BULLWHIP REDUCTION AND WIN-WIN STRATEGIES

Robert N. Boute[1], Stephen M. Disney[2], Marc R. Lambrecht[1],
Wim Van de Velde[3] and Benny Van Houdt[4]

[1] Department of Applied Economics, Katholieke Universiteit Leuven, Naamsestraat 69, 3000 Leuven, Belgium.
E-mail: robert.boute@econ.kuleuven.ac.be, marc.lambrecht@econ.kuleuven.ac.be.

[2] Logistics Systems Dynamics Group, Cardiff Business School, Cardiff University, Aberconway Building, Colum Drive, Cardiff, CF10 3EU, UK. E-mail: disneysm@cardiff.ac.uk.

[3] Procter and Gamble Services Company NV, Temselaan 100, 1853 Strombeek-Bever, Belgium.
E-mail: vandevelde.wm@pg.com.

[4] Department of Mathematics and Computer Science, University of Antwerp, Middelheimlaan 1, 2020 Antwerpen, Belgium. Email: benny.vanhoudt@ua.ac.be.

ABSTRACT
We consider a two echelon supply chain with a single retailer and a single manufacturer. Inventory replenishment policies at the retailer level transmit customer demand variability to the manufacturer, sometimes even in an amplified form (known as the bullwhip effect). When production is inflexible at the manufacturing level, significant costs may be incurred by ramping up and down production levels frequently. In this paper we focus on an inventory replenishment rule that reduces the variation of upstream orders and generates a smooth ordering pattern. This is beneficial for the manufacturer. However, dampening the variability in orders inflates the retailer's safety stock requirements due to the increased variance of the retailer's inventory levels. We can turn this conflicting issue into a win-win situation for both supply chain echelons when we treat the lead time as an endogenous variable. A less variable order pattern generates shorter and less variable lead times, creating a compensating effect on the retailer's safety stock.

Keywords: Bullwhip effect, order smoothing, supply chain collaboration

INTRODUCTION

Distorted information throughout the supply chain can lead to inefficiencies: excessive inventory investment, poor customer service, lost revenues, misguided capacity plans, ineffective transportation, and missed production schedules. Lee et al. (1997a, 1997b) describe a problem frequently encountered in supply chains, called the bullwhip effect: demand variability increases as one moves up the supply chain. Even when demand variability is not amplified but merely transmitted to the upstream manufacturer, this order variability can have large upstream cost repercussions. Production may be inflexible and significant costs may be incurred by ramping up and down production levels frequently. Upstream manufacturers therefore prefer minimal variability in the replenishment orders from the (downstream) retailers. Balakrishnan et al. (2004) emphasize opportunities to reduce supply chain costs by dampening order variability. This has led to the creation of new replenishment rules that are able to generate smooth ordering patterns, which we call 'smoothing replenishment rules'. Smoothing is a well-known method to reduce variability.

A number of production level smoothing rules were developed in the 1950's and 1960's. The more recent work on smoothing replenishment rules can be found in Dejonckheere et al. (2003) and Balakrishnan et al. (2004).

We have to be careful not to focus only on one side of the production smoothing 'coin'. The manufacturer does benefit from smooth production, but retailers, driven by the goal of reducing inventory (holding and shortage/backlog) costs, prefer to use replenishment policies that chase demand rather than dampen consumer demand variability. Dampening variability in orders may have a negative impact on the retailer's customer service due to inventory variance increases (Disney and Towill 2003). Inventory acts as a buffer, absorbing increases or decreases in demand while production remains relatively steady (Allen 1997). Disney et al. (2004, 2005) quantify the variance of the net stock and compute the required safety stock as a function of the smoothing intensity. Their main conclusion is that when customer demand is i.i.d., order smoothing comes at a price: in order to guarantee the same fill rate, more investment in safety stock is required.

