ELSEVIER

European Journal of Operational Research 141 (2002) 113–132

EUROPEAN JOURNAL OF OPERATIONAL RESEARCH

www.elsevier.com/locate/dsw

Production, Manufacturing and Logistics

Optimal and near-optimal inventory control policies for a make-to-order inventory–production system

Q.-M. He [a,*], E.M. Jewkes [b], J. Buzacott [c]

[a] *Department of Industrial Engineering, Dalhousie University, Halifax, Nova Scotia, Canada B3J 2X4*
[b] *Department of Management Sciences, University of Waterloo, Waterloo, Ont., Canada N2L 3G1*
[c] *Schulich School of Business, York University, North York, Ont., Canada M3J 1P3*

Received 22 November 1999; accepted 29 June 2001

Abstract

This paper examines several inventory replenishment policies for a make-to-order inventory–production system that consists of a production workshop and a warehouse. Demands arrive to the production workshop according to a Poisson process, and are processed in an FCFS manner. The production workshop requires that the warehouse provides, as needed, raw materials for use in the production process. The warehouse inventory is replenished according to an inventory replenishment policy. The optimal replenishment policy, which minimizes the average total cost per product, is derived using a Markov decision process approach. The structure of the optimal replenishment policy is explored. Simple "order-up-to", "myopic", and heuristic replenishment policies are introduced. The myopic and heuristic replenishment policies are easy to compute, and yet perform almost as well as the optimal replenishment policy.

Keywords: Inventory theory; Production system; Markov decision processes; Optimal policy; Queueing theory

1. Introduction

In the manufacturing and distribution sectors, the importance of effective inventory control policies has long been recognized and has been studied extensively. Recently, interest in supply chain management has further increased awareness of the importance of inventory control throughout various links in a supply chain. In this paper, we are concerned with a small supply chain (see Fig. 1) that is made up of a production workshop and a warehouse of raw materials. The production workshop manufactures products on

* Corresponding author. Tel.: +1-902-494-6141; fax: +1-902-420-7858.
E-mail address: qi-ming.he@dal.ca (Q.-M. He).

PII: S0377-2217(01)00257-0

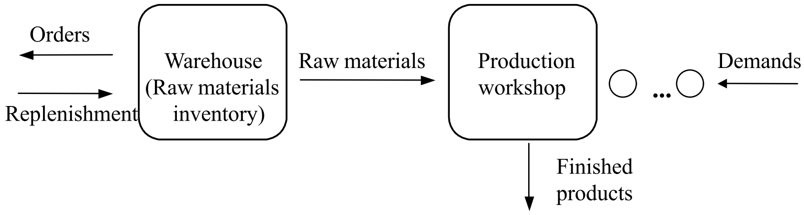

Fig. 1. The inventory–production system.

a make-to-order basis to meet customer demands. The warehouse provides raw materials to the production workshop. We explore inventory replenishment policies that trade off the cost of ordering with the cost of holding inventory of raw materials. The unique aspect of our model is the optimization of inventory decisions by taking explicit account of what is happening in the production workshop.

Our model is a special two-echelon model (see [6]) in which a raw material inventory exists in the warehouse, none exists in the production workshop, and production is make-to-order (i.e., the production workshop carries no finished goods inventory). Thus, inventory control is necessary only for raw materials replenishment in the warehouse. Studies of the optimal or near-optimal replenishment policies for multi-echelon systems are extensive (see [1,4,6,8]). For instance, Axsäter [1] looked at the $(S-1,S)$ policy for inventory systems with slow moving items. Chen and Zheng [4] conducted a performance analysis on production/inventory systems with (R,nQ) type policies. For the models considered in these papers, inventories are allowed at every stage of the echelon system and inventory control is based on the inventory levels (or inventory positions) at various stages. Our model is unique compared to other two-echelon systems since there is a delay between the arrival of a demand and the time raw materials are actually needed to produce a product for that demand. This is a result of explicitly modeling the production process as a queue. Consequently, inventory control in the warehouse will depend on both the inventory level in the warehouse and the number of demands queued in the production workshop, instead of just the inventory level (or inventory position) as in classical inventory control models (see [2,6,7]). Veatch and Wein [18] considered the optimal control of a two-station tandem production–inventory system. Their focus is on the control policy of the service rates at various production stages, not the inventory control policy.

Hariharan and Zipkin [8] studied an inventory model in which replenishment decisions are made based on the current inventory level (position) as well as some customer-order information. The idea of using customer-order information in inventory control is similar to that of using the number of waiting demands in raw materials replenishment in our model. But we use the exact queue length of waiting demands in making replenishment decisions, while they use information about the arrival of the next demand.

In this paper, we show the existence of the optimal replenishment policy, analyze the structure of the optimal replenishment policy, construct simpler heuristic replenishment policies, and compare these replenishment policies. Although the model considered in this paper is simple, the results obtained give insights into more complex supply chain models. Note that since the focus of this paper is on the optimal and near-optimal replenishment policies, issues related to performance analysis and the value of inventory control are not discussed. Interested readers are referred to [9–11].

The rest of the paper is organized as follows. Section 2 defines the inventory–production model explicitly. Section 3 introduces a Markov decision process for the optimal replenishment policy problem. The optimality equation for the optimal stationary replenishment policy is established. The structure of the optimal replenishment policy is explored in Section 4. Based on the results obtained in Section 4, Section 5 introduces four simple replenishment policies as approximations to the optimal policy. In Section 6, numerical examples are presented to show how well the simple replenishment policies perform. Finally, Section 7 summarizes the paper and gives future research directions.

2. The inventory–production model

The supply chain depicted in Fig. 1 is modeled as an integrated inventory–production system that consists of a workshop and a warehouse (see Fig. 2). The workshop manufactures a single type make-to-order product based on customer demands. Customer demands arrive one at a time to the workshop according to a Poisson process with parameter λ and are processed in a first-come-first-served manner. All demands are processed in the workshop by a single machine in batch sizes of one. Production (or processing) times of products have a common exponential distribution with parameter μ. The demand arrival process and the production times are assumed to be independent.

The warehouse stores raw materials for production in the workshop. The warehouse places orders to an outside supplier for new raw materials. The raw materials replenishment leadtimes are zero so that all orders to the supplier are filled immediately. There is a fixed ordering cost K associated with each order of raw materials, regardless of the order size. The holding cost is C_h per unit of raw materials held per unit time. The inventory level, defined as the number of units of raw materials, is reviewed continuously.

The warehouse and the workshop co-ordinate their operations in the following way: when the workshop starts to process a customer demand, a call for a unit of raw materials is sent to the warehouse. If the warehouse is not empty, a unit of raw materials is sent immediately to the workshop and production on that unit begins. If the warehouse is empty, an order for more raw materials is placed. The order is filled immediately, a unit of raw materials is sent to the workshop, and production on that unit begins. The transportation time between the warehouse and the workshop is assumed negligible so that there is no production delay in the workshop.

According to the definition above, the inventory–production system can be decomposed into two subsystems: an M/M/1 queue (the workshop) and an inventory system (the warehouse). The workshop can be modeled as a classical M/M/1 queue [5] since no shortage of raw materials is allowed, i.e., the queueing process is not influenced by the replenishment process. The warehouse can be modeled as a stochastic inventory system with zero leadtimes and demands that occur at the epochs when the workshop begins to produce a new product.

The status of the M/M/1 queue at time t can be represented by the number of customers in the workshop (i.e., unfilled demands, or the queue length), which we denote by $q(t)$. Throughout this paper, we assume that the traffic intensity $\rho = \lambda/\mu < 1$ so that the M/M/1 queue is stable. The status of the inventory system can be represented by the number of units of raw materials in the warehouse (i.e., the inventory level) at time t. We denote by $I(t)$ the total number of units of raw materials in the system, i.e., the inventory in the warehouse plus the unit of raw materials in the workshop (if any). Thus, the overall status of the inventory–production system at time t can be represented by $(q(t), I(t))$.

The raw material replenishment policy determines *when and how much* raw material should be ordered from the supplier. In this paper, we consider stationary replenishment policies that are based on the system status $(q(t), I(t))$. Since leadtimes are zero, raw materials *should not be ordered* when $I(t)$ is positive or $q(t)$ is zero. That observation implies that a stationary replenishment policy can be simply represented by a vector

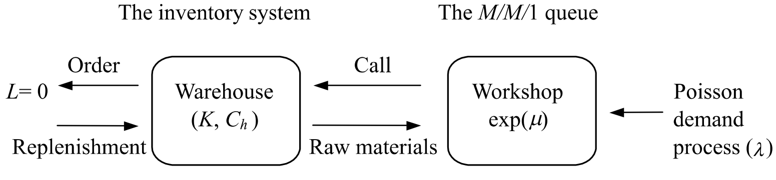

Fig. 2. The inventory–production system with zero leadtimes.

$\boldsymbol{\pi} = (\pi(0), \pi(1), \pi(2), \ldots)$, where $\pi(q)$ is the order size when the inventory level is zero and the number of unfilled demands is q. Therefore, the set of replenishment policies under consideration is given by

$$\boldsymbol{\Pi} = \{\boldsymbol{\pi} : \pi(0) = 0 \text{ and } \pi(q) \geqslant 1 \text{ for } q \geqslant 1\}. \tag{2.1}$$

Throughout this paper, we only consider the stationary replenishment policies in $\boldsymbol{\Pi}$. For brevity, we may drop the word "stationary" in the sequel.

