipation of future demand. Much of the queueing literature that incorporates impatience is focused on performance evaluation and not optimal control. Moreover, the optimal control problem in a queueing system is simpler as there is typically only a decision about whether or not to admit a customer.

The rest of the paper is organized as follows. In Section 2, we formulate the problem. In Section 3, we characterize the structure of the optimal policy. In Section 4, we describe a performance evaluation model. In Section 5, we present numerical results. In Section 6, we offer a summary and some concluding comments.

2 Problem Formulation

We consider a system where a single product is produced at a single facility to fulfill demand from customers who place orders continuously over time according to a Poisson process with rate λ. Items are produced one unit at a time with exponentially-distributed production times with mean $1/\mu$. The production facility can produce ahead of demand in a make-to-stock fashion. However, items in inventory incur a holding cost h per unit per unit time. Upon arrival, an order is either fulfilled from inventory, if any is available, backordered, or rejected. If an order is rejected, the system incurs a rejection cost r. If an order is backordered, the system incurs no immediate cost. However, customers are impatient and may decide to cancel their orders if their waiting time in backlog exceeds a patience time. If a customer cancels her order, the system incurs a cancellation cost c. We assume that the rejecting cost is larger than the cancellation cost ($r \leq c$). Otherwise, it is optimal for the customers to accept all orders. The rejection cost can be viewed as a lost sale cost (e.g., the opportunity cost of generating revenue from the sale of one unit), while the cancellation cost can be viewed as the sum of a lost sale cost and a penalty for backlogging the order and not fulfilling it within the customer's patience time. We assume there is no other cost to backordering, although it is possible to impose an additional cost that increases with the amount of time an order stays in backlog. Customer patience times are independent and exponentially distributed with mean $1/\gamma$. This means that customers are willing to wait for an amount of time that is exponentially distributed for their orders to be fulfilled; otherwise, they cancel their orders. We assume that there is a finite upper bound M on the number of orders that can be on backorder at any time. "This assumption, which is made for mathematical tractability, is however not restrictive as we allow this upper to be arbitrarily large.

At any point in time, the system manager must decide whether or not to produce an item. We assume that preemption is possible, so that deciding not to produce could mean interrupting the production of a unit that was previously initiated. If interruption occurs, we assume it can be resumed the next time production is initiated (because of the memoryless property of the exponential distribution, resuming production from where it was interrupted is equivalent to initiating it from scratch). We assume that there are no costs associated with interrupting production. This conforms to earlier treatment of production-inventory systems in the literature; see, for example, Ha (1997a, 1997b). This assumption is not restrictive since, as we show in Theorem 1, it turns out that, generally, it is not optimal to interrupt production of an item once it has been initiated. At any point in time, the system manager must also decide on how to handle incoming orders. In particular, should an order arise and there is no inventory on-hand, a decision must be made on whether to backorder it or to reject it.

In our model, we assume that demand is Poisson and both production times and patience times are exponentially distributed. These assumptions are made in part for mathematical tractability as they allow us to formulate the control problem as an MDP and enable us to describe the structure of an optimal policy. They are also useful in approximating the behavior of systems where variability is high. The assumptions of Poisson demand and exponential production times are consistent with previous treatments of production-inventory systems; see for example, Buzacott and Shanthikumar (Chapter 4, 1993), Ha (1997a, 1997b), Zipkin (2000), and de Véricourt et al. (2002), among others. In Section 6, we discuss how these assumptions may be partially relaxed. The assumption of exponentially distributed patience times has been widely used in modeling customer impatience in queueing systems; see for example Mandelbaum and Zeltyn (2005) and Garnet et al. (2002). It captures the realistic feature that customers' willingness to wait decreases quickly with time, leading to a setting where most customers are willing to wait for only a relatively short period, with few (the most loyal customers) willing to wait for an extensive length of time.

The state of the system at time t can be described by net inventory $X(t)$, where $X(t)^+ = \max[0, X(t)]$ corresponds to on-hand inventory, and $X(t)^- = -\min[0, X(t)]$ to backorder level (the number of orders that are still waiting to be fulfilled). Note that because of the possibility of interrupting production, it is not necessary to include in the state description whether an item is currently being produced or not. Furthermore, because both order inter-arrival times and production times are exponentially distributed, the system is memoryless and decision epochs can be restricted to only times when the state changes (i.e., the completion of an item, the arrival of an order, or the cancellation of an order due to customer impatience). The memoryless property allows us to formulate the problem as an MDP and to restrict our attention to the class of Markovian policies for which actions taken at a particular decision epoch depend only on the current state of the system. In each state, the system manager makes two types of decisions, one regarding production and the other regarding order fulfillment. A policy d specifies for each state x whether production should be initiated or not and should an order arise whether it should be fulfilled from on-hand inventory, backordered or rejected (if there is inventory on-hand, it is trivial to show that it is always optimal to fulfill it).

Let $R(t)$ denote the number of orders that have been rejected up to time t and $N(t)$ the number of orders that have been cancelled by customers due to impatience up to time t. Then the expected discounted cost (the sum of inventory holding, order cancellation, and lost sales costs) over an infinite planning horizon obtained under a policy d and a starting state x can be written as:

$$v^d(x) = E_x^d \left[\int_0^\infty e^{-\alpha t} h X^+(t) dt + \int_0^\infty e^{-\alpha t} r dR(t) + \int_0^\infty e^{-\alpha t} c dN(t) \right],$$

where $\alpha > 0$ is the discount rate (extending the analysis to the case where the objective is to minimize average cost is straightforward and is briefly described at the end of Section 3). Our objective is to choose a policy d^* that minimizes the expected discounted cost. We refer to the optimal cost function as v^* where $v^* = v^{d^*}$. Following Lippman (1975), we work with a uniformized version of the problem in which the transition rate in each state under any action is $\beta = \lambda + \mu + M\gamma$ so that the transition times between decision epochs form a sequence of i.i.d. exponential random variables, each with mean $1/\beta$. The introduction of the uniform transition rate allows us to transform the continuous time decision process into a discrete time decision process, simplifying the analysis considerably. To further simplify the analysis, and without loss of generality, we also rescale time by letting $\alpha + \beta = 1$. The optimal cost function can now be shown to satisfy the following optimality equation:

$$v^*(x) = hx^+ + \lambda T_{arr} v^*(x) + \mu T_{prod} v^*(x) + \gamma T_{imp} v^*(x), \tag{1}$$

where the operators T_{prod}, T_{arr}, and T_{imp} are defined as follows,

$$T_{prod} v(x) = \min(v(x), v(x+1)), \tag{2}$$

$$T_{arr} v(x) = \begin{cases} \min(v(x-1), v(x)+r) & \text{if} \quad x > -M \\ v(x)+r, & \text{if} \quad x = -M, \text{ and} \end{cases} \tag{3}$$

$$T_{imp} v(x) = \begin{cases} -x[v(x+1)+c] + (M+x)v(x) & \text{if} \quad -M \leq x \leq -1 \\ Mv(x) & \text{if} \quad x \geq 0. \end{cases} \tag{4}$$

Operator T_{prod} is associated with the production decision: a decision to produce would increase the inventory level by one unit once production is completed while a decision not to produce will leave the inventory level unchanged. Operator T_{arr} is associated with the handling of the arrival of an order: fulfilling the order from on-hand inventory or backordering it reduce inventory level by one unit while rejecting it leaves the inventory level unchanged but leads the system to incur cost r. (Note that when the backorder level reaches M an incoming order is always rejected and the cost r is incurred.) Operator T_{imp} is associated with customers canceling their orders due to impatience. If the number of orders backordered is k, then the transition rate out of this state due to cancellations is $k\gamma$. To maintain a uniform transition rate, we use the standard approach of allowing for fictitious transitions from the state to itself with rate $(M-k)\gamma$, so that the overall transition rate is always $M\gamma$.

3 The Structure of the Optimal Policy

In this section, we characterize the structure of the optimal policy. In order to do so, we show that the optimal value function $v^*(x)$ for all states x satisfies certain properties as specified in Definition 1 below. We then show that these properties imply a specific rule for the optimal action in each state.

