A win-win solution for the bullwhip problem

ROBERT N. BOUTE * [1,2], STEPHEN M. DISNEY [3],
MARC R. LAMBRECHT [2] and BENNY VAN HOUDT [4]

[1] *Operations & Technology Management Center, Vlerick Leuven Gent Management School.*
Vlamingenstraat 83, 3000 Leuven, Belgium.

[2] *Research Center for Operations Management, Katholieke Universiteit Leuven.*
Naamsestraat 69, 3000 Leuven, Belgium.

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

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

Abstract: An important supply chain research problem is the bullwhip effect where demand variability increases as one moves up the supply chain. This distorted information may lead to inefficiencies. In this paper we suggest a remedy to reduce the bullwhip effect. We focus on an inventory replenishment rule that reduces the variability of upstream orders and generates a *smooth* order pattern. However, dampening the order variability has a negative impact on customer service due to an increased inventory variance. We resolve this conflicting issue by taking the impact of the replenishment rule on lead times into account. A smooth order pattern generates shorter and less variable (production/replenishment) lead times, introducing a compensating effect on the inventory levels. We show that by including *endogenous* lead times in our analysis, the order pattern can be smoothed to a considerable extent without increasing stock levels, resulting in a win-win solution for both supply chain echelons. Finally we discuss several order smoothing approaches from an industrial perspective and comment how our results may influence these cases.

Keywords: Supply chain control, supply chain collaboration, bullwhip effect

---

* Corresponding author: robert.boute@vlerick.be

1   Introduction: the bullwhip problem

A major cause of supply chain deficiencies is the bullwhip problem, which refers to the tendency of replenishment orders to increase in variability as it moves up a supply chain. Jay Forrester (1961) was among the first researchers to describe this phenomenon, then called 'Demand Amplification'. Procter and Gamble first coined the phrase *bullwhip effect* to describe the ordering behaviour witnessed between customers and suppliers of Pampers diapers. While diapers enjoy a fairly constant consumption rate, P&G found that wholesale orders tended to fluctuate considerably over time. They observed further amplification of the oscillations of orders placed to their suppliers of raw material.

A number of researchers designed games to illustrate the bullwhip effect. The most famous game is the "Beer Distribution Game". This game has a rich history: growing out of the industrial dynamics work of Forrester and others at MIT, it is later on developed by Sterman in 1989. The Beer Game is by far the most popular simulation and the most widely used game in many business schools, supply chain electives and executive seminars. Simchi-Levi et al. (1998) developed a computerized version of the beer game, and several versions of the beer game are nowadays available, ranging from manual to computerized and even web-based versions (e.g. Chen and Samroengraja 2000, Jacobs 2000).

This bullwhip effect throughout the supply chain can lead to tremendous inefficiencies; excessive inventory investment, poor customer service, lost revenues, misguided capacity plans, ineffective transportation, and missed production schedules (Lee et al. 1997a). Lee et al. (1997b) identify five major operational causes of the bullwhip effect; the use of 'demand signal processing', non-zero lead times, order batching, supply shortages and price fluctuations. Our focus is on the issue of demand

signal processing, which refers to the practice of adjusting the parameters of the inventory replenishment rule. These *rational* adjustments may cause over-reactions to short-term fluctuations and lead to variance amplification. In other words, the replenishment rule used by the members of the chain may be a contributory factor to the bullwhip effect. Following the same line of argument it can be seen that the replenishment policy can also be used to reduce or *tame* the bullwhip effect. This is exactly what we aim to do in this contribution.

The remainder of the paper is organized as follows. In the next section we describe our model and introduce notation. In section 3 we propose a replenishment policy that is able to dampen the order variability. This reduces the bullwhip effect in an effective manner. However, as will be explained in section 3, dampening the order variability may have a negative impact on customer service. We do find a win-win solution when we include the impact of the replenishment rule on the manufacturer's lead times. This is done in section 4 where we show that a smooth order pattern generates shorter and less variable (production/replenishment) lead times, introducing a compensating effect on the safety stock. Section 5 numerically illustrates our findings. In section 6 we discuss other techniques to reduce order variability and illustrate with a practical application in industry. Section 7 concludes.

