IIE Transactions

Taylor & Francis
Taylor & Francis Group

ISSN: 0740-817X (Print) 1545-8830 (Online) Journal homepage: https://www.tandfonline.com/loi/uiie20

Integrating advance order information in make-to-stock production systems

FIKRI KARAESMEN , JOHN A. BUZACOTT & YVES DALLERY

To cite this article: FIKRI KARAESMEN , JOHN A. BUZACOTT & YVES DALLERY (2002) Integrating advance order information in make-to-stock production systems, IIE Transactions, 34:8, 649-662, DOI: 10.1080/07408170208928902

To link to this article: https://doi.org/10.1080/07408170208928902

Published online: 17 Apr 2007.

IIE Transactions (2002) **34**, 649–662

Integrating advance order information in make-to-stock production systems

FIKRI KARAESMEN[1], JOHN A. BUZACOTT[2] and YVES DALLERY[1]

[1]*Laboratoire Productique Logistique, Ecole Centrale Paris, Grande Voie des Vignes 92295 Châtenay-Malabry Cedex, France*
E-mail: fikri@pl.ecp.fr or dallery@pl.ecp.fr
[2]*Schulich School of Business, York University, North York, Ontario, Canada*
E-mail: jbuzacot@schulich.yorku.ca

Received January 2001 and accepted January 2002

Increased cooperation between supply chain partners and information technology are enabling the availability of advance order information for contract suppliers. Control mechanisms that take into account this availability are necessary in order to achieve the potential improvements in performance. We investigate the structure of optimal control policies for a discrete-time make-to-stock queue with advance order information. Since the optimal policy does not have a simple structure, we then propose a heuristic policy which is an extension of the base stock system that incorporates advance order information through a release lead time parameter. In order to quantify the benefits due to advance order information, we investigate the performance of the proposed mechanism and benchmark it against against the optimal control policy.

1. Introduction

In most manufacturing and distribution environments, safety stocks constitute the principal measure in order to cope with the uncertainties in demand as well as in production and transportation. An alternative measure of protection against randomness would be to reduce demand uncertainty through increased information sharing between the partners of a supply chain. In practice, some information on future demands is usually available for each partner under varying forms ranging from forecasts to supply contracts. When used effectively, advance demand information serves to reduce the uncertainty in future demands thereby enabling better inventory and service performance.

In this paper, we consider a particular type of advance demand information that we also refer to as *advance order information*: early commitments on orders from customers (which may themselves be downstream partners of the same supply chain). As outsourcing of production becomes more prevalent in several industries, suppliers are facing a new class of problems where external demand is still uncertain in the long-run but is not completely random in the short-term because of early commitments which are specified by contractual agreements.

In order to motivate the issues that will be investigated, let us consider two firms, A (an Original Equipment Manufacturer - OEM) and B (a Contract Manufacturer-CM) that are partners of the same supply chain (Firm B is a contract supplier to Firm A). The focus of our investigation will be on the production planning problem that arises for Firm B. In an outsourcing/contracting type relationship, collaboration is expected to improve supply chain performance and is strategically desirable. In such a case, the OEM may be willing to share early order information with the CM in return for better service levels or price discounts.

There are several ways through which OEMs could transmit early order information to its contract manufacturer. In certain industries, the OEM may itself be receiving advance orders from its external customers. This is prevalent in electronic retailing where end-clients purchase items over the internet and reveal advance demand information by accepting future delivery dates. In the auto industry, Firm B may be a contract supplier that delivers components to the final assembly line of the OEM (Firm A), an auto manufacturer. In this case, Firm A optimizes and freezes its assembly schedule several days in advance and can, in principle, share this information with Firm B.

Despite the contracted relationship, the advance information exchange does not always take place smoothly in practice. The reason seems to be that both parties have difficulty in evaluating the benefits that the exchange would bring. In this context, neither party has a strong willingness to push forward for the sharing of this information. Without an assessment of the expected benefits and costs due to advance order information, it is difficult to achieve cooperation at an operational level.

Placed within this context, our objective is two-fold: first, investigating how the supplier (Firm B) should use advance order information in an effective manner and second, evaluating the value of this information. A key question that we would like to answer is the following one: what are the potential benefits for Firm B if Firm A ordered earlier? In order to respond to these questions, we first investigate the structure of optimal production control policies for the supplier and then propose a simple control mechanism that makes use of advance order information by integrating a release timing mechanism with a base stock type inventory control policy. In order to gain insights into the performance of the proposed mechanism and into the potential benefits of advance order information, we provide a detailed analysis of a single-stage discrete-time production/inventory system (i.e., a make-to-stock queue).

An important first step in analyzing how advance demand information improves inventory related costs is understanding how safety stocks relate to controlled (safety) lead times. Lambrecht *et al.* (1984) describe how safety stocks induce safety lead times in multi-stage inventory systems. Yano (1987) considers the problem of optimizing planned release times. Gong *et al.* (1994) show that the problem of planning optimal release times is the equivalent of the problem of setting optimal safety stock levels in serial systems. In a single-period make-to-stock environment, Milgrom and Roberts (1988) show that advance demand information is a substitute for safety stocks in the single-period newsvendor model.

In the operations management literature, integrating advance demand information in inventory control mechanisms has been studied from several different perspectives in recent years. Most of this research focuses on models with exogeneous replenishment-type models which do not model limited production capacity. Hariharan and Zipkin (1995) study the effects of incorporating advance demand information in inventory control policies. Their results are highly intuitive: early demand information reduces inventory control costs and this reduction becomes more important as future uncertainty is resolved earlier. Axsater and Rosling (1993, 1994), show that by incorporating a timing mechanism into classical multi-stage inventory control systems (such as installation stock and echelon stock control), one can obtain a new class of control policies that actually dominate the performance of classical systems. Finally, Gallego and Özer (2000, 2001) have recently investigated optimal replenishment policies for single-echelon and multi-echelon inventory systems with advance demand information. The above papers present interesting and important results on the integration of benefits of advance demand information for uncapacitated systems. While our analysis of the capacitated system manifests certain parallels to the above papers, limited production capacity and variability play a central role in our analysis and in the results obtained.

There are fewer papers that discuss or analyze the integration of advance demand information in a manufacturing setting where limited production capacity cannot be ignored. A general framework that comprises an advance order information component in this setting is the multi-stage PAC (Production Authorization Card) mechanism described in detail by Buzacott and Shanthikumar (1992, 1993). The PAC framework takes into account a particular kind of advance demand information by the inclusion of a mechanism for delaying the release of orders.

Buzacott and Shanthikumar (1994) analyze an analytical model of a single-stage make-to-stock queueing system with advance demand information coming from a downstream MRP system. Their analysis provides many interesting insights for a particular release policy. It is, however, difficult to extend this framework to investigate whether other, more effective, control policies can be designed or not. In fact, the model of this paper can be viewed as the discrete-time version of the model in Buzacott and Shanthikumar (1994). The discrete-time framework enables us not only to provide numerical performance measures for a given mechanism but also to address *optimality* issues within a dynamic control framework. The optimal control policy is partially characterized and the optimality gap (when a simple release policy is used instead of the optimal one) is quantified.

In other related work on make-to-stock queues with demand forecast evolution, Güllü (1996) studies a (capacitated) production/inventory system evolving in discrete time within a model that integrates the forecast evolution process. He shows that the forecast information can be used to reduce safety stocks and overall inventory related costs. In a recent paper, Toktay and Wein (2001) also study a discrete-time production/inventory system along with a similar demand forecast process. Their asymptotic analysis reveals that advance information in terms of demands forecasts enables a reduction in safety stocks that depends both on the average load on the system and on demand or forecast variability. Our model is simpler in its representation of demand, production capacity and the information evolution structure than these two papers. In particular, we focus exclusively on the uncertainty in the timing of demands (and neglect the uncertainty in quantities). This enables us to obtain several explicit results and to investigate the important issue of the value of advance order information as a function of the duration of the information horizon. We also confirm, through an exact analysis, some of the earlier results of Güllü (1996) and Toktay and Wein (2001) for the model considered here.

Finally, a relevant line of research investigates issues of how lead times should be quoted to customers. For instance, Duenyas and Hopp (1995) study the problem of optimizing the quoted due-dates to clients for a manufacturing firm. These issues are not considered in our framework since due-dates are externally set by contracts (or by the planning system).

The principal contribution of the paper is the comparison of a simple policy that integrates advance order information with the overall optimal policy that uses detailed advance order information. We show, numerically, that the simple policy can be surprisingly effective and conjecture that it is optimal in certain cases that can be identified. The analysis reveals the importance of setting planned release times and underlines the strong interaction between planned release times and order information availability.

The paper is structured as follows: Section 2 describes the problem, introduces an analytical model and presents results on the structure of the optimal policy with advance order information. Section 3 presents the proposed control policy followed by its analysis and optimization. Section 4 reports the numerical results and the observations. Our conclusions are presented in Section 5.

