ASB and ACORS 1999 Conference
Dalhousie University, Halifax

Qi-Ming HE
Department of Industrial Engineering
DalTech, Dalhousie University

Elizabeth M. Jewkes
Department Management Science
University of Waterloo

John Buzacott
Schulich School of Business, York University

ANALYSIS OF THE VALUE OF INFORMATION USED IN INVENTORY CONTROL OF AN INVENTORY-PRODUCTION SYSTEM

This paper studies the value of information in an inventory-production system consisting of a warehouse and a workshop. The focal point is the use of information about the production status of the workshop in inventory control in the warehouse. Systems with various information levels are analyzed to determine what information is most valuable. Examples are presented to demonstrate how to identify the most important information about production status.

1. Introduction

This paper investigates a simple supply chain in an attempt to gain a better understanding of the benefits of information in making inventory replenishment decisions. The supply chain consists of a supplier, a raw materials warehouse, and a manufacturer (see Figure 1.1). The manufacturer makes products to order based on customer demand, and uses raw materials as inputs to the production process. Raw materials are replenished by the supplier.

In making decisions about replenishing the warehouse inventory, it would be normal to use information about the current level of inventory in the warehouse. However, it would be useful to have up-to-date information on the number of outstanding unfilled orders in the workshop. But such information may not be readily available and it may be difficult and expensive to collect, particularly as the level of detail required increases. To assess the value of obtaining such information, this paper looks at the extent to which information on the outstanding orders benefits inventory control at the warehouse.

The supply chain shown in Figure 1.1 is a special two-echelon model (Federgruen, 1993). Studies of the optimal or nearly optimal replenishment policies for two echelon models are extensive (Axsäter, 1993; Federgruen, 1993; Veatch and Wein, 1994; and references therein). In particular, for the inventory-production model studied in this paper, HE, Jewkes, and Buzacott (1997, 1999) developed algorithms for computing the optimal replenishment policies. HE and Jewkes (1999) developed efficient algorithms for computing the average total cost per product and other performance measures. The results obtained in the last three papers are used in this paper to study the value of information about the production status to replenishment decisions.

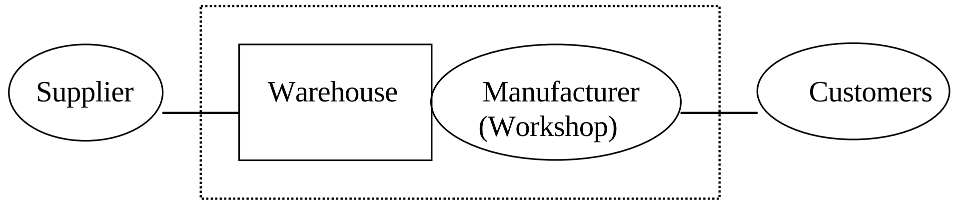

Figure 1.1 The supply chain of interest

The sutdy of information available to manufacturers and their suppliers for decision making has attracted attention from researchers and practitioners. While the importance of information is well noted (Davis, 1993; Federgruen, 1993), further work needs to be done on what information is useful (Lacity, Willcocks, and Feeny, 1996; Meyer and Zack, 1996) and the value of the information. The wide use of computers and information technologies has made more information available for production management, but not all of the information is as useful as expected (Earl, 1996). Thus, it is important to identify useful information so that it can be utilized properly.

The rest of the paper is organized as follows. Section 2 defines the inventory-production model explicitly. The concept of information level is introduced in terms of available information about the number of unfilled orders. Section 3 briefly discusses the optimal replenishment policies for the inventory-production system with various levels of information. Section 4 presents numerical examples to show which information levels are important, how to identify the important information levels, and the relationship between information levels and system parameters. Section 5 summarizes our results and indicates future research directions.

2. The Inventory-Production System

The inventory-production system of interest in this paper is defined as follows. The workshop manufactures make-to-order products based on customer demands. Customer demands arrive one at a time to the workshop according to a Poisson process with parameter λ. All demands are processed in the workshop by a single machine in batch sizes of one. Production (or processing) times have a common exponential distribution with parameter μ. When the workshop is ready to process a customer order, a call for a unit of raw materials is sent to the warehouse. A unit of raw materials is immediately sent to the workshop and production on that unit begins. The transportation time between the warehouse and the workshop is assumed negligible. Raw materials in the warehouse are replenished from a supplier according to a continuous review replenishment policy. Replenishment leadtimes of raw materials are assumed to be zero so that production in the workshop can occur whenever there are demands. The replenishment policy for raw materials is such that it does not allow raw material shortages, i.e., production will occur whenever there are demands in the workshop.