2 Model description

We consider a two echelon supply chain with a single retailer and a single manufacturer. Every period, the retailer observes customer demand, denoted by D_t, representing 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 period. The order quantity O_t is determined by the retailer's replenishment policy. We assume that the manufacturer does not hold a finished goods inventory, but instead produces on a make-to-order basis. The replenishment orders of size O_t enter the production facility where they are processed on a first-come-first-served basis. Orders that arrive at a busy production facility must wait in a queue. 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 *production* or *replenishment* lead time, denoted by T_p. A schematic of our model is shown in figure 1.
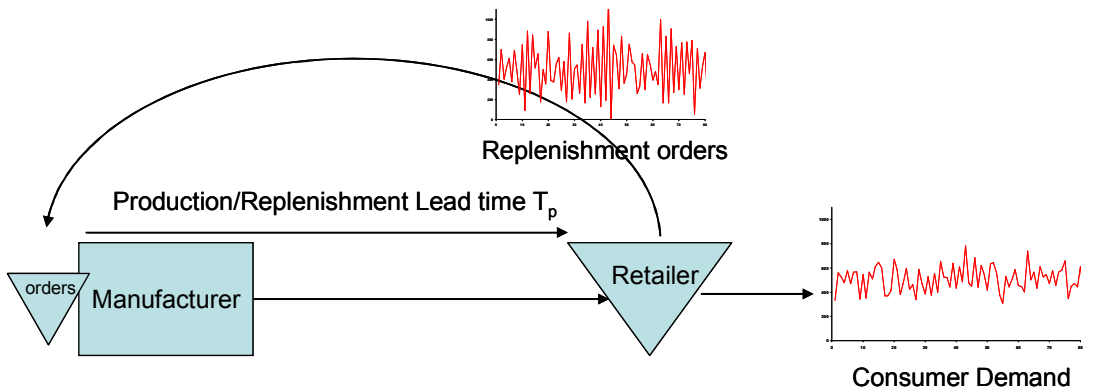

Figure 1: A two echelon supply chain modeled as a production/inventory system

3 Taming the bullwhip: order smoothing

Due to the bullwhip effect, the retailer's orders O_t to the manufacturer tend to have a larger variance than the consumer demand D_t that triggers the orders. This order variability can have large upstream cost repercussions. The upstream manufacturer aims to smooth production and therefore he prefers minimal variability in the replenishment orders from the (downstream) retailer. The manufacturer not only prefers a level production schedule, the smoothed demand also allows him to minimize his raw materials inventory cost. Balakrishnan et al. (2004) emphasize the 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* order 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 1950s and 1960s (e.g., Simon 1952, Magee 1958). The more recent work on smoothing replenishment rules can be found in Dejonckheere et al. (2003) and Balakrishnan et al. (2004).

3.1 A smoothing replenishment policy

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. 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 (items immediately available to meet demand) and the inventory on order (items ordered but not yet arrived due to the

lead time) minus the backlog (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;

$$O_t = S - IP_t. \tag{1}$$

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)). Magee (1958) and Forrester (1961) propose to order only a fraction of the inventory deficit, resulting in the following ordering policy (see also Warburton 2004):

$$O_t = \beta \cdot (S - IP_t). \tag{2}$$

Forrester (1961) refers to $1/\beta$ as the 'adjustment time' and hence explicitly acknowledges that the deficit recovery should be spread out over time.

When customer demand is i.i.d., the base-stock level S is a fixed constant. Boute et al. (2007) show that Eqn. (2) gives rise to an autocorrelated order pattern, given by

$$O_t = (1-\beta) \cdot O_{t-1} + \beta \cdot D_t. \tag{3}$$

Hence, the retailer's replenishment orders are not statistically independent, because from Eq. (3) we derive that the correlation between the orders is equal to

$$corr(O_t, O_{t-x}) = (1 - \beta)^x. \tag{4}$$

Moreover, Boute et al. (2007) demonstrate that the base-stock level S in Eqn. (2) 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

$$S = SS + [E(T_p) + 1] \cdot E(D) + (1-\beta)/\beta \cdot E(D), \qquad (5)$$

where SS denotes the safety stock and $E(T_p)$ and $E(D)$ represent resp. the average lead time and average demand.

