Int. J. Production Economics 115 (2008) 433–438

Contents lists available at ScienceDirect

Int. J. Production Economics

journal homepage: www.elsevier.com/locate/ijpe

Periodic review inventory models with stochastic supplier's visit intervals

Chi Chiang

Department of Management Science, National Chiao Tung University, Hsinchu, Taiwan, ROC

ARTICLE INFO

Article history:
Received 30 September 2006
Accepted 20 March 2008
Available online 22 June 2008

Keywords:
Inventory model
Periodic review
Supply chain
Lost-sales
Dynamic programming

ABSTRACT

Periodic review inventory models are widely used in practice, especially for inventory systems in which many different items are purchased from the same supplier. However, all periodic review models have assumed a fixed length of the review periods. In practice, it is possible that the review periods are of a variable length. Such periodic systems result mainly from supply uncertainties. For example, the supplier visits the downstream retailers and replenishes inventories for them, but does not always come in constant intervals. This may be because retailers are geographically dispersed in the supply chain, the supplier is in a relatively more powerful position, or the supplier simply does not have a reliable visit schedule. In such situations, the replenishment cycle length is random in nature. In this paper, we use dynamic programming to model such institutional contexts. We assume that the supplier's visit intervals are independently and identically distributed. With a suitable transformation, the back-logged periodic review model derived becomes a standard discrete-time model. The computation shows that ignoring the variability of the supplier's visit intervals can incur extremely large losses, especially if shortage is costly, demand variability is low, and/or the replenishment lead-time is short.

1. Introduction

Though the use of computer systems has made continuous review models more attractive, periodic review models are still applied in many situations (e.g., Prasad et al., 2005; Silver et al., 1998), especially for inventory systems in which the coordination of ordering and transportation for different items is important (which is especially true if these items are purchased from the same supplier). Also, as Porteus (1985) observes, continuous review systems that keep inventory records current, but order periodically, are equivalent to periodic review systems. Often, periodic systems have the review periods that are possibly longer than the supply lead-times.

One fundamental assumption about periodic systems is that the review periods are of a fixed length. In practice, however, the review periods may be of a variable length. Such periodic systems result mainly from supply uncertainties. For example, many supermarkets have suppliers who come to visit regularly and replenish the inventory of various items (and even sell) for them. However, the supplier does not always come in constant (say, 10 days) intervals. Depending on her visit plans or work schedules and loads, she often arrives at a particular supermarket one or few hours (or days) early or late. The elapsed time between two consecutive visits varies basically. Ertogral and Rahim (2005) also observed institutional settings or constraints that are internal to the supply chain, in which the supplier is strategically dominant, in a relatively more powerful position, and/or the retailers are located in a geographically disadvantageous remote location, so that the supplier decides when to visit and replenish the retailers' inventories. In general, for such situations, the

E-mail address: cchiang@mail.nctu.edu.tw

0925-5273/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.ijpe.2008.03.012

replenishment epochs are not under retailers' control; rather, they are under the supplier's control. Hence, if the supplier arrives at a particular retailer in irregular intervals, the replenishment cycle length for that retailer is random in nature.

To our knowledge, the possibility of stochastic review periods or replenishment intervals has not been investigated in the inventory literature, though there are some works on inventory models with supply uncertainties (e.g., Lee et al., 1997; Mohebbi, 2004; Ozekici and Parlar, 1999). It was studied only recently by Ertogral and Rahim (2005) who derived the expected profit per replenishment cycle by assuming independently and identically distributed (i.i.d.) replenishment intervals, constant demand and zero replenishment lead-time.