The cost components of the inventory-production system are as follows: there is a fixed ordering cost K associated with each order from the warehouse to the supplier, regardless of the order size, and the holding cost is C_h per unit of raw materials held in the system per unit time.

According to the above definition, leadtimes at both the supplier and the warehouse are zero. No shortage of raw materials is allowed at the supplier, the warehouse, or the workshop, and all demands will be processed. Thus, the inventory-production system can be decomposed into two subsystems: an *M/M/*1 queue (the workshop) and an inventory system (the warehouse) (Figure 2.1). The workshop is modeled as an *M/M/*1 queue (Cohen, 1982) since no shortage of raw materials is allowed and the queueing process is not influenced by raw material replenishment. Without loss of generality, it is assumed that all demands are processed on a first-come-first-served basis. The warehouse is modeled as an inventory system with zero replenishment leadtimes and demands that come from the workshop when the workshop begins to produce new products.

The status of the *M/M/*1 queue at time t is represented by the number of customers (the number of unfilled demands or the queue length) in the workshop, denoted by $q(t)$. It is assumed that $\rho = \lambda/\mu < 1$ so that the *M/M/*1 queue is ergodic. Let $I(t)$ be the total number of units of raw materials in the inventory-production system at time t, i.e., the number of units of raw materials in the warehouse plus the number of units of raw materials in the workshop, if the workshop is busy. Then the status of the inventory-production system can be represented by $(q(t), I(t))$ at time t.

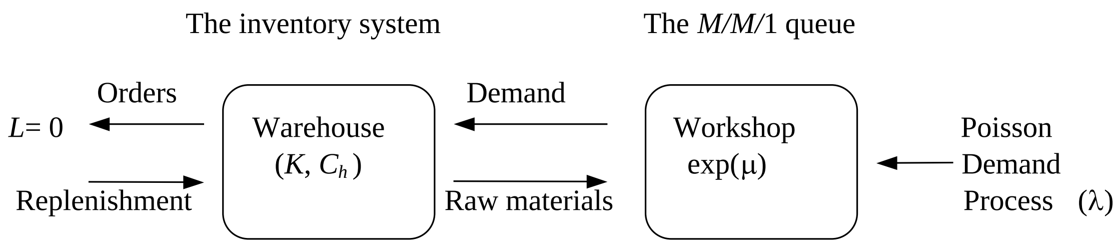

Figure 2.1 The inventory-production system with zero leadtimes

Raw material replenishment in the warehouse is controlled according to a specific replenishment policy. This replenishment policy determines when and how much raw material to order from the outside supplier. In this paper, only replenishment policies based on system status $(q(t), I(t))$ are considered. This implies that both the queue length $q(t)$ and the inventory level $I(t)$ are reviewed continuously. Since leadtimes of raw materials are zero, it makes no sense to order raw materials when $I(t)$ is positive. Thus, a replenishment policy is a function of the number of customers $q(t)$ only and can be represented as a vector $\pi = (\pi(0), \pi(1), \pi(2), ...)$, where $\pi(q)$ is the order size when *the inventory level is zero* and the number of customers is q. At time t, if $I(t) = 0$ and $\pi(q(t))>0$, an order of size $\pi(q(t))$ is issued and filled; 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., $I(t+) = \pi(q(t))$.

The *average total cost per product* is considered as the fundamental performance measurement by which replenishment policies are evaluated. The optimal replenishment policy considered in this paper is the one which minimizes the average total cost per product given the information available on the number of unfilled demands in the system when orders are placed.

While many other types of information regarding production status are useful in inventory control, we shall focus on *the number of unfilled demands in the workshop*, $q(t)$, as the information of potential use in making replenishment decisions since it has a major impact on inventory control in the warehouse. To quantify the benefits of information used in inventory replenishment, the concept of *information level* is introduced. Level l (≥ 0) information implies that when $q(t) \leq l$, the exact queue length is available to the raw materials replenishment decision maker; otherwise, only the fact that the $q(t) > l$ is known. For instance, level –1 means that no information about the queue length is available in making replenishment decisions. Level 0 means that information about whether or not the workshop is busy is available, but the exact queue length is not available. When information of level ∞ is available, the decision maker can find the exact queue length whenever a replenishment decision must be made.