It is notable that the replenishment rule described by Eqn. (3) 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 this case given by

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

Hence, if $\beta = 1$, these expressions reduce to the standard base-stock policy, where $O_t = D_t$; we chase sales and thus there is no variance amplification. For $1 < \beta < 2$ we create bullwhip, i.e. the order variance is amplified compared to the demand variance. This tendency is often observed in reality, or when playing the Beer Distribution Game (Dejonckheere et al. 2003). For $0 < \beta < 1$ we find that this replenishment rule generates a smooth replenishment pattern, i.e. it dampens the order variability. Under a fixed lead

time assumption such a smoothing 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 the production capacity is fixed and lead times result from a single server queueing system (as in the model described in this paper), this replenishment rule enables to smooth the manufacturer's production, resulting in shorter order-to-delivery times and more balanced, peak shaving production schedules, which are beneficial for the manufacturer. Besides the benefits realized through a smoother planning, the manufacturer also realizes cost savings on its own raw materials and/or component inventories (see Balakrishnan et al. 2004).

3.2 Is smooth smart?

Since the bullwhip effect has a number of highly undesirable cost implications, taming the bullwhip, or dampening the order variability, seems to be a dominating operations strategy. We have to be careful not to focus only on one side of the production smoothing 'coin' however. In developing a replenishment rule one has to consider the impact on the inventory variance as well. The manufacturer does benefit from smooth production, but dampening variability in orders may have a negative impact on the retailer's customer service due to inventory variance increases (Bertrand 1986, Disney and Towill 2003).

Disney et al. (2006) 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. As a consequence, 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. This leads to a tension between the preferred order variability of retailers and manufacturers.

However, we can model a two echelon supply chain as a production-inventory system, as illustrated in Figure 1. This implies that a replenishment order generated by the retailer's inventory results in an arrival of a production order at the manufacturer. Hence the choice of the retailer's replenishment policy (amplifying or dampening customer demand variability in the replenishment orders) determines the arrival process of production orders at the manufacturer's production queue and as such it affects the distribution of the production lead times. According to the laws of factory physics (Hopp and Spearman 2001), a smooth order pattern will give rise to shorter and less variable lead times. This in turn exercises a downward effect on the retailer's inventory level, which may compensate the increase in inventory variance. The quest for a win-win solution (smooth production *and* lower inventory levels) is the topic of the next section.

4   In search of a win-win solution

4.1   Impact of order variance dampening on lead times

Most inventory models proposed in the literature take the replenishment lead time T_p as a fixed constant or as an exogenous variable with a given probability distribution (for example see Kim et al. 2006). However, the replenishment orders do in fact load the production facilities. The nature of this loading process relative to the available capacity and the variability it creates are the primary determinants of lead times in the production 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 (Hopp and Spearman 2001).
Zipkin (2000, p.246) states: "to understand the overall inventory system, we need to understand the supply system. For this purpose we can and do apply the results of queueing theory".

It is essential to extend pure inventory systems with *exogenous* lead times to production-inventory systems with *endogenous* lead times. After all, inventory influences production by initiating orders, and production influences inventory by completing and delivering those orders to inventory. In Figure 2 the interaction between the retailer's replenishment policy and the manufacturer's production system is illustrated: the replenishment orders constitute the arrival process at the manufacturer's queue. The time until the order is produced (the sojourn time in the queueing system) is the time to replenish the order. This replenishment lead time is a prime determinant in setting the safety stock requirements for the retailer.
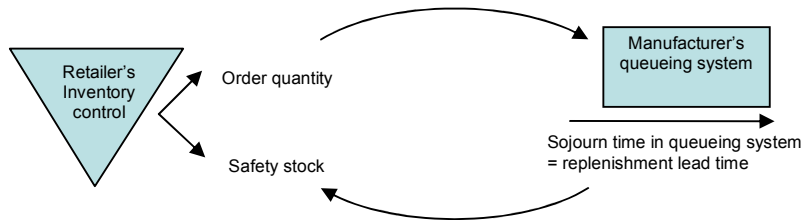
Figure 2: Interaction between retailer's inventory system and manufacturer's production system