Definition 1 *Let $\mathcal{U}$ be a set of real valued functions defined on the set of integers $\mathbf{Z}$, such that if $v \in \mathcal{U}$, then:*

Property P1 $\Delta v(x) \geq -c$, *for all x,*

Property P2 $\Delta^2 v(x) \geq 0$, *for all x,*

Property P3 $\Delta v(x) \geq -r$, *for all $x \geq 0$, and*

Property P4 $\Delta v(x) \leq 0$, *for all $x < 0$,* where $\Delta v(x) = v(x+1) - v(x)$ and $\Delta^2 v(x) = \Delta v(x+1) - \Delta v(x)$. Therefore, convexity of $v(x)$ is equivalent to $\Delta^2 v(x) \geq 0$.

Lemma 1 *If $v \in \mathcal{U}$, then $Tv \in \mathcal{U}$ where $Tv(x) = hx^+ + \lambda T_{arr} v(x) + \mu T_{prod} v(x) + \gamma T_{imp} v(x)$. Furthermore, the optimal cost function v^* is an element of $\mathcal{U}$. That is, $v^* \in \mathcal{U}$,*

The proof of Lemma 1 and all other subsequent results can be found in the Appendix. In the proof, we first show that the operator T preserves properties P1-P4, which together with the convergence of value iteration, allows us to conclude that the optimal cost function v^* satisfies properties P1-P4. Applied to v^*, property P1 indicates that it is never desirable to have cancellations due to impatience. Property P2 implies that the marginal cost difference due to increasing net inventory is non-increasing. That is, the optimal value function is convex. Property P3 indicates that it is more preferable to fulfill orders from on-hand inventory, if there is any, than to reject them. Property P4 implies that it is optimal to produce whenever there is a backlog.

In order to describe the optimal policy implied by the above properties of the value function, we first define the following two threshold parameters:

$$s^* \equiv \min(x : \Delta v^*(x) \geq 0), \quad (5)$$

$$\text{and } w^* \equiv \max\{-M, \min(x : \Delta v^*(x) + r \geq 0)\}. \quad (6)$$

We are now ready to characterize the optimal policy.

Theorem 1 *There exists an optimal policy that can be specified by thresholds s^* and w^* as follows. The optimal production policy is a base-stock policy with base-stock level s^*, such that it is optimal to produce if $x < s^*$ and not to produce otherwise. The optimal order fulfillment policy is a limited admission policy with admission threshold w^*, such that it is optimal to accept an order if $x > w^*$ and to reject it otherwise. An admitted order is fulfilled from on-hand inventory if there is any and is backordered otherwise. Moreover, we have the following:*

- *It is always optimal to produce if there are any backorders; that is, $s^* \geq 0$.*
- *It is always optimal to accept orders if there is on-hand inventory; that is, $w^* \leq 0$.*
- *If $s^* > 0$, then it is never optimal to preempt production once it has been initiated.*

In contrast to common pure lost sales and pure backorder policies, the optimal policy allows for both backordering and order rejection. In doing so, the policy limits both inventory and backorder levels (inventory is costly because of holding costs and backordering is costly because it increases the chance of customers canceling their orders due to impatience). Note that as long as $s^* > 0$, it is never optimal to preempt production once it is initiated. In particular, if it is optimal to produce in state $0 \leq x < s^*$, then it continues, by virtue of the fact that the production policy is a base-stock policy with base-stock level s^*, to be optimal to produce in state $x - 1$ if an order arrives and we decide not to reject it (of course it continues to be optimal if an order arrives and we decide to reject it, leaving the system in state x). Similarly, if it is optimal to produce in state $w^* \leq x < 0$, then it continues to be optimal to produce in both states $x + 1$, corresponding to an order cancellation, and state $x - 1$, corresponding to an order arrival."

The only scenario under which preemption is possible is when $s^* = 0$ (i.e., inventory is never held and we produce only if there is a backorder). However, even in this case, preemption is optimal only if the system is in state $x = -1$ and an order is cancelled, moving the system to state $x = 0$.

The structure of the optimal policy in Theorem 1 can be shown to continue to hold for several variants of the problem, including systems where there is a one time backordering cost per unit backordered, a linear production cost and a convex holding cost. It also continues to hold in the case where the optimization criterion is the average cost per unit time instead of the expected discounted cost. The existence of an optimal policy for the average cost, and for this average cost to be finite and independent of the starting state, can be proven via an argument involving taking the limit as $\alpha \to 0$ in the discounted cost problem (see for example Cavazos-Cadena and Sennott (1992) and Weber and Stidham (1987)).

In the following theorem, we further characterize the structure of the optimal policy by characterizing the impact of various system parameters on the base-stock level and the admission threshold.

Theorem 2 *The optimal base-stock level s^* is non-increasing in h, μ, and M and is non-decreasing in c, r, and λ. The optimal admission threshold w^* is non-increasing in h, μ, r, and M and is non-decreasing in c and λ.*

The proof of Theorem 2 involves defining various super- and sub-modularity properties and showing that these are satisfied by the optimal cost function, v^*, which then implies the monotonicity results described in the theorem. Full details of the proof can be found in the Appendix. Theorem 2 also pertains to the average-cost case. In Figure 1, we provide representative numerical results, for the average cost criterion, that illustrate the impact of system parameters on s^* and w^*. These results show that both s^* and w^* can be quite sensitive to changes in system parameter values. For example, when $c \to r$, $w^* \to -\infty$ and it becomes optimal never to reject any order. When $c \to \infty$, $w^* \to 0$ and it becomes optimal to always reject orders when there is no inventory on-hand. When λ is much larger than μ, it also becomes optimal to always reject when there is no on-hand inventory (In this case, there is not sufficient capacity to fulfill demand and a fraction of total demand must always be rejected). Note that we were not able to establish monotonicity results with respect to the impatience parameter γ. In Section 5, we present numerical results regarding the effect of γ and examine the impact of customer impatience on system performance.