In this paper, we use dynamic programming to model the supply chain situations where the supplier's visit intervals (i.e., replenishment intervals) are random. We will assume that the supplier's visit intervals are i.i.d., as in Ertogral and Rahim (2005). However, unlike Ertogral and Rahim, we will assume stochastic demand which is usually true in the real world; also, we will allow the replenishment lead-time to be positive (i.e., it may take a positive time to replenish inventories after the supplier arrives at a retailer and reviews his inventories). We will develop both the backlogged and lost-sales periodic review inventory models. With a suitable transformation, the backlogged model derived becomes a standard discrete-time model. Thus, an order-up-to policy is known to be optimal for the infinite horizon problem. This is also true of the lost-sales problem with zero lead-time (due to a result from the inventory literature). For the lost-sales problem with positive lead-time, we suggest a simple heuristic policy in Hadley and Whitin (1963).

The computation shows that ignoring the variability of the supplier's visit intervals can incur unnecessary large costs, especially if shortage is costly, the replenishment lead-time is short, and/or demand variability is not high. It is thus important for a retailer to incorporate this variability into inventory models when the supplier does not visit in constant intervals. It would be better if the retailer can have the supplier to visit in more regular intervals (i.e., the visit interval has a smaller variability) so that his cost can be reduced, as shown in the computation. This may not be an easy task, since the institutional constraints are perhaps difficult to change in the short run (for example, the supplier is in a relatively more powerful position as described above). For such institutional contexts, we suggest that the retailer should somehow persuade the supplier to visit more regularly. The retailer should at least communicate with the supplier often so that she understands the consequence of the irregular visit intervals and hopefully, she will continue to improve on her visit schedule in terms of the stability/reliability in the future.

Of course, it is possible that the supplier completely fixes the visit interval after the retailer's persuasion. The supplier and the retailer may even cooperate closely in the supply chain, or form a strategic alliance in the long run. Then the supplier will also visit the retailer and replenish his inventories more often (not only more regularly) so as to further reduce his costs, and in return, the retailer could negotiate a long-term supply contract or purchase other products from the supplier, for example. All of these are certainly a significant change of status-quo, i.e., a breakthrough of the supply chain. Note that we are not saying that it is not good to have the replenishment epochs under the supplier's control; it may be one of the most efficient ways of operating the supply chain (in terms of replenishing the downstream retailers' inventories), especially for the institutional settings described above. We simply say that cooperation between the supplier and the retailer could result in a win–win situation. If indeed the supplier no longer visits the retailer in irregular intervals, then the ordinary periodic review models found in textbooks can be used, i.e., the periodic review models derived in this paper need not be used.

2. Backlogged periodic review inventory models

We first assume that all demand not filled immediately is backlogged. Let c denote the unit item cost and L the (deterministic) lead-time. Demand is stochastic with mean rate μ per unit time, and is assumed to be non-negative and independently distributed in disjoint time intervals. Let T be the period length, i.e., the supplier's visit interval. Successive T's are assumed to be i.i.d. random variables. Let $\phi(\cdot)$ be the probability density function (pdf) of T and D the demand during T. Also let $g(\cdot|\tau)$ be the conditional pdf of demand during a time interval of length τ. Thus, $g(\cdot|T)$ is the conditional pdf of D.

Let α be the discount rate, y the inventory position (i.e., inventory on hand minus backorder plus inventory on order) after an order is placed at a review epoch (i.e., upon the supplier's visit), and H the expected one-period inventory holding and shortage cost (H is a function of y). Given time 0 at a review epoch, H is charged for the time interval $[L, T{+}L)$. Denote $V_n(x)$ as the expected discounted cost with n periods remaining until the end of the planning horizon when the starting inventory position is x and an optimal ordering policy is used at every review epoch. $V_n(x)$ satisfies the functional equation:

$$V_n(x) = \min_{x\leqslant y}\{e^{-aL}[cy + H(y)] + E_T[e^{-aT}E_{D|T}[V_{n-1}(y-D)]]\} - e^{-aL}cx, \quad (1)$$