2. The model and its analysis

In this section, we describe a model of a contract supplier that satisfies the demands of an equipment manufacturer that places orders in advance. The general modeling framework is described in Section 2.1 and is further specified in Section 2.2 which includes the main results on the structure of optimal policies. Section 2.3 focuses on a benchmark case of no advance demand information.

2.1. *The modeling framework*

We consider a single-stage production system, consisting of a manufacturing stage and a finished goods inventory. The system operates in a make-to-stock setting and the output of the manufacturing stage is placed in the finished goods inventory. In particular, our focus is on a contracted component supplier to an equipment manufacturer which may provide advance order information. Since information about future demands is available, the time at which a demand is claimed (due-date of the order) may be later than the arrival time of the order and production may be initiated by firm orders rather than in anticipation of expected future demands. More precisely, we assume that the supplier receives all orders exactly H periods in advance of their due-date. Under this assumption, at time t, the supplier has exact information on the timing of all demands to satisfy between the time window $(t, t+H]$. We refer to H as the *horizon of visibility* since the supplier has perfect visibility of demand within this time window.

In order to simplify the analysis, we assume that setup costs and setup times are negligible so that production lot sizing issues can be avoided. We assume unit order arrivals and assume that orders that are not fulfilled at their required due-dates are fully backlogged. The production control problem is to find a part release policy (i.e., to decide when to release parts to the manufacturing stage) in order to minimize the sum of the (average) holding and penalty costs. Further analysis requires the specification of a model for the order arrival and the production processes. This will be done in the next section.

2.2. *A discrete-time make-to-stock queue with advance order information*

Consider the following single stage make-to-stock system: time is divided into equal length intervals and $X(t)$ $(t = 1, 2, \ldots)$ denotes the finished inventory level (where negative values of $X(t)$ represent backlogs) at the beginning of interval t. In each time interval, the probability of an order arrival is q (note that orders arrive in single units) and the probability of a production completion is p (whenever the facility is allowed to produce).

Since order inter-arrival times and processing times are geometrically distributed, the basic model is a discrete-time make-to-stock queue with geometric processing times. Even though this model is too simple to represent classical periodic review inventory dynamics (where demand and production processes can be more general), it can be viewed as the approximation of a (Markovian) continuous-time make-to-stock queue. In this regard, it has an important feature in that, it is well-suited to an optimal control type formulation (unlike the continuous-time make-to-stock queue) when advance order information is integrated. Moreover, just like the M/M/1 queue in continuous time, the Geo/Geo/1 queue captures the most significant effects that arise due to congestion (see Pujolle *et al.* (1986) for a discussion of the similarities). We exploit these features in the following development.

Based on our previous assumption, advance order information is in the form of firm orders placed exactly H periods in advance. The consequence of this assumption is that at time t, all demands that have to be satisfied in the periods $t+1, t+2, \ldots, t+H$ are known with certainty. This information will be summarized by the vector $\mathbf{D}(t) = (D_1(t), D_2(t), \ldots, D_H(t))$ where an element $D_h(t)$ of the vector gives the number of demands that will claim finished items at the end of the period $t+h-1$. The system state will then be $(X(t), \mathbf{D}(t))$. We denote by x a particular realization of the random variable $X(t)$ and by $\mathbf{d}$, a particular realization of the random vector $\mathbf{D}(t)$.

The assumption of H periods in advance order information combined with geometrically distributed order inter-arrival times significantly simplifies the structure of the vector $\mathbf{D}(t)$. First, under these assumptions, all components of the vector $\mathbf{D}(t)$ are either zero or one. Second, the demand inter-arrival time process which is a translation by H units of the order process and is itself geometrically distributed. The order queue in front of the manufacturing facility is then precisely the Geo/Geo/1 queue in discrete-time and can be viewed as the limiting version of the M/M/1 queue as the unit time interval $\Delta t \to 0$ and the transition probabilities $p \to (1 - e^{-\mu\Delta t})$ and $q \to (1 - e^{-\lambda\Delta t})$.

Let $y(t)$ be a Bernoulli random variable corresponding to an order arrival at time t where

$$y(t) = \begin{cases} 1 & \text{with probability } q, \\ 0 & \text{otherwise.} \end{cases}$$

The evolution of the advance demand information vector then becomes:

$$\mathbf{D}(t+1) = (D_2(t), D_3(t), \dots, D_H(t), y(t)).$$

Note that this is a particular case of advance demand information evolution. More general advance demand information evolution structures can be represented by using a vector summarizing future demand or forecast information over the planning horizon (see Güllü (1996), Gallego and Özer (2001) and Toktay and Wein (2001) for more general representations). The advance demand information vector can also be viewed as a vector representing the state of the demand process as in Song and Zipkin (1993).

In order to describe the inventory level dynamics of the system, let us describe the sequence of events that take place within a period: the controller observes the inventory level and the information vector $(X(t), \mathbf{D}(t))$ at the beginning of the period and decides whether to produce or not. If the "produce" decision is taken, a production completion takes place with probability p at the end of the period. As for the demands, exactly $D_1(t)$ demands have to be fulfilled at the end of the period. At the same time, a new demand order arrives to the system with probability q and

$$D_H(t+1) = \begin{cases} 1 & \text{with probability } q, \\ 0 & \text{otherwise.} \end{cases}$$

Let $a(t)$ denote the production decision where $a(t) = 1(0)$ corresponds to production (non-production) at time t. $m(t)$ is Bernoulli random variable corresponding to a production completion (with probability p). The evolution of the stock level, $X(t)$, can be expressed by:

$$X(t+1) = X(t) - D_1(t) + a(t)m(t).$$

For optimization purposes, let us consider the classical formulation, with a linear holding cost, h, and a linear backorder cost, b. We would then like to obtain the control policy that minimizes the infinite horizon discounted costs (with discount factor α, $0 \le \alpha < 1$):

$$\lim_{T\to\infty} E\left[\sum_{t=0}^{T} \alpha^t c(X(t), \mathbf{D}(t))\right], \tag{1}$$

where $c(x, \mathbf{d}) = hx$ when $x \ge 0$ and $c(x, \mathbf{d}) = -bx$ when $x < 0$.

The objective is then to find the optimal production control policy (a proxy for the optimal release policy) in order to minimize (1). To facilitate the ensuing development, let us denote by $\mathbf{e}_k$ a vector with one in the kth entry and zeroes elsewhere, and define the "shifted" vector $\mathbf{d}^+ = (d_2, d_3, \dots, d_{H-1}, 0)$. The functional equation for this problem is:

$$\begin{aligned} V(x, \mathbf{d}) = c(x) + \alpha \min\{ & pqV(x - d_1 + 1, \mathbf{d}^+ + \mathbf{e}_H) \\ & + p(1-q)V(x - d_1 + 1, \mathbf{d}^+) + (1-p) \\ & \times qV(x - d_1, \mathbf{d}^+ + \mathbf{e}_H) + (1-p)(1-q) \\ & \times V(x - d_1, \mathbf{d}^+), qV(x - d_1, \mathbf{d}^+ + \mathbf{e}_H) \\ & + (1-q)V(x - d_1, \mathbf{d}^+)\} \end{aligned}$$

where $V(\ )$ is the value function of the dynamic program.

We will now show that, the optimal production policy is of threshold type for each demand vector $\mathbf{d}$.

Proposition 1. *For each demand vector,* $\mathbf{d}$, $V(x, \mathbf{d})$ *is convex in* x *and the optimal control policy is of threshold type with a threshold level* $S_{\mathbf{d}}$ *corresponding to the vector* $\mathbf{d}$.

Proof. See Appendix A1.

Proposition 1 establishes that optimal policies have to be of the threshold type in inventory levels for each advance information vector. It is also reasonable that these threshold levels exhibit certain monotonicity properties as a function of the advance demand information vector. In particular, more demands in the foreseeable future should intuitively imply increased production in advance through higher optimal threshold levels. In order to strengthen Proposition 1 in this direction, let us define a partial order of the demand vectors. We will say that $\mathbf{d} \succeq \mathbf{d}'$ if $d_i \ge d'_i$, $\forall i$, $i = 1, 2, \dots, H$ or if $\mathbf{d} = \mathbf{d}' + d'_k\mathbf{e}_j - d'_k\mathbf{e}_k$ $(2 \le j \le k \le H)$. This order is slightly weaker than a lexicographic order but it has the property that it is preserved under the shift operation: if $\mathbf{d} \succeq \mathbf{d}'$ then $\mathbf{d}^+ \succeq (\mathbf{d}')^+$.

Proposition 2. *If it is optimal to produce in state* $(x, \mathbf{d})$, *then it is optimal to produce at state* $(x, \mathbf{d}')$ *where* $\mathbf{d} \succeq \mathbf{d}'$. *Consequently, if* $\mathbf{d} \succeq \mathbf{d}'$ *then the optimal threshold levels are such that* $S_{\mathbf{d}} \ge S_{\mathbf{d}'}$.