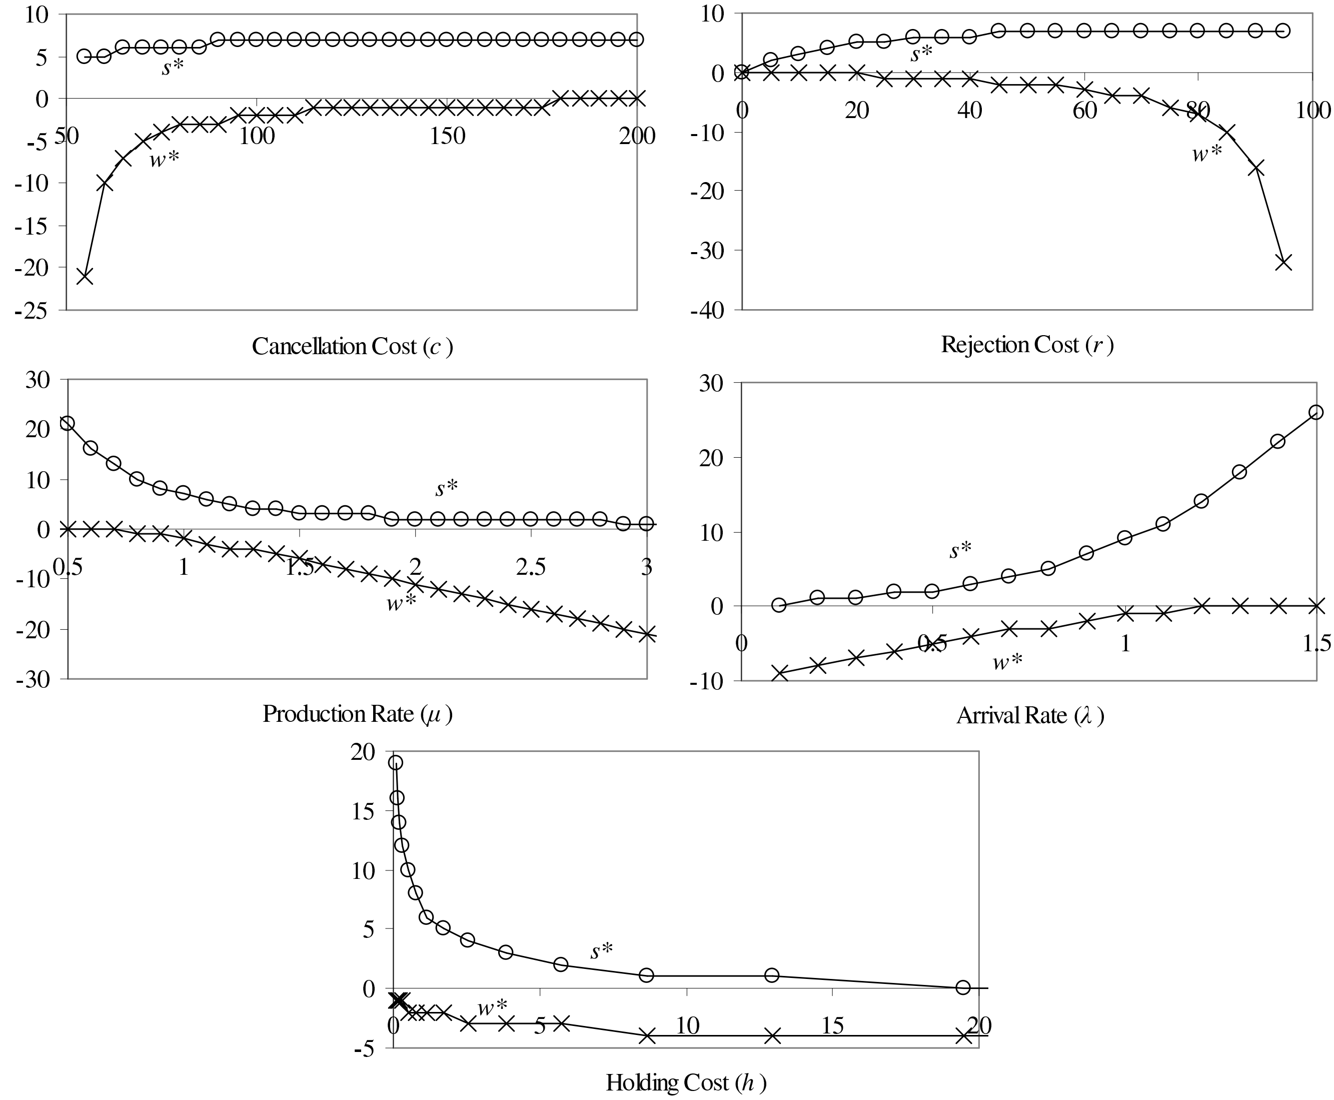

Figure 1: Impact of system parameters on w^* and s^*. Unless they are being varied, following parameter values are used: $\gamma = 0.1, \lambda = 0.9, \mu = 1, c = 100, r = 50, h = 1$

4 A Performance Evaluation Model

In this section, we use knowledge of the structure of the optimal policy to construct a performance evaluation model to compute efficiently the optimal base-stock level and the optimal admission threshold under the average cost criterion. Having such a model eliminates the need to use dynamic programming to carry out computations. Moreover, a dynamic programming algorithm, depending on problem parameter values, may require truncation of the state space, making the corresponding results approximate.

The approach we take follows from the recognition that a system operating under a control policy specified by a fixed base-stock level s and an admission threshold w can be modeled as a Markov chain. In particular, the net inventory level, $X(t)$, evolves as a continuous-time Markov chain with transition rates from state j to state k, q_{jk}, given by

$$q_{jk} = \begin{cases} \lambda & \text{if } k = j-1, w < j \leq s, \\ \mu + \gamma j^- & \text{if } k = j+1, w \leq j < s, \\ 0 & \text{otherwise,} \end{cases}$$

where $j^- = -\min[0, j]$. This Markov chain is graphically illustrated in Figure 2.

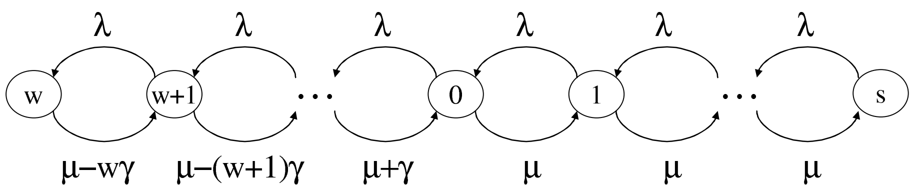

Figure 2: The transition diagram for the Markov chain model

The stationary probabilities, $\pi_j = \lim_{t\to\infty} P(X(t) = j)$, for the Markov chain can be shown to be given by the following:

$$\pi_j = \begin{cases} \rho^{s-j}\pi_s & \text{if } 0 \leq j \leq s, \\ \left(\prod_{k=1}^{-j} \frac{\lambda}{\mu+\gamma k}\right) \rho^s \pi_s & \text{if } w \leq j < 0, \\ 0 & \text{otherwise,} \end{cases} \tag{7}$$

where $\rho = \frac{\lambda}{\mu}$ and

$$\pi_s = \left[\sum_{j=0}^{s} \rho^j + \sum_{j=w}^{-1} \left(\prod_{k=1}^{-j} \frac{\lambda}{\mu + \gamma k} \right) \rho^s \right]^{-1}. \quad (8)$$

For given s and w, the expected cost, which we denote by $V(s, w)$, can now be obtained as:

$$V(s, w) = hE(X^+) + c\gamma E(X^-) + r\lambda\pi_w = h \sum_{j=1}^{s} j\pi_j + c\gamma \sum_{j=w}^{-1} -j\pi_j + r\lambda\pi_w. \quad (9)$$

The above expression involves the sum of finite terms and, therefore, can be computed efficiently. The optimal values for s and w can be obtained via an exhaustive search over a large enough range of s and w (unfortunately, the function $V(s, w)$ is not jointly convex in s and w). The computational effort for carrying out this search is generally modest. For example, a search over a 1,000 by 1,000 grid takes only few seconds on a standard personal computer. The computations can be further expedited by noting that the optimal base-stock level, s^*, has an upper bound given by the optimal base-stock level, $\hat{s}^*$, of a system where backorders are never allowed and where items are always rejected when they cannot be fulfilled from on-hand inventory. In the appendix, we show how an upper bound on $\hat{s}^*$ and, therefore, also on s^* can be obtained in closed form.

5 Some Numerical Results

In this section, we briefly provide some numerical results that illustrate the impact of customer impatience on optimal average cost and that examine the sensitivity of the base-stock level and admission threshold to system parameters. We also provide numerical results that compare the optimal policy to other commonly used policies.

In Figure 3, we present results that show the impact of varying the patience time parameter γ on optimal average cost. As we can see, customer impatience can have a significant impact on the cost. Cost is increasing in a roughly concave fashion in the impatience parameter, γ. These results highlight the importance of carefully accounting for customer impatience, as under or over-estimating customers' willingness to wait can lead to significantly under or over-estimating the true cost. They can also lead to significant errors in selecting values for the base-stock level and the admission thresholds (Figure 4).

In Figure 5, we compare the performance of the optimal policy with four other policies, that are perhaps simpler to implement as they all involve a single control parameter, but that either ignore or do not fully account for the impact of customer impatience. The policies are as follows:

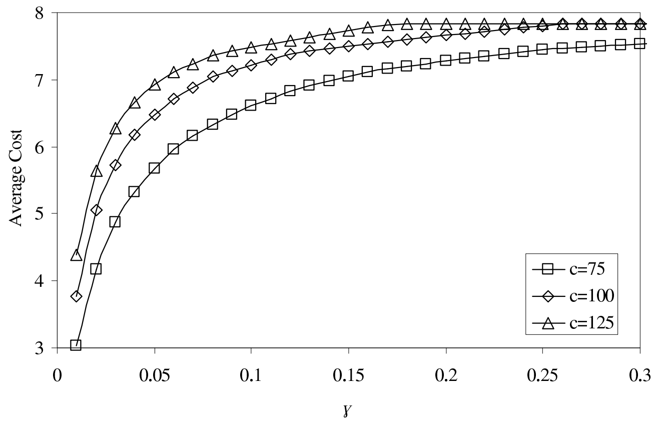

Figure 3: Impact of impatience rate on the average cost when $\lambda = 0.9, \mu = 1, r = 50, h = 1$