However, we can model a two echelon supply chain as a production-inventory system, where the retailer's inventory replenishment lead times are endogenously determined by the manufacturer's production facility. As a consequence the choice of the retailer's replenishment policy (transmitting demand variability or smoothing orders) determines the arrival process at the manufacturer's production facility and as such it affects the distribution of replenishment lead times. We expect that a smooth order pattern gives rise to a shorter lead time (Hopp and Spearman 2001), which may then have a compensating effect on the safety stock.

In this paper we consider a smoothing replenishment policy and we treat the replenishment lead time as an endogenous variable. More specifically, we estimate the lead time distribution given the order pattern generated by our smoothing replenishment rule. We then focus on the resulting impact of order smoothing on the safety stock requirements to provide a given service level.

The remainder of the paper is organized as follows. First we describe our model and we compare a smoothing replenishment rule with the standard base-stock policy. After that we highlight the methodology to estimate the lead time distribution and to compute the impact on customer service and safety stock. Finally we illustrate with a numerical example.

MODEL DESCRIPTION

We consider a two echelon supply chain with a single retailer and a single manufacturer. Every period, the retailer observes the customer demand, denoted by D_t, that represents a finite number of items that customers buy from the retailer. We assume that customer demand D_t is identically and independently distributed (i.i.d.) over time. If there is enough on-hand inventory available, the demand is immediately satisfied. If not, the shortage is backlogged.

To maintain an appropriate amount of inventory on hand, the retailer places a replenishment order with the manufacturer at the end of every review period. The order quantity O_t is determined by the retailer's replenishment policy. We assume that the manufacturer produces on a make-to-order basis and consequently does not hold a finished goods inventory. The replenishment orders of size O_t enter the production facility where they are processed on a first-come-first-served basis. When the production facility is busy, they join a queue of unprocessed orders. We assume that the production times for a single product are i.i.d. random variables and to ensure stability (of the queue), we assume that the utilization of the production facility (average batch production time divided by average batch inter-arrival time) is strictly smaller than one.

Once the complete batch (equal to the replenishment order) is produced, it is immediately sent to the retailer. The time from the moment the order arrives at the production system to the point that the production of the entire batch is finished, is the *manufacturing* lead time or response time T_r. Note that the manufacturing lead time is not necessarily an integer number of periods. Since in our inventory model events occur on a discrete time basis with a time unit equal to one period, the *replenishment* lead time, denoted by T_p, has to be expressed in terms of an integer number of periods. We therefore rely on the sequence of events in a period. We assume that the retailer first receives goods from the manufacturer, then he observes and satisfies customer demand and finally, he places a replenishment order with the manufacturer.

In this sequence of events, the retailer is always able to satisfy demand after the receipt of products from the manufacturer. For instance, the retailer places an order at the end of period t, and it turns out that the manufacturing lead time is 0.8 periods. This order quantity will be added to the inventory in period t + 1, and will be used to satisfy demand in period t + 1. Therefore the replenishment lead time is 0 periods. A manufacturing lead time of 1.4 periods will be added to the inventory in period t + 2. Consequently we will treat the 1.4 period manufacturing lead time as an integer 1 period replenishment lead time. We therefore round the response time T_r down to the nearest integer T_p (i.e., setting $T_p = \lfloor T_r \rfloor$) to obtain the (discrete) replenishment lead time. A schematic of our model is shown in figure 1.