By finding the optimal replenishment policy for level *l* information and its associated average total cost per product for $l = -1, 0, 1, \ldots$, we can determine the value of increasing levels of information. The value of level *l* information is the reduction in average total cost per product which comes about by having level *l* information as opposed to level *l*-1 information. The rest of the paper finds the value of information and the relationship between the optimal replenishment policy, the value of information, and system parameters.

3. The Inventory-Production System with Different Information Levels

In this section, we briefly discuss the optimal replenishment policy and its computation. The discussion is divided into four parts according to the level of information available: 1) *no information*; 2) *busy/idle status information*; 3) *partial information*; 4) *full information*.

<u>3.1 Systems with no information ($l = -1$)</u> When no information about the queue length is available, the order size is independent of the queue length (i.e., the order size is the same for all queue lengths). Since the penalty cost of raw materials shortage is infinite, an order must be issued whenever the inventory level becomes zero. Assume that *n* units of raw materials are ordered whenever the inventory level hits zero. The replenishment policy $\pi = (n, n, ...)$, is an "order up to" policy. Due to its simple structure, explicit results can be obtained for the average total cost per product. It has been proved in HE and Jewkes (1999) that the average total cost per product, $g_{-1}(n)$, when the order size is *n*, is given by:

$$g_{-1}(n) = \frac{K}{n} + \frac{(n+1)C_h}{2\lambda}. \tag{3.1}$$

Then the average total cost per product $g_{-1}(n)$ is minimized at EOQ(λ) (Bartmann and Beckmann, 1992):

$$\mathrm{EOQ}(\lambda) = \arg\min_n \{\frac{K}{n} + \frac{(n+1)C_h}{2\lambda}\}. \tag{3.2}$$

Therefore, the optimal replenishment policy with no information is given by π^*_{-1} = (EOQ(λ), EOQ(λ), ...).

<u>3.2 Workshop busy/idle status information (l=0)</u> In this situation, information about whether the workshop is busy or idle is available, but the exact queue length is not. Since leadtimes are zero, no order is placed when q=0. When q>0, a replenishment decision has to be made without further details about the queue length. Thus, the order size for q>0 is independent of the exact queue length if the queue length is positive. Therefore, the replenishment policies with level 0 information will have the form $\pi = (0, n, n, ...)$, for some positive integer *n*. It can be proved (HE and Jewkes, 1999) that the average total cost per product is given by

$$g_0(n) = \frac{K}{n} + \frac{(n+1)C_h}{2\lambda} - (1-\rho)\frac{C_h}{\lambda} = g_{-1}(n) - (1-\rho)\frac{C_h}{\lambda}. \tag{3.3}$$

The average total cost per product is again minimized at EOQ(λ). Compared to the no information situation, the total average cost per product has decreased by $(1-\rho)C_h/\lambda$. The reason is that no order is placed when the queue length is zero, so that the holding costs during such idle periods are eliminated. It will be shown that this cost savings is the most significant cost savings achievable.

3.3 Partial information ($0<l<\infty$) No explicit results exist for the optimal replenishment policy with partial information. What can be said, though, is that for level l information, the order size may vary according to the queue length up to $q = l$, and that order sizes for queue lengths larger than l will be the same. The common order size will be referred to as the *tail order size*. The reason for this special form is that the exact queue length is not available when $q(t)>l$ so that a common order size must be chosen. Thus, the policies under consideration when information of the level l is available are

$$\prod[l] \equiv \{\pi = (\pi(0), \pi(1), \ldots, \pi(l), \pi(\infty), \pi(\infty), \ldots): \ \pi(q)\geq 0,\ 0\leq q\leq l,\ \pi(\infty)>0\}. \qquad (3.4)$$

No explicit formula has been found for the minimal total cost per product g^*_l. Computationally, however, g^*_l can be found - the first step is to find the optimal policy in $\prod[l]$, and then to use an algorithm to calculate the average total cost per product. The methodology is explained below.