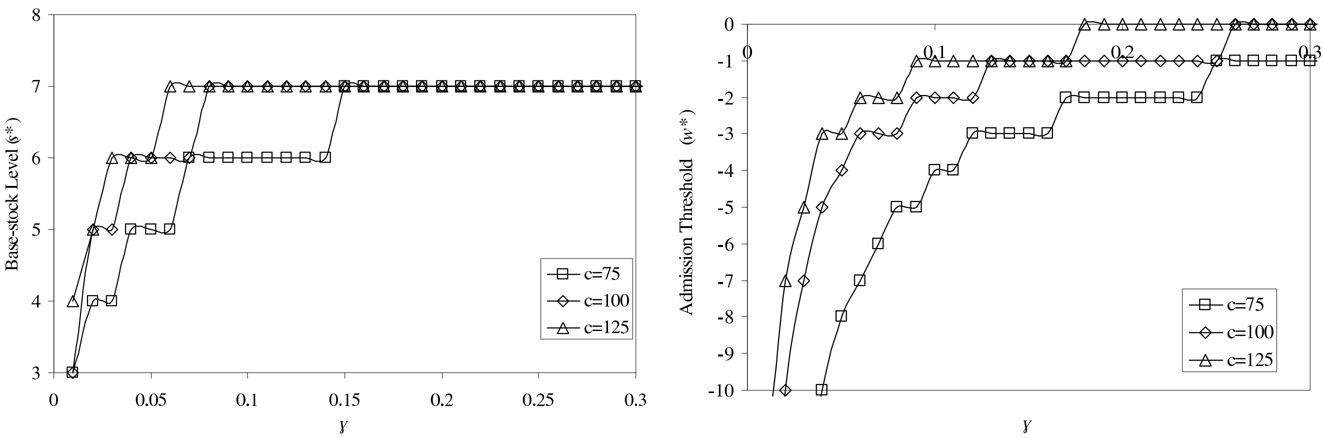

Figure 4: Impact of impatience rate on s^* and w^* when $\lambda = 0.9, \mu = 1, r = 50, h = 1$

Policy H_1: Orders are never rejected and are always backordered. Production is managed according to a base-stock policy with a fixed base-stock level.

Policy H_2: Orders that cannot be fulfilled from on-hand inventory are always rejected. Production is managed according to a base-stock policy with a fixed base-stock level.

Policy H_3: No inventory is held in anticipation of future demand and orders are always backordered as long as the backorder level does not exceed a specified threshold.

Policy H_4: No inventory is held in anticipation of future demand and orders are always backordered when they arrive.

The above policies can all be viewed as special cases of an (s, w) policy, where production is managed using a base-stock policy with base-stock level s and order fulfillment is managed using an admission policy with an admission threshold w. In the case of H_1, $s \geq 0$ and $w = -\infty$; for H_2, $s \geq 0$ and $w = 0$; for H_3, $s = 0$ and $w \leq 0$; and for H_4, $s = 0$ and $w = -\infty$. To allow for a fair comparison against the optimal policy, the parameters of the four policies are always chosen optimally. Figure 5 shows the percentage cost difference between the cost of the optimal policy and the optimal cost of each policy. The percentage cost difference, δ_i, for policy H_i is computed as $\delta_i = (C_i^* - C^*)/C^* \times 100\%$, where C_i^* is the optimal average cost under policy H_i and C^* is the average cost under the optimal policy.

As we can see from Figure 5, all four policies can perform poorly. In general, policies that do not allow for rejection ($w = -\infty$) perform poorly when customers are very impatient (γ is high), cancellation cost is high (c is high), rejection cost is low (r is low), or when the utilization of the production facility is high (the ratio λ/μ is high). On the other hand, policies that always reject orders when they cannot be fulfilled from on-hand inventory ($w = 0$) perform poorly when customers are patient (γ is low), impatience cost is low (c is low), rejection cost is high (r is high), or when the utilization of the production facility is low (the ratio λ/μ is low). There are of course settings where each of the four policies performs reasonably well. However, in most settings when impatience matters, either because of a low customer patience time or a high cost of cancellation, there are significant benefits to using the optimal or, to a lesser degree, a policy that limits the number of backorders, such as policy H_2 or H_3.

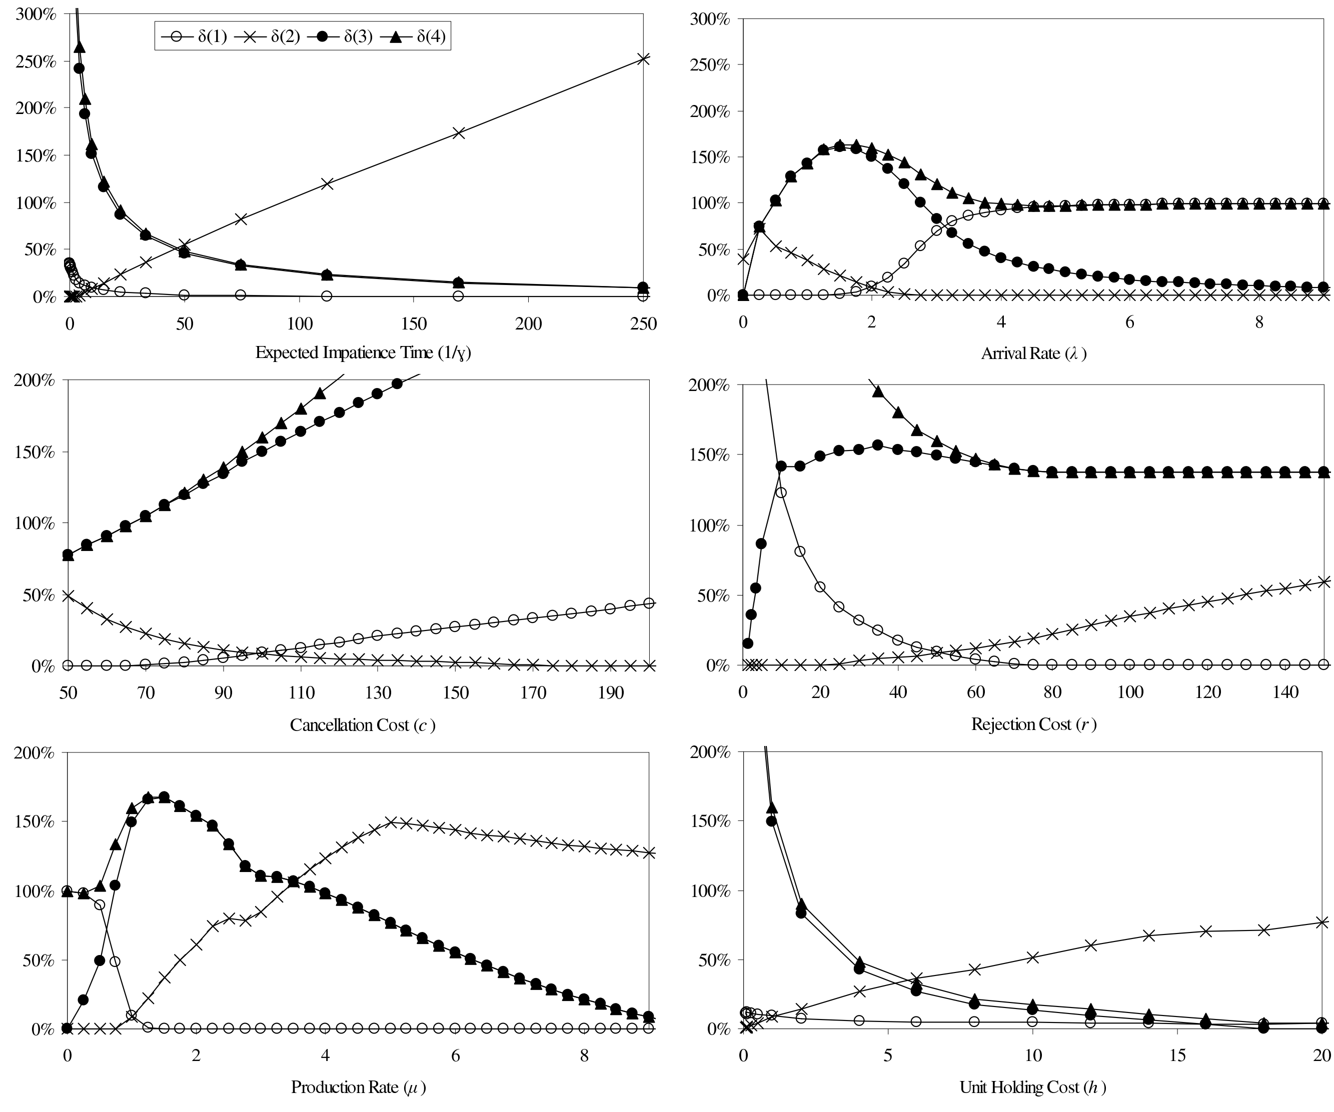

Figure 5: Impact of system parameters on δ_i. Unless they are being varied, following parameter values are used: $\gamma = 0.1, \lambda = 0.9, \mu = 1, c = 100, r = 50, h = 1$

6 Conclusions and Future Research

In this paper, we have analyzed a production-inventory system with impatient customers. We characterized the structure of the optimal policy and studied the sensitivity of the optimal policy to various operating parameters. Given the structure of the optimal policy, we described a performance evaluation model that allows for computing control parameters of the optimal policy efficiently. From numerical results, we investigated the impact of customer impatience on total system cost and compared the performance of the optimal policy against the performance of alternative control policies that ignore or do not fully account for customer impatience.