At time t, if $I(t) = 0$ and $\pi(q(t)) > 0$, an order of size $\pi(q(t))$ is issued and filled immediately; otherwise, no action takes place. If an order of the size $\pi(q(t))$ is filled at time t, the inventory level becomes $\pi(q(t))$, i.e., $I(t+) = \pi(q(t))$. It is easy to see that for any replenishment policy $\boldsymbol{\pi} \in \boldsymbol{\Pi}$, the corresponding stochastic process $(q(t), I(t))$ is a continuous time Markov chain.

In order to evaluate replenishment policies, the *average total cost per unit of raw materials delivered to the workshop* (or equivalently the average total cost per product) is chosen as the fundamental measurement. The primary objectives of this paper are to find the optimal replenishment policy in $\boldsymbol{\Pi}$ that minimizes the average total cost per product, and to examine some simple, near-optimal replenishment policies.

3. The existence of the optimal stationary replenishment policy

In this section, a Markov decision process approach is utilized to establish an optimality equation for the optimal stationary replenishment policy. For that purpose, we need to introduce the set of feasible policies, the Markov chains associated with each feasible policy, the action sets associated with each state, and the immediate average cost associated with each state. For information about Markov decision processes, we refer to [3,13,16,17]. First, we state the following useful result.

Lemma 3.1. *For the inventory–production system introduced in Section 2, to minimize the average total cost per product, the order size must be less than* $1 + K\mu/C_{\text{h}}$, *regardless of the queue length at the decision epoch.*

Proof. See Appendix A. □

Lemma 3.1 implies that the replenishment policy which minimizes the average total cost per product must be in the following set of stationary policies:

$$\boldsymbol{\Pi}_b = \{\boldsymbol{\pi} : \pi(0) = 0 \text{ and } 1 \leqslant \pi(q) \leqslant \lfloor 1 + K\mu/C_{\text{h}} \rfloor \text{ for } q \geqslant 1\}, \tag{3.1}$$

where $\lfloor 1 + K\mu/C_{\text{h}} \rfloor$ represents the largest integer that is smaller than or equal to $1 + K\mu/C_{\text{h}}$. Thus, from now on, we shall only consider stationary replenishment policies in $\boldsymbol{\Pi}_b$ with the understanding that the optimal stationary replenishment policy in $\boldsymbol{\Pi}_b$ (if it exists) is also optimal in the set $\boldsymbol{\Pi}$.

Let t_n be the departure time of the nth customer (i.e., the completion time of the nth product), $n \geqslant 0$. We assume that a unit of production is completed at time zero, i.e., $t_0 = 0$. Let $q_n = q(t_n+)$ and $I_n = I(t_n+)$, $n \geqslant 0$. Apparently, the set $\{t_n, n \geqslant 0\}$ consists of the epochs at which a replenishment decision must be made. Note that a replenishment decision is made right after the departure of a customer. Since leadtimes are zero and there is no production delay, $\{t_n, n \geqslant 0\}$ are the departure epochs of the M/M/1 queue, which are the same for all replenishment policies in $\boldsymbol{\Pi}_b$. It is easy to see that the process $\{(q_n, I_n), n \geqslant 0\}$ is an embedded Markov chain. For each stationary policy $\boldsymbol{\pi}$ in $\boldsymbol{\Pi}_b$, the corresponding probability transition matrix of $\{(q_n, I_n), n \geqslant 0\}$ can be constructed explicitly (see [9]).

We point out that if $q_n = I_n = 0$, the replenishment decision must be "no action" since no raw materials are needed until the next customer arrives. That implies that, if $q_n = I_n = 0$, an order must be placed when the next demand arrives. However, the set $\{t_n, n \geqslant 0\}$ does not include any arrival epoch. To avoid this inconvenience, we assume that, if $q_n = I_n = 0$, the replenishment decision at the next demand arrival epoch is made at t_n, but the order is filled at the next demand arrival epoch. Such a change in the decision process does not change the stochastic process under consideration. With that modification, $\{t_n, n \geqslant 0\}$ contains all the decision epochs at which a replenishment decision must be made.

Let $A(q,i)$ be the set of actions associated with the state (q,i). It is easy to see that

$$A(q,i) = \begin{cases} \{\text{no order}\}, & q \geqslant 0 \text{ and } i \geqslant 1, \\ \{1,2,\ldots,\lfloor 1+K\mu/C_{\mathrm{h}}\rfloor\}, & q \geqslant 0 \text{ and } i = 0. \end{cases} \tag{3.2}$$

Note that $A(0,0) = \{1,2,\ldots,\lfloor 1+K\mu/C_{\mathrm{h}}\rfloor\}$ is really the action set at the next arrival epoch (at which the state is $(1,0)$) because of the modification of the decision in the state $(0,0)$.

Let $c(q,i,i_{\mathrm{a}})$ be the (average) immediate cost associated with the state (q,i) and action $i_{\mathrm{a}} \in A(q,i)$, i.e., $c(q,i,i_{\mathrm{a}})$ is the average total cost incurred between two consecutive departures, given that the state right after the first departure is (q,i) and the action i_{a} is taken right after the first departure. The sequence of state change and decision making is: state change first, decision making second, and implementation of the decision last. It can be shown that

$$c(q,i,i_{\mathrm{a}}) = \begin{cases} K + i_{\mathrm{a}}C_{\mathrm{h}}/\mu, & q \geqslant 0,\ i = 0,\ i_{\mathrm{a}} \geqslant 1, \\ iC_{\mathrm{h}}/\lambda + iC_{\mathrm{h}}/\mu, & q = 0,\ i \geqslant 1,\ i_{\mathrm{a}} = \text{no order}, \\ iC_{\mathrm{h}}/\mu, & q \geqslant 1,\ i \geqslant 1,\ i_{\mathrm{a}} = \text{no order}. \end{cases} \tag{3.3}$$

If $q = i = 0$ and the action is to order i_{a}, according to the modification, the order is filled when the next customer arrives. Thus, the total ordering cost is K and the expected total holding cost is $i_{\mathrm{a}}C_{\mathrm{h}}/\mu$ before the next departure, which leads to a total cost $K + i_{\mathrm{a}}C_{\mathrm{h}}/\mu$. Similarly, we can prove Eq. (3.3) for the case with $q \geqslant 1$, $i = 0$, and action i_{a}. If $q = 0$ and $i \geqslant 1$, these i units of raw materials must be held during the idle period of the M/M/1 queue and during the production time of the next unit. The mean length of the idle period in the M/M/1 queue is $1/\lambda$ and the mean production time is $1/\mu$. Thus, the immediate cost incurred until the next departure is $iC_{\mathrm{h}}(1/\lambda + 1/\mu)$. The last case $\{q \geqslant 1, i \geqslant 1\}$ can be proven easily.

In summary, the Markov decision process of interest can be defined with the following four components:

1. A set of stationary replenishment policy $\mathbf{\Pi}_b$ given in Eq. (3.1).
2. The Markov chain $\{(q_n, I_n), n \geqslant 0\}$ with a state space: $\{(q,i) : q \geqslant 0, 0 \leqslant i \leqslant \lfloor 1+K\mu/C_{\mathrm{h}}\rfloor\}$.
3. Action sets $\{A(q,i)\}$ defined by Eq. (3.2).
4. A set of immediate costs $\{c(q,i,i_{\mathrm{a}})\}$ (given in Eq. (3.3)) associated with individual states in the state space.

Clearly, we are dealing with a Markov decision process with a countable state space, finite decision sets, and bounded immediate costs. To find the optimal stationary replenishment policy $\boldsymbol{\pi}$ in $\mathbf{\Pi}_b$, we need to prove its existence first. Next, we show the following lemma about the Markov chain $\{(q_n, I_n), n \geqslant 0\}$, which leads to the existence of the optimal replenishment policy in $\mathbf{\Pi}_b$.

Lemma 3.2. *For any replenishment policy* $\boldsymbol{\pi}$ *in* $\mathbf{\Pi}_b$, *the corresponding Markov chain* $\{(q_n, I_n), n \geqslant 0\}$ *has only one closed set of states and is positive recurrent when it is restricted to the closed set. The state* $(0,0)$ *is always in the closed set. Let* $\{x(q,i) : q \geqslant 0, 0 \leqslant i \leqslant \lfloor 1+K\mu/C_{\mathrm{h}}\rfloor\}$ *be the stationary distribution of* $\{(q_n, I_n), n \geqslant 0\}$ (*if the state* (q,i) *is not in the closed set, then* $x(q,i) = 0$). *Then*

$$x(q,0) + x(q,1) + \cdots + x(q, \lfloor 1+K\mu/C_{\mathrm{h}}\rfloor) = (1-\rho)\rho^q \quad \textit{for } q \geqslant 0.$$

(*Note that* $\rho = \lambda/\mu$.)

Proof. See Appendix B. □

With Lemma 3.2, we are ready to show that there exists an optimal stationary replenishment policy in $\mathbf{\Pi}$.

Theorem 3.3. *For the inventory–production system defined in Section 2, if $\rho = \lambda/\mu < 1$, there exists a stationary replenishment policy in $\boldsymbol{\Pi}$ that has the minimal average total cost per product. We call that policy the optimal replenishment policy. The optimal replenishment policy can be found by solving a set of optimality equations (when restricted to $\boldsymbol{\Pi}_b$).*

Proof. See Appendix C. □

Based on Theorem 3.3, we can establish the optimality equation for computing the optimal replenishment policy. Define, for $0 \leqslant i,\ q \leqslant n$ and $n \geqslant 1$:

$V^*(q,i,n) =$ The minimal average total cost to produce n products, given that there are q demands (customers) and i units of raw materials in the system initially.