To estimate the lead time distribution we develop a discrete time queueing model. By analyzing the characteristics of the replenishment orders, we implicitly analyse the characteristics of the production orders that arrive to the production system. 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. The supply system is *bulk queue*, which tends to be difficult to analyse (Chaudry and

Templeton 1983). Moreover as we can see from Eqn. (3), the batch sizes generated by our smoothing rule are not i.i.d., rather they are autocorrelated. Therefore the resulting queueing model is substantially different from the M/M/1 make-to-stock queue, as considered by, for example, Karaesmen et al. (2004).

The analysis of our queueing model can be solved using *matrix analytic methods* (Neuts 1981, Latouche and Ramaswami 1999). These methods are popular as modeling tools because they can be used to construct and analyse a wide class of stochastic models. They are applied in several areas, of which the performance analysis of telecommunication systems is one of the most notable (Latouche and Ramaswami 1999). In a separate paper, the authors of this paper discuss the solution procedure of this queueing model (see Boute et al. 2007). The results confirm our expectation that a smooth order pattern generates shorter and less variable lead times.

4.2   Resulting 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 from 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, more safety stock is required to buffer the difference between usage and supply. A reduction of order variations comes with the cost of an increase in inventory variability (Bertrand 1986).

When lead times are endogenously determined, however, dampening variability in orders affects the replenishment lead time distribution as well. A smooth order pattern generates shorter and less variable lead times, introducing a compensating effect on the required safety stock. The aim is to find values for the smoothing parameter $0 < \beta < 1$

11

where the decrease in lead times compensates the increase in inventory variance. In that case we can smooth production without having to increase inventory levels to provide the same customer service.

To do so, we characterize the inventory random variable and use it to find the safety stock requirements for the system. Since the inventory is controlled by stochastic lead times, the inventory is not necessarily replenished every period and we do not know exactly when a replenishment occurs. Moreover, the queueing analysis implies that it takes a longer time to produce (and consequently replenish) a larger order quantity. Hence the order quantity and its replenishment lead time are correlated, which has an impact on the calculation of the inventory distribution. Therefore, if we want to determine the inventory distribution and the corresponding safety stock requirements in an exact way, we need to take this correlation into account.

We measure customer service with the fill rate, which is the proportion of the demand that can immediately be delivered from the inventory on hand (Zipkin 2000)

$$\text{Fill rate} = 1 - \frac{\text{expected number of backorders}}{\text{expected demand}}. \qquad (7)$$

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

The inventory on hand or net stock NS_t is then equal to the initial inventory on hand plus all replenishment orders received so far minus total observed customer demand. At the end of period t, the order O_{t-k} is in service, and the orders placed more than k periods ago, i.e., O_{t-i}, $i \geq k+1$, are already delivered in inventory, while customer demand is satisfied up to the current period t. Assuming the initial inventory level is equal to the base stock level S, the net stock after satisfying demand in period t is equal to

$$NS_t = S + \sum_{i=k+1}^{t-1} O_{t-i} - \sum_{i=0}^{t-1} D_{t-i} \; . \qquad (8)$$

Substituting (3) into (8), we obtain

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

Boute et al. (2007) evaluate the steady state distribution of NS_t. Some care must be taken when evaluating (9), however, as the value of D_{t-k} influences the age k of the order in service: 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 customer demand, the demand terms D_{t-i}, $i=k+1,...,t$ also influence the order's age k.

From the steady state distribution of the inventory variable NS, we can easily deduct the expected number of backorders $E(NS^-)$, where $NS^- = \max\{0,-NS\}$, and the corresponding fill rate realised with a given base-stock level S. In practice, decision makers often determine the minimal base-stock level that is required to achieve a target fill rate. From this base-stock level S, we then find the corresponding safety stock using Eqn. (5).