Proof. See Appendix A2.

Propositions 1 and 2 shed light onto the structure of the optimal policy. The optimal policy has a threshold-type structure but the optimal threshold levels fluctuate as a function of actual advance demand information. More (and closer) demands in the future imply higher threshold levels in correspondence with the intuition that future peaks may be absorbed with building inventories in advance. Indeed, the optimal policy seems to confirm most basic intuitive properties but the exact structure is difficult to express in a few parameters. In particular, the optimal threshold level depends not only on the number of demands within the visible horizon but also on the timing of future demands.

2.3. *The case without advance order information*

This section focuses on a special case of the model where there is no advance order information ($H = 0$). This case constitutes a benchmark in order to quantify the gains that can be obtained through advance order information.

Since $H = 0$, orders correspond to demands to be fulfilled immediately. Hence, $y(t)$ now represents the random variable representing the demand arrival (with probability q). The system then evolves according to the equation:

$$X(t+1) = X(t) + a(t)m(t) - y(t). \tag{2}$$

By Proposition 1, a base stock (i.e., single threshold) policy is optimal. The facility should produce when the inventory level is below the base stock level S and should stop otherwise.

The only step that remains is the computation of the optimal base stock level. This computation can be done using standard methods and is presented in Appendix B. It turns out that, the optimal base stock level is:

$$S^* = \left\lfloor \ln\left[\frac{h}{(h+b)}\frac{(1-\beta)}{\kappa}\right] \Big/ \ln[\beta] \right\rfloor.$$

3. A base stock policy with a release lead time parameter

In Section 2 some results on the structure of the optimal control policy under advance order information were presented. In an outsourcing context, this structure sheds certain light onto how the contract manufacturer (Firm B) should coordinate its production decisions when advance order information is available. There are, however, several difficulties in following the exact optimal policy. For instance, production decisions depend on very detailed order information comprising not only the number of future orders but also their precise timings. It is of interest to investigate whether a simpler control policy can be sufficiently effective.

In this section, we describe and analyze a dynamic control policy which is an extension of the well-known base stock control policy that integrates demand lead time information in a simple and natural way. Section 3.1 introduces the policy and Section 3.2 presents the analysis.

3.1. *Description of the policy*

For a single-stage system, a base stock policy is completely described by a single parameter S (which corresponds to the produce-up-to-level). In order to incorporate advance demand information in this policy, we associate with the production stage an additional parameter L. The control parameter L is called the *release lead time* since it will be used to regulate the timing of material release into the manufacturing stage. We start with a general description of the policy before describing its implications for the fixed demand lead time model of Section 2. Consider the instant when the nth demand d_n arrives to the system. Under a standard base stock policy, this arrival would trigger the release of a part to the manufacturing stage at its time of arrival t_n. The objective in the case where advance demand information is available is then to drive the release mechanism of parts using demand lead time information rather than arrival times (or due-dates) of demands.

In order to understand how the release lead time parameter L regulates the release timing, let us assume first that a demand arrives with a sufficient demand lead time, i.e., $l_n \geq L$. Releasing a part too early in the manufacturing stage leads to excess finished goods inventory (we are assuming that early deliveries to customers are not authorized) and is costly. It is then logical to delay this release by $l_n - L$ units of time (so that the part is released exactly L units of time in advance). Alternatively, the demand lead time of an arrival may not be sufficient with respect to the parameter L (i.e., $l_n < L$). In this case, the part should be released immediately into the manufacturing stage without a further delay. Within this logic, the timing of part releases into manufacturing is entirely driven by the respective values of L and l_n. The requested release time of a part to the stage is, hence, $t_n + \max(0, l_n - L)$. Note that, except in special cases, demand n is satisfied by the existing (or in-process) inventories and not by the corresponding release. In general, the parameter L regulates the inventory replenishment process but some randomness remains and safety stocks are still needed. Note that the release lead time L clearly has to be related to the production lead time, W. The exact relationship between L and W, however, is not immediately obvious since W is a random variable that is endogenous to the system.

Figure 1 depicts the base stock mechanism with a lead time parameter. In the figure, MF1 is the manufacturing stage, and P_1 and P_2 are the buffers corresponding respectively to raw materials (where infinite supply is assumed) and finished items. D_1 and D_2 correspond to demands that pull items respectively from P_1 and P_2. The round nodes that regulate the input into D_1 and D_2 are delay nodes through which demand information is transmitted. Note also the synchronization of buffers P_1 with D_1 and P_2 with D_2. A part is delivered to a customer if and only if there is a demand in D_2 and a finished part available in P_2. Similarly, a raw part is released into the manufacturing stage if and only if there is a demand in D_1 and a raw part available in P_1 (note that this is a special case; since unlimited raw part supply is assumed, a raw

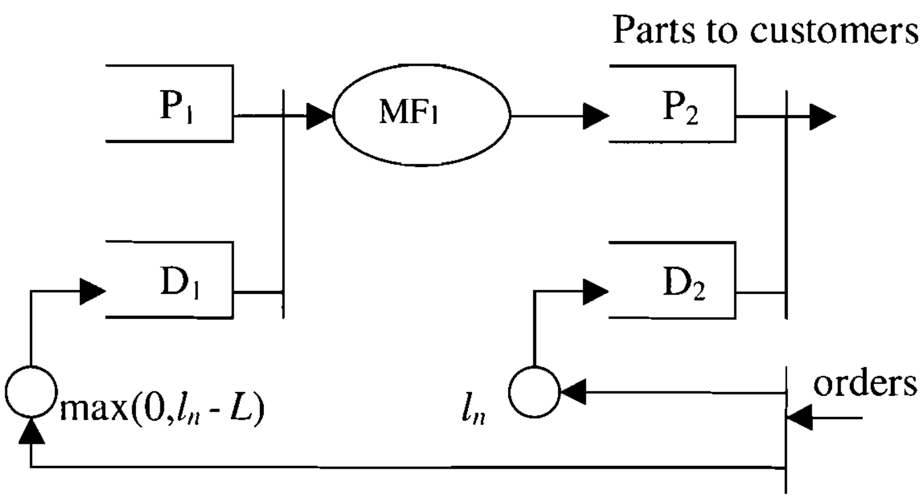

Fig. 1. The single stage base stock system with a release lead time parameter.

part is always released as soon as there is an arrival to D_1).

Initially, there are S finished items in P_2, and both D_1 and D_2 are empty (P_1 is assumed to have an infinite supply of raw material). Consider the arrival of the nth demand to the system with a demand lead time of l_n. This information is transmitted to buffer D_2 after a delay of l_n. The delay node before D_2 hence ensures that a finished part will be claimed exactly at the due-date of the arrival. The delay node before D_1 drives the release of parts. This node ensures that part release is regulated through the parameter L. A part will be released immediately if $l_n \leq L$ and will be delayed by the required amount if $l_n > L$.

In the case where there is no advance demand information (i.e., $l_n = 0$, $\forall n$), the delay nodes of Fig. 1 become inactive and a new part is released at the instant of each demand arrival. The two-parameter policy thus reduces to the classical base stock policy with parameter S.

The focus of Section 2 was in the special case where all demand lead times are constant (i.e., $l_n = H$, $\forall n$). In this case, the above policy reduces to the MRP interpretation of the PAC system as described in detail by Buzacott and Shanthikumar (1993). In the PAC system arriving demand signals are delayed by an amount of H ($H \geq 0$). Recall that in our system, demand signals are delayed by an amount of $\max\{H - L, 0\}$. Apart from its generality (under more general advance demand information structures), the description above emphasizes the distinction between the two parameters L and H. L is an internally set parameter whereas H is the visibility available and H acts a constraint on L.

Policies of similar nature have also been proposed by Hariharan and Zipkin (1995) for uncapacitated supply systems. In fact, Hariharan and Zipkin show that for uncapacitated supply systems these simple policies can sometimes be optimal. Our results in Section 2 indicate that, with endogenous supply lead times, the optimal policies are more complicated.

3.2. *Analysis of the base stock policy with a release lead time parameter*

The two-parameter-per-stage policy introduced in Section 3.1 is conceptually fairly simple. It would be interesting to verify whether this simple way of integrating advance demand information in the coordination of the system leads to significant cost savings. This section focuses on the optimization and performance evaluation of the two-parameter policy for the Geometric/Geometric/1 make-to-stock queue.

Let us denote by $o(t)$, the orders that arrive in period t (which is, in general, different than $d(t)$, the demands that have their due-dates in period t). Once again, our focus is on the case where the demand lead time is a constant H ($l_n = H$, $\forall n$). Letting $X(t)$ the inventory position at the beginning of period t, the dynamics of the system can be described by:

$$X(t+1) = X(t) + a(t)m(t) - o(t-H),$$

where $a(t)$ is the indicator variable and

$$a(t) = \begin{cases} 1 & \text{if production takes place at time } t, \\ 0 & \text{otherwise.} \end{cases}$$

Note however that $a(t)$ is indirectly controlled by the part release policy and $a(t) = 1$ if there are items waiting to be processed in the manufacturing queue.