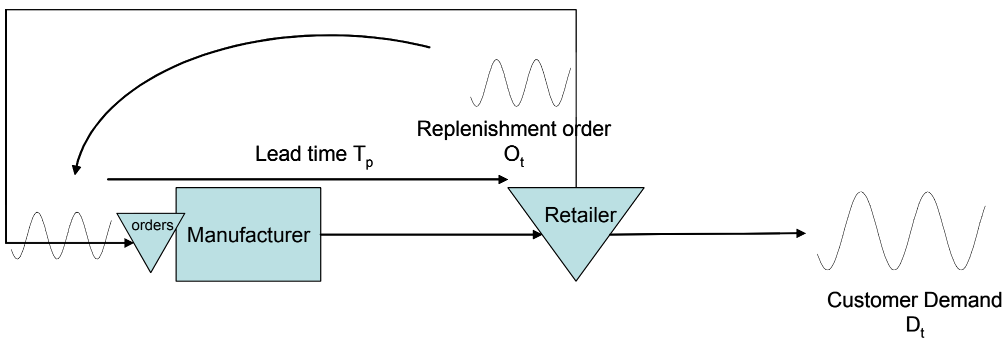

Figure 1: A two echelon supply chain modelled as a production/inventory system

TAMING THE BULLWHIP: ORDER SMOOTHING

There are many different types of replenishment policies, of which two are commonly used: the periodic review, replenishment interval, order-up-to policy and the continuous review, reorder point, order quantity model. Given the common practice in retailing to replenish inventories frequently (e.g. daily) and the tendency of manufacturers to produce to demand, we will focus our analysis on periodic review, base-stock or order-up-to replenishment policies.

Standard base-stock control policy

The standard periodic review base-stock replenishment policy is the (R,S) replenishment policy (Silver et al. 1998). At the end of every review period R, the retailer tracks his inventory position IP_t, which is the sum of the inventory on hand (that is, items immediately available to meet demand) and the inventory on order (that is, items ordered but not yet arrived due to the lead time) minus the backlog (that is, demand that could not be fulfilled and still has to be delivered). A replenishment order is then placed to raise the inventory position to an order-up-to or base-stock level S, which determines the order quantity O_t:

$$O_t = S - IP_t. \qquad (1)$$

The base-stock level S is the inventory required to ensure a given customer service level. Orders are placed every R periods. After an order is placed, it takes T_p periods for the replenishment to arrive, with T_p being the random replenishment lead time. Hence the risk period (the time between placing a replenishment order until receiving the subsequent replenishment order) is equal to the review period plus the replenishment lead time R + T_p. Since the lead time and the demand size are both random i.i.d. variables independent of each other, the demand during this risk period is a sum of a random number of random i.i.d. variables. Consequently, the base-stock level equals

$$S = [E(T_{p}) + R] \cdot E(D) + SS, \quad (2)$$

with E(D) the average demand and SS denoting the safety stock.

When the review period R is equal to one base period (that is, we place a replenishment order every period t), the order pattern generated by Eqn. (1) is simply equal to the demand pattern: $O_t = D_t$. Hence the variability of the replenishment orders is the same as the variability of the original demand.

We now introduce a smoothing replenishment policy that reduces the variability of the replenishment orders. In the remainder of this paper we will assume that the review period is one base period, so that we place an order at the end of every period t.

Smoothing replenishment policy

A smoothing replenishment policy is a policy where the decision maker does not recover the entire deficit between the base-stock level and the inventory position in one time period (contrary to what happens in Eqn. (1)). Forrester (1961) and Magee (1958) for example propose to order only a fraction of the inventory deficit, resulting in the following ordering policy:

$$O_t = \beta \cdot (S - IP_t). \quad (3)$$

Forrester (1961) refers to 1/β as the "adjustment time" and hence explicitly acknowledges that the deficit recovery should be spread out over time. Such a policy is justified when production (or ordering) and holding costs are convex or when there is a cost of changing the level of production (Veinott 1966).

When customer demand is i.i.d., the base-stock level S is a fixed constant. It can be proved (see Boute et al. 2005) that Eqn. (3) gives rise to an autocorrelated order pattern, given by

$$O_t = (1-\beta) \cdot O_{t-1} + \beta \cdot D_t. \quad (4)$$

The correlation between the orders is equal to

$$\mathrm{corr}(O_t, O_{t-x}) = (1 - \beta)^x. \quad (5)$$

This result is remarkable. Although we explicitly include endogenous lead times in our inventory model, the ordering decision is independent of the lead time (distribution). This implies that the lead time distribution does not affect the ordering decision, which simplifies the analysis of the combined production-inventory system considerably.