4.3 The bullwhip effect results in a lose-lose situation

Note that the discussion above considers the situation where we smooth the replenishment orders, which implies a replenishment parameter β smaller than one in Eqn. (2). We may extend the analysis, however, to the case where the replenishment parameter β is larger than one, which implies an overreaction to the inventory deficit. This policy is often observed in reality and leads to order variance amplification, or equivalently, induces the bullwhip effect.

When $1 < \beta < 2$, the order pattern generated by the replenishment policy (2) is negatively correlated and it may generate negative order quantities. Since in our model it is not possible to send negative orders to production, we have to preclude the possibility of negative orders. The following restriction on β given the minimum demand D_{min} and the maximum demand D_{max} ensures that $O_t \geq 1$ (we refer to Boute (2006) for a proof):

$$D_{min} + (1 - \beta) \cdot D_{max} \geq 2 - \beta. \tag{10}$$

What is the impact of the bullwhip effect on the performance in the supply chain? First of all, Disney et al. (2006) prove that the inventory variance increases as we either smooth the order pattern ($\beta < 1$) or amplify the orders ($\beta > 1$), compared to a pure chase sales policy where $\beta = 1$. This increased inventory variability inflates the safety stock requirements at the retailer.

Moreover, this replenishment decision has an impact on the distribution of the lead times. More specifically, order variance amplification increases the variability at the production queue, which increases the lead times as a consequence. This leads to higher safety stocks. In other words, the bullwhip effect leads to an increased inventory

variance, and additionally, it generates longer lead times, reinforcing the inflated safety stock requirements. This is clearly a lose-lose situation.

5   Numerical example

To illustrate our findings, we set up a numerical experiment where a retailer observes an i.i.d. random customer demand on a daily basis of between 11 and 30 products with an average of 20.5 units per day and a coefficient of variation of 0.135. The retailer satisfies this demand from his inventory on hand and replenishes with the smoothing replenishment rule given by Eqn. (2). We assume that the manufacturer's production operates 24 hours per day and the production time of a single unit is geometrically distributed with an average of 64 minutes per unit. Hence the average production load is $(20.5 \times 64) / (24 \times 60) = 0.91$.

The retailer has to determine the parameter β to control his inventory. When $\beta = 1$, the retailer places orders equal to demand and hence the variability in demand is transmitted to the manufacturer. This policy results in an average lead time of 0.67 periods and a variance of 0.44. The safety stock required to provide a 98% fill rate is equal to 36.95 units.

Suppose that the retailer chooses to smooth his orders with a parameter $\beta = 0.4$. This results in an order pattern which is four times less variable ($\text{Var}(O)/\text{Var}(D) = 0.4/(2-0.4) = 0.25$). When we maintain the same lead time distribution, 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. This is clearly a *win-lose* situation: the manufacturer can smooth production, but at the expense of an increase in the retailer's inventory.

However, working with the same lead times, is incomplete. When the retailer smoothes his orders, he sends a less variable pattern to the manufacturer. 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 $\beta = 0.4$ to the manufacturer's production, we observe a lower and less variable lead time distribution. The average lead time decreases to 0.49 and its variance equals 0.36. This in turn introduces a compensating effect on the required safety stock. We find that a safety stock of 36.41 is sufficient to provide a 98% fill rate, which is slightly lower than when we do not smooth the orders ($\beta = 1$).
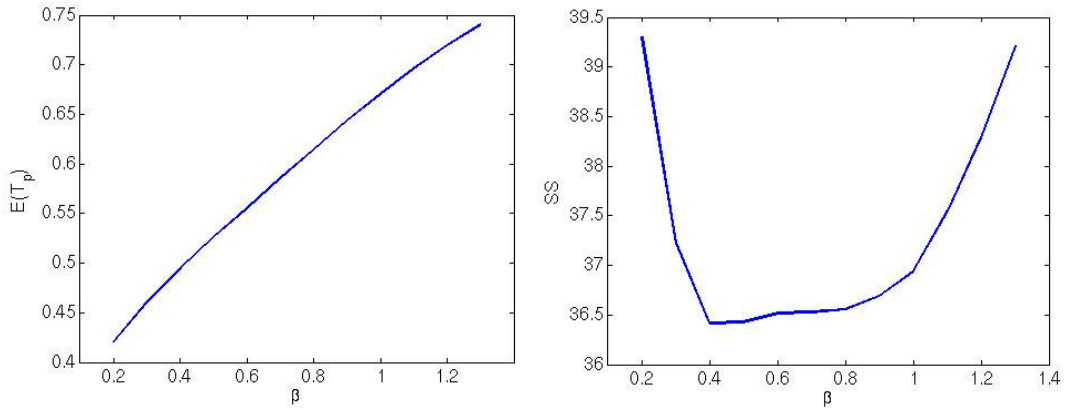