The objective is then to select the values of the parameters L and S (where $S, L \geq 0$) in order to optimize the inventory-related cost of the system. This will be done in two steps. First, let us assume that the horizon of visibility, H, is infinite. This implies that all orders are known in advance with certainty. Under this assumption, let us denote the corresponding release lead time parameter by L_∞ where $0 \leq L_\infty \leq \infty$ (note that under the assumption of infinite visibility, there is no restriction on the planned release time). In the first step, we focus on the optimization of the pair (S, L_∞). It will be seen later that, once the optimal value of L_∞ is obtained, the passage to any finite horizon of visibility – a more realistic assumption – is not difficult.

Let us denote by $R(t)$ and $J(t)$ respectively, the number of items produced (i.e., delivered to buffer P_2 of Fig. 1) up to time t, and the number of demands that have due-dates (number of demands transmitted to buffer D_2 of Fig. 1) before t. We can then define the underlying (unrestricted) queueing-type process $N^*(t) = R(t) - J(t)$. Note that $N^*(t)$ can also be expressed in terms of the inventory process $X(t)$ ($N^*(t) = S - X(t)$). It should be remarked that, due to overshoots, $X(t)$ can exceed the base stock level S and the process $N^*(t)$ can take negative values.

The ensuing analysis then follows along the lines of Buzacott and Shanthikumar (1994) for a continuous-time make-to-stock queue. Let us denote by $N(t) = S - X(t)$ the difference between cumulative production and demand for a system without advance information. The stationary distribution of $N(t)$ is presented in Appendix B.

Note then that:

$$N^*(t) = N(t - L_\infty) - J(t - L_\infty, t),$$

where $J(t - L_\infty, t)$ denotes the number of departures from the queue during the interval $(t - L_\infty, t]$.

We can then write:

$$P\{N^*(t) = n\} = \sum_{k=n}^{\infty} P\{J(t - L_\infty, t) = k - n | N(t - L_\infty) = k\} \times P\{N(t - L_\infty) = k\},$$

when $n \geq 1$.

Because processing times are geometrically distributed:

$$P\{J(t - L_\infty, t) = k - n | N(t - L_\infty) = k\} = \binom{L_\infty}{k-n} p^{k-n}(1-p)^{L_\infty - (k-n)},$$

and

$$P\{N^*(t) = n\} = \sum_{k=n}^{\infty} \binom{L_\infty}{k-n} p^{k-n}(1-p)^{L_\infty - (k-n)} \times P\{N(t - L_\infty) = k\}.$$

Denoting by $\pi^*(n)$ and $\pi(n)$ the respective stationary distributions of $N^*(t)$, and $N(t)$, we get:

$$\pi^*(n) = \sum_{k=n}^{\infty} \binom{L_\infty}{k-n} p^{k-n}(1-p)^{L_\infty-(k-n)} \pi(k) \text{ for } n \geq 1.$$

Replacing the stationary probability $\pi(k)$ in right-hand side of the above equation using the expression in (A2) in Appendix B, we obtain after some manipulation:

$$\pi^*(n) = \left(\frac{1-p}{1-q}\right)^{L_\infty} \pi(n) \text{ for } n \geq 1.$$

Let us denote the multiplicative factor on the right-hand side of the above expression by γ^{L_∞} ($\gamma \equiv (1-p)/(1-q)$). In other words, as in Buzacott and Shanthikumar (1994), the advance release mechanism causes the stationary distribution of the underlying queue to be rescaled by the factor γ^{L_∞} (when $n \geq 1$).

For given values of S and L_∞, the expected backorder level $E[B]$ and the expected inventory level $E[I]$ can be obtained as follows:

$$E[B] = \kappa\gamma^{L_\infty} \frac{\beta^{S+1}}{1-\beta}, \tag{3}$$

and

$$E[I] = S + E[B] + qL_\infty - \frac{q(1-q)}{p-q}. \tag{4}$$

Let us denote by $C(S, L_\infty)$ the expected optimal cost for given values of S and L_∞. Using (3) and (4), $C(S, L_\infty)$ can be expressed as:

$$C(S, L_\infty) = h\left(S + qL_\infty - \frac{q(1-q)}{p-q}\right) + (h+b)\kappa\gamma^{L_\infty} \frac{\beta^{S+1}}{1-\beta}.$$

It turns out then that for a fixed L_∞, the optimal base stock level is given by:

$$S^*(L_\infty) = \left\lfloor \left[\ln\left(\frac{h}{h+b} \frac{(1-\beta)}{\kappa}\right) \Big/ \ln\beta + \frac{L_\infty \ln\gamma}{\ln\beta} \right] \right\rfloor, \tag{5}$$

provided that $S^*(L_\infty) \geq 0$. The second term of (5) is negative and characterizes the reduction of the base stock level as a function of the release lead time, L_∞. It is interesting to note that this reduction is approximately linear (if the integrality correction is ignored). In other words, base stock levels can be reduced at a linear rate as a function of L_∞ as long as they stay non-negative.

In addition, minimization of the cost function $C(S^*(L_\infty), L_\infty)$ as a function of L_∞ yields that the optimal value of L_∞ is given by:

$$L_\infty^* = \left\lceil \left[\ln\left(\frac{h}{h+b}\right) \Big/ \ln\left(\frac{1-p}{1-q}\right) \right] \right\rceil. \tag{6}$$

Equation (6) expresses, in a concise manner, the optimal release lead time. Since the corresponding value of S^* is zero, if the horizon of visibility were infinite, it would be optimal to operate the system in a make-to-order mode releasing parts exactly L_∞^* periods in advance of their due-date. Another important remark is that visibility beyond L_∞^* is unnecessary since releasing parts earlier does not improve performance. L_∞^* hence defines the *desired horizon of visibility*, H^*.

Note that L_∞^* (and hence H^*) is increasing in the average load, q/p, and in the arrival/processing time variability for a fixed value of q/p. In other words, in order to optimize performance, more visibility is necessary as the average load and the variability increase. A second interesting consequence of (6) is that the optimal release lead time (and the desired horizon of visibility) cannot be zero. In other words, if advance order information is available, it is always optimal to use it for early release of parts.

Finally, let us address the relationship between L^*_∞ and W, the random variable representing the actual production lead time. If we let $F_W(\ )$ denote the cumulative distribution function of the production lead time, then L^*_∞ is such that $F_W(L^*_\infty) = b/(b+h)$. In other words, the desired horizon of visibility depends on the tail distribution of the production lead time (as well as the holding/backorder costs).

Up to this point, we have assumed an infinite horizon of visibility. In order to extend the analysis to any finite horizon H, it remains to relate the unconstrained value of the release lead time, L_∞ to its constrained counterpart L. Note that, when visibility is unconstrained, the choice of L_∞ causes the release of a part exactly L_∞ periods in advance of its due-date. When the visibility H is greater than L_∞, the same release policy can be implemented for the constrained problem by setting $L = L_\infty$. On the other hand, when H is less than L_∞, the unconstrained release policy has to be modified taking into account the visibility constraint. Nevertheless, the optimal unconstrained solution can easily be transformed into an optimal solution for the constrained problem as the next proposition states:

Proposition 3. *For a finite horizon of visibility H, the optimal values of the control parameters (S, L) are given by:*

1. $S = 0$ *and* $L = L^*_\infty$ *(if* $H \geq L^*_\infty$*)*;
2. $S = S^*(H)$ *and* $L = L^*_\infty$ *(if* $H < L^*_\infty$*)*.

Proof. The proof follows by checking the properties of the expected cost function (5). In Part 1, the visibility H is greater than the desired release lead time L^*_∞. It can be verified that the cost function $C(S^*(L), L)$ is increasing in L when $L > L^*_\infty$ and is decreasing in L when $L \leq L^*_\infty$. It is, hence, optimal to set the parameter as $L = L^*_\infty$. The corresponding optimal value of S is given by: $S = S^*(L^*_\infty) = 0$.

In part 2, the visibility H is less than the desired release time. However, since $C(S^*(L), L)$ is decreasing in L when $L \leq L^*_\infty$, the optimal choice of the release lead time parameter is any L that is greater than or equal to H so that the release takes place exactly H periods in advance. Even though the optimal value of the release lead time parameter is not unique, in the statement of the proof, we select $L = L^*_\infty$ to underline the fact that the optimal L is determined by the desired release lead time. The corresponding optimal S parameter is then given by $S = S^*(H)$. ■

Remark. An alternative statement of Proposition 3 would be the following. For a finite horizon of visibility H, the optimal values of the control parameters (S, L) are given by $L = \min(H, L^*_\infty)$ and $S = S^*(L)$.