The optimal replenishment policy for a system with level l information can be found using a two stage approach. In the first stage, the order size for a queue length larger than l is fixed. A Markov decision process approach (Puterman, 1994; Tijms, 1990) is applied to find the optimal replenishment policy among all the policies with the same fixed tail order size. In the second stage, the tail order size is enumerated over all its possible values to find the optimal replenishment level l policy. The procedure is as follows. First, the tail order size $\pi(\infty)$ is fixed. Define, for $0\leq i,\ q\leq n$ and $n\geq 1$:

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

In steady state ($\rho=\lambda/\mu<1$), the average total cost per product is defined as $g^*_l = \lim_{n\to\infty} \frac{V^*(i,q,n)}{n}$. Note that g^*_l is independent of the initial state, is finite, and it is *not* the average total cost between decision (or transition) epochs but the average total cost between product completion epochs (i.e., average total cost per product). Define the relative cost function as:

$$h(i, q, n) = V^*(i, q, n) - V^*(1, 1, n), \quad i\geq 0,\ q\geq 0,\ n\geq 1. \qquad (3.5)$$

By Theorem 3.1 in Tijms (1990), the limit of $\{h(i, q, n), n\geq 1\}$ exists and is finite for each state (i, q). Define, for $i\geq 0$ and $q\geq 0$,

$$h(i, q) = \lim_{n\to\infty} h(i, q, n) = \lim_{n\to\infty} [V^*(i, q, n) - V^*(1, 1, n)] \qquad (3.6)$$

$$h^*(q) = \begin{cases} \min_{i\geq 1}\{K + h(i,q)\}, & 1 \leq q \leq l; \\ K + h(\pi(\infty), q)\}, & l < q. \end{cases} \qquad (3.7)$$

The optimal replenishment policy π^*_l is given by

$$\pi_l^*(q) = \begin{cases} \arg\min_{i\geq 1}\{h(i,q)\}, & 1 \leq q \leq l; \\ \pi(\infty), & l < q. \end{cases} \tag{3.8}$$

It is easy to see that when $\lambda<\mu$, the stationary optimal replenishment policy π_l^* exists. This is intuitive since $\lambda<\mu$ ensures a stable *M/M/*1 queue. The functions $h(i, q)$ and $h^*(q)$ are computed using functional equations given in the Appendix.

It can be proved that the optimal order size never exceeds $2\sqrt{K\mu / C_h}+2$ (HE, Jewkes, and Buzacott, 1997). Thus, let the tail order size $\pi(\infty)$ vary from 1 to $2\sqrt{K\mu / C_h}+2$ to find the optimal replenishment policy.

3.4 Full information case ($l=\infty$) When full information is available, the order size varies according to the queue length. When $l = \infty$, the optimal replenishment policy π_∞^* has been obtained (HE, Jewkes, and Buzacott, 1997). The optimal replenishment policy with full information has the structure $\pi_\infty^* = (0, \pi^*(1), \pi^*(2), ..., \pi^*(q^*-1), \text{EOQ}(\mu), \text{EOQ}(\mu), ...)$, where q^* is an integer and EOQ(μ) is defined by equation (3.2). Thus, only level q^* (an integer) information is needed for implementing the optimal replenishment policy with full information.

4. The Value of Information: A Numerical Example

This section compares the optimal replenishment policies for various information levels and their corresponding average total costs per product so as to analyze the value of information. The base cases for comparison purposes are the no information case (i.e. level -1 information) and the busy/idle workshop information case (i.e. level 0 information). The comparisons take several forms: 1) comparison of the optimal average total cost per product for level l information; 2) the relative cost savings obtainable, e.g., $(g^*_{-1} - g^*_\infty)/ g^*_{-1}$; and 3) a comparison of the optimal replenishment policies for level l information.

Example 4.1 Consider an inventory-production system with C_h=0.2, K=10, and μ=1 and demand rates of λ=0.1, 0.4, 0.618, and 0.95. Columns 2 to 4 in Table 4.1 present the average total costs per product for information levels from -1 to 8 for each demand rate.