The results of the paper highlight the importance of incorporating customer impatience in the management of inventory systems. The results also highlight the inadequacy of existing inventory models which tend to assume that orders that cannot be immediately fulfilled from on-hand inventory are either all backordered (pure backorder systems) or all rejected (pure lost sales systems) and illustrates the need for models that allow for both backordering and rejection.

This paper is obviously only a first step toward a more comprehensive modeling and analysis of inventory systems with impatience. Avenues for future research are many. It will be useful to consider systems with different demand, production time, and patience time distributions. For example, it is possible to substitute the exponential distribution by Phase-type distributions which can be constructed to approximate other more general distributions. Phase-type distributions retains the Markovian property of the system and continues to allow the formulation of the problem as an MDP. It will also be useful to extend the analysis to systems with multiple demand classes with different patient time parameters and different rejection and cancellation costs. This would give rise to new types of decisions regarding how on-hand inventory should be allocated and how fulfillment priorities should be assigned to orders that are in backlog.

References

[1] M. Armony, E. Plambeck, and S. Seshadri. Sensitivity of optimal capacity to customer impatience in an unobservable $M/M/S$ queue (Why you shouldn't shout at the DMV). *Manufacturing and Service Operations Management*, 2007.

[2] J.A. Buzacott and G.J. Shanthikumar. *Stochastic models of manufacturing systems.* Prentice Hall, 1993.

[3] R. Cavazos-Cadena and LI Sennott. Comparing recent assumptions for the existence of average optimal stationary policies. *Operations research letters*, 11(1):33–37, 1992.

[4] E.B. Çil, E.L. Örmeci, and F. Karaesmen. Effects of system parameters on the optimal policy structure in a class of queueing control problems. *Queueing Systems*, 61(4):273–304, 2009.

[5] C. Das. The $(S-1, S)$ Inventory Model under Time Limit on Backorders. *Operations Research*, 25:835–850, 1977.

[6] F. de Véricourt, F. Karaesmen, and Y. Dallery. Optimal Stock Allocation for a Capacitated Supply System. *Management Science*, 48:1486–1501, 2002.

[7] N. Gans, G. Koole, and A. Mandelbaum. Telephone call centers: Tutorial, review, and research prospects. *Manufacturing and Service Operations Management*, 5(2):79–141, 2003.

[8] O. Garnet, A. Mandelbaum, and M. Reiman. Designing a call center with impatient customers. *Manufacturing & Service Operations Management*, 4(3):208–227, 2002.

[9] A.Y. Ha. Inventory Rationing in a Make-to-Stock Production System with Several Demand Classes and Lost Sales. *Management Science*, 43(8):1093–1103, 1997a.

[10] A.Y. Ha. Stock-rationing policy for a make-to-stock production system with two priority classes and backordering. *Naval Research Logistics*, 44(5):457–472, 1997b.

[11] O. Jouini and Y. Dallery. Analysis of a multiple priority queue with impatient customers. *Working Paper, Ecole Centrale Paris*, 2007b.

[12] O. Jouini, A. Pot, Y. Dallery, and G. Koole. Real-time dynamic scheduling policies for multiclass call centers with impatient customers. *Working Paper, Ecole Centrale Paris*, 2007a.

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

[14] A. Mandelbaum and S. Zeltyn. The Palm/Erlang-A Queue, with Applications to Call Centers. *Working Paper, Tel Aviv University*, 2005.

[15] K. Moinzadeh. Operating characteristics of the (S-1, S) inventory system with partial backorders and constant resupply times. *Management Science*, pages 472–477, 1989.

[16] D.C. Montgomery, MS Bazaraa, and A.K. Keswani. Inventory models with a mixture of backorders and lost sales. *Naval Research Logistics Quarterly*, 20(2), 1973.

[17] S. Nahmias and S.A. Smith. Optimizing inventory levels in a two-echelon retailer system with partial lost sales. *Management Science*, pages 582–596, 1994.

[18] M.J.M. Posner and B. Yansouni. A class of inventory models with customer impatience. *Naval Research Logistics Quarterly*, 19(3), 1972.

[19] M. Puterman. *Markov Decision Processes.* John Wiley and Sons Inc., New York, 1994.

[20] E. Smeitink. A Note on "Operating Characteristics of the (S-1, S) Inventory System with Partial Backorders and Constant Resupply Times". *Management Science*, pages 1413–1414, 1990.

[21] A.R. Ward and S. Kumar. Asymptotically optimal admission control of a queue with impatient customers. *Mathematics of Operations Research*, 33(1):167–202, 2008.

[22] R.R. Weber and S. Stidham Jr. Optimal control of service rates in networks of queues. *Advances in Applied Probability*, 19(1):202–218, 1987.

[23] P.H. Zipkin. *Foundations of inventory management.* McGraw-Hill Boston, 2000.

A Online Appendix

A.1 Proof of Lemma 1

Throughout this proof, we assume that the value function v belongs to $\mathcal{U}$. In order to prove Lemma 1, we will prove that Tv satisfies properties P1, P2, P3 and P4.

Preliminaries

As v satisfies Property P2, we can define the thresholds s and w:

$$s = \min(x : \Delta v(x) \geq 0), \tag{10}$$

$$w = \max\{-M, \min(x : \Delta v(x) + r \geq 0)\}, \tag{11}$$

and for $x \geq -M$, we have:

$$\Delta v(x) \geq 0 \text{ if and only if } x \geq s, \tag{12}$$

$$\Delta v(x) + r \geq 0 \text{ if and only if } x \geq w. \tag{13}$$

Using the definitions of s and w, we can rewrite operators T_{arr} and T_{prod} as follows:

$$T_{arr}v(x) = \begin{cases} v(x-1) & \text{if} \quad x > w \\ v(x) + r & \text{if} \quad x \leq w, \end{cases} \tag{14}$$

$$T_{prod}v(x) = \begin{cases} v(x+1) & \text{if} \quad x < s \\ v(x) & \text{if} \quad x \geq s. \end{cases} \tag{15}$$

The first order differences for the operators T_{arr}, T_{prod} and T_{imp} can be written as follows:

$$\Delta T_{arr}v(x) = \begin{cases} \Delta v(x) & \text{if} \quad x \le w-1 \\ -r & \text{if} \quad x = w \\ \Delta v(x-1) & \text{if} \quad x > w, \end{cases} \tag{16}$$

$$\Delta T_{prod}v(x) = \begin{cases} \Delta v(x+1) & \text{if} \quad x+1 < s \\ 0 & \text{if} \quad x+1 = s \\ \Delta v(x) & \text{if} \quad x+1 > s, \end{cases} \tag{17}$$

$$\Delta T_{imp}v(x) = \begin{cases} -c-(x+1)\Delta v(x+1)+(M+x)\Delta v(x) & \text{if} \quad x \le -1 \\ M\Delta v(x) & \text{if} \quad x \ge 0, \end{cases} \tag{18}$$

and the second order differences are:

$$\Delta^2 T_{arr}v(x) = \begin{cases} \Delta^2 v(x-1) & \text{if} \quad x > w \\ \Delta v(x)+r & \text{if} \quad x = w \\ -r-\Delta v(x) & \text{if} \quad x+1 = w \\ \Delta^2 v(x) & \text{if} \quad x+2 \le w, \end{cases} \tag{19}$$

$$\Delta^2 T_{prod}v(x) = \begin{cases} \Delta^2 v(x+1) & \text{if} \quad x < s-2 \\ -\Delta v(x+1) & \text{if} \quad x = s-2 \\ \Delta v(x+1) & \text{if} \quad x = s-1 \\ \Delta^2 v(x) & \text{if} \quad x \ge s, \end{cases} \tag{20}$$

$$\Delta^2 T_{imp}v(x) = \begin{cases} -(x+2)\Delta^2 v(x+1)+(M+x)\Delta^2 v(x) & \text{if} \quad -M \le x \le -2 \\ M\Delta^2 v(x)+\Delta v(x)+c & \text{if} \quad x = -1 \\ M\Delta^2 v(x) & \text{if} \quad x \ge 0. \end{cases} \tag{21}$$

Since we defined the first and second order differences, we are now ready to prove that properties P1-P4 hold for Tv.