Figure 3 - left: Average lead times in function of the replenishment parameter β
right: Safety stock required to ensure a 98% fill rate with endogenous lead times

In Figure 3 we show the effect of order smoothing on the (average) lead times and safety stocks for a smoothing parameter $\beta = 0.2$ to $\beta = 1.3$. As β decreases, the average lead times decrease as well (Figure 3 – left). This is intuitively clear, as the order variability decreases with a smaller β, leading to a less variable arrival pattern at the queue and consequently decreasing lead times. When we include this effect of order smoothing on lead times, the safety stock becomes a U-shaped function of the smoothing intensity. We can smooth the replenishment orders to some extent without having to increase the safety stock, whilst maintaining customer service at the same

16

target level. Moreover, we can even decrease our safety stock when we smooth the order pattern (up to $\beta = 0.35$).

As such we may obtain 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. This Pareto-improving policy may require contractual arrangements between the supply chain partners so that the lead time reduction is effectively implemented (Tsay 1999).

However, as of a certain point (around $\beta = 0.4$) the safety stock increases sharply. When β approaches zero, the lead time reduction cannot compensate the increase in inventory variability anymore and the safety stock exceeds the safety stock that is required when the orders are not smoothed ($\beta = 1$).

When $\beta > 1$, we observe that lead times increase further together with the safety stocks. Obviously, this is a *lose-lose* situation and needs to be avoided.

This numerical example well illustrates the dynamics resulting from the retailer's inventory decision on the lead times and safety stocks. Obviously, the degree to which we should smooth and the exact amount of safety stock decrease depend on the observed demand pattern.

6   Some practical examples of reducing order variability

Order smoothing combined with endogenous lead times may create a win-win situation for both the retailer and the manufacturer. In order to effectively implement such a policy, the supply chain partners have to align their replenishment policies, i.e. the type of replenishment rule used and the setting of the 'best' parameter value (β). It is important to notice that such a strategy goes far beyond 'information sharing'. In a practical setting, however, other coordination schemes may be used to achieve the same

objective. We therefore briefly discuss a range of other order variance reduction tools and add real life examples where applicable. An excellent overview can be found in Holweg et al. (2005).

In a traditional supply chain, each level in the supply chain issues production orders and replenishes stock without considering the situation at either up- or downstream tiers of the supply chain. This is how most supply chains still operate, no formal collaboration between the retailer and supplier. Collaboration on the other hand can be installed through a wide range of concepts such as Collaborative Forecasting Planning and Replenishment (CPFR), Information Sharing, Vendor Managed Inventory (VMI, including Continuous Replenishment). A more drastic solution can be obtained by a redesign of the supply chain by eliminating echelons. Let us first focus on VMI.

VMI eliminates one decision point and merges the replenishment decision with the production and materials planning of the supplier. Here, the supplier takes charge of the customer's inventory replenishment on the operational level, and uses this visibility in planning his own supply operations (e.g. more efficient production schedules and transportation planning). With VMI, multi-echelon supply chains can act in the same way, dynamically, as a single echelon of a supply chain. VMI often results in more frequent replenishments and consequently the order quantity variance is reduced. Economies in transportation can be obtained through an optimization of the route planning and with methods such as joint replenishment and inventory routing techniques. VMI is in other words an alternative to the smoothing replenishment policy proposed in this paper.