It is also notable that the replenishment rule described by Eqn. (4) is exactly the same as the exponential smoothing policy proposed by Balakrishnan et al. (2004) to decrease order variability. To examine the variability in orders created by our smoothing rule, we look at the ratio of the variance of the orders over the variance of demand (in the literature this variance ratio is commonly used as a measure for the bullwhip effect), which is in our case given by

$$\frac{\mathrm{Var(O)}}{\mathrm{Var(D)}} = \frac{\beta}{2 - \beta}. \quad (6)$$

From Eqn. (6) we find that, if we do not smooth, i.e. if β = 1, these expressions reduce to the standard base-stock policy, where $O_t = D_t$ and we chase sales and thus there is no variance amplification. For 1 < β < 2 we create bullwhip (variance amplification) and for 0 < β < 1 we generate a smooth replenishment pattern (dampening order variability).

Moreover, it can be proved (see Boute et al. 2005) that the base-stock level S in Eqn. (3) is not only affected by lead time demand, as in the standard base-stock policy, but it also contains an additional 'smoothing' component. More specifically, the base-stock level is given by

$$S = SS + [E(T_p)+ R] \cdot E(D) + (1-\beta)/\beta \cdot E(D), \qquad (7)$$

so that the base-stock level S increases the more we smooth (as β decreases towards zero). This can intuitively be understood. As depicted by Eqn. (3), the order quantities comprise only a fraction of the gap between inventory position and base-stock level. Since the average order quantity has to cover the average demand, the base-stock level has to increase when the fraction β decreases towards zero.

IMPACT OF ORDER SMOOTHING ON LEAD TIMES AND CUSTOMER SERVICE

Most replenishment rules proposed in the inventory literature take the replenishment lead time as a fixed constant or as an exogenous variable with a given probability distribution. However, the replenishment orders in fact load the production facilities, and the nature of this loading process relative to available capacity and the variability it creates are the primary determinants of lead times in the facility. Therefore the inventory control system should work with a lead time which is a good estimate of the real lead time, depending on the production load, the interarrival rate of orders and the variability of the production system. Zipkin (2000, p. 246) states "to understand the overall inventory system, we need to understand the supply system. For this purposes we can and do apply the results of queueing theory".

Methodology to estimate lead times

To estimate the lead time distribution we develop a discrete time queueing model. We explicitly model the two echelon supply chain described before as a production-inventory system, where the retailer's inventory replenishment lead times are endogenously determined by the manufacturer's production system. This implies that the choice of the retailer's replenishment policy determines the arrival process at the manufacturer's production facility. In a periodic review base-stock policy, the arrival pattern consists of batch arrivals with a fixed interarrival time (equal to the review period, R = 1) and with variable batch sizes. Moreover as we can see from Eqn. (4), the batch sizes generated by our smoothing rule are not i.i.d., rather they are autocorrelated.

The production time M of a single product is i.i.d. We slightly reduce the complexity of our queueing model by fitting a discrete phase-type (PH) distribution to the mean E(M) and variance Var(M) of the service time of a single product. The key idea behind PH distributions is to exploit the Markovian structure of the distribution to simplify the queueing analysis. Moreover, any general discrete distribution can be approximated in sufficient detail by means of a PH distribution.

In order to estimate the lead time distribution we define the following random variables:

- $a(n)$: the arrival time of the order in service at time t_n,
- B_n : the age of the order in service at time t_n,
- $O_{a(n)}$: the order quantity of the order in service at time t_n,
- C_n : the number of items part of the order in service that still need to start or complete service at time t_n,
- S_n : the service phase at time t_n,

where t_n denotes the time of the n-th observation point, which we define as the n-th time epoch during which the server is busy. Then, $(B_n, O_{a(n)}, C_n, S_n)$ forms a discrete time Markov process. Note that we need to keep track of the order quantity $O_{a(n)}$ in order to determine the order size of the next batch arrival (see Eqn. (4)).

The analysis of this Markov process can be solved using matrix analytic methods (Neuts 1981, Latouche and Ramaswami 1999). In this paper we will not discuss the solution procedure of this Markov process however. We refer the interested reader to Boute et al. (2005) for more details on this procedure as it is very complex. However, the result is clear. As expected, smooth order patterns generate shorter and less variable lead times.