Property P1

First of all, we have $\Delta v(x) \ge -c$ for all x since v satisfies P1. Secondly, $-r \ge -c$ by assumption. Using (12)-(13) and (16)-(18), we show that, for all x, $\Delta T_{arr}v(x) \ge -c$, $\Delta T_{prod}v(x) \ge -c$ and $\Delta T_{arr}v(x) \ge -Mc$.

We can now conclude that Property P1 holds for Tv, since, for all x:

$$\Delta Tv(x) = h + \mu\Delta T_{prod}v(x) + \lambda\Delta T_{arr}v(x) + \gamma\Delta T_{imp}v(x) \ge -(\mu+\lambda+\gamma M)c \ge -c.$$

The last inequality results from the rescaling of time assumption given in Section 2 which implies that $\mu + \lambda + \gamma M < 1$.

Property P2

Using (12)-(13), (19)-(21) and properties satisfied by v (in particular P1, P2 and P3), we show that $\Delta^2 T_{arr}v(x)$, $\Delta^2 T_{prod}v(x)$ and $\Delta^2 T_{imp}v(x)$ are non-negative, for all x. Therefore, for all x, we have the following inequality:

$$\Delta^2 Tv(x) = \mu\Delta^2 T_{prod}v(x) + \lambda\Delta^2 T_{arr}v(x) + \gamma\Delta^2 T_{imp}v(x) \geq 0,$$

and we conclude that Tv satisfies Property P2.

Property P3

We have $\Delta v(x) \geq -r$ when $x \geq 0$ since v satisfies P3. As a result, equations (16)-(18) imply that $\Delta T_{arr}v(x) \geq -r$, $\Delta T_{prod}v(x) \geq -r$ and $\Delta T_{imp}v(x) \geq -Mr$ if $x \geq 0$. Therefore, when $x \geq 0$, we have the following inequality:

$$\Delta Tv(x) = h + \mu\Delta T_{prod}v(x) + \lambda\Delta T_{arr}v(x) + \gamma\Delta T_{imp}v(x) \geq -(\mu + \lambda + \gamma M)r \geq -r,$$

and Tv satisfies P3.

Property P4

We have $\Delta v(x) \leq 0$ when $x < 0$ since v satisfies P4. As a result, equations (16)-(18) imply that $\Delta T_{arr}v(x) \leq 0$, $\Delta T_{prod}v(x) \leq 0$ and $\Delta T_{imp}v(x) \leq 0$ if $x < 0$. Finally, when $x < 0$, we have the following inequality:

$$\Delta Tv(x) = \mu\Delta T_{prod}v(x) + \lambda\Delta T_{arr}v(x) + \gamma\Delta T_{imp}v(x) \leq 0,$$

and Tv satisfies P4.

Conclusion

We have shown that $v \in \mathcal{U}$ implies that $Tv \in \mathcal{U}$. The fixed point theorem (Puterman 1994) ensures that the sequence of functions $v_{n+1} = Tv_n$ converges to v^* for any v_0. If we take $v_0(x) = 0$ for all x, we have $v_0 \in \mathcal{U}$ and we conclude by induction that $v^* \in \mathcal{U}$.

A.2 Proof of Theorem 1

From Lemma 1, we know that $v^* \in \mathcal{U}$. Property P2 guarantees the existence of threshold levels s^* and w^*. Furthermore it implies that it is optimal to produce if $x < s^*$ and not to produce otherwise and to accept an order if $x > w^*$ and to reject it otherwise. Property P3 states that it is optimal to accept orders if there is on-hand inventory and implies that $w^* \leq 0$. Property P4 states that it is optimal to produce if $x < 0$ and implies that $s^* \geq 0$.

A.3 Proof of Theorem 2

In order to show monotonicity properties for s^* and w^* with respect to various system parameters, we adapt the following approach. We compare the optimal value functions of two systems that are identical except for the value of one system parameter, denoted by p. For short, we write $p = \lambda$ when demand rate is varied, $p = r$ when rejecting cost is varied and so on. The optimal base-stock level, admission threshold and value function corresponding to a given system parameter p will be represented by s_p^*, w_p^* and $v_p^*(x)$ respectively, where p belongs to the set of system parameters $\{\lambda, \mu, h, c, r, M\}$.

We state that a function v_p is submodular in x and p (denoted by $SubM(x, p)$), if and only if,

$$\Delta v_p(x) \geq \Delta v_{p+\epsilon}(x), \forall x \ \geq -M, \forall p \neq M, \forall \epsilon \geq 0.$$

The supermodularity in x and p , denoted by $SuperM(x, p)$, is the opposite inequality ($\Delta v_p(x) \leq \Delta v_{p+\epsilon}(x)$). These definitions can be used when $p \in \{r, c, h, \lambda, \mu\}$. However, when $p = M$, p is discrete and the state space depends on M. In this case, we state that v is $SubM(x, M)$ if and only if the following inequality holds:

$$\Delta v_M(x) \geq \Delta v_{M+1}(x), \forall x \geq -M, \forall M \in \mathbb{N}.$$

$SuperM(x, M)$ is the same inequality in opposite direction.

Let us start our proof by defining the necessary properties to be satisfied in order to show the results given in Theorem 2. Firstly, we define $\mathcal{V}$ a set of real valued functions, with the following properties:

Definition 2 *If $v \in \mathcal{V}$, then:*

Property Q1 $v \in \mathcal{U}$,

Property Q2 $\forall p \in \{\mu, h, M\}$, *$v$ is $SuperM(x, p)$ and* $\forall p \in \{\lambda, r, c\}$, *$v$ is $SubM(x, p)$,*

Property Q3 $\Delta v_{r+\epsilon}(x) + \epsilon \geq \Delta v_r(x), \forall r \geq 0, \forall \epsilon \geq 0.$

If we prove that $v^* \in \mathcal{V}$, then, we obtain the monotonicity results for s^* and w^* as described in Theorem 2.

The uniformization rate depends on $\{\lambda, \mu, M\}$ and needs to be constant for two systems to be comparable. We rescale the time using a uniformization rate δ which is sufficiently larger than $(\alpha+\lambda+\mu+M\gamma)$ to have the same uniformization rate with parameter values p or $p+\epsilon$. Therefore, the optimality equations are redefined by adding a new operator, T_{unif}, to maintain a constant uniformization rate. T_{unif} is a fictitious event operator that transfers the system into the same state. Optimality equations can be rewritten as follows:

$$\begin{aligned} & v_p^*(x) = Tv_p^*(x), \forall x, \\ \text{with} \quad & Tv_p(x) = \frac{1}{\delta}\left[hx^+ + \lambda T_{arr}v_p(x) + \mu T_{prod}v_p(x) + \gamma T_{imp}v_p(x) + T_{unif}v_p(x)\right], \\ \text{and} \quad & T_{unif}v_p(x) = (\delta - \alpha - \lambda - \mu - M\gamma)v_p(x). \end{aligned}$$

Operators T_{arr}, T_{prod} and T_{imp} are defined as previously. Note that all operators may depend on p. For instance, if $p = \lambda$, we have the following optimal operator T when the arrival rate equals $\lambda + \epsilon$:

$$\begin{aligned} Tv_{\lambda+\epsilon}(x) = & \frac{1}{\delta}[hx^+ + (\lambda+\epsilon)T_{arr}v_{\lambda+\epsilon}(x) + \mu T_{prod}v_{\lambda+\epsilon}(x) + \gamma T_{imp}v_{\lambda+\epsilon}(x) \\ & + (\delta - \alpha - \lambda - \epsilon - \mu - M\gamma)v_{\lambda+\epsilon}(x)]. \end{aligned}$$

In order to prove that T preserves properties of the set $\mathcal{V}$, we are going to prove two lemmas. In Lemma 2, we show that individual operators $T_{prod}, T_{imp}, T_{arr}, T_{unif}$ preserve some monotonicity properties.