Proposition 3 implies that there are two different regimes for the determination of the optimal decisions: If the visibility is not sufficient ($H < L^*_\infty$) it is optimal not to delay the release and to compensate the remaining variability by holding a positive base stock level. Alternatively, if the visibility is sufficient ($H \geq L^*_\infty$), the optimal base stock level is zero (the system operates in a make-to-order fashion) and delaying the release of parts reduces unnecessary holding costs. In this case, the parameter L should be used to regulate the release such that the release takes place exactly L^*_∞ periods in advance.

In conclusion, the desired release lead time L^*_∞ subsumes all the information that is required in order to set the optimal release parameter L. Further demand visibility than L^*_∞ does not have any additional value since it will not be used. With a different interpretation, L^*_∞ determines the optimal planning (or forecast) horizon for the class of (S, L) policies. When these policies are used, information beyond a horizon of L^*_∞ is not used for planning purposes.

4. Numerical examples: a comparison of the optimal policy with the proposed policy

The purpose of the numerical examples that follow is two-fold. First, we would like to quantify the cost reduction due to the use of advance order information. Second, we would like to evaluate the performance of (S, L) policies with respect to the optimal policy.

In the discrete-time make-to-stock queue that we consider here, there are two inherent factors affecting system performance. The first (and stronger) factor is the system load summarized by the ratio q/p and the second one is the overall variability summarized by the value of q (or p) for a fixed load q/p. In order to capture both effects we vary q/p for two different values of q corresponding to a system with highly regular production times ($p = 0.9$) and moderately variable production times ($p = 0.5$).

As for the cost parameters, we take $h = 1$, and $b = 10$ and 100. Finally, we vary the horizon of visibility, H, between 0 and 9. For each of these values, we computed the average cost of the optimal policy using dynamic programming and the cost of the optimal (S, L) policy using the results of the previous section. Tables 1–4 report these results.

An overview of the Tables 1–4 leads to two important observations. First, integrating advance order information decreases costs in a significant way in general. Second, (S, L) policies are surprisingly effective. We will focus on each of the two issues and their implications in more detail.

4.1. *Benefits of advance order information and the performance of the (S, L) policy*

Figure 2 compares the relative percentage cost reduction due to using the optimal policy ((C(no. info)− C(opt))/

Table 1. Cost of the optimal policy and the percentage suboptimality of the optimal (S, L) policy (reported in parentheses) if different than the optimal policy, for $h = 1$, $p = 0.5$, $b = 10$

	Visibility (H)										
q/p	0	1	2	3	4	5	6	7	8	9	L^*_∞
0.1	0.956	0.556	0.316	0.214	0.184	0.184	0.181 (1.66)	0.179 (2.79)	0.179 (2.79)	0.179 (2.79)	4
0.2	1.05	1.028	0.739	0.499	0.411	0.406	0.396 (2.53)	0.387 (4.91)	0.385 (5.45)	0.384 (5.73)	5
0.3	1.343	1.202	1.180	0.901	0.716	0.668	0.668	0.644 (3.73)	0.635 (5.20)	0.631 (5.86)	5
0.4	1.833	1.583	1.440	1.425	1.162	1.026	1.016	0.985 (3.15)	0.954 (6.50)	0.949 (7.06)	6
0.5	2.167	2.111	1.972	1.815	1.793	1.586	1.474	1.474	1.411 (4.46)	1.380 (6.81)	6
0.6	2.859	2.765	2.632	2.623	2.439	2.370	2.284	2.146	2.133	2.062 (3.44)	8
0.7	3.886	3.837	3.725	3.719	3.577	3.536	3.445	3.365	3.329	3.207	10
0.8	5.918	5.897	5.814	5.812	5.715	5.696	5.620	5.583	5.529	5.474	14
0.9	11.950	11.944	11.900	11.899	11.850	11.845	11.800	11.791	11.751	11.737	26

Table 2. Cost of the optimal policy and the percentage suboptimality of the optimal (S, L) policy (reported in parentheses) if different than the optimal policy, for $h = 1$, $p = 0.9$, $b = 10$

	Visibility (H)										
q/p	0	1	2	3	4	5	6	7	8	9	L^*_∞
0.1	0.911	0.111	0.0923	0.0923	0.0923	0.0923	0.0923	0.0923	0.0923	0.0923	2
0.2	0.850	0.250	0.189	0.189	0.189	0.189	0.189	0.189	0.189	0.189	2
0.3	0.829	0.429	0.292	0.292	0.292	0.292	0.292	0.292	0.292	0.292	2
0.4	0.867	0.667	0.408	0.408	0.408	0.408	0.408	0.406 (0.4)	0.406 (0.4)	0.406 (0.4)	2
0.5	1.000	1.000	0.550	0.550	0.550	0.550	0.544 (1.10)	0.540 (1.85)	0.540 (1.85)	0.540 (1.85)	2
0.6	1.300	1.065	0.749	0.749	0.749	0.742 (0.94)	0.716 (4.61)	0.713 (5.05)	0.713 (5.05)	0.711 (5.34)	2
0.7	1.477	1.252	1.090	1.090	1.060 (2.83)	0.993 (9.77)	0.972 (12.14)	0.972 (12.14)	0.969 (12.49)	0.964 (13.07)	2
0.8	1.886	1.857	1.769	1.601	1.521 (5.26)	1.495 (7.09)	1.495 (7.09)	1.481 (8.10)	1.448 (10.57)	1.417 (12.99)	3
0.9	3.237	3.151	3.079	3.019	2.973	2.948 (0.84)	2.944 (0.99)	2.907 (2.27)	2.834 (4.90)	2.781 (6.90)	4

Table 3. Cost of the optimal policy and the percentage suboptimality of the optimal (S, L) policy (reported in parentheses) if different than the optimal policy, for $h = 1$, $p = 0.5$, $b = 100$

	Visibility (H)										
q/p	0	1	2	3	4	5	6	7	8	9	L^*_∞
0.1	1.456	1.240	1.150	1.126	0.913	0.575	0.421	0.364	0.357	0.356 (0.28)	8
0.2	2.056	2.031	1.754	1.508	1.416	1.409	1.043	0.846	0.781	0.781	8
0.3	2.781	2.460	2.332	2.319	2.013	1.843	1.805	1.582	1.363	1.296	9
0.4	3.308	3.193	3.182	2.889	2.780	2.751	2.469	2.368	2.321	2.051	10
0.5	4.185	4.124	3.997	3.831	3.804	3.608	3.489	3.478	3.235	3.157	12
0.6	5.483	5.345	5.332	5.154	5.081	5.002	4.859	4.842	4.670	4.593	14
0.7	7.494	7.380	7.373	7.233	7.191	7.101	7.020	6.986	6.862	6.848	18
0.8	11.402	11.336	11.320	11.233	11.228	11.134	11.112	11.040	11.000	10.950	26
0.9	22.999	22.998	22.949	22.944	22.899	22.890	22.850	22.836	22.802	22.783	49

$C(\text{no. info}) \times 100)$ and using the optimal (S, L) policy $((C(\text{no. info}) - C(S^*, L^*))/C(\text{no. info}) \times 100)$ for different values of the system load when advance order information is abundant ($H = 9$). The figure shows that the relative cost reduction can be very significant when the load is low but diminishes rapidly as the load increases. The figure also shows that (S, L) policies attain almost all of the potential cost reduction due to advance order information. Both of these observations are valid throughout Tables 1–4. The cost reduction appears in a sharper manner when the production times are less variable or when the backorder costs are higher but the qualitative behavior is the same. Similarly, even though the relative suboptimality of (S, L) policies increases for less variable production times or increased backorder costs, a significant part of the cost reduction with respect to the no-information base case can be obtained through these policies.