Table 4.1 The average total cost per product for various information levels

Level	λ=0.1	λ=0.4	λ=0.618	λ=0.95
–1	7.333333	3.416667	2.706311	2.157895
0	5.533333	3.116667	2.582686	2.147368
1	5.528434	3.103620	2.571827	2.144377
2	5.527970	3.100114	2.569436	2.144112
3	5.527926	3.099178	2.568092	2.143818
4	(≈)5.527922	3.099035	2.568029	2.143818
5	(≈)5.527922	3.099035	2.568029	2.143818
6	(≈)5.527922	3.099028	2.568029	2.143818
7	(≈)5.527922	3.099028	2.568029	2.143818
8	(≈)5.527922	3.099026	2.568029	2.143818

Naturally, the average total cost per product decreases when more information is available, but the reductions diminish with higher information levels. For λ=0.10 and λ=0.95, the average total costs per product decline to a limiting value for low information levels. For moderate demand rates, λ=0.40 and λ=0.618, the decline is evident, but continues for higher information levels. To gain more insights, we look at the relative cost savings and the optimal replenishment policies.

Table 4.2 Relative cost savings obtainable for Example 4.1

	λ=0.1	λ=0.4	λ=0.618	λ=0.95
$(g^*_{-1} - g^*_{\infty})/ g^*_{-1}$	24.62%	9.30%	5.11%	0.65%

First, the relative cost savings obtainable, $(g^*_{-1} - g^*_{\infty})/ g^*_{-1}$ can be generated from Table 4.1. Table 4.2 and other computational experience indicate that the relative cost savings obtainable with full information decreases with respect to the traffic intensity. There is a straightforward explanation for this behavior: if the workshop is empty (q = 0), an order can be delayed until the next demand arrives. This will have the long term effect of reducing inventory holding costs during idle periods. If the workshop has few outstanding orders, a small replenishment order will be placed, thus avoiding holding costs. Without any information about the queue length, these cost savings cannot be taken advantage of and the impact will be large when the utilization rate of the system is small. As the traffic intensity is increased, the impact on the potential cost savings is reduced as the idle periods are shorter and less frequent. Computational experience has shown that as the ordering cost or holding cost increases, the benefits of the lower levels of information (especially 0 and 1) are sustained at higher utilization rates, as one might expect. Next, we look at the structure of the optimal replenishment policies.

λ = 0.1. For a demand rate of 0.1, the optimal replenishment policies with information levels 0, 1, 5, 10, and full are given in Figure 4.1.

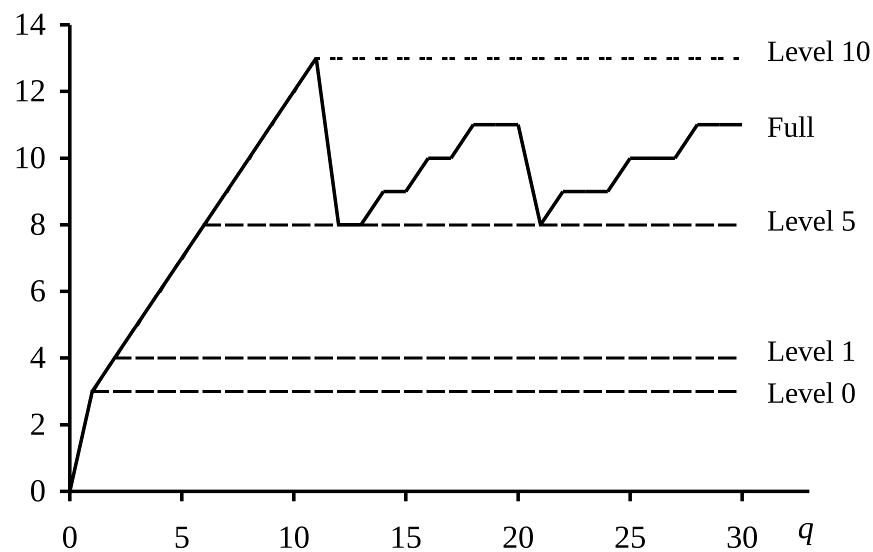

Figure 4.1 The optimal replenishment policies when λ =0.1

Figure 4.1 shows that the optimal replenishment policies for different information levels are dramatically different, though as the information level increases, the policy converges to that of full information. Reference to Table 4.1 indicates that despite these differences in policies, the corresponding average total costs per product converges to the minimum average total cost per product quickly. This is due to the low traffic intensity - the likelihood of having a queue longer than three is very small. It is worth to point out that for the optimal replenishment policy π_l^*, its order size $\pi_l^*(q)$ does not have to be a monotone function of the queue length q (e.g., the full information case in Figure 4.1).