The minimal average total cost per product can be defined as

$$g = \lim_{n\to\infty} \frac{V^*(q,i,n)}{n}. \tag{3.4}$$

According to Theorem 3.3, the limit in Eq. (3.4) exists, and g is finite (since g is upper bounded by $K + C_{\mathrm{h}}/\mu$) and independent of the initial state. We choose the state $(0,0)$ as a *distinguished* state [17]. Define the relative cost functions as

$$h(q,i,n) = V^*(q,i,n) - V^*(0,0,n), \quad i \geqslant 0,\ \ q \geqslant 0,\ \ n \geqslant 1. \tag{3.5}$$

Let, for $i \geqslant 0$ and $q \geqslant 0$

$$h(q,i) = \lim_{n\to\infty} h(q,i,n) = \lim_{n\to\infty} [V^*(q,i,n) - V^*(0,0,n)]. \tag{3.6}$$

By Theorem 3.3, the limits in Eq. (3.6) exist. The following equations for the optimal replenishment policy can be established:

$$\begin{aligned}
g + h(0,0) &= \min_{1 \leqslant i \leqslant 1+\lfloor K\mu/C_{\mathrm{h}} \rfloor} \left\{ K + \frac{iC_{\mathrm{h}}}{\mu} + \sum_{n=0}^{\infty} \omega(1-\omega)^n h(n, i-1) \right\}, \quad q = i = 0,\\
g + h(q,0) &= \min_{1 \leqslant i \leqslant 1+\lfloor K\mu/C_{\mathrm{h}} \rfloor} \left\{ K + \frac{iC_{\mathrm{h}}}{\mu} + \sum_{n=0}^{\infty} \omega(1-\omega)^n h(q-1+n, i-1) \right\}, \quad q \geqslant 1,\ \ i = 0,\\
g + h(0,i) &= \frac{iC_{\mathrm{h}}}{\lambda} + \frac{iC_{\mathrm{h}}}{\mu} + \sum_{n=0}^{\infty} \omega(1-\omega)^n h(n, i-1), \quad q = 0,\ \ 1 \leqslant i \leqslant 1 + \lfloor K\mu/C_{\mathrm{h}} \rfloor,\\
g + h(q,i) &= \frac{iC_{\mathrm{h}}}{\mu} + \sum_{n=0}^{\infty} \omega(1-\omega)^n h(q-1+n, i-1), \quad q \geqslant 1,\ \ 1 \leqslant i \leqslant 1 + \lfloor K\mu/C_{\mathrm{h}} \rfloor,
\end{aligned} \tag{3.7}$$

where $\omega = \mu/(\lambda+\mu)$, the probability that a production completes before the next arrival occurs. The relationships in Eq. (3.7) can be proven easily. For instance, given the current state (q,i), the Markov chain can be in the state $(q-1+n, i-1)$ after the next transition with probability $\omega(1-\omega)^n$, for $n \geqslant 1$. Note that for the state (q,i) with $i > 0$, the optimal action is "no order" since the leadtimes are zero. Also note that in the state $(0,0)$ (i.e., $q = i = 0$), the optimal replenishment decision in the state $(0,0)$ is for the optimal decision in the state $(1,0)$ and it is indeed equivalent to the optimal decision in the state $(1,0)$.

Based on the solutions to Eq. (3.7) and the understanding of the replenishment decisions in the state (0, 0), the optimal replenishment policy $\boldsymbol{\pi}^*$ can be obtained by

$$\pi^*(q) = \begin{cases} 0, & q = 0, \\ \arg \min\limits_{1 \leqslant i \leqslant 1+\lfloor K\mu/C_h \rfloor} \{h(q,i)\}, & q \geqslant 1. \end{cases} \tag{3.8}$$

With the optimality equation (3.7), we are now able to analyze the optimal stationary replenishment policy.

4. Structure of the optimal stationary replenishment policy

In this section, we analyze the structure of the optimal replenishment policy and explain how and why the optimal replenishment policy fluctuates. In addition, we make a remark on the computation of the optimal replenishment policy. First, we introduce the following example to show numerically how a typical optimal replenishment policy behaves.

Example 4.1. Consider an inventory–production system with the following parameters: $C_h = 1$, $K = 30$, $\lambda = 0.3$, and $\mu = 1$. The optimal replenishment policy for this system is plotted in Fig. 3, in which the horizontal axis represents the queue length, the vertical axis represents the order size and

$$\mathrm{EOQ}(\mu) = \arg \min_{1 \leqslant i < \infty} \left\{ \frac{K}{i} + \frac{(i+1)C_h}{2\mu} \right\}. \tag{4.1}$$

$\mathrm{EOQ}(\mu)$ can be interpreted as the optimal order size in the warehouse if the demands to the warehouse follow a Poisson process with parameter μ [2].

As Fig. 3 illustrates, the optimal order size is strongly dependent on the queue length. Intuitively, when the queue length is near zero, the demand rate to the warehouse is small so that the order size of raw materials should be small in order to avoid unnecessary inventory build-up during possible idle periods in the workshop. As the queue length increases, the demand for raw materials in the workshop increases and so does the optimal order size. When the queue length is moderate, the optimal order size fluctuates around the $\mathrm{EOQ}(\mu)$ in order to reduce inventory holding costs during possible idle periods in the workshop. Lastly, when the queue length is large, the optimal order size converges to $\mathrm{EOQ}(\mu)$ since the demand rate to the warehouse will be μ (approximately) for a long period of time. In general, knowing the queue length at a decision epoch makes it possible to predict the length of the current busy period, or equivalently the average demand rate to the warehouse in the near future. Then the order size can be adjusted accordingly in order to reduce the total inventory costs.

From Example 4.1 and many other numerical examples, we observed the following important properties of the optimal replenishment policy.

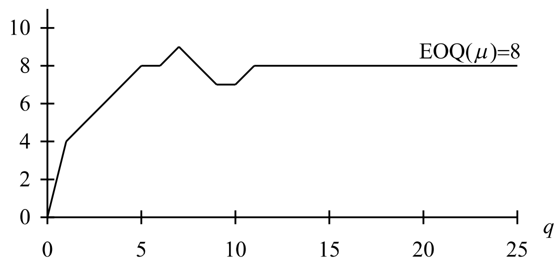

Fig. 3. The optimal replenishment policy.

Property 4.1. *For the inventory–production system introduced in Section 2, (a) the optimal order size $\pi^*(q)$ has a upper bound, and (b) $\pi^*(q) = \mathrm{EOQ}(\mu)$ for large q.*

Proof. Part (a) has been proved in Lemma 3.1. The basic argument for part (b) has been made already: when the queue length is large, the demand rate for raw materials will be μ for a long time. In such a case, the warehouse can be considered approximately as a stochastic inventory system with a Poisson demand process with parameter μ, which implies that $\mathrm{EOQ}(\mu)$ is the optimal choice for the order size. Details of the proof of part (b) are given in Appendix D. □

Remark 4.1. Property 4.1 provides important information about the structure of the optimal replenishment policy and can be used to develop algorithms for computing the optimal replenishment policy. For $m \geqslant 1$, define the following subsets of stationary replenishment policies:

$$\mathbf{\Pi}_b(m) = \{\boldsymbol{\pi} : \pi(0) = 0, 1 \leqslant \pi(q) \leqslant \lfloor 1 + K\mu/C_\mathrm{h} \rfloor \text{ for } 1 \leqslant q < m, \text{ and } \pi(q) = \mathrm{EOQ}(\mu), q \geqslant m\}. \quad (4.2)$$

According to Property 4.1, when m is large enough, the optimal replenishment policy in $\mathbf{\Pi}_b(m)$ is equivalent to the optimal replenishment policy in $\mathbf{\Pi}$. Thus, the optimal replenishment policy can be obtained by finding the optimal replenishment policy in $\mathbf{\Pi}_b(m)$ when m is large enough. According to He et al. [11], the problem of finding the optimal replenishment policy in $\mathbf{\Pi}_b(m)$ can be transformed into a Markov decision process with finite states and finite actions. Then some existing algorithms for computing the optimal policy of finite Markov decision processes can be utilized for computing the optimal replenishment policy efficiently. We refer to [11] and [15] for details of that computational approach.

Fluctuation in the optimal order size is strongly related to the systems parameters, particularly the traffic intensity $\rho = \lambda/\mu$ and the cost ratio K/C_h. We use the following example to demonstrate that relationship, which is useful in developing simple replenishment policies (Section 5) and analyzing those simple policies (Section 6).

Example 4.2. Consider an inventory–production system with system parameters $K = 10$, $C_\mathrm{h} = 0.2$, and $\mu = 1$. The optimal replenishment policies for $\lambda = 0.1$, 0.4, 0.618, and 0.95 are plotted in Fig. 4. The horizontal axis represents the queue length and the vertical axis represents the order size.