where the procurement cost $c(y–x)$ is paid upon delivery. The above dynamic program is an inventory problem with discrete but random epochs. Let $\beta = E_T[e^{-\alpha T}]$. Using a standard approach in semi-Markov decision processes (e.g., Puterman, 1994, p. 542) and defining $\varphi(\cdot)\equiv E_T[e^{-\alpha T}g(\cdot|T)]/\beta$, i.e., $\beta\varphi(\cdot) = E_T[e^{-\alpha T}g(\cdot|T)]$ is the discount density of D and $\varphi(\cdot)$ is the "normalized" pdf of D, we can express $E_T[e^{-\alpha T}E_{D|T}[V_{n-1}(y–D)]]$ by

$$\begin{aligned}&E_T\left[e^{-\alpha T}E_{D|T}[V_{n-1}(y-D)]\right]\\&= \int_0^\infty e^{-\alpha T}\left(\int_0^\infty V_{n-1}(y-D)g(D|T)\mathrm{d}D\right)\phi(T)\mathrm{d}T\\&= \int_0^\infty V_{n-1}(y-D)\left(\int_0^\infty e^{-\alpha T}g(D|T)\phi(T)\mathrm{d}T\right)\mathrm{d}D\end{aligned}$$

$$= \beta \int_0^\infty V_{n-1}(y-D)\left\{\int_0^\infty e^{-\alpha T} g(D|T)\phi(T)\,dT)/\beta\right\} dD$$
$$= \beta \int_0^\infty V_{n-1}(y-D)\varphi(D)\,dD$$
$$= \beta E_D[V_{n-1}(y-D)], \quad (2)$$

where the expectation E_D is taken over the pdf $\varphi(\cdot)$. After the above transformation, $V_n(x)$ is written as

$$V_n(x) = \min_{x\leqslant y}\{e^{-aL}[cy + H(y)] + \beta E_D[V_{n-1}(y-D)]\} - e^{-aL}cx. \quad (3)$$

The original problem in (1) is now a standard discrete-time model.

Next, we give an expression for $H(y)$. Let h be the holding cost per unit held per unit time and p the shortage cost per unit. Assume that an order when arriving is almost always sufficient to meet any outstanding backorders (see Hadley and Whitin, 1963, p. 239, for a detailed discussion). Thus, backorders that occur during the time interval $[L, L+T)$ can be computed for the interval $[0, L+T)$ and the expected on-hand inventory immediately after the arrival of an order is $y-L\mu$. If $D\leqslant y-L\mu$, as the expected on-hand inventory just prior to the arrival of the next order is $y-L\mu-D$, the average holding cost over the interval $[L, L+T)$ is $hT(y-L\mu-0.5D)$. On the other hand, if $D>y-L\mu$, the expected on-hand inventory falls to zero at some time during the interval $[L, L+T)$. Assuming that the expected inventory decreases linearly with time (e.g., Hadley and Whitin, 1963, p. 238), the expected on-hand inventory falls to zero at time $L+T(y-L\mu)/D$; since the average on-hand inventory over the interval $[L, L+T(y-L\mu)/D)$ is $0.5(y-L\mu)$, the average holding cost over this interval, which is also the average holding cost over the interval $[L, L+T)$, is $0.5hT(y-L\mu)^2/D$. As D appears in the denominator, this exact expression will complicate the subsequent analysis. We thus use a lower bound $hT(y-L\mu-0.5D)$ for this expression, i.e., the same one as when $D\leqslant y-L\mu$ (see Chiang (2003) for a similar approach in the two-supply-mode setting). This approximation is similar to the one in the ordinary periodic model where the expected on-hand inventory approximately equals the expected net inventory (Hadley and Whitin, 1963, pp. 237–239). Hence, the average holding cost over the interval $[L, L+T)$, after taking the expectation of D, is $hT(y-L\mu-0.5T\mu)$. It follows that if T is constant,

$$H(y) = hT(y - L\mu - 0.5\,T\mu) + \int_y^\infty p(\zeta - y) g(\zeta|T+L)\,dz \quad (4)$$