λ=0.4. The optimal replenishment policy with full information is reached when the information level is 8. The optimal replenishment policies with various information levels are shown in Table 4.3. Each row represents an optimal replenishment policy with a given level of information.

For cells with bold numbers in Table 4.3, the corresponding order size can be adjusted individually in order to minimize the average total cost per product. Compared to λ=0.1, the fluctuations in the optimal replenishment policies are reduced, and information levels up to 8 produce cost savings. Since the traffic intensity is moderate, being able to adjust the order sizes for low queue lengths decreases the holding costs, even when the queue length is as large as 8.

Table 4.3 The optimal replenishment policies when λ=0.4

Level	q=0	q=1	q=2	q=3	q=4	q=5	q=6	q=7	q≥8
–1	6	6	6	6	6	6	6	6	6
0	**0**	6	6	6	6	6	6	6	6
1	**0**	**6**	8	8	8	8	8	8	8
2	**0**	**6**	**7**	8	8	8	8	8	8
3	**0**	**6**	**7**	**8**	9	9	9	9	9
4	**0**	**6**	**7**	**8**	**9**	10	10	10	10
5	**0**	**6**	**7**	**8**	**9**	**10**	10	10	10
6	**0**	**6**	**7**	**8**	**9**	**10**	**10**	11	11
7	**0**	**6**	**7**	**8**	**9**	**10**	**10**	**11**	11
(≥)8	**0**	**6**	**7**	**8**	**9**	**10**	**10**	**11**	**11**

λ=0.618. For this arrival rate, the busy periods are longer as are queue lengths. Thus, the impact of avoiding holding inventory during idle and low demand periods is reduced. The optimal policy with full information takes effect for smaller queue lengths (see also Table 4.1). Table 4.5 shows that the optimal order size is 10 for a queue length larger than 4 and hence only the information levels from 0 to 4 produce cost reductions.

Table 4.5 The optimal replenishment policies when λ=0.618

level	q=0	q=1	q=2	q=3	q=4	q≥5
–1	8	8	8	8	8	8
0	**0**	8	8	8	8	8
1	**0**	**7**	9	9	9	9
2	**0**	**7**	**8**	9	9	9
3	**0**	**7**	**8**	**9**	10	10
(≥)4	**0**	**7**	**8**	**9**	**9**	10

λ=0.95. Compared to λ=0.1, the optimal inventory policy is much simpler. As are shown in Tables 4.1 and 4.6, only information up to level 3 is valuable. Once again this is due to the combined effect of shorter idle periods and a faster convergence to the optimal replenishment policy with full information.

Table 4.6 The optimal replenishment policies when λ=0.95

level	q=0	q=1	q=2	q=3	q≥4
–1	10	10	10	10	10
0	**0**	10	10	10	10
1	**0**	**8**	10	10	10
2	**0**	**9**	**9**	10	10
(≥)3	**0**	**8**	**9**	**9**	10

From Figure 4.1 and Tables 4.4 to 4.6, it is clear that once the optimal replenishment policy with full information is known, one can determine which information levels are valuable, i.e., information up to the level $q^* = \min\{q: \pi_\infty^*(n)=\text{EOQ}(\mu), \text{ for } n>q\}$. However, the structure of the optimal replenishment policy does not show the magnitude of the cost savings at each level. The relative cost savings percentage function can be used to determine the cost savings.

5. Conclusions

In this paper, we introduced the concept of information level in an integrated inventory-production system. The value of information about the queue length present at the workshop when a replenishment order is made was analyzed and discussed. The main conclusions are given as follows.

1. In general, queue length information is useful in reducing the average total cost per product, especially when the traffic intensity of the system is low or moderate.
2. Lower level information, especially level 0 information, is more important than higher level information. This is particularly true when the holding cost is high, as level 0 information will allow the decision maker to postpone an order during idle periods until the next demand arrives. Higher level information is useful when the traffic intensity is moderate and the ratio of the ordering cost per order and the holding cost per unit per unit time is moderate as well.
3. The value of information on workshop status is closely related to the system parameters. In general, as the utilization rate increases, holding costs or ordering cost decreases, the value of information decreases.