Table 4. Cost of the optimal policy and the percentage suboptimality of the optimal (S, L) policy (reported in parentheses) if different than the optimal policy, for $h = 1$, $p = 0.9$, $b = 100$

	Visibility (H)										
q/p	0	1	2	3	4	5	6	7	8	9	L_∞^*
0.1	1.011	1.001	0.202	0.182	0.182	0.182	0.182	0.182	0.182	0.182	3
0.2	1.300	1.036	0.463	0.373	0.373	0.373	0.373	0.372	0.372 (0.27)	0.372 (0.27)	3
0.3	1.741	1.135	0.820	0.578	0.578	0.578	0.576 (0.35)	0.573 (0.87)	0.573 (0.87)	0.573 (0.87)	3
0.4	1.742	1.354	1.345	0.818	0.818	0.818	0.800 (2.25)	0.796 (2.76)	0.796 (2.76)	0.795 (2.89)	3
0.5	1.909	1.818	1.517	1.134	1.134	1.120 (1.25)	1.063 (6.68)	1.061 (6.68)	1.058 (7.18)	1.056 (7.18)	3
0.6	2.465	2.108	1.820	1.646	1.626	1.462 (11.22)	1.415 (14.91)	1.415 (14.91)	1.406 (15.65)	1.386 (17.32)	4
0.7	2.727	2.610	2.602	2.352	2.122	2.037 (4.17)	2.037 (4.17)	1.988 (6.74)	1.905 (11.391)	1.871 (13.41)	4
0.8	3.634	3.542	3.498	3.461	3.286	3.137	3.039 (3.22)	3.001 (4.53)	2.994 (4.78)	2.905 (7.99)	5
0.9	6.214	6.127	6.051	5.988	5.939	5.905	5.889	5.875	5.787	5.710 (1.35)	8

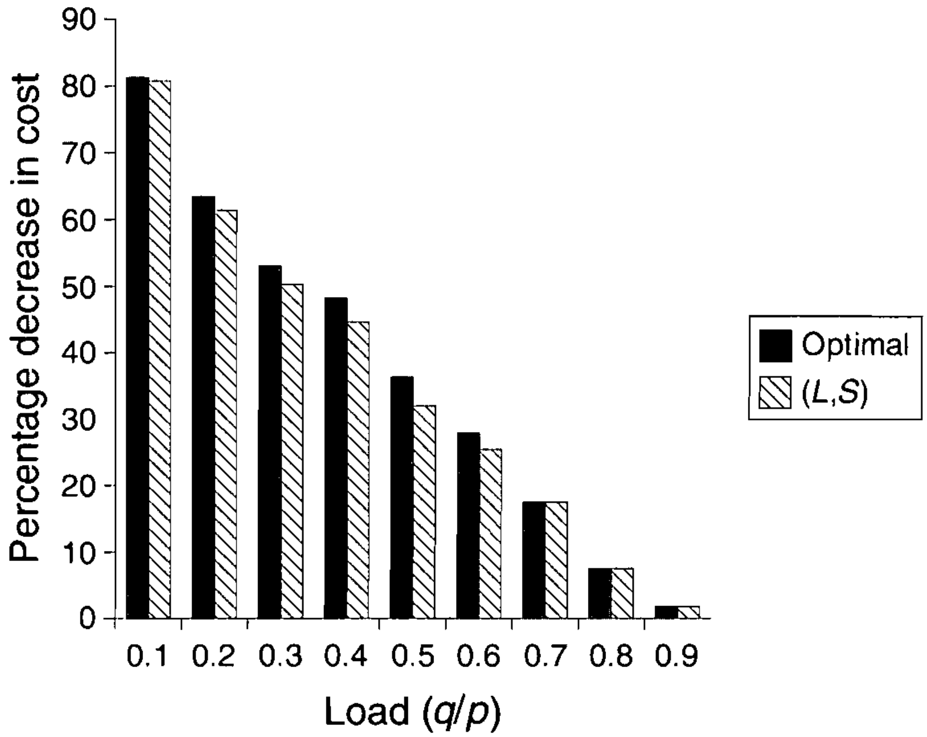

Fig. 2. Percentage cost reduction due to advance order information through the optimal policy and the (S, L) policy ($p = 0.5$, $h = 1$, $b = 10$).

We will hence focus on this class, in order to characterize the potential cost reductions due to advance order information. Let $C(S^*(0), 0)$ denote the optimal cost per unit time of a system that does not use advance order information and $C(0, L_\infty^*)$ be the optimal cost of a system that has sufficient advance demand information so that it can set its lead time parameter using the desired release lead time L_∞^* (as in Proposition 3). One possible measure of the benefits of using advance demand information is then: $[C(S^*(0), 0) - C(0, L_\infty^*)]/C(S^*(0), 0)$, the relative cost improvement obtained when advance demand information is incorporated through the desired release lead time L_∞^*.

4.2. *The effect of the system load*

Figure 3 displays the relative improvement in cost (in terms of the measure described above) as a function of the system load q/p for three different values of p (with $h = 1$ and $b = 10$). Once again, the significant impact of the system load is evident from the figure. The second-order effect of variability is also distinguishable. As the

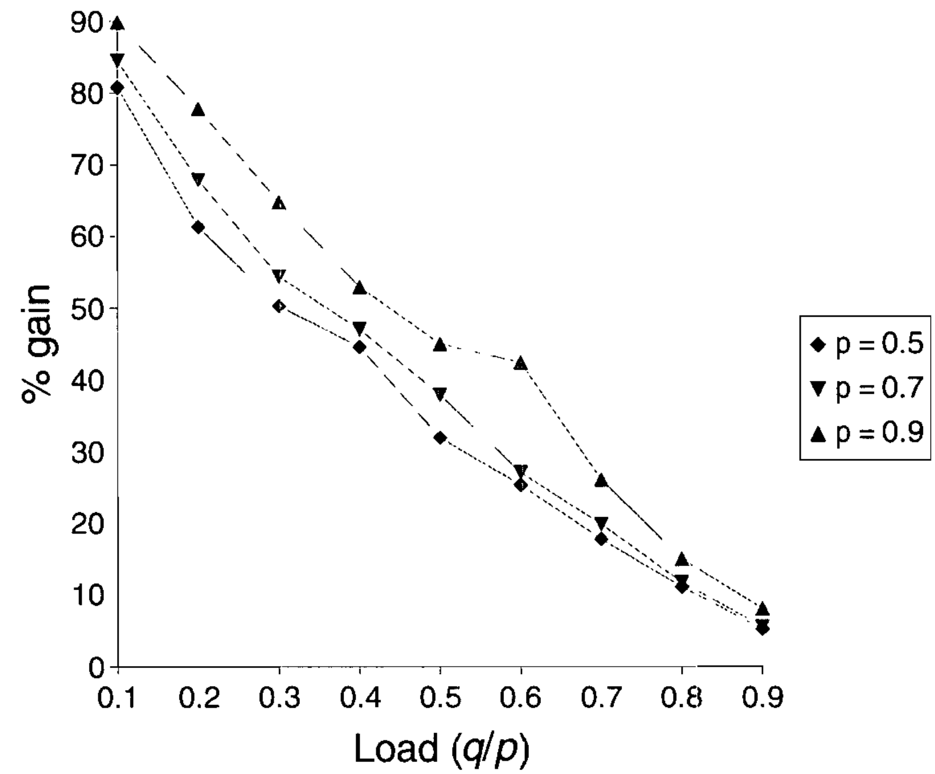

Fig. 3. Percentage relative benefit of advance order information.

system becomes less variable (p increases for a given q/p), the relative cost improvement increases.

There is also a more subtle effect of the system load that we can observe from Figs. 2 and 3. We remark that, in both figures, the relative benefit of advance order information disappears as the system load increases. In addition, Tables 1–4 indicate that at the other extreme (light system loads), even though the relative benefits are significant, the absolute benefits are small. In fact, when the trade-off between absolute versus relative gains is considered, advance order information is especially desirable for systems that are not in extreme load situations (say in the regime $0.4 \leq q/p \leq 0.9$). Figure 2 shows an additional observation in the same general direction. In situations of extreme load (light or heavy), the average performance is not very sensitive to the precise control policy and the optimal (S, L) policy is almost as good as the overall optimal policy. Alternatively, in medium load, the best (S, L) policy can be somewhat inferior to the overall optimal policy. It seems that in this medium-load range, detailed demand state information has more relative value.

4.3. *The rate of cost savings as a function of visibility*

A final observation from Tables 1–4 is that the marginal cost reductions are diminishing as the horizon of visibility increases. In fact, we can observe that the marginal cost reduction approaches zero as expected. It seems likely that there are finite forecast horizons beyond which advance information does not provide any additional cost reduction. Although our numerical results cannot provide an answer to this important issue, we can investigate the existence of approximate forecast horizons beyond which advance order information brings little additional value. Tables 1–4 follow a similar pattern for the optimal costs as a function of visibility. After an initial steep decrease, the optimal cost function becomes rather flat after a threshold visibility level. In fact, it can be observed that there are two distinct regimes for the cost reduction. In the first regime (for example visibility between zero and four for the first four rows in Table 1), costs decrease rapidly as advance information becomes available earlier until reaching a threshold horizon. In the second regime (visibility between five and nine for the first four rows in Table 1), further visibility is of little additional value.

4.4. *Characterizing approximate planning horizons*

An analytical characterization of exact or approximate forecast horizons is difficult when optimal policies are considered. On the other hand, from the analysis in the previous section, we have seen that, when (S, L) policies are considered, the desired release lead time L^*_∞ determines the optimal planning horizon. In addition, Tables 1–4 indicate that (S, L) policies coincide frequently with the optimal policies. More precisely, it seems that (S, L) policies are optimal until a threshold horizon of visibility.

In fact, a more careful investigation of Tables 1–4 reveals that (S, L) policies are optimal whenever the horizon of visibility is shorter than or equal to L^*_∞. It is highly likely that this observation reflects a general property that is not restricted to the examples considered above. The conjecture below states this property.