Impact on customer service and safety stock

When demand is probabilistic, there is a definite chance of not being able to satisfy some of the demand directly out of stock. Therefore, a buffer or safety stock is required to meet unexpected fluctuations in demand. The goal is to reduce inventory without diminishing the level of service provided to customers. When the retailer faces (and satisfies) a variable customer demand, but replenishes through a smooth order pattern, this means that the variability in customer demand is transmitted to the retailer's inventory. In this case the inventory acts as a buffer, absorbing increases or decreases in demand while production remains relatively steady (Allen 1997). As a consequence, in order to provide the same service level a larger safety stock will be needed than the traditional standard base-stock replenishment policy where orders have the same variability as customer demand (Disney and Towill 2003, Disney et al. 2005).

However, since lead times are endogenously determined, dampening variability in orders affects the replenishment lead time distribution. A smoother order pattern generates a shorter and a less variable lead time. This creates a compensating effect on the required safety stock.

We characterize the inventory random variable and use it to find the safety stock requirements for the system. We first determine the optimal value of the base-stock level S. From this we can find the corresponding safety stock SS using Eqn. (7). We study the fill rate, which is a popular metric of customer service. It measures the proportion of the demand that can immediately be delivered from the inventory on hand:

$$\text{Fill rate} = 1 - \frac{\text{expected number of backorders}}{\text{expected demand}}. \qquad (8)$$

To calculate the fill rate, we monitor the inventory on hand after the customer demand is observed and we retain the number of shortages when a stock-out occurs. To do so, we observe the system at the end of every period t, after the demand is satisfied and after a replenishment order has been placed with the manufacturer (and before a possible order delivery at the retailer). At that time there may be $k \geq 0$ orders waiting in the production queue and there is always 1 order in service (since the observation moment is immediately after an order placement) which is placed k periods ago. Although k is a function of t, we write k as opposed to k(t) to simplify the notation. At this time instant the inventory on hand is equal to

$$NS_t = S - \sum_{i=0}^{k-1} D_{t-i} - \sum_{i=k}^{t} (1-\beta)^{i-k} D_{t-i}, \qquad (9)$$

and the number of backorders at the end of period t is given by

$$B_t = NS_t^-, \qquad (10)$$

where $NS_t^- = \max\{0, -NS_t\}$.

Some care must be taken when evaluating Eqn. (9) as the value of D_{t-k} influences the age of the current order in service k: the larger the demand size, the larger the order size and consequently the longer it takes to produce the order. Moreover, since the order quantity is also affected by previously realised demand terms (see Eqn. (12)), the demand terms D_{t-i}, $i = k + 1, \ldots, t$ also influence the current order's age k. Now, the age of the current order k determines the number of demand terms in the summations in Eqn. (9). Hence there is correlation between the different terms that make up NS_t.

At first sight the correlation between the different terms of NS_t seems to necessitate some kind of approximation. However, the Markov process (B_n, $O_{a(n)}$, C_n, S_n) used to estimate the lead times, retains the age k of the current order in service and the order quantity O_{t-k}. According to Eqn. (4) this order quantity is equal to

$$O_{t-k} = (1-\beta) \cdot O_{t-k-1} + \beta \cdot D_{t-k}, \tag{11}$$

or after backward substitution,

$$O_{t-k} = \sum_{i=k}^{t} \beta(1-\beta)^{i-k} D_{t-i}. \tag{12}$$

Dividing Eqn. (12) by β and substituting into Eqn. (9) gives

$$NS_t = S - \sum_{i=0}^{k-1} D_{t-i} - \frac{O_{t-k}}{\beta}, \tag{13}$$

where both the order quantity O_{t-k} and the age of the current order in service k are measures that can be obtained from the Markov process (B_n, $O_{a(n)}$, C_n, S_n). Hence, from the steady state analysis of this Markov process, we can determine the joint distribution of k and O_{t-k}, so that we can determine the steady state distribution of the net stock NS_t exactly. Hence the fill rate can be calculated as

$$\text{Fill rate} = 1 - \frac{E\left(\left(\sum_{i=0}^{k-1} D_{t-i} + \frac{O_{t-k}}{\beta} - S\right)^{+}\right)}{E(D)}. \tag{14}$$

In practice, decision makers often have to find the minimum safety stock that is required to achieve a given fill rate. From Eqn. (14) we can compute the minimal base-stock level S that is required such that an imposed fill rate is met. The corresponding safety stock can be found using Eqn. (7).