Lemma 2 *If $v \in \mathcal{V}$, then:*

	$SuperM(x,p)$	$SubM(x,p)$
hx^+	$\forall p \in \{\mu, h, M\}$	$\forall p \in \{\lambda, c, r\}$
$T_{prod}v$	$\forall p \in \{\mu, h, M\}$	$\forall p \in \{\lambda, c, r\}$
$T_{imp}v$	$\forall p \in \{\mu, h\}$	$\forall p \in \{\lambda, c, r\}$
$T_{arr}v$	$\forall p \in \{\mu, h, M\}$	$\forall p \in \{\lambda, c, r\}$
$T_{unif}v$	$\forall p \in \{\mu, h\}$	$\forall p \in \{\lambda, c, r\}$

Table 1: Preservation of submodularity and supermodularity by the operators

Table 1 can be interpreted as follows. For instance, if $v \in \mathcal{V}$ then $T_{imp}v$ is $SuperM(x,p)$ for all $p \in \{\mu, h\}$ and $SubM(x,p)$ for all $p \in \{\lambda, c, r\}$.

Proof: We assume in all this proof that $v \in \mathcal{V}$. It is clear that hx^+ is $SuperM(x,p)$ for $p \in \{\mu, h, M\}$ and $SubM(x,p)$ for $p \in \{\lambda, c, r\}$.

If we adapt results from Çil et al. (2009), we have $T_{prod}v$ is $SuperM(x,p)$ for $p \in \{\mu, h\}$ and $SubM(x,p)$

for $p \in \{\lambda, c, r\}$. When $p = M$, the state space depends on M. However, the argument used by Çil et al. still holds, since there is no transition from state $(-M)$ to state $(-M-1)$ associated with production.

For operator T_{imp}, Equation (18) implies that $T_{imp}v$ is $SuperM(x,p)$ for $p \in \{\mu, h\}$ and $SubM(x,p)$ for $p \in \{\lambda, c, r\}$.

For operator T_{arr}, Çil et al. (2009) show that $T_{arr}v$ is $SuperM(x,p)$ for $p \in \{\mu, h\}$ and $SubM(x,p)$ for $p \in \{\lambda, c\}$. Consider the case $p = r$, it is slightly more complex since operator T_{arr} depends on r. We have:

$$\Delta T_{arr} v_r(x) = \begin{cases} \Delta v_r(x) \leq -r & \text{if } x < w_r \\ -r & \text{if } x = w_r \\ \Delta v_r(x-1) \geq -r & \text{if } x > w_r. \end{cases} \tag{22}$$

Above equation implies that

$$\Delta T_{arr} v_r(x) \geq -r \text{ if and only if } x \geq w_r. \tag{23}$$

As v satisfies Q3, we have $w_{r+\epsilon} \leq w_r$ and the following cases:

$$\begin{aligned} &\Delta T_{arr} v_{r+\epsilon}(x) - \Delta T_{arr} v_r(x) \\ &= \begin{cases} \Delta v_{r+\epsilon}(x) - \Delta v_r(x) \leq 0 & \text{if } x < w_{r+\epsilon} \leq w_r \\ -(r+\epsilon) - \Delta v_r(x) \leq -(r+\epsilon) - \Delta v_{r+\epsilon}(x) \leq 0 & \text{if } x = w_{r+\epsilon} < w_r \\ \Delta v_{r+\epsilon}(x-1) - \Delta v_r(x) \leq \Delta v_{r+\epsilon}(x) - \Delta v_r(x) \leq 0 & \text{if } w_{r+\epsilon} < x < w_r \\ -\epsilon \leq 0 & \text{if } x = w_{r+\epsilon} = w_r \\ \Delta v_{r+\epsilon}(x-1) + r \leq \Delta v_r(x-1) + r \leq 0 & \text{if } w_{r+\epsilon} < x = w_r \\ \Delta v_{r+\epsilon}(x-1) - \Delta v_r(x-1) \leq 0 & \text{if } w_{r+\epsilon} \leq w_r < x. \end{cases} \end{aligned} \tag{24}$$

Inequalities in (24) are based on property P2 (i.e. $\Delta v_r(x) \leq \Delta v_r(x+1)$), on $SubM(x,r)$ of v (i.e. $\Delta v_{r+\epsilon}(x) \leq \Delta v_r(x)$) and on definition of w_r which ensures that $\Delta v_r(x) + r \geq 0$ if and only if $x \geq w_r$.

Consider now the case where $p = M$. As v is $SuperM(x,M)$, we have $\Delta v_M(x) \leq \Delta v_{M+1}(x)$ when $x \geq -M$. Moreover, $w_{M+1} \leq w_M$. We have:

$$\begin{aligned} &\Delta T_{arr} v_{M+1}(x) - \Delta T_{arr} v_M(x) \\ &= \begin{cases} \Delta v_{M+1}(x) - \Delta v_M(x) \geq 0 & \text{if } -M \leq x < w_{M+1} \\ \underbrace{\Delta T_{arr} v_{M+1}(x)}_{\geq -r} - \underbrace{\Delta T_{arr} v_M(x)}_{\leq -r} \geq 0 & \text{if } w_{M+1} \leq x \leq w_M, x \geq -M \\ \Delta v_M(x-1) - \Delta v_{M+1}(x-1) \geq 0 & \text{if } x > w_M \geq -M. \end{cases} \end{aligned} \tag{25}$$

Therefore, we have $\Delta T_{arr}v_{M+1}(x) - \Delta T_{arr}v_M(x) \geq 0$ when $x \geq -M$ and $T_{arr}v$ is $SuperM(x, M)$.

Finally, it is clear that $T_{unif}v$ is $SuperM(x,p)$ for $p = h$ and $SubM(x,p)$ for $p \in \{c, r\}$. ■

The following lemma will be used to prove that T preserves $SuperM(x, \mu)$ and $SubM(x, \lambda)$. We have omitted p since these results hold independent of p.

Lemma 3 *If $v \in \mathcal{U}$, then*

1. $\Delta[T_{prod}v(x) - v(x)] \geq 0$*, for all* $x \geq -M$
2. $\Delta[T_{arr}v(x) - v(x)] \leq 0$*, for all* $x \geq -M$

Proof: Assume that $v \in \mathcal{U}$. From Çil et al. (2009), we know that convexity of v in x implies that $(T_{prod}v - v)$ is non-decreasing in x.

We will now show that $(T_{arr}v - v)$ is non-increasing in x.

$$\Delta[T_{arr}v(x) - v(x)] = \begin{cases} 0 & \text{if } x < w \\ -r - \Delta v(x) \leq 0 & \text{if } x = w \\ \Delta v(x-1) - \Delta v(x) \leq 0 & \text{if } x > w. \end{cases} \tag{26}$$

Inequalities in (26) come from the definition of w and convexity of $v \in \mathcal{U}$. ■

The last lemma, based on lemmas 2 and 3, show that T preserves properties of $\mathcal{V}$. It also implies, by value iteration, that $v*$ belongs to $\mathcal{V}$ and thus Theorem 2.

Lemma 4

If $v \in \mathcal{V}$ then $Tv \in \mathcal{V}$.

Proof: Assume that $v \in \mathcal{V}$.

Property Q1

By Lemma 1, $Tv \in \mathcal{U}$ and hence satisfies Q1.

Property Q2

When p belongs to $\{c, r, h\}$, Tv is a linear combination of hx^+, $T_{arr}v$, $T_{prod}v$, $T_{imp}v$, $T_{unif}v$ where the coefficients of the combination do not depend on h, c, r. Therefore, Lemma 2 implies that T preserves $SubM(x,p)$ for $p \in \{c, r\}$ and preserves $SuperM(x,p)$ for $p = h$.

When p belongs to $\{\lambda, \mu, M\}$, the transition rates depend on λ, μ and M and we can not apply above argument. Çil et al. (2009) have shown how lemmas 2 and 3 imply that Tv is $SubM(x, \lambda)$ and $SuperM(x, \mu)$.