The implications of these results for the decision maker are: (1) information on workshop status should be made available; (2) the most valuable information is whether the workshop is busy or idle at the decision epoch. While more information on the queue length is useful, its marginal value is less. Higher levels of information can result in substantial savings but these savings need to be traded off against the cost of maintaining and supplying the information. The methodology in this paper enables the benefits of more information to be evaluated.

A generalization of this paper would be to consider models with non-zero leadtimes. For instance, in HE, Jewkes, and Buzacott (1999), an inventory-production model with exponential leadtimes was investigated. In principle, the main conclusions (1 to 3) obtained for the zero leadtime case hold for cases with an exponential leadtime, but the analysis for the later is much more complicated.

Appendix. Functional Equations for the Optimal Replenishment Policy

For brevity, the derivation is given for the full information case. The basis for the functional equation is the principle of optimality (Puterman, 1994).

$$\begin{aligned}
& g_\infty^* + h(1,1) = \frac{C_h}{\mu} + \omega h^*(1) + \sum_{j=1}^{\infty} \omega(1-\omega)^j h^*(j); \\
q \geq 2, \quad & g_\infty^* + h(1,q) = \frac{C_h}{\mu} + \sum_{j=0}^{\infty} \omega(1-\omega)^j h^*(q-1+j); \\
i \geq 2, \quad & g_\infty^* + h(i,1) = \frac{iC_h}{\mu} + \omega[\frac{(i-1)C_h}{\lambda} + h(i-1,1)] + \sum_{j=1}^{\infty} \omega(1-\omega)^j h(i-1,j); \\
i,q \geq 2, \quad & g_\infty^* + h(i,q) = \frac{iC_h}{\mu} + \sum_{j=0}^{\infty} \omega(1-\omega)^j h(i-1,q-1+j).
\end{aligned} \tag{A.1}$$

The above method is used to compute the best replenishment policy with a fixed tail order size $\pi(\infty)$. To find the optimal replenishment policy for inventory-production systems with information of the level *l*, one needs to search through all possible values of $\pi(\infty)$. The search range of $\pi(\infty)$ is bounded and is from 1 to $2\sqrt{K\mu / C_h}$ +2.

References

Axsäter, S., Continuous Review Policies for Multi-level Inventory Systems with Stochastic Demand, *Logistics of Production and Inventory*, Handbooks in Operations Research and Management Science, North-Holland, **Vol 4**, 195-198, 1993.

Bartmann, D. and Beckmann, M.J., *Inventory Control*: *Models and Methods*, Springer-Verlag, 1992.

Cohen, J.W., *The Single Server Queue*, North-Holland series in Applied Mathematics and Mechanics, North-Holland, 1982.

Davis, T., Effective supply chain management, *The Solan Management Review/ Summer*, 35-46, 1993.

Earl, M.J., The risks of outsourcing IT, *The Sloan Management Review*, **Vol 37**, No.3, 26-32, 1996.

Federgruen, A., Centralized planning models for multi-echelon inventory systems under uncertainty, *Logistics of Production and Inventory*, Handbooks in Operations Research and Management Science, North-Holland, **Vol 4**, 133-174, 1993.

HE, Qi-Ming and E.M. Jewkes, Performance measures of a make-to-order inventory-production system, *IIE Transaction*, 1999 (accepted).

HE, Qi-Ming and E.M. Jewkes, and J. Buzacott, Inventory control policies for an inventory-production system, 1997 (submitted for publication).

HE, Qi-Ming and E.M. Jewkes, and J. Buzacott, The value of information used in a make-to-order inventory-production system, 1999 (submitted to *IIE Transaction* for publication).

Lacity, M.C., Willcocks, L.P. and Feeny, D.F., The value of selective IT sourcing, *The Sloan Management Review*, **Vol 37**, No.3, 13-25, 1996.

Meyer, M.H. and Zack, M.H., The design and development of information products, *The Sloan Management Review*, **Vol 37**, No.3, 43-59, 1996.

Puterman, M.L., *Markov Decision Processes*, John Wiley & Sons, Inc., 1994.

Tijms, H., *Stochastic Modelling and Analysis: A Computational Approach*, John Wiley & Sons, Inc., 1990.

Veatch, M.H. and Wein, L.M., Optimal control of a two-station tandem production/inventory system, *Operations Research*, **Vol 42**, 337-350, 1994