NUMERICAL EXAMPLE

To illustrate this procedure, we consider an example where the retailer daily observes a random customer demand between 1 and 20 products with an average of 5 and a standard deviation of 2.82. Its (discrete) probability distribution is shown in figure 2.

The retailer who satisfies this customer demand has to determine the parameter β to control his inventory. In the case of $\beta = 1$, the retailer places orders equal to demand and hence the variability in demand is transmitted to the manufacturer. We find that this policy results in an average replenishment lead time of 1.25 periods and a standard deviation of 1.42 periods. The lead time distribution is plotted in figure 4. In order to obtain a fill rate of 95%, the base-stock level should equal 33.23 resulting in a safety stock of 22 items.

Suppose that the retailer chooses to smooth his orders with a parameter $\beta = 0.4$. This results in an order pattern with less variability. In figure 2 we plot the order pattern resulting from this smoothing decision together with the observed customer demand pattern. Both probability distributions are centered around the same average value, but the order pattern is clearly less variable. Recall that when $\beta \neq 1$, the order pattern is correlated, while customer demand is i.i.d.

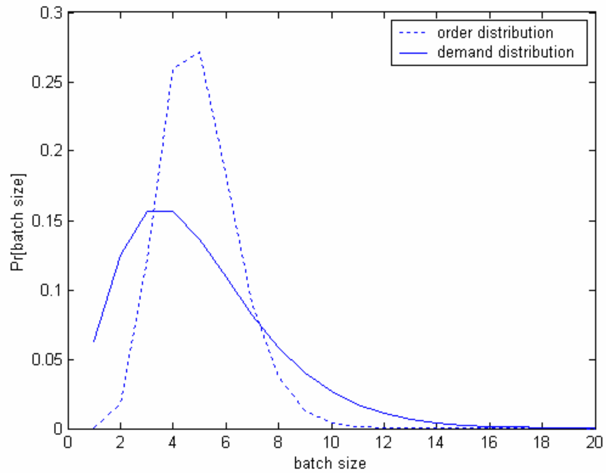

Figure 2: Demand pattern and smoothed order pattern with β = 0.4

When we consider an exogenous lead time distribution (e.g., the lead time distribution when we do not smooth the orders), this smoothing decision would lead to an increase in inventory variance, since inventory absorbs the variability in demand while the replenishments are relatively steady. As a consequence a higher safety stock has to be kept in order to maintain the same fill rate. In our example, smoothing with β = 0.4 would lead to a base stock level of 49.62 and a corresponding safety stock of 30.89. As an illustration we plot in figure 3 the safety stock in function of the smoothing parameter β when we use an exogenous lead time distribution. As expected, the safety stock increases the more we smooth.

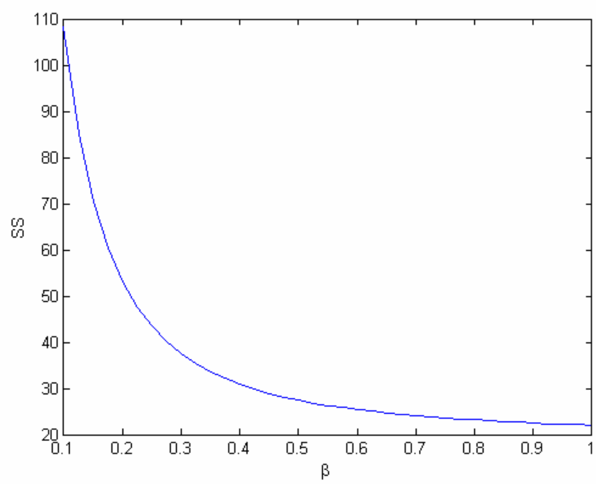

*Figure 3: Safety stock required to ensure a 95% fill rate with **exogenous** lead times*