We remind their argument for $p = \lambda$. We have the following inequality:

$$
\begin{aligned}
&\Delta T v_{\lambda+\epsilon}(x) \leq \Delta T v_\lambda(x) \\
&\Leftrightarrow \begin{pmatrix} \mu \Delta T_{prod} v_{\lambda+\epsilon}(x) \\ +\lambda \Delta T_{arr} v_{\lambda+\epsilon}(x) \\ +\gamma \Delta T_{imp} v_{\lambda+\epsilon}(x) \\ +(\delta - \alpha - \lambda - \mu - M\gamma)\Delta v_{\lambda+\epsilon}(x) \\ +\epsilon \Delta[T_{arr} v_{\lambda+\epsilon}(x) - v_{\lambda+\epsilon}(x)] \end{pmatrix} \leq \begin{pmatrix} \mu \Delta T_{prod} v_\lambda(x) \\ +\lambda \Delta T_{arr} v_\lambda(x) \\ +\gamma \Delta T_{imp} v_\lambda(x) \\ +(\delta - \alpha - \lambda - \mu - M\lambda)\Delta v_\lambda(x) \end{pmatrix}.
\end{aligned} \tag{27}
$$

The first four lines of (27) satisfy the inequality ($\leq$) since T_{prod}, T_{arr} and T_{imp} preserve $SubM(x, \lambda)$ (Lemma 2). Moreover, $\Delta[T_{arr} v_{\lambda+\epsilon}(x) - v_{\lambda+\epsilon}(x)] \leq 0$ by Lemma 3. Therefore, inequality (27) holds and Tv is $SubM(x, \lambda)$. The arguments are similar for proving supermodularity of Tv in (x, μ).

Now, consider $SuperM(x, M)$. We have:

$$
\begin{aligned}
Tv_M(x) =& \frac{1}{\delta}[hx^+ + \mu T_{prod} v_M(x) + \lambda T_{arr} v_M(x) + \gamma T_{imp} v_M(x) \\
&+ (\delta - \alpha - \lambda - \mu - M\gamma) v_M(x), \\
Tv_{M+1}(x) =& \frac{1}{\delta}[hx^+ + \mu T_{prod} v_{M+1}(x) + \lambda T_{arr} v_{M+1}(x) + \gamma T_{imp} v_{M+1}(x) \\
&+ (\delta - \alpha - \lambda - \mu - (M+1)\gamma) v_{M+1}(x)],
\end{aligned}
$$

and the following inequality, when $x \geq -M$,

$$
\begin{aligned}
&\Delta T v_{M+1}(x) \geq \Delta T v_M(x) \\
&\Leftrightarrow \begin{pmatrix} \mu \Delta T_{prod} v_{M+1}(x) \\ +\lambda \Delta T_{arr} v_{M+1}(x) \\ +\gamma[\Delta T_{imp} v_{M+1}(x) - \Delta v_{M+1}(x)] \\ +(\delta - \alpha - \lambda - \mu - M\gamma)\Delta v_{M+1}(x) \end{pmatrix} \geq \begin{pmatrix} \mu \Delta T_{prod} v_M(x) \\ +\lambda \Delta T_{arr} v_M(x) \\ +\gamma \Delta T_{imp} v_M(x) \\ +(\delta - \alpha - \lambda - \mu - M\gamma)\Delta v_M(x) \end{pmatrix}.
\end{aligned} \tag{28}
$$

Inequalities in lines 1, 2 and 4 of Equation (28) hold from Lemma 2. Remains to prove that $\Delta T_{imp} v_{M+1}(x) - \Delta v_{M+1}(x) \geq \Delta T_{imp} v_M(x)$, for $x \geq -M$. We have:

$$
\Delta T_{imp} v_M(x) = \begin{cases} -c - (x+1)\Delta v_M(x+1) + (M+x)\Delta v_M(x) & \text{if } -M \leq x \leq -1 \\ M \Delta v_M(x) & \text{if } x \geq 0, \end{cases} \tag{29}
$$

and

$$\begin{aligned}
&\Delta T_{imp} v_{M+1}(x) - \Delta v_{M+1}(x) \\
&= \begin{cases} -c - (x+1)\Delta v_{M+1}(x+1) + (M+x)\Delta v_{M+1}(x) & \text{if } -M \leq x \leq -1 \\ M\Delta v_{M+1}(x) & \text{if } x \geq 0. \end{cases}
\end{aligned} \tag{30}$$

Since v is $SuperM(x, M)$, we have $\Delta v_{M+1}(x) \geq \Delta v_M(x)$. If we compare equations (29)-(30), it follows that $[\Delta T_{imp} v_{M+1}(x) - \Delta v_{M+1}(x)] \geq \Delta T_{imp} v_M(x)$. We conclude that inequality (28) holds and Tv is $SuperM(x, M)$. Finally, Tv satisfies Q2.

Property Q3
Here, we have:

$$\Delta T v_r(x) + r = \frac{1}{\delta}\left\{h + \mu[\Delta T_{prod} v_r(x) + r] + \lambda[\Delta T_{arr} v_r(x) + r] + \gamma[\Delta T_{imp} v_r(x) + r] + [\Delta T_{unif} v_r(x) + r]\right\}. \tag{31}$$

We can use similar arguments to those used for property Q2 to show that $(\Delta T_{prod} v + r)$, $(\Delta T_{imp} v + r)$ and $(\Delta T_{unif} v + r)$ are non-decreasing in r.

As v satisfies Property Q3, we have $\Delta v_{r+\epsilon}(x) + \epsilon \geq \Delta v_r(x)$ and $w_{r+\epsilon} \leq w_r$. Therefore,

$$\begin{aligned}
&\Delta T_{arr} v_{r+\epsilon}(x) + \epsilon - \Delta T_{arr} v_r(x) \\
&= \begin{cases} \Delta v_{r+\epsilon}(x) + \epsilon - \Delta v_r(x) \geq 0 & \text{if } x < w_{r+\epsilon} \leq w_r \\ \underbrace{\Delta T_{arr} v_{r+\epsilon}(x)}_{\geq -(r+\epsilon)} + \epsilon - \underbrace{\Delta T_{arr} v_r(x)}_{\leq -r} \geq 0 & \text{if } w_{r+\epsilon} \leq x \leq w_r \\ \Delta v_{r+\epsilon}(x-1) + \epsilon - \Delta v_r(x-1) \geq 0 & \text{if } w_{r+\epsilon} \leq w_r < x. \end{cases}
\end{aligned} \tag{32}$$

Inequalities $(\leq -(r+\epsilon))$ and $(\geq -r)$ in (32) come from (23). Finally, we conclude that $\Delta T v_r(x) + r$ is non-decreasing in r and Tv satisfies Q3. ■

A.4 Bounds on the Optimal Base-stock Level

Theorem 3

$$0 \leq s^* \leq s_u \quad \textit{with} \quad s_u = \begin{cases} \sqrt{2\lambda r/h} & \textit{if } \rho \leq 1 \\ (\rho-1)/(h'\rho) + 1/\ln\rho & \textit{if } \rho > 1 \end{cases} \quad \textit{and} \quad h' = h/(\lambda r)$$

Proof: The lower bound $0 \leq s^*$ is a direct consequence of Theorem 2.

From Theorem 2, we know that $s^*(M)$ is non-increasing in M and hence we have $s^*(M) \leq s^*(M = 0)$. When $\rho > 1$, Ha (1997a) has established that $s^*(M = 0)$ is bounded by the smallest nonnegative integer larger than $(\rho - 1)/(h'\rho) + 1/\ln\rho - 1$ with $h' = h/(\lambda r)$. When $\rho \leq 1$, Ha shows that the average cost $V(s, 0)$ is convex in s.

We complete the analysis of Ha by providing a simple upper-bound on the optimal base-stock level when $\rho \leq 1$ (or $\lambda \leq \mu$). We consider the limiting case where $\mu = \lambda$ and $M = 0$ where the stationary probabilities simplify to $\pi_i = \frac{1}{s+1}$. The average cost has a simple expression: $V(s, 0) = hs/2 + \lambda r/(s + 1)$. Since $V(s, 0)$ is convex in s, the optimal base-stock equals $\lfloor\sqrt{2\lambda r/h}\rfloor$. From Theorem 2, the optimal base-stock level is non-increasing in μ and hence $\sqrt{2\lambda r/h}$ is an upper-bound on s^* when $\lambda \leq \mu$. ■