Fig. 4 shows that the structure of the optimal replenishment policy $\boldsymbol{\pi}$ depends strongly on the traffic intensity $\rho = \lambda/\mu$. If ρ is close to zero, the optimal order size fluctuates with the queue length dramatically

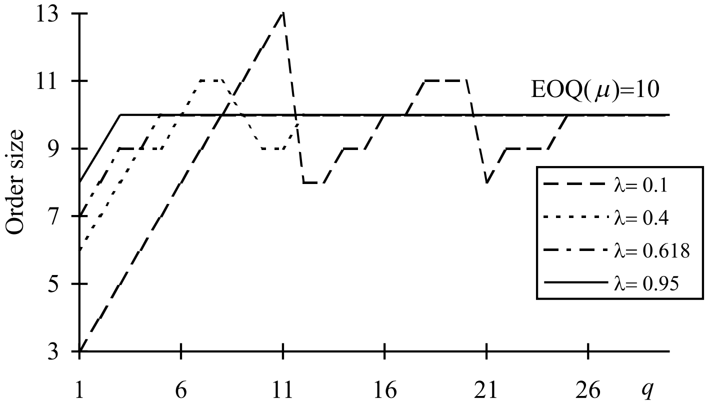

Fig. 4. The optimal policies for Example 4.2.

(see the curve for $\lambda = 0.1$ in Fig. 4). If ρ is close to 1, the optimal order size is EOQ(μ), even for a small queue length (see the curve for $\lambda = 0.95$ in Fig. 4). Thus, when the traffic intensity ρ increases from 0 to 1, the optimal replenishment policies becomes less variable and the optimal order size settles at EOQ(μ) quickly.

For the cost ratio K/C_h, numerical examples show that the optimal order size tends to be constant if K/C_h is small and is more variable if K/C_h is large. This is intuitively explained by Lemma 3.1.

To end this section, we briefly discuss why the optimal order size fluctuates with respect to the queue length. Consider a modified inventory–production system obtained by ignoring the holding costs incurred in all the idle periods of the workshop. The only difference between the modified system and the original system is how the holding costs incurred in idle periods are calculated. For the modified inventory–production system, the optimal replenishment policy can be found similarly. In fact, all the concepts, results of existence, and equations introduced in Section 3 can be used, except that the holding costs (iC_h/λ) incurred in the idle periods are removed from Eq. (3.7). It can be proven that, for the modified inventory–production system, the optimal order size is EOQ(μ) for all queue lengths except zero (details are omitted). That implies that, for the inventory–production system introduced in Section 2, the optimal order size fluctuates because of the holding cost incurred during idle periods. In other words, the optimal order size is adjusted according to the queue length so as to reduce the holding costs during possible idle periods.

5. Simple stationary replenishment policies

Property 4.1 shows that the optimal order size converges to EOQ(μ) as the queue length increases. However, prior to the convergence, the optimal order size fluctuates and can be somewhat complex. The complexity of the optimal replenishment policy makes it difficult to study theoretically, and it also makes it difficult to implement in practice. To overcome such difficulties, an alternative approach is to look for simple, but effective policies. Four simple policies, two "order-up-to" policies, a myopic policy (see [12] for more discussion on myopic policies), and a heuristic policy, are introduced as follows.

An order-up-to policy is defined as a replenishment policy $\boldsymbol{\pi} = (l, l, \ldots)$, i.e., $\pi(q) = l$ for all $q \geqslant 0$, where l is a positive integer. For such a replenishment policy, if the inventory level becomes zero, an order of the size l is placed immediately (regardless of the queue length). Order-up-to replenishment policies are simple to analyze (see [9]) and straightforward to implement. In fact, it has been proven in [9] that the average total costs per product of a replenishment policy $\boldsymbol{\pi} = (l, l, \ldots)$ is given by

$$g_\pi(l) = \frac{K}{l} + \frac{(l+1)C_h}{2\lambda}. \tag{5.1}$$

It is easy to see that expression (5.1) is minimized at EOQ(λ) (see Eq. (4.1) for definition).

Two order-up-to policies are of special interest: the EOQ(λ) policy and the EOQ(μ) policy, which are defined as the order-up-to policies with order sizes EOQ(λ) and EOQ(μ), respectively. The EOQ(λ) policy is interesting since it is the best order-up-to policy. The EOQ(μ) policy is interesting because EOQ(μ) is the order size of the optimal replenishment policy when the queue length is large.

The *myopic* policy is defined as the replenishment policy that minimizes the average total cost per product for the current production run (i.e., until the current raw materials inventory is depleted). No future cost on or after the next replenishment epoch is considered in making the current order size decision. The difference between the myopic policy and the optimal replenishment policy is how the costs after the next replenishment epoch are dealt with. Since the myopic policy ignores the influence of future costs on the current decision, it is simpler and an explicit description of the structure of it can be obtained. Let:

$V(q,i)$: The average total cost of producing exactly i products, given that the queue length is q and the inventory level is i initially.

The following recursive formulas for $\{V(q,i)\}$ can be easily derived:

$$\begin{aligned}
&V(q,0) = 0, \quad q \geqslant 0, \; i = 0,\\
&V(0,1) = \frac{C_{\mathrm{h}}}{\lambda} + \frac{C_{\mathrm{h}}}{\mu}, \quad q = 0, \; i = 1,\\
&V(q,1) = \frac{C_{\mathrm{h}}}{\mu}, \quad q \geqslant 0, \; i = 1,\\
&V(0,i) = \frac{iC_{\mathrm{h}}}{\lambda} + \frac{iC_{\mathrm{h}}}{\mu} + \sum_{n=0}^{\infty} \omega(1-\omega)^n V(n, i-1), \quad q = 0, \; i > 1,\\
&V(q,i) = \frac{iC_{\mathrm{h}}}{\mu} + \sum_{n=0}^{\infty} \omega(1-\omega)^n V(n+q-1, i-1), \quad q \geqslant 1, \; i > 1.
\end{aligned} \tag{5.2}$$

The myopic policy, denoted as $\boldsymbol{\pi}^m$, is defined as

$$\pi^m(q) = \begin{cases} 0, & q = 0,\\ \arg\min\limits_{i \geqslant 1} \left\{ \dfrac{K + V(q,i)}{i} \right\}, & q \geqslant 1. \end{cases} \tag{5.3}$$

The following result shows that the myopic order size is a monotone function of the queue length and it is no larger than $\mathrm{EOQ}(\mu)$. These properties are useful in determining upper bounds on i and q in computing the myopic policy.

Property 5.1. *For the myopic policy* $\boldsymbol{\pi}^m$, $\pi^m(q) \leqslant \mathrm{EOQ}(\mu)$ *for* $q \geqslant 0$, *and* $\pi^m(q) = \mathrm{EOQ}(\mu)$ *for* $q \geqslant \mathrm{EOQ}(\mu)$. $\pi^m(q)$ *is nondecreasing in* q.

Proof. See Appendix E. □

To find the order sizes of the myopic policy, Property 5.1 implies that the search range for both i and q are from zero to $\mathrm{EOQ}(\mu)$. Based on this result, the following efficient algorithm is developed for computing the myopic policy.

Step 1. Find $\mathrm{EOQ}(\mu)$ using Eq. (4.1).

Step 2. Calculate $V(q,i)$ recursively by using Eq. (5.2) for $0 \leqslant i$, $q \leqslant \mathrm{EOQ}(\mu)$.

Step 3. Find the optimal order size $\pi^m(q)$ using Eq. (5.3).

The myopic policy is much simpler than the optimal replenishment policy (see Table 1). However, the order size of the myopic policy for small queue lengths is not a simple function of the queue length. Thus, we consider a policy with an even simpler structure.

A Heuristic policy: For the optimal replenishment policy shown in Fig. 3, $\pi^*(1)$ is the smallest order size. As the queue length increases, the optimal order size first increases linearly, and then fluctuates around $\mathrm{EOQ}(\mu)$ until finally it coincides with $\mathrm{EOQ}(\mu)$. The heuristic policy approximates the optimal policy by removing the fluctuations in the optimal order size.

The heuristic policy $\boldsymbol{\pi}^h$ is constructed as follows. When $\pi^h(1) \leqslant \mathrm{EOQ}(\mu)$, the heuristic policy takes the following form:

$$\boldsymbol{\pi}^h = (0, \pi^h(1), \pi^h(1)+1, \ldots, \mathrm{EOQ}(\mu)-1, \mathrm{EOQ}(\mu), \mathrm{EOQ}(\mu), \ldots), \tag{5.4}$$

where $\pi^h(1)$ is an approximation of the optimal order size $\pi^*(1)$. When $\pi^h(1) > \mathrm{EOQ}(\mu)$, the heuristic policy has a slightly different form:

$$\boldsymbol{\pi}^h = (0, \pi^h(1), \pi^h(1) - 1, \ldots, \mathrm{EOQ}(\mu) + 1, \mathrm{EOQ}(\mu), \mathrm{EOQ}(\mu), \ldots). \tag{5.5}$$

A key in the construction of the heuristic policy is to find a good approximation $\pi^h(1)$ for $\pi^*(1)$ since there is no explicit solution of $\pi^*(1)$. Based on some properties about the optimal order size $\pi^*(1)$, the following $\pi^h(1)$ is proposed:

$$\pi^h(1) = \max\left\{ \frac{\min_{i \geqslant 1}\{(K/i) + ((i+1)C_\mathrm{h}/2\mu)\} - (C_\mathrm{h}/\mu)}{(C_\mathrm{h}/\mu) + \omega(C_\mathrm{h}/\lambda)}, \sqrt{\frac{2K}{C_h}\frac{1}{((1/\mu) + \omega(1/\lambda))}} \right\}. \tag{5.6}$$

The first term on the right-hand side of Eq. (5.6) comes from the inequality $(g - C_\mathrm{h}\mu)/(C_\mathrm{h}/\mu + \omega C_\mathrm{h}/\lambda) \leqslant \pi^*(1)$, which can be proved by using Eq. (3.7). Note that $\min_{i \geqslant 1}\{K/i + (i+1)C_\mathrm{h}/(2\mu)\}$ is used in Eq. (5.6) as an approximation (upper bound) of g. The second term on the right hand side of Eq. (5.6) is (approximately) $\mathrm{EOQ}(1/(1/\mu + \omega/\lambda))$ (see Eq. (4.1)), where the average arrival rate is $1/(1/\mu + \omega/\lambda)$. Note that $1/\mu + \omega/\lambda$ is an approximation of the mean time between two consecutive demands to the warehouse if the queue length is one initially.