(Hadley and Whitin, 1963, p. 240). For the present model in which T is a variable, $H(y)$ is given by

$$H(y) = E_T[hT(y - L\mu - 0.5\,T\mu)] + \int_y^\infty p(\zeta - y) g^*(\zeta|L)\,dz, \quad (5)$$

where $g^*(\cdot|L)\equiv E_T[g(\cdot|T+L)]$. Since $H(y)$ is convex, $V_n(x)$ in (3) is a convex function (by induction and Proposition B-4 of Heyman and Sobel, 1984). Hence, a stationary order-up-to policy (i.e., base-stock policy) is known to be optimal for the infinite horizon problem. To obtain the optimal order-up-to level y^*, we minimize the following myopic function:

$$J(y) = cy(1-\beta) + H(y) \quad (6)$$

(e.g., Veinott and Wagner, 1965, p. 527). Thus, we set the first derivative of $J(y)$ to zero:

$$c(1-\beta) + hE[T] - \int_y^\infty pg^*(\zeta|L)d\zeta = 0$$

or

$$\int_y^\infty g^*(\zeta|L)\,d\zeta = \{c(1-\beta) + hE[T]\}/p \quad (7)$$

and solving for the optimal y^*. J is basically the expected cost of the upcoming period. Notice that the constant scaling factor $e^{-\alpha L}$ (discounted to the present time) and procurement cost $e^{-\alpha L}c\beta E_D[D]$ are not included in (6) for simplicity. As the ratio $\{c(1-\beta)+hE[T]\}/p$ should be less than 1 (since the average backorder level is assumed to be small [Hadley and Whitin, 1963, p. 241]), y^* is guaranteed to be obtained. Let y' be the optimal y found if T is fixed.

3. Lost-sales periodic review inventory models

Suppose now that demand not satisfied at once is lost. Assume that L is less than the minimum T (i.e., there is at most one order outstanding). Let D_1 be the demand during the lead-time L and D_2 the demand during the time interval $[L, T)$ (thus $D = D_1+D_2$). Also, let z be the order quantity placed at a review epoch and redefine x as the starting on-hand inventory. Let $(\cdot)^+\equiv\max\{\cdot, 0\}$. Then, $V_n(x)$ satisfies the recursive equation:

$$V_n(x) = \min_{z\geqslant 0}\{e^{-\alpha L}(cz + E_{D_1}[H((x-D_1)^+ + z)]) + E_T\left[e^{-aT}E_{D_1,D_2|T}[V_{n-1}(((x-D_1)^+ + z - D_2)^+)]\right]\} \quad (8)$$

Consider first the simplest case of $L = 0$. Then (8) reduces to

$$V_n(x) = \min_{z\geqslant 0}\{cz + H(x+z) + E_T\left[e^{-\alpha T}E_{D|T}[V_{n-1}((x+z-D)^+)]\right]\}$$
$$= \min_{x\leqslant y}\{cy + H(y) + E_T\left[e^{-\alpha T}E_{D|T}[V_{n-1}((y-D)^+)]\right]\} - cx. \quad (9)$$

If T is constant, (9) simplifies to

$$V_n(x) = \min_{x\leqslant y}\{cy + H(y) + e^{-\alpha T}E_D[V_{n-1}((y-D)^+)]\} - cx. \quad (10)$$

Veinott and Wagner (1965, p. 528) showed that the lost-sales model in (10) could be viewed as a backlog model in which a credit of $e^{-\alpha T}c$ is given to each unit of demand actually backlogged (an order-up-to policy is thus optimal). For the model in (9) where T is stochastic, we can transform (9) to (as in the backlogged model)

$$V_n(x) = \min_{x\leqslant y}\{cy + H(y) + \beta E_D[V_{n-1}((y-D)^+)]\} - cx, \quad (11)$$

where the expectation E_D is taken over the pdf $\varphi(\cdot)$. Hence, an order-up-to policy is optimal for the infinite horizon problem and the optimal order-up-to level y^* is obtained by minimizing

$$J(y) = cy(1-\beta) + H(y) - \beta c \int_y^\infty (D-y)\varphi(D)\,\mathrm{d}D, \qquad (12)$$

where the last term (i.e., the difference between (6) and (12)) is the credit given to demand not satisfied and the exact one-period holding and shortage cost is given by

$$H(y) = E_T\left[0.5hT\left(y + \int_0^y (y-D)g(D|T)\mathrm{d}D\right) + \int_y^\infty p(D-y)g(D|T)\mathrm{d}D\right] \qquad (13)$$