Working with exogenous lead times is incomplete however. When the retailer smoothes his orders, he sends a less variable pattern to the manufacturer (see figure 2). This inevitably results in a different lead time distribution. Indeed, when we estimate the lead time distribution when we send a smooth order pattern with β = 0.4 to the manufacturer's production, we observe that the average lead time is equal to only 1.05 periods with a standard deviation of 1.31. Clearly, smoothing leads to a lower and less variable lead time. The resulting lead time distribution when we smooth our orders with β = 0.4 is plotted in figure 4. For comparison purposes, we add the lead time distribution in the case we do not smooth (dotted line).

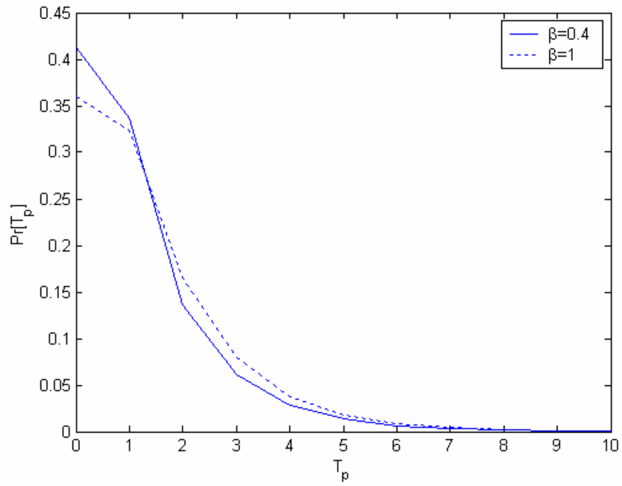

Figure 4: Lead time distribution when β = 1 (no smoothing) and when β = 0.4 (smoothing)

When we include this lead time distribution in our model, we find that a base-stock level of 38.91 is required to ensure a fill rate of 95%. Moreover this corresponds to a safety stock of only 21.12 items. This is even less safety stock then when we would not smooth our orders (in which case SS = 22). This implies that we can smooth orders without having to increase the safety stock in order to maintain customer service at the same target level. Moreover, we can even decrease our safety stock when we smooth the order pattern. This is clearly a win-win situation for both the retailer and the manufacturer. The manufacturer receives a less variable order pattern and the retailer can decrease his safety stock while maintaining the same fill rate.

Let us gradually decrease the smoothing parameter from β = 1 to β = 0.1. The lead time decreases the more we smooth. This introduces a compensating effect on the safety stock.

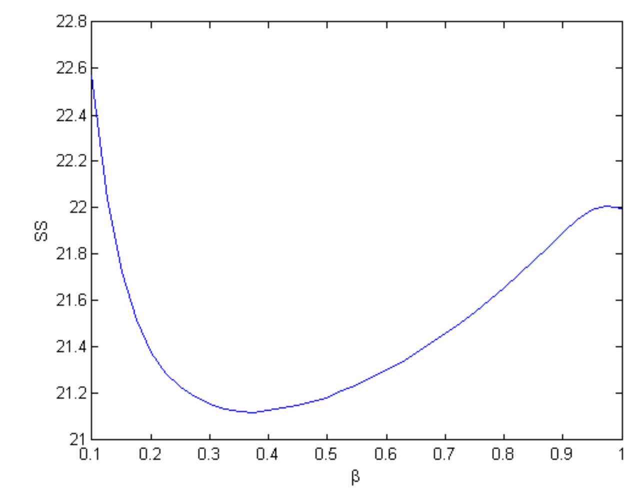

*Figure 5: Safety stock required to ensure a 95% fill rate with **endogenous** lead times*

As can be seen in figure 5, the safety stock is a U-shaped function of the smoothing intensity. We can smooth the replenishment orders to some extent without having to increase the safety stock. In our example, we can smooth up to approximately β = 0.375 while decreasing the safety stock. However, from that point on the safety stock increases exponentially. When β < 0.15, the safety stock is even larger than the safety stock when we do not smooth the order pattern. In that case the decrease in lead time cannot compensate the increase in safety stock any more.