6. Numerical examples

The replenishment policies introduced in Section 5 are simple, easy to compute, and easy to implement. However, they must perform well compared to the optimal replenishment policy to be effective substitutes. In this section, we compare the performance of these simple policies with that of the optimal policy and analyze numerically when these simple policies perform well.

Example 6.1 (Example 4.1 continued). Consider the inventory–production system with system parameters $C_\mathrm{h} = 1$, $K = 30$, $\lambda = 0.3$, and $\mu = 1$. The optimal replenishment policy, the myopic policy, the heuristic policy, and their corresponding average total cost per product (g) are given in Table 1.

According to Table 1, the difference in the average total costs per product of the optimal policy and the two order-up-to policies is significant. One of the reasons is that the order-up-to policies place an order even when the queue length is zero. Thus, these two policies do not perform well compared to the optimal policy in general. Nonetheless, numerical experiments do show that the performance of the two polices can be close to that of the optimal policy under some conditions.

First, both the $\mathrm{EOQ}(\lambda)$ and $\mathrm{EOQ}(\mu)$ policies perform well if the system traffic intensity ($\rho = \lambda/\mu$) is close to one. This is intuitive since, in this case, the optimal order size is $\mathrm{EOQ}(\mu)$ for almost all queue lengths and the idle periods are much shorter than the busy periods. Second, both $\mathrm{EOQ}(\lambda)$ and $\mathrm{EOQ}(\mu)$ can perform well if the cost ratio K/C_h is small. When K/C_h is small, the optimal order size will be close to one and the holding costs incurred during the idle periods account for a small portion in the total inventory costs. Note that the $\mathrm{EOQ}(\lambda)$ policy always outperforms the $\mathrm{EOQ}(\mu)$ policy since the $\mathrm{EOQ}(\lambda)$ policy is the best order-up-to policy.

Table 1
Comparison of replenishment policies

$q(t)$	0	1	2	3	4	5	6	7	8	9	10	$q \geqslant 11$	g
Optimal	0	4	5	6	7	8	8	9	8	7	7	8	15.64039
Myopic	0	4	5	5	6	6	7	7	8	8	8	8	15.65638
Heuristic	0	4	5	6	7	8	8	8	8	8	8	8	15.64044
$\mathrm{EOQ}(\lambda)$	4	4	4	4	4	4	4	4	4	4	4	4	15.83333
$\mathrm{EOQ}(\mu)$	8	8	8	8	8	8	8	8	8	8	8	8	18.75000

Table 1 also shows that the myopic and heuristic policies perform very close to that of the optimal policy. In fact, our numerical experiments show that these two simple policies perform extremely well for most of the cases we examined. We present the next example to demonstrate more specifically when the myopic and heuristic policies perform well.

Example 6.2 (Example 4.2 continued). We consider the inventory–production system with fixed $C_h = 0.2$ and $\mu = 1$ and changing K and λ. The corresponding average total costs per product of the three replenishment policies are shown in Table 2.

The computational results in Table 2 indicate that two simple inventory replenishment policies perform well in general. The reason is that both the myopic and heuristic policies choose order sizes close to the optimal ones, while removing the fluctuation in the optimal order sizes. Nonetheless, Table 2 does show the difference between the three policies. In fact, similar to the order-up-to policy case, Table 2 shows that the performance of two simpler policies is related to the traffic intensity $\rho = \lambda/\mu$ and the cost ratio K/C_h.

First, the two simple policies work extremely well when the traffic intensity is close to one. In this case, the optimal order size is $EOQ(\mu)$ for almost all the queue lengths (except zero). Thus, the three policies behave similarly. Second, the two simple policies work extremely well when the traffic intensity is close to zero. In this case, the optimal policy usually fluctuates dramatically. However, since the traffic intensity is small, the queue is small most of the time. Therefore, only the order sizes corresponding to small queue length matter. Since the myopic and heuristic policies match the optimal policy when the queue length is small, their performance is close to the optimal policy. Third, when the traffic intensity is moderate, the myopic or heuristic policy can achieve good performance but the difference between these simple policies and the optimal one can be significant.

In terms of the cost ratio K/C_h, Table 2 shows that the difference between the simple policies and the optimal policy becomes larger when the cost ratio is larger. If the cost ratio K/C_h is small, the optimal order size is small (Lemma 3.1). The two simple policies are close to the optimal policy with respect to the order sizes. If the ratio K/C_h is large, more choices are available for the optimal order size (which can fluctuate), while the two simple policies have fewer choices of order size because of their structure. Thus, the difference between the optimal policy and the two simple policies can be significant.

Table 2
Cost comparison for systems with $C_h = 0.2$ and $\mu = 1$

		$K = 0.1$	$K = 0.5$	$K = 1$	$K = 10$
$\lambda = 0.1$	Optimal	0.3	0.697409	1.192844	5.527922
	Myopic	0.3	0.697450	1.192859	5.528331
	Heuristic	0.3	0.697450	1.192859	5.527922
$\lambda = 0.4$	Optimal	0.3	0.668832	0.940697	3.099032
	Myopic	0.3	0.668953	0.943844	3.106538
	Heuristic	0.3	0.668953	0.940697	3.099035
$\lambda = 0.618$	Optimal	0.3	0.611812	0.835806	2.568029
	Myopic	0.3	0.611812	0.835806	2.576088
	Heuristic	0.3	0.611812	0.835806	2.578051
$\lambda = 0.95$	Optimal	0.3	0.555262	0.743858	2.143805
	Myopic	0.3	0.555262	0.743858	2.144449
	Heuristic	0.3	0.555262	0.743858	2.147364

7. Conclusions and future research

Several inventory replenishment policies are examined in this paper for the make-to-order inventory–production system depicted in Fig. 1. It was shown that the optimal stationary replenishment policy exists and can be obtained by solving an optimality equation. It was found that the optimal order size, though fluctuating, is bounded and converges to EOQ(μ) when the queue length increases. The characterization of the optimal replenishment policy, though not complete, provides good insights into the structure of the optimal policy. These insights led to the introduction of the myopic and heuristic policies. The myopic and heuristic policies have a simpler structure than the optimal one and yet they perform extremely close to the optimal policy. This is useful in practice where a simple control scheme is more feasible to implement.

The current work can be generalized in two directions. The first one includes investigating the value of queue length information in the replenishment decision. Some results in this direction have been reported in [10]. The second research direction is to consider an inventory–production system with nonzero leadtimes. In such cases, orders might be issued when the inventory level is still positive, i.e., the optimal order size will depend on the inventory level as well as the queue length. The problem becomes much more challenging than the current one.

Appendix A. Proof of Lemma 3.1

We prove Lemma 3.1 by comparing the following two ordering schemes for $i+1$ units of raw materials. The first scheme is to place an order of the size $i+1$. The second scheme is to place two orders of the size i and 1 (one) respectively. In both cases, the $i+1$ units of raw materials are used to produce the next $i+1$ products.

Now, we compare the average total costs associated with the two schemes. In terms of the ordering costs, it is clear that the second scheme leads to an extra cost of the amount K. For the holding costs, we consider the average total holding costs associated with each of the $i+1$ units. It is easy to see that the average total holding costs associated with the first i units used in production are the same for both schemes. But the average total holding costs associated with the last unit of raw materials (i.e., the $(i+1)$th unit) is different for the two schemes. Since the last unit is in the system from the very beginning for the first scheme, holding costs are incurred during the production of the first i units and any possible idle periods between these productions. That extra holding cost (associated with the $(i+1)$th unit of raw materials) is at least iC_{h}/μ. Note that the mean production time of a product is $1/\mu$. The extra average total holding cost associated with the $(i+1)$th units can be larger than iC_{h}/μ since the holding cost incurred during possible idle periods of the workshop (when $q(t)=0$) is not included in the expression.

If $i \geqslant 1+K\mu/C_{\mathrm{h}}$, it is easy to verify that $iC_{\mathrm{h}}/\mu > K$, which implies that the first scheme has a larger average total cost. Thus, the first scheme has a larger average total cost per product as well. Therefore, in order to reduce the total inventory cost per product, the order size must be less than $1+K\mu/C_{\mathrm{h}}$, regardless of the queue length at the decision epoch. This completes the proof of Lemma 3.1. □

Appendix B. Proof of Lemma 3.2

For any replenishment policy $\boldsymbol{\pi}$ in $\boldsymbol{\Pi}_b$, we organize the states of the corresponding Markov chain $\{(q_n, I_n), n \geqslant 0\}$ according to its first coordinate. We call the set of states $\{(q,i), 0 \leqslant i \leqslant \lfloor 1+K\mu/C_{\mathrm{h}} \rfloor\}$ the level q of states for $q \geqslant 0$.