Conjecture 1. *If $H \leq L^*_\infty$, the overall optimal policy is an (S, L) policy whose parameters are: $S = S^*(H)$ and $L = L^*_\infty$.*

5. Conclusion

Advances in information technology are significantly facilitating data transmission between the partners of a supply chain. In parallel, long-term contracting relationships between various partners generates a relatively cooperative environment. This underlines two fundamental issues: how should the additional information be used in order to improve performance and what is the potential benefit that can be expected. We have investigated these issues in the context of advance order information. In regard to the the first issue, we presented a production control policy that incorporates advance order information within a base stock mechanism. In regards to the second issue, we analyzed the potential benefits due to integration of advance order information in a single-stage model. The analysis provides qualitative insights into the benefits that contract suppliers can expect through downstream order information. Our results indicate that advance order information can be expected to improve performance significantly if production capacity is sufficient. In general, it is likely that a joint investment in information technology and in additional capacity is necessary to reap the full benefits of increased information sharing.

The control mechanism proposed is simple and natural which constitutes a significant advantage for its implementation. Interestingly, even though it is a genuine dynamic control rule, the mechanism follows an MRP-type logic that combines safety stocks and planned release times. It should be emphasized however that, in the dynamic vision, planned lead release times are parameters that have to be optimized unlike in most MRP implementations.

The analysis of the model also yields other interesting results that should be investigated under more general conditions. For instance, the simple MRP-type logic is surprisingly effective in a single-stage setting if the parameters are chosen optimally. Indeed, the parameter selection issue is critical and non-trivial in general. The analysis reveals properties that are of interest in this sense as well. In particular, there are critical planning horizons (which may be used to determine release time parameters) that can be obtained as functions of production lead times. Longer and more variable production lead times necessitate longer planning horizons and the benefits that can be expected are strongly related to the effective visibility and the critical planning horizon.

In order to obtain explicit analytical and quantitative results, we have employed a simple model. It is likely, however, that most qualitative results obtained in this paper continue to hold under more general assumptions. A recent evidence in this direction is the work of Toktay and Wein (2001) which considers a more general framework encompassing forecast evolutions but reaches similar conclusions on the benefits of advance demand information through heavy traffic approximations.

There seem to be several interesting directions of future research. First of all, the findings of the single-stage model should be investigated under more general conditions. Second, the multi-stage case poses several challenges. For instance, it is known that base stock policies are not optimal for capacitated systems and that work-in-process inventories may have to be bounded. This motivates further investigation of Generalized Kanban (or PAC) type mechanisms (Buzacott and Shanthikumar, (1993); Dallery and Liberopoulos, 2000). Finally, parameter optimization problems are critical and should have a major impact on the implementation level.

Acknowledgements

The authors would like to thank Beril Toktay and George Liberopoulos for many helpful discussions on the topic and for their comments on previous versions of the paper.

References

Axsater, S. and Rosling, K. (1993) Installation vs. echelon stock policies for multi-level inventory control. *Management Science*, **39**, 1274–1280.

Axsater, S. and Rosling, K. (1994) Multi-level production-inventory control: material requirements planning or reorder point policies. *European Journal of Operational Research*, **75**, 405–412.

Buzacott, J.A. and Shanthikumar, J.G. (1992) A general approach for coordinating production in multiple-cell manufacturing systems. *Production and Operations Management*, **1**, 34–52.

Buzacott, J.A. and Shanthikumar, J.G. (1993) *Stochastic Models of Manufacturing Systems*, Prentice-Hall, Englewood Cliffs, NJ, 1993.

Buzacott, J.A. and Shanthikumar, J.G. (1994) Safety stock versus safety time in MRP controlled production systems. *Management Science*, **40**, 1678–1689.

Dallery, Y. and Liberopoulos, G. (2000) Extended kanban control system: combining kanban and base stock. *IIE Transactions*, **32**, 369–386.

Duenyas, I. and Hopp, W.J. (1995) Quoting customer lead times. *Management Science*, **41**, 43–57.

Gallego, G. and Özer, Ö. (2000) Optimal replenishment policies for multi-echelon inventory problems under advance order information. *Working paper*, Columbia University, 2000.

Gallego, G. and Özer, Ö. (2001) Integrating replenishment decisions with advance order information. *Management Science*, **47**, 1344–1360.

Gong, L., de Kok, T. and Ding, J. (1994) Optimal leadtimes planning in a serial production system. *Management Science*, **40**, 629–632.
Güllü, R. (1996) On the value of information in dynamic production/inventory problems under forecast evolution. *Naval Research Logistics*, **43**, 289–303.
Hariharan, R. and Zipkin, P. (1995) Customer-order information, leadtimes and inventories. *Management Science*, **41**, 1599–1607.
Lambrecht, M.R., Muckstadt, J.A. and Luyten, R. (1984) Protective stocks in multi-stage manufacturing systems. *International Journal of Production Research*, **22**, 1001–1025.
Milgrom, P. and Roberts, J. (1988) Communication and inventory as substitutes in organizing production. *Scandinavian Journal of Economics*, **90**, 275–289.
Pujolle, G., Claude, J.P. and Seret, D. (1986) A discrete queueing system with a product form solution, in *Computer Networking and Performance Evaluation*, Hasegawa, T., Takagi, H. and Takahashi, Y. (eds), Elsevier Science Publishers, New York, pp. 139–147.
Song, J.-S. and Zipkin, P. (1993) Inventory control in a fluctuating demand environment. *Operations Research*, **41**, 351–370.
Toktay, B.L. and Wein, L.M. (2001) Analysis of a forecasting-production-inventory system with stationary demand. *Management Science*, **47**, 1268–1281.
Weber, R.R. and Stidham, S. (1987) Optimal control of service rates in networks of queues. *Advances in Applied Probability*, **19**, 202–218.
Yano, C. (1987) Setting planned leadtimes in serial production systems with tardiness costs. *Management Science*, **33**, 95–106.

Appendices

Appendix A: Proofs of Propositions 1 and 2

In order to prove Propositions 1 and 2, we use standard induction-based arguments of stochastic dynamic programming. As frequently is the case in dynamic programming problems, for structural results, it is easier to argue inductively through the discounted formulation (1). Even though the formulation of the problem and the proofs below are based on the discounted version of the problem, it is important to note that, for the make-to-stock queue, the structure under a discounted formulation will be preserved under average cost optimization (see for example Weber and Stidham (1987).)

A1: Proof of Proposition 1 Let $\Delta(x,\mathbf{d}) = V(x,\mathbf{d}) - V(x-1,\mathbf{d})$. In order to show that $\Delta(x+1,\mathbf{d}) \geq \Delta(x,\mathbf{d})$, we argue that the dynamic programming operator preserves convexity.

A base stock policy is ensured if the difference:

$$\begin{aligned}&[pqV(x-d_1+1,\mathbf{d}^+ + \mathbf{e}_H) + p(1-q)V(x-d_1+1,\mathbf{d}^+)\\&+(1-p)qV(x-d_1,\mathbf{d}^+ + \mathbf{e}_H) + (1-p)(1-q)V(x-d_1,\mathbf{d}^+)]\\&-[qV(x-d_1,\mathbf{d}^+ + \mathbf{e}_H) + (1-q)V(x-d_1,\mathbf{d}^+)],\end{aligned}$$

is increasing in x.

The above difference can be expressed as:

$$pq\Delta(x+1-d_1,\mathbf{d}^+ + \mathbf{e}_H) + p(1-q)\Delta(x+1-d_1,\mathbf{d}^+),$$

which is increasing in x as long as $\Delta(x,\mathbf{d})$ is increasing in x for a fixed $\mathbf{d}$.

It remains then to check the preservation of convexity when the dynamic programming operator is applied on a function that initially satisfies the condition: $\Delta(x,\mathbf{d})$ is increasing in x for a fixed $\mathbf{d}$. Let us take three states: $\{V(x+1,\mathbf{d}), V(x,\mathbf{d}), V(x-1,\mathbf{d})\}$, with associated controls $\{a(x+1,\mathbf{d}), a(x,\mathbf{d}), a(x-1,\mathbf{d})\}$. Due to monotonicity (induction assumption), the controls $(a(x+1,\mathbf{d}), a(x,\mathbf{d}), a(x-1,\mathbf{d}))$ belong to the set $\{(0,0,0),(1,1,1),(0,1,1),(0,0,1)\}$. We can then distinguish three cases:

Case 1: $(a(x+1,\mathbf{d}),\ a(x,\mathbf{d}),\ a(x-1,\mathbf{d})) \in \{(0,0,0),(1,1,1)\}$.