(which is convex in y). The holding cost of (13) is derived based on the average of beginning and ending inventories of a period. Also, the shortage penalty p has a different meaning in the lost-sales model and it usually includes the unit sales revenue. As $J(y)$ is convex, y^* is found by solving its first-order condition, i.e.,

$$c(1-\beta) + E_T\left[0.5hT\left(1 + \int_o^y g(D|T)\mathrm{d}D\right) - \int_y^\infty pg(D|T)\mathrm{d}D\right] + cE_T\left[\mathrm{e}^{-\alpha T}\int_y^\infty g(D|T)\mathrm{d}D\right] = 0$$

(noting that $\beta\varphi(\cdot) = E_T[\mathrm{e}^{-\alpha T}g(\cdot|T)]$), or

$$c(1-\beta) + E_T\left[0.5hT(1 + 1 - \int_y^\infty g(D|T)\mathrm{d}D) - \int_y^\infty pg(D|T)\mathrm{d}D\right] + cE_T\left[\mathrm{e}^{-\alpha T}\int_y^\infty g(D|T)\mathrm{d}D\right] = 0,$$

i.e.,

$$E_T\left[(0.5hT + p - \mathrm{e}^{-\alpha T}c)\int_y^\infty g(D|T)\mathrm{d}D\right] = c(1-\beta) + hE[T]. \qquad (14)$$

If L is positive (deterministic), the lost-sales model given by (8) is difficult to solve. If T is fixed, see, e.g., Hadley and Whitin (1963, p. 285), Morton (1971), and Zipkin (2000, pp. 411–413) (note that Chiang (2006) recently develops optimal ordering policies for the case of $L \leqslant T$). For the present model in (8) where T is variable, heuristic approaches need to be used. We suggest that one uses the order-up-to policy in Hadley and Whitin (1963, pp. 240–242). Using our notation, the expected undiscounted cost of the upcoming time interval [L, T+L), if T is fixed, is expressed by

$$H(y) = hT(y - \mu L - 0.5T\mu) + (0.5hT + p - c)\int_y^\infty (\zeta - y)g(\zeta|T+L)\mathrm{d}\zeta. \qquad (15)$$

As a note, there is an error in expressions (5–11) of Hadley and Whitin. The integral should be multiplied by a factor of 0.5. Since the base-stock policy orders filled demands just as it does in a backlogged model (except that now lost sales are not counted), the average cycle stock (i.e., order quantity) over the interval [L, T+L) is given by

$$0.5\left\{T\mu - \int_y^\infty (\zeta - y)g(\zeta|T+L)\mathrm{d}\zeta\right\} \qquad (16)$$

instead of $0.5T\mu$. Note that (16) is an approximation, since it ignores the effects of lost sales that can occur between the time an order is placed and the time it arrives. Adding (16) to the safety stock, given by (5–10) of Hadley and Whitin, would yield the holding cost component of (15). For the present model in which T is a variable, the expected undiscounted cost of the upcoming time interval [L, T+L) is written by

$$H(y) = E_T[hT(y - \mu L - 0.5T\mu)] + E_T\left[(0.5hT + p - c)\int_y^\infty (\zeta - y)g(\zeta|T+L)\mathrm{d}\zeta\right]. \qquad (17)$$

The optimal order-up-to level y^* is then obtained by minimizing (17), i.e., setting the first derivative of (17) to zero:

$$hE[T] - E_T\left[(0.5hT + p - c)\int_y^\infty g(\zeta|T+L)\mathrm{d}\zeta\right] = 0,$$

or

$$E_T\left[(0.5hT + p - c)\int_y^\infty g(\zeta|T+L)\mathrm{d}\zeta\right] = hE[T] \qquad (18)$$

and solving for the optimal y^*. Let y' be the optimal level found from (14) or (18) when T is fixed.