For the Markov chain $\{(q_n, I_n), n \geqslant 0\}$, the state $(0,0)$ can be reached from any other states for two reasons: (1) since the M/M/1 queue is not influenced by inventory control, q_n will become zero after a finite number of transitions, which implies that the level 0 is always reachable from any state; (2) since the number of products produced in a busy period can be any positive number, it is possible that the inventory level becomes zero at the end of a busy period. Consequently, the Markov chain $\{(q_n, I_n), n \geqslant 0\}$ has only one closed set.

Since the M/M/1 queue is not influenced by inventory control, the limit of the probability $\mathbf{P}\{q_n = 0\}$ as n goes to infinity exists and is $1 - \rho$, i.e., the probability that the queueing system is empty (see [5]). Since

$$\mathbf{P}\{q_n = 0\} = \mathbf{P}\{q_n = 0, I_n = 0\} + \mathbf{P}\{q_n = 0, I_n = 1\} + \cdots + \mathbf{P}\{q_n = 0, I_n = 1 + \lfloor 1 + K\mu/C_\mathrm{h}\rfloor\},$$

$\lim_{\{n\to\infty\}} \mathbf{P}\{q_n = 0\} = 1 - \rho$, and the state $(0,0)$ is reachable from any other state, we must have $x(0,0) = \lim_{\{n\to\infty\}} \mathbf{P}\{q_n = 0, I_n = 0\} > 0$. That implies that the Markov chain $\{(q_n, I_n), n \geqslant 0\}$ is positive recurrent if it is restricted to its only closed set. Finally, we must have

$$x(q,0) + x(q,1) + \cdots + x(q, \lfloor 1 + K\mu/C_\mathrm{h}\rfloor) = \lim_{\{n\to\infty\}} \mathbf{P}\{q_n = q\} = (1-\rho)\rho^q \quad \text{for } q \geqslant 0.$$

This completes the proof of Lemma 3.2. □

Appendix C. Proof of Theorem 3.3

We apply Theorem 10.4 in [13] to prove the theorem. In order to do so, we need to verify two conditions:

(1) for any stationary replenishment policy $\boldsymbol{\pi}$ in $\mathbf{\Pi}_b$, the corresponding Markov chain $\{(q_n, I_n), n > 0\}$ is recurrent;

(2) the stationary distributions $\{x(q,i) : q \geqslant 0, 0 \leqslant i \leqslant \lfloor 1 + K\mu/C_\mathrm{h}\rfloor\}$ of all replenishment policies in $\mathbf{\Pi}_b$ is tight.

First, condition (1) is satisfied because of Lemma 3.2. Second, condition (2) is satisfied because, according to Lemma 3.2,

$$x(q,0) + x(q,1) + \cdots + x(q, \lfloor 1 + K\mu/C_\mathrm{h}\rfloor) = (1-\rho)\rho^q \quad \text{for } q \geqslant 0$$

for any replenishment policy in $\mathbf{\Pi}_b$. Since the Markov decision process has a countable state space, finite decision sets, and bounded immediate costs, we conclude that all the conditions of in Theorem 10.4 [13] are satisfied. Therefore, an optimal replenishment policy exists in $\mathbf{\Pi}_b$. Furthermore, the optimality equations hold and can be used to find the optimal replenishment policy. Finally, according to Lemma 3.1, the optimal policy in $\mathbf{\Pi}_b$ is also optimal in $\mathbf{\Pi}$. This completes the proof of Theorem 3.3. □

Appendix D. Proof of Property 4.1

To prove Property 4.1, we first show that the sequence $\{h(q,i) - h(q+1,i), q \geqslant 0\}$ converge to a common limit for all $0 \leqslant i \leqslant \lfloor 1 + K\mu/C_\mathrm{h}\rfloor$. That implies that $\{h(q,i), 0 \leqslant i \leqslant \lfloor 1 + K\mu/C_\mathrm{h}\rfloor\}$ is minimized at a common order size for large q. Then we show that $\{h(q,i), 0 \leqslant i \leqslant \lfloor 1 + K\mu/C_\mathrm{h}\rfloor\}$ is minimized at EOQ(μ) for large q. First, the following preliminary results are needed.

Lemma D.1. *The relative functions $\{h(q,i)\}$ introduced in Section 3 satisfy*

$$0 \leqslant h(q,i) - h(q+1,i) \leqslant \frac{(1 + \lfloor K\mu/C_\mathrm{h}\rfloor)C_\mathrm{h}}{\lambda}. \tag{D.1}$$

Proof. We consider the cost functions $V^*(q,i,n)$ introduced in Section 3. We first estimate the difference $V^*(q,i,n) - V^*(q+1,i,n)$ by considering two systems starting with (q,i) and $(q+1,i)$, respectively. If the system starting in (q, i) places orders using the optimal order sizes of the system starting in $(q+1,i)$ until the queue length goes to zero for the first time, the average total cost incurred is higher than that of using its own optimal order sizes. When the queue length of the system starting in (q, i) reaches zero, the workshop stops working until the next demand arrives so that the system becomes $(1, i')$, where i' is the inventory level right after the epoch when the queue length becomes zero. The solid line in Fig. 5 shows the change in $q(t)$ when $q(0) = q$. It is not difficult to see that $(1, i')$ is the state of the system starting in $(q, i+1)$ at the epoch when the other system becomes $(0, i')$. The dashed line in Fig. 5 shows the change in $q(t)$ when $q(0) = q+1$. If the system with $q(0) = q$ does not reach queue length zero before finishing n products, then the two systems has the same total costs (but the system with $q(0) = q$ is not operating optimally).

Now, it is clear that the costs incurred in both systems after their corresponding $(1, i')$ epochs are the same, but the system starting in (q, i) has paid an extra holding cost $i'C_\mathrm{h}/\lambda$, but it is not operating optimally. Therefore, it must be that

$$V^*(q,i,n) \leqslant V^*(q+1,i,n) + i'C_\mathrm{h}/\lambda \leqslant V^*(q+1,i,n) + (1 + \lfloor K\mu/C_\mathrm{h} \rfloor)C_\mathrm{h}/\lambda.$$

By Eqs. (3.5) and (3.6), we obtain

$$h(q,i,) \leqslant h(q+1,i) + (1 + \lfloor K\mu/C_\mathrm{h} \rfloor)C_\mathrm{h}/\lambda.$$

On the other hand, if both systems using the optimal replenishment policy of the system starting with (q, i), the same argument leads to

$$V^*(q,i,n) \geqslant V^*(q+1,i,n) + i'C_\mathrm{h}/\lambda \geqslant V^*(q+1,i,n),$$

which leads to $h(q,i) \geqslant h(q+1,i)$. This completes the proof of Lemma D.1. □

This result implies that the decrease in cost with respect to the increases in queue length is bounded uniformly, i.e., the increase in queue length by one can only reduce cost by a limited amount. The next result shows that the magnitude of the decrease converges when the queue length goes to infinity.

Lemma D.2. *For the relative functions defined by Eq.* (3.6), *we have, for* $0 \leqslant i \leqslant \lfloor 1 + K\mu/C_\mathrm{h} \rfloor$,

$$\lim_{q\to\infty} [h(q,i) - h(q+1,i)] = \eta_\infty, \tag{D.2}$$

where η_∞ *is a finite nonnegative constant.*

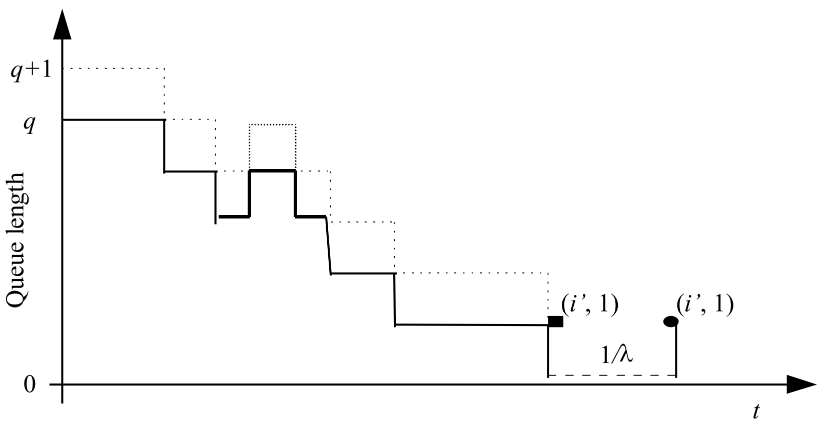

Fig. 5. The sample paths of the queue length.