The difference in this case is:

$$\begin{aligned}\Delta(x+1,\mathbf{d}) - \Delta(x,\mathbf{d}) &= c(x+1) - 2c(x) + c(x-1)\\&\quad + pq(\Delta(x+1-d_1,\mathbf{d}^+ + \mathbf{e}_H)\\&\quad - \Delta(x-d_1,\mathbf{d}^+ + \mathbf{e}_H))\\&\quad + p(1-q)(\Delta(x+1-d_1,\mathbf{d}^+)\\&\quad - \Delta(x-d_1,\mathbf{d}^+)),\end{aligned}$$

which is positive by the assumption.

Case 2: $(a(x+1,\mathbf{d}), a(x,\mathbf{d}), a(x-1,\mathbf{d})) = \{(0,1,1)\}$.

First note that

$$\begin{aligned}\Delta(x+1,\mathbf{d}) &\geq pq\Delta(x-d_1+1,\mathbf{d}^+ + \mathbf{e}_H)\\&\quad + p(1-q)\Delta(x-d_1+1,\mathbf{d}^+)\\&\quad + (1-p)q\Delta(x+1-d_1,\mathbf{d}^+ + \mathbf{e}_H)\\&\quad + (1-p)(1-q)\Delta(x-d_1+1,\mathbf{d}^+),\end{aligned}$$

which in turn implies that:

$$\begin{aligned}\Delta(x+1,\mathbf{d}) - \Delta(x,\mathbf{d}) &\geq c(x+1) - 2c(x) + c(x-1)\\&\quad + pq(\Delta(x+1-d_1,\mathbf{d}^+ + \mathbf{e}_H)\\&\quad - \Delta(x-d_1,\mathbf{d}^+ + \mathbf{e}_H))\\&\quad + p(1-q)(\Delta(x+1-d_1,\mathbf{d}^+)\\&\quad - \Delta(x-d_1,\mathbf{d}^+)),\end{aligned}$$

where the right-hand side is non-negative by the assumption.

Case 3: $(a(x+1,\mathbf{d}), a(x,\mathbf{d}), a(x-1,\mathbf{d})) = \{(0,0,1)\}$.

As in case 2

$$\begin{aligned}\Delta(x,\mathbf{d}) &\leq pq\Delta(x-d_1+1,\mathbf{d}^+ + \mathbf{e}_H)\\&\quad + p(1-q)\Delta(x-d_1+1,\mathbf{d}^+)\\&\quad + (1-p)q\Delta(x-d_1,\mathbf{d}^+ + \mathbf{e}_H)\\&\quad + (1-p)(1-q)\Delta(x-d_1,\mathbf{d}^+),\end{aligned}$$

which in turn implies that:

$$\begin{aligned}\Delta(x+1,\mathbf{d}) - \Delta(x,\mathbf{d}) &\geq c(x+1) - 2c(x) + c(x-1)\\&\quad + (1-p)q(\Delta(x+1-d_1,\mathbf{d}^+ + \mathbf{e}_H)\\&\quad - \Delta(x-d_1,\mathbf{d}^+ + \mathbf{e}_H))\\&\quad + p(1-q)(\Delta(x+1-d_1,\mathbf{d}^+)\\&\quad - \Delta(x-d_1,\mathbf{d}^+)),\end{aligned}$$

where the right-hand side is once again non-negative. ■

A2: Proof of Proposition 2 We have to show that $\Delta(x,\mathbf{d}) \geq \Delta(x,\mathbf{d}')$. The convexity condition and the induction assumption require that:

$$(a(x+1,\mathbf{d}), a(x,\mathbf{d}), a(x+1,\mathbf{d}'), a(x,\mathbf{d}')) \in \{(0,0,0,0),(0,0,0,1),(0,0,1,1),(0,1,0,1),(0,1,1,1),(1,1,1,1)\}.$$

For most cases, the required inequality directly follows when the differences are expressed explicitly through the optimality equation. In fact there are only two non-trivial cases to consider: $(a(x+1,\mathbf{d}), a(x,\mathbf{d}), a(x+1,\mathbf{d}'), a(x,\mathbf{d}')) = (0,0,0,1)$ and $(a(x+1,\mathbf{d}), a(x,\mathbf{d}), a(x+1,\mathbf{d}'), a(x,\mathbf{d}')) = (0,1,1,1)$. We focus on the first one (i.e., $(0,0,0,1)$). The difference can be expressed as:

$$\Delta(x,\mathbf{d}) - \Delta(x,\mathbf{d}') = q\Delta(x-d_1,\mathbf{d}^+ + \mathbf{e}_H) + (1-q)\Delta(x-d_1,\mathbf{d}^+) - ((1-p)q\Delta(x-d_1,\mathbf{d}'^+ + \mathbf{e}_H) + (1-p)(1-q)\Delta(x-d_1,\mathbf{d}'^+)). \quad \text{(A1)}$$

Regrouping the first and third terms of (A1), we note that:

$$q\Delta(x-d_1,\mathbf{d}^+ + \mathbf{e}_H) - (1-p)q\Delta(x-d_1,\mathbf{d}'^+ + \mathbf{e}_H) \geq 0,$$

by the induction assumption. Similarly, regrouping the second and fourth terms of (A1), we obtain, by the induction assumption, that:

$$(1-q)\Delta(x-d_1,\mathbf{d}^+) - (1-p)(1-q)\Delta(x-d_1,\mathbf{d}'^+) \geq 0.$$

This establishes that $\Delta(x,\mathbf{d}) - \Delta(x,\mathbf{d}') \geq 0$. All the other cases follow in a straight-forward manner and are omitted. ■

Appendix B: The optimal base stock policy without advance demand information

In this section, we present the details of the computation of the optimal base stock policy without advance demand information for a Geo/Geo/1 queue in discrete-time.

Note that under a base stock policy with base stock level S (≥ 0), the process $N(t) = S - X(t)$ is a Geometric/Geometric/1 queue. Let $\pi(n)$ denote the stationary probability that there are n customers in this queue. It turns out that (see Pujolle *et al.* (1986) for example):

$$\pi(0) = 1 - (q/p), \qquad \pi(i) = \left[1 - \frac{q}{p}\right]\frac{1}{1-p}\beta^i \quad i \geq 1 \quad \text{(A2)}$$

where: $\beta = [q(1-p)]/[(1-q)p]$.

Let F_N denote the cumulative distribution function of the stationary random variable N. The optimal base stock level is given by S, $S = \min\{n : F_N(n) > b/(h+b)\}$. Defining the parameter: $\kappa = (p-q)/(p(1-p))$, we obtain the optimal base stock level, S^*, as:

$$S^* = \left\lfloor \ln\left[\frac{h}{(h+b)}\frac{(1-\beta)}{\kappa}\right] \Big/ \ln[\beta] \right\rfloor.$$

As expected, S^* is increasing in the average load q/p. Furthermore, keeping q/p constant and varying p (and q) we can observe the effects of processing (and order interarrival) time variability. S^* can be seen to be decreasing in p (or q) for a fixed value of q/p. This reflects the two significant factors affecting optimal safety stocks: the capacity utilization and the variability in processing or order interarrival times.

Biographies

Yves Dallery received his Ph.D. and the degree of Habilitation a Diriger des Recherches from the Institut National Polytechnique de Grenoble (INPG) in 1984 and 1989, respectively. He is currently Professor of Manufacturing and Logistics at the Ecole Centrale de Paris. Before that, he was Directeur de Recherche at the Centre National de la Recherche Scientifique (CNRS). In 1984–1985, he was a post-doctoral fellow at Harvard University. In 1991–1992, he was a visiting scientist at MIT and in 1992–1993 he was an Associate Professor of Manufacturing Engineering at Boston University. His research interests are in stochastic models, production management, supply chain management and service operations management.

John Buzacott was born in Sydney, Australia. He obtained a B.Sc. in Physics and a B.E. with First Class Honours in Electrical Engineering from the University of Sydney. He then went to the UK where he worked in industry for a couple of years, obtained an M.Sc. in Operational Research from the University of Birmingham, worked in industry again, and then returned to Birmingham to obtain a Ph.D. He moved to Canada in 1967 and has taught at the Universities of Toronto and Waterloo and, since 1991, at York University. He has been President of the Canadian Operational Research Society, Chair of the ORSA Technical Section on Manufacturing and Operations Management, and President of the Production and Operations Management Society. He is a Senior Editor of *M&SOM*, an Associate Editor of *Management Science*, *Operations Research*, *Queueing Systems: Theory and Applications*, and on the Editorial Board of the *International Journal of Flexible Manufacturing Systems*. His research interests are in developing models to provide insight into operational and strategic issues in manufacturing and services, in particular models that help understand the impact of variability and uncertainty. He is co-author of two books and about 80 papers in refereed journals. In 2001 the Technical University of Eindhoven in the Netherlands awarded him the degree of Doctor Honoris Causa.

Fikri Karaesmen received his Ph.D. from Northeastern University in 1996. In 1997–1999, he was a visiting researcher at the Universite Paris 6. He is currently an Associate Professor of Manufacturing and Logistics at Ecole Centrale de Paris. His research interests are in stochastic models of production/inventory systems and supply chains.