4. Computational results

We investigate the effect of a variable T on the expected cost, if a retailer fails to incorporate it when developing inventory policies. In the following experiments, we assume that demand is normal with mean $\mu\tau$ and variance $\sigma^2\tau$ for a time interval of length τ. The common data used are $\mu = 10$/day (unit time is one day), $h = \$0.1$, $\alpha = 0$, and $E[T] = 10$ days. Also, T is either triangularly or uniformly (discretely) distributed. In the former case, $Pr(T=8) = Pr(T=12) = 1/9$, $Pr(T=9) = Pr(T=11) = 2/9$, and $Pr(T=10) = 1/3$; in the latter case, $Pr(T=8) = Pr(T=9) = Pr(T=10) = Pr(T=11) = Pr(T=12) = 1/5$. Thus, T has a larger variability if it is uniformly distributed.

First consider the backlogged model with $L = 0$. Table 1 reports computational results as p and σ^2 are varied. Three observations can be made from Table 1. First, ignoring T's variability, i.e., using y' of the ordinary periodic model when in fact T is stochastic, can incur unnecessary large costs, especially if T is uniformly distributed (i.e., has a larger variability). Second, it appears that as the unit shortage cost p is higher, a firm incurs larger losses. This result can be seen by comparing (5) to (4). T's variability affects $J(y)$ only through its shortage cost component. As p is higher, T's variability has a larger effect on $J(y)$. Third (and probably most importantly), as σ is smaller, T's variability has a larger impact on $J(y)$. In the extreme case where $\sigma = 0$ (i.e., deterministic demand), ignoring T's variability can increase the cost by more than 300 percent! This is possibly because the variability of demand during T plus L includes T's variability and demand variability; the introduction of T's variability into an inventory model increases the overall variability of demand during T plus L more significantly when demand variability is smaller.

Table 1
Effect of variable T on the expected cost (backlogged model with $L=0$)

σ	p	y'	(A) T is triangularly distributed				(B) T is uniformly distributed			
			y^*	$J(y^*)$	$J(y')$	% Off	y^*	$J(y^*)$	$J(y')$	% Off
0	10	100	120	69.3	93.8	35.4	120	69.0	109.0	56.0
	20	100	120	69.3	138.2	99.4	120	69.0	169.0	144.9
	40	100	120	69.3	227.1	227.7	120	69.0	289.0	318.8
2	10	109	118	72.5	78.5	8.3	121	74.8	87.1	16.4
	20	111	122	75.9	91.4	20.4	125	77.9	106.4	36.6
	40	113	128	79.3	109.1	37.6	128	80.5	134.2	66.7
4	10	117	123	80.1	81.7	2.0	125	82.7	86.5	4.6
	20	121	129	85.5	90.1	5.4	132	88.1	97.3	10.4
	40	125	134	90.2	98.5	9.2	137	92.8	108.8	17.2
8	10	133	136	99.4	99.7	0.3	138	101.6	102.3	0.7
	20	142	147	108.5	109.5	0.9	149	111.0	113.1	1.9
	40	150	156	116.7	118.4	1.5	159	119.5	123.2	3.1

Table 2
Effect of variable T on expected cost (backlogged model with $L=6$)