Proof. Since $\{h(q,i,)-h(q+1,i), q\geqslant 0, 0\leqslant i\leqslant \lfloor 1+K\mu/C_h\rfloor\}$ are all bounded, then any subsequence has at least one converging subsequence. We shall use these converging subsequences to prove Property 4.1. Let, for $0\leqslant i\leqslant \lfloor 1+K\mu/C_h\rfloor$,

$$\bar{\eta}(i)=\lim_{q\to\infty}\sup[h(q,i)-h(q+1,i)] \quad \text{and} \quad \underline{\eta}(i)=\lim_{q\to\infty}\inf[h(q,i)-h(q+1,i)]. \tag{D.3}$$

For $q>i$, there is no idle period before the next ordering epoch. Then, by Eq. (3.7) and Eq. (E.1), for large q ($>i$),

$$\begin{aligned} h(q,i)-h(q+1,i) &= \sum_{j=0}^{\infty}\omega(1-\omega)^j[h(q-1+j,i-1)-h(q+j,i-1)] \\ &= \sum_{j=0}^{\infty}\left(P^i_{(q,j)}-P^i_{(q+1,j)}\right)h(j,0)=\sum_{j=0}^{\infty}P^i_{(q,j)}[h(j,0)-h(j+1,0)], \end{aligned} \tag{D.4}$$

where $P^i_{(q,j)}$ is the (q,j)th element of the matrix P^i. Note that the qth row of the matrix P^i is obtained by shifting the $(q+1)$th row of the matrix P^i to the left by one (when $q>i$). Since P is a stochastic matrix, so is P^i. Then Eq. (D.4) leads to, using Fatou Lemma,

$$\underline{\eta}(0)\leqslant\underline{\eta}(1)\leqslant\cdots\leqslant\underline{\eta}(1+\lfloor K\mu/C_h\rfloor)\leqslant\bar{\eta}(1+\lfloor K\mu/C_h\rfloor)\leqslant\cdots\leqslant\bar{\eta}(1)\leqslant\bar{\eta}(0). \tag{D.5}$$

By Lemma D.1, we must have $0\leqslant\bar{\eta}(0)\leqslant\lfloor 1+K\mu/C_h\rfloor C_h/\lambda$. Let $\aleph(\eta)=\{j: h(j,0)-h(j+1,0)<\eta\}$ for some $\eta<\bar{\eta}(0)$. Suppose that a subsequence $\{q_n\}$ leads $\{h(q_n,0)-h(q_n+1,0)\}$ to $\bar{\eta}(0)$. Then we have

$$\begin{aligned} h(q_n,\pi^*(q_n+1))-h(q_n+1,\pi^*(q_n+1)) &= \sum_{j=0}^{\infty}p^{\pi^*(q_n+1)}_{(q_n,j)}[h(q_n+j-\pi^*(q_n+1),0)-h(q_n+1+j-\pi^*(q_n+1),0)] \\ &\leqslant \sum_{j=0:j\notin\aleph(\eta)}^{\infty}p^{\pi^*(q_n+1)}_{(q_n,j)}(\bar{\eta}(0)+\varepsilon)+\sum_{j=0:j\in\aleph(\eta)}^{\infty}p^{\pi^*(q_n+1)}_{(q_n,j)}\eta \\ &\equiv [1-S(q_n,\aleph(\eta))](\bar{\eta}(0)+\varepsilon)+S(q_n,\aleph(\eta))\eta. \end{aligned} \tag{D.6}$$

On the other hand,

$$h(q_n,\pi^*(q_n+1))-h(q_n+1,\pi^*(q_n+1))\geqslant h(q_n,0)-h(q_n+1,0)\geqslant\bar{\eta}(0)-\varepsilon. \tag{D.7}$$

Then Eqs. (D.6) and (D.7) yield

$$\bar{\eta}(0)-\varepsilon\leqslant[1-S(q_n,\aleph(\eta))](\bar{\eta}(0)+\varepsilon)+S(q_n,\aleph(\eta))\eta, \quad \text{i.e.,} \quad S(q_n,\aleph(\eta))\leqslant\frac{2\varepsilon}{\bar{\eta}(0)+\varepsilon-\eta}. \tag{D.8}$$

Eq. (D.8) leads to $S(q_n,\aleph(\eta))\to 0$ when $q_n\to\infty$. That implies that $h(q,0)-h(q+1,0)$ is approximately as large as η when q is large enough. Since that conclusion holds for any η that is less than $\bar{\eta}(0)$, we can conclude that any convergent sequence of $\{h(q,0)-h(q+1,0)\}$ converges to $\bar{\eta}(0)$. Therefore, $\bar{\eta}(0)=\underline{\eta}(0)$ and the sequence $\{h(q,0)-h(q+1,0)\}$ converges. By Eq. (D.5), all the sequences $\{h(q,i)-h(q+1,i)\}$ converge and the limits are the same. □

Using Lemma D.2, we prove the following useful result.

Lemma D.3. *Let*

$$\Delta(q)=K+h(q,1)-h(q,0) \quad \textit{for } q\geqslant 1. \tag{D.9}$$

Then $\Delta^* = \lim_{q\to\infty} \Delta(q)$ *exists and is finite and nonnegative.*

Proof. By definition, $K + h(q,1) \geqslant h(q,0)$ so that $\Delta(q)$ is nonnegative. If $h(q,0) < h(q,1)$, it has $K + h(q,\pi^*(q)) < h(q,1)$. Then it is optimal to order at $(q,1)$ to bring down the cost, which is a contradiction. Thus $h(q,0) \geqslant h(q,1)$ and $\Delta(q)$ is upper bounded by *K*. Then $\{\Delta(q), q \geqslant 0\}$ must have a convergence subsequence. Since

$$\Delta(q) - \Delta(q+1) = h(q,1) - h(q+1,1) - h(q,0) + h(q+1,0) \to 0, \quad q \to \infty,$$

it is easy to see, by Lemma D.2, that there is only one limit. The limit is nonnegative and less than *K*. This completes the proof of Lemma D.3. □

Using Lemmas D.2 and D.3, the following property shows that the optimal order size converges to $\mathrm{EOQ}(\mu)$, which is Property 4.1.

Property D.4. $\lim_{\{q\to\infty\}} \pi^*(q) = \mathrm{EOQ}(\mu)$ *and* $\Delta^* = K - C_{\mathrm{EOQ}}(\mathrm{EOQ}(\mu)) + C_{\mathrm{h}}/\mu$, *where* $C_{\mathrm{EOQ}}(i) = K/i + (i+1)C_{\mathrm{h}}/(2\mu)$.

Proof. $\forall \varepsilon > 0$, q_ε satisfies for all $q > q_\varepsilon$, $|\Delta(q) - \Delta^*| < \varepsilon$. By Lemma D.3, it must have

$$\frac{iC_{\mathrm{h}}}{\mu} - K + \Delta^* - \varepsilon < h(q,i+1) - h(q,i) < \frac{iC_{\mathrm{h}}}{\mu} - K + \Delta^* + \varepsilon. \tag{D.10}$$

Taking the summation for *i* from 1 to *i* in the above inequality gives, for $i > 1$,

$$\begin{aligned} h(q,i) - h(q,1) &> \frac{i(i-1)C_{\mathrm{h}}}{2\mu} - (i-1)K + (i-1)\Delta^* - (i-1)\varepsilon, \\ h(q,i) - h(q,1) &< \frac{i(i-1)C_{\mathrm{h}}}{2\mu} - (i-1)K + (i-1)\Delta^* + (i-1)\varepsilon. \end{aligned} \tag{D.11}$$

The above inequalities further imply that

$$\begin{aligned} h(q,0) - (K + h(q,1)) &> \min_{1\leqslant i\leqslant \lfloor 1+K\mu/C_{\mathrm{h}}\rfloor} \left\{ \frac{i(i-1)C_{\mathrm{h}}}{2\mu} - (i-1)(K-\Delta^*) - (i-1)\varepsilon \right\}, \\ h(q,0) - (K + h(q,1)) &< \min_{1\leqslant i\leqslant \lfloor 1+K\mu/C_{\mathrm{h}}\rfloor} \left\{ \frac{i(i-1)C_{\mathrm{h}}}{2\mu} - (i-1)(K-\Delta^*) + (i-1)\varepsilon \right\}. \end{aligned} \tag{D.12}$$

Letting $q \to \infty$ and $\varepsilon \to 0$, by the definition of Δ^*, gives

$$\Delta^* = -\min_{1\leqslant i\leqslant \lfloor 1+K\mu/C_{\mathrm{h}}\rfloor} \left\{ \frac{i(i-1)C_{\mathrm{h}}}{2\mu} - (i-1)(K-\Delta^*) \right\} \Rightarrow 0 = \min_{1\leqslant i\leqslant \lfloor 1+K\mu/C_{\mathrm{h}}\rfloor} \left\{ \frac{i(i-1)C_{\mathrm{h}}}{2\mu} - (i-1)K + i\Delta^* \right\}.$$

Dividing both sides of the last equation by *i* yields

$$0 = \min_{1\leqslant i\leqslant \lfloor 1+K\mu/C_{\mathrm{h}}\rfloor} \left\{ \frac{(i-1)C_{\mathrm{h}}}{2\mu} + \frac{K}{i} + \Delta^* - K \right\} \Rightarrow \frac{C_{\mathrm{h}}}{\mu} + K - \Delta^* = \min_{1\leqslant i\leqslant \lfloor 1+K\mu/C_{\mathrm{h}}\rfloor} \left\{ \frac{(i+1)C_{\mathrm{h}}}{2\mu} + \frac{K}{i} \right\}.$$

The right-hand side of the last equation is $C_{\mathrm{EOQ}}(\mathrm{EOQ}(\mu))$, which gives the value of Δ^*.