CONCLUSIONS

Disney and Towill (2003) question "to what extent can production rates be smoothed in order to minimise production adaptation costs without adversely increasing inventory costs". This is an important trade-off because if a perfectly level production rate is used then large inventory deviations are found and hence large inventory costs are incurred. Conversely, if inventory deviations are minimised (by "passing on orders"), highly variable production schedules are generated and hence production adaptation costs are incurred. We have shown that by treating the lead time as endogenous variable, we can turn this conflicting situation into a win-win situation. A smooth order pattern gives rise to shorter and less variable lead times. This introduces a compensating effect on the retailer's inventory level. In this paper we showed that we can smooth the order pattern to a considerable extent without increasing stock levels. This may motivate the retailer to generate a smooth ordering pattern, resulting in a win-win situation for both parties.

ACKNOWLEDGMENTS

This research contribution is supported by contract grant G.0051.03 from the Research Programme of the Fund for Scientific Research - Flanders (Belgium) (F.W.O.-Vlaanderen). Benny Van Houdt is a postdoctoral Fellow of F.W.O.-Vlaanderen.

REFERENCES

Allen, D.S. (1997). Do inventories moderate fluctuations in output? *Review, Federal Reserve Bank of St. Louis* (July/August), pp 39-50.

Balakrishnan, A. , Geunes, J. and Pangburn, M. (2004). Coordinating supply chains by controlling upstream variability propagation. *Manufacturing & Service Operations Management*, 6(2), pp 163-183.

Boute, R.N., Disney, S.M., Lambrecht, M.R. and Van Houdt, B. (2005). A smoothing replenishment rule with endogenous lead times. *Working paper, Department of Applied Economics, K.U.Leuven, Belgium.*

Dejonckheere, J., Disney, S.M., Lambrecht, M.R. and Towill, D.R. (2003). Measuring and avoiding the bullwhip effect: A control theoretic approach. *European Journal of Operational Research,* 147, pp 567-590.

Disney, S.M. and Towill, D.R. (2003). On the bullwhip and inventory variance produced by an ordering policy. *Omega*, 31, pp 157-167.

Disney, S.M., Farasyn, I., Lambrecht, M.R., Towill, D.R. and Van de Velde, W. (2004). Taming the bullwhip whilst watching customer service in a single supply chain echelon. *European Journal of Operational Research*. To appear.

Disney, S.M., Farasyn, I., Lambrecht, M.R., Towill, D.R. and Van de Velde, W. (2005). Dampening variability by using smoothing replenishment rules. *Research Report 0502, Department of Applied Economics, K.U.Leuven, Belgium.*

Forrester, J. (1961). *Industrial Dynamics*. MIT Press, Cambridge MA.

Hopp, W.J. and Spearman, M.L. (2001). *Factory Physics.* 2nd edn. Irwin, McGraw-Hill.

Latouche, G. and Ramaswami, V. (1999). *Introduction to matrix analytic methods and stochastic modeling.* SIAM. Philadelphia.

Lee, H.L., Padmanabhan, V. and Whang, S. (1997a). The bullwhip effect in supply chains. *Sloan Management Review*, Spring 38(3), pp 93-102.

Lee, H.L., Padmanabhan, V. and Whang, S. (1997b). Information distortion in a supply chain: the bullwhip effect. *Management Science*, 43(4), pp 546-558.

Neuts, M. (1981). *Matrix-geometric solutions in stochastic models, an algorithmic approach.* John Hopkins University Press.

Magee, J. (1958). *Production Planning and Inventory Control.* London, McGraw-Hill.

Silver, E.A., Pyke, D.F.and Peterson, R. (1998). *Inventory management and production planning and scheduling.* John Wiley & Sons, New York.

Veinott, A.F. (1966). The status of mathematical inventory theory. *Management Science*, 12(11), pp 745-777.

Zipkin, P.H. (2000). *Foundations of inventory management*. McGraw-Hill. New York.