σ	p	y'	(A) T is triangularly distributed				(B) T is uniformly distributed			
			y^*	$J(y^*)$	$J(y')$	% Off	y^*	$J(y^*)$	$J(y')$	% Off
0	10	160	180	69.3	93.8	35.4	180	69.0	109.0	56.0
	20	160	180	69.3	138.2	99.4	180	69.0	169.0	144.9
	40	160	180	69.3	227.1	227.7	180	69.0	289.0	318.8
2	10	171	179	74.0	78.3	5.8	182	76.5	85.5	11.8
	20	174	183	77.9	87.6	12.5	186	80.1	99.0	23.5
	40	176	187	81.2	101.8	25.4	190	83.1	120.6	45.1
4	10	181	186	84.5	85.5	1.2	188	86.9	89.3	2.8
	20	187	193	90.7	93.0	2.5	196	93.2	98.0	5.2
	40	192	199	96.2	100.2	4.2	202	98.7	107.0	8.4
8	10	202	204	109.9	110.0	0.1	205	111.6	111.8	0.2
	20	213	217	120.8	121.1	0.3	219	122.8	123.7	0.7
	40	223	228	130.5	131.3	0.6	230	132.8	134.6	1.4

Table 3
Effect of variable T on expected cost (lost-sales model with $L=0$)

σ	p	y'	(A) T is triangularly distributed				(B) T is uniformly distributed			
			y^*	$J(y^*)$	$J(y')$	% Off	y^*	$J(y^*)$	$J(y')$	% Off
0	110	100	120	69.3	96.3	39.0	120	69.0	112.5	63.0
	120	100	120	69.3	140.8	103.2	120	69.0	172.5	150.0
	140	100	120	69.3	239.7	245.9	120	69.0	292.5	323.9
2	110	109	118	72.8	79.7	9.5	122	75.2	88.8	18.1
	120	111	122	76.1	92.3	21.3	125	78.0	107.8	38.2
	140	113	128	79.3	109.8	38.4	128	80.5	135.3	68.1
4	110	117	124	80.6	82.6	2.5	126	83.2	88.8	6.7
	120	121	129	85.7	90.6	5.7	133	88.3	98.1	11.1
	140	125	135	90.3	98.9	9.5	138	92.8	109.3	17.8
8	110	133	137	100.1	100.4	0.3	139	102.4	103.1	0.7
	120	142	147	108.8	110.0	1.1	149	111.4	113.7	2.1
	140	150	157	116.9	118.7	1.5	159	119.6	123.5	3.3

Table 4
Effect of variable T on expected cost (lost-sales model with $L=6$)

σ	p	y'	(A) T is triangularly distributed				(B) T is uniformly distributed			
			y^*	$J(y^*)$	$J(y')$	% Off	y^*	$J(y^*)$	$J(y')$	% Off
0	110	160	180	69.3	96.3	39.0	180	69.0	112.5	63.0
	120	160	180	69.3	140.8	103.2	180	69.0	172.5	150.0
	140	160	180	69.3	239.7	245.9	180	69.0	292.5	323.9
2	110	171	179	74.4	79.4	6.7	182	76.8	87.0	13.3
	120	174	184	78.1	88.3	13.1	187	80.3	100.1	24.7
	140	176	188	81.3	102.4	26.0	190	83.2	121.4	45.9
4	110	181	186	85.0	86.4	1.6	189	87.5	90.4	3.3
	120	187	193	90.9	93.4	2.8	196	93.5	98.7	5.6
	140	192	199	96.3	100.4	4.3	202	98.8	107.4	8.7
8	110	202	204	110.8	111.0	0.2	206	112.5	113.0	0.4
	120	213	217	121.1	121.6	0.4	219	123.2	124.3	0.9
	140	223	228	130.7	131.6	0.7	230	133.0	134.9	1.4

Next, consider the backlogged model with a positive L (which equals, for example, 6 days). Comparing Table 2 to 1, we see that except for the case of deterministic demand, a positive L dilutes the effect of a variable T on $J(y)$. In fact, as L is larger (other things being equal), ignoring T's variability incurs smaller losses (more experiments are available from the author). This result can be explained by the same reason above: as the demand during T plus L becomes more volatile, the introduction of T's variability into an inventory model has a less significant effect on the cost.