Replacing Δ^* in Eq. (D.12) results in, for $q > q_\varepsilon$,

$$\left| \frac{(i+1)C_{\mathrm{h}}}{2\mu} - C_{\mathrm{EOQ}}(\mathrm{EOQ}(\mu)) - [h(q,i+1) - h(q,i)] \right| < \varepsilon. \tag{D.13}$$

This implies that $\{h(q,i+1)-h(q,i)\}$ (when q is large enough) lies between two linear functions (with respect to i) which are independent of q. The two linear functions are parallel to each other and could be arbitrarily close to each other. Then $\{h(i,q)\}$ is minimized at a common point when ε is small enough. This common point is given by

$$\min_{1\leqslant i\leqslant \lfloor 1+K\mu/C_{\rm h}\rfloor} \arg\left\{\frac{(i+1)C_{\rm h}}{\mu}-C_{\rm EOQ}({\rm EOQ}(\mu)):\ \text{closest to zero}\right\}$$
$$=\min_{1\leqslant i\leqslant \lfloor 1+K\mu/C_{\rm h}\rfloor} \arg\left\{\min\left\{\frac{(i+1)C_{\rm h}}{2\mu}-\frac{K}{i}\right\}:\ \text{closest to zero}\right\}={\rm EOQ}(\mu). \tag{D.14}$$

The last equality is true since

$$[{\rm EOQ}(\mu)+1]C_{\rm h}/(2\mu)>K/{\rm EOQ}(\mu)>[{\rm EOQ}(\mu)-1]C_{\rm h}/(2\mu),$$

and, when $i>{\rm EOQ}(\mu)$, it must have

$$(i+1)C_{\rm h}/(2\mu)-K/i>[{\rm EOQ}(\mu)+1]C_{\rm h}/(2\mu)-K/{\rm EOQ}(\mu)>0;$$

when $i<{\rm EOQ}(\mu)$, it has

$$(i+1)C_{\rm h}/(2\mu)-K/i<[{\rm EOQ}(\mu)-1]C_{\rm h}/(2\mu)+C_{\rm h}/(2\mu)-K/i<[{\rm EOQ}(\mu)-1]C_{\rm h}/(2\mu)-K/{\rm EOQ}(\mu)<0.$$

This completes the proof of Property D.4. □

Appendix E. Proof of Property 5.1

To prove Property 5.1, we need the concepts of majorization and monotonicity [14]. For vectors $\mathbf{x}$ and $\mathbf{y}$ in $\Re^\infty$, we say that $\mathbf{x}$ is *majorized* by $\mathbf{y}$ if and only if $\sum_{i=1}^n x_i\leqslant\sum_{i=1}^n y_i$ for $n\geqslant 1$. Denote this relationship as $\mathbf{x}\prec\mathbf{y}$. Let X be a matrix with rows $\{\mathbf{x}_n\}$. The matrix X is called *monotone* if $\mathbf{x}_1\succ\mathbf{x}_2\succ\mathbf{x}_3\succ\cdots$. We note that the concept of "majorization" and "monotonicity" defined here is slightly different from that in some classical literature [14]. Let

$$P=\begin{pmatrix}\omega & \omega(1-\omega) & \omega(1-\omega)^2 & \cdots\\ \omega & \omega(1-\omega) & \omega(1-\omega)^2 & \cdots\\ & \omega & \omega(1-\omega) & \ddots\\ & & \omega & \ddots\\ & & & \ddots\end{pmatrix}. \tag{E.1}$$

By definition, it is easy to see that the matrix P^j is monotone for $j\geqslant 1$. Therefore, $(P^j)_{q,0}\geqslant(P^j)_{q+1,0}$ for $q\geqslant 0$ and $j\geqslant 0$. Proofs of these conclusions can be found in [14].

Let $\mathbf{V}(i)=(V(0,i),V(1,i),\ldots)^{\rm T}$, where "T" represents the transpose of matrix. Eq. (5.2) can be rewritten as

$$V(i)=\frac{i(i+1)C_{\rm h}}{2\mu}\mathbf{e}+\sum_{j=1}^{i}P^{i-j}\begin{pmatrix} jC_{\rm h}/\lambda\\ 0\\ \vdots\end{pmatrix},\quad i\geqslant 1. \tag{E.2}$$

Eq. (E.2) gives the cost functions $\{V(q,i)\}$ explicitly. By Eqs. (5.2) and (E.2), we have

$$\begin{aligned}
\frac{K+V(q,i+1)}{i+1}-\frac{K+V(q,i)}{i} &= \frac{iV(q,i+1)-(i+1)V(q,i)-K}{i(i+1)} \\
&= \frac{1}{i(i+1)}\left\{ i\left[\frac{(i+1)(i+2)C_\mathrm{h}}{2\mu}+\sum_{j=1}^{i+1}(P^{i+1-j})_{q,0}\frac{jC_\mathrm{h}}{\lambda}\right]\right. \\
&\quad \left. -(i+1)\left[\frac{i(i+1)C_\mathrm{h}}{2\mu}+\sum_{j=1}^{i}(P^{i-j})_{q,0}\frac{jC_\mathrm{h}}{\lambda}\right]-K\right\} \\
&= \frac{1}{i(i+1)}\left[\frac{i(i+1)C_\mathrm{h}}{2\mu}+\frac{C_\mathrm{h}}{\lambda}\sum_{j=0}^{i}j(P^j)_{q,0}-K\right] \\
&\geqslant \frac{1}{i(i+1)}\left[\frac{i(i+1)C_\mathrm{h}}{2\mu}+\frac{C_\mathrm{h}}{\lambda}\sum_{j=0}^{i}j(P^j)_{q+1,0}-K\right] \\
&= \frac{K+V(q+1,i+1)}{i+1}-\frac{K+V(q+1,i)}{i}.
\end{aligned} \tag{E.3}$$

In Eq. (E.3), we used the fact that $(P^j)_{q,0} > (P^j)_{q+1,0}$ since P^j is monotone. Note that $(P^j)_{q,0}$ can be interpreted as the probability that the M/M/1 queue has zero customer after j service completions, given that there are q customers in the queue initially. Therefore

$$\frac{K+V(q,i+j)}{i+j}-\frac{K+V(q,i)}{i} \geqslant \frac{K+V(q+1,i+j)}{i+j}-\frac{K+V(q+1,i)}{i}, \tag{E.4}$$

which implies that $\pi^m(q) \leqslant \pi^m(q+1)$. The second equality in (E.3) implies that

$$\left[\frac{K+V(q,i+1)}{i+1}-\frac{K+V(q,i)}{i}\right] \geqslant \frac{1}{i(i+1)}\left[\frac{i(i+1)C_\mathrm{h}}{2\mu}-K\right] = \frac{1}{(i+1)}\left[\frac{(i+1)C_\mathrm{h}}{2\mu}-\frac{K}{i}\right]. \tag{E.5}$$

For $i > \mathrm{EOQ}(\mu)$, the right-hand side of Eq. (E.5) is positive and so is the left-hand side. Therefore, i ($>\mathrm{EOQ}(\mu)$) is not optimal. When $q \geqslant \mathrm{EOQ}(\mu)$, optimal order size must be $\mathrm{EOQ}(\mu)$ since it is where the average total cost per product is minimized under that condition. This completes the proof of Property 5.1. □

References

[1] S. Axsäter, Continuous review policies for multi-level inventory systems with stochastic demand, in: Logistics of Production and Inventory, Handbooks in Operations Research and Management Science, vol. 4, North-Holland, Amsterdam, 1993, pp. 195–198.

[2] D. Bartmann, M.J. Beckmann, Inventory Control: Models and Methods, Springer, Berlin, 1992.

[3] R. Bellman, Dynamic Programming, Princeton University Press, Princeton, NJ, 1957.

[4] F. Chen, Y.S. Zheng, Evaluating Echelon stock (R, nQ) policies in serial production/inventory systems with stochastic demand, Management Science 40 (1994) 1262–1275.

[5] J.W. Cohen, The Single Server Queue, in: North-Holland Series in Applied Mathematics and Mechanics, North-Holland, Amsterdam, 1982.

[6] A. Federgruen, Centralized planning models for multi-echelon inventory systems under uncertainty, in: Logistics of Production and Inventory, Handbooks in Operations Research and Management Science, vol. 4, North-Holland, Amsterdam, 1993, pp. 133–174.

[7] G. Hadley, T.M. Whitin, Analysis of Inventory Systems, Prentice-Hall, Englewood Cliffs, NJ, 1963.

[8] R. Hariharan, P. Zipkin, Customer-order information, leadtimes, and inventories, Management Science 41 (10) (1995) 1599–1607.

[9] Q.-M. He, E.M. Jewkes, Performance measures of a make-to-order inventory–production system, IIE Transactions 32 (2000) 409–419.

[10] Q.-M. He, E.M. Jewkes, J. Buzacott, Analysis of the value of information used in inventory control of an inventory–production system, in: Proceedings of ABS/ACORS 99, Halifax, Canada, 1999.

[11] Q.-M. He, E.M. Jewkes, J. Buzacott, An efficient algorithm for computing the optimal replenishment policy for an inventory–production system, in: A.S. Alfa, S. Chakravarthy (Eds.), Advances in Matrix-Analytic Methods for Stochastic Models, Proceedings of the Second International Conference on Matrix-Analytic Methods, Notable Publications, Inc, New Jersey, 1998, pp. 381–401.

[12] D.P. Heyman, M.J. Sobel, Stochastic Models in Operations Research, vol. II, McGraw-Hill, New York, 1984.

[13] A. Hordijk, Dynamic Programming and Markov Potential Theory, vol. 51, Mathematical Centre Tracts, Amsterdam, 1974.

[14] A.W. Marshall, I. Olkin, Inequalities: Theory of Majorization and its Applications, Academic Press, New York, 1979.

[15] M.F. Neuts, Matrix-Geometric Solutions in Stochastic Models: An algorithmic Approach, Johns Hopkins University Press, Baltimore, MD, 1981.

[16] M.L. Puterman, Markov Decision Processes, Wiley, New York, 1994.

[17] L.I. Sennott, Stochastic Dynamic Programming and the Control of Queueing Systems, Wiley, New York, 1999.

[18] M.H. Veatch, L.M. Wein, Optimal control of a two-station tandem production/inventory system, Operations Research 42 (1994) 337–350.