In Tables 3 and 4, we consider the lost-sales models with zero and positive lead-times, respectively. The unit cost c is assumed to be \$100 and p is varied such that $(p–c)$ is the same as p in the backlogged models. As we see from these two tables, similar results to those in Tables 1 and 2 are observed, though the percentage loss of ignoring T's variability seems to be a little higher in general.

Notice that if T is fixed (equal to 10 days) rather than stochastic, the cost per period is significantly reduced (by using the ordinary periodic review models found in textbooks). For example, if $\sigma=0$, the optimal order-up-to level is 100 and the cost per period is only \$50, as compared to \$69.3 in the above tables. This illustrates the importance of a fixed visit schedule on the part of the supplier. As a note, the cost per period is only \$12.5 if T is fixed and cut down to 5 days (and $\sigma=0$). This certainly signals a message: the retailer should somehow have the supplier to come more often and regularly, as mentioned in Section 1.

5. Conclusions

In this paper, we considered periodic inventory models with stochastic supplier's visit intervals. We assumed that the supplier's visit intervals were independently and identically distributed. With a suitable transformation, the backlogged dynamic programming model derived became a standard discrete-time model. In addition, we suggested a simple order-up-to policy for the lost-sales periodic problem with positive lead-time. The periodic review policies developed in this paper were thus easy to implement.

The computation showed that ignoring the variability of the supplier's visit intervals could incur large losses if shortage was costly. It also showed that a retailer was more vulnerable to this variability if the replenishment lead-time was short and/or demand variability was small. This was because the introduction of T's variability into an inventory model increased the overall variability of demand during T plus L more significantly when lead-time was shorter and/or demand variability was smaller. In the extreme case where demand was deterministic, ignoring T's variability could increase the retailer's cost by more than 300 percent!

Moreover, the computation showed that a retailer could avoid some losses by reducing the variability of the supplier's visit intervals (e.g., from a uniform to triangular distribution). This might not be an easy task, because the replenishment epochs were under the supplier's control and such institutional constraints were perhaps difficult to change in the short run. However, the retailer could discuss this issue and explain its effect on his cost with the suppler. The retailer should at least communicate with the supplier often so that she would come to visit and replenish inventories more punctually (or even in constant intervals) in the future.

References

Chiang, C., 2003. Optimal replenishment for a periodic review inventory system with two supply modes. European Journal of Operational Research 149, 229–244.

Chiang, C., 2006. Optimal ordering policies for periodic-review systems with replenishment cycles. European Journal of Operational Research 170, 44–56.

Ertogral, K., Rahim, M.A., 2005. Replenish-up-to inventory control policy with random replenishment intervals. International Journal of Production Economics 93–94, 399–405.

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

Heyman, D.P., Sobel, M.J., 1984. Stochastic Models in Operations Research. McGraw-Hill, New York.

Lee, H.L., Padmanabhan, V., Whang, S., 1997. Information distortion in a supply chain: the bullwhip effect. Management Science 43, 546–558.

Mohebbi, E., 2004. A replenishment model for the supplier-uncertainty problem. International Journal of Production Economics 87, 25–37.

Morton, T.E., 1971. The near-myopic nature of the lagged-proportional-cost inventory problem with lost sales. Operations Research 19, 1708–1716.

Ozekici, S., Parlar, M., 1999. Inventory models with unreliable suppliers in a random environment. Annals of Operations Research 91, 123–136.

Porteus, E., 1985. Numerical comparisons of inventory policies for periodic review systems. Operations Research 33, 134–152.

Prasad, S., Tata, J., Madan, M., 2005. Build to order supply chains in developed and developing countries. Journal of Operations Management 23, 551–568.

Puterman, M., 1994. Markov Decision Processes. Wiley, New York.

Silver, E.A., Pyke, D.F., Peterson, R., 1998. Inventory Management and Production Planning and Scheduling. Wiley, New York.

Veinott Jr., A.F., Wagner, H.M., 1965. Computing optimal (s, S) inventory policies. Management Science 11, 525–552.

Zipkin, P.H., 2000. Foundations of Inventory Management. McGraw-Hill, New York.