We discuss two cases to illustrate the benefits of dampening the order variance. First, we analyse the ordering pattern of a bakery company focusing on authentic specialties in the biscuit and cake world: caramelized biscuits, waffles, frangipane, and

cake specialties among others. For certain products, a make-to-order policy is employed and the assumptions used in this paper are largely satisfied. In 2002, the firm introduced a VMI program implemented in the SAP software, referred to as "Customer Replenishment Planning" (CRP). In Figure 4 we show a graph of the shipments from the production facility to the distribution centre of a retailer (for one specific product) in the pre-CRP period (2001-mid 2002) and the shipments in the post-CRP period (mid 2002-2005). The coefficient of variation of the shipment quantities went down from 1.14 to 0.45 (a number observed for other products as well). We were also able to collect (post-CRP) data on the shipments from the distribution centre of the retailer to the different retail outlets. For the specific product discussed above, we obtain a coefficient of variation of 0.40. The company now benefits from a higher flexibility in its production planning and reduced its transportation costs considerably. Moreover, inventories decreased, improving the freshness of the products of the end consumer.
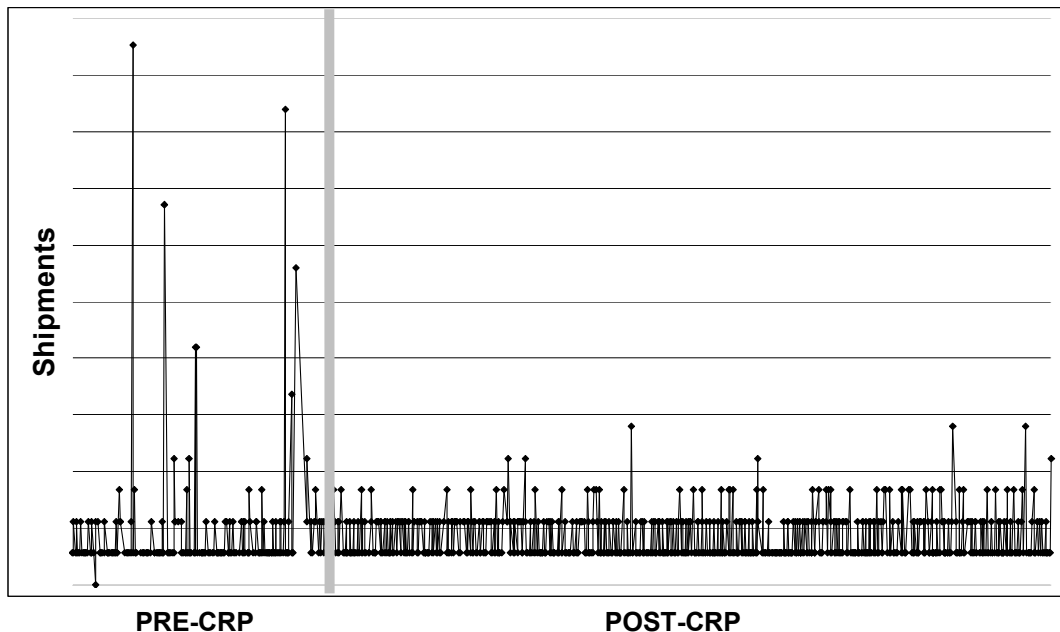

PRE-CRP   **POST-CRP**

Figure 4: The impact of VMI on the order variability for a selected product

19

The second case deals with a more traditional example of order smoothing of a UK grocery retailer. Here we were asked to identify and reduce the cause of workload variability in their own warehousing (cross docking) and transportation activities. We discovered that the replenishment algorithms that maintained stock levels at individual stores were the source of a bullwhip effect. There were several different replenishment algorithms in use, and we where able to introduce a proportional controller into half of them. These modified algorithms controlled 65% of the sales volume, but only 35% of product lines. In general these were the higher volume products. Figure 5 illustrates the performance of the system with a "before" and "after" simulation of a single product using real demand data for a single product from a single store. The company had identified that this modification had allowed a very significant reduction in contract staff in distribution centre and amount of third party logistics costs to meet the peak demand on certain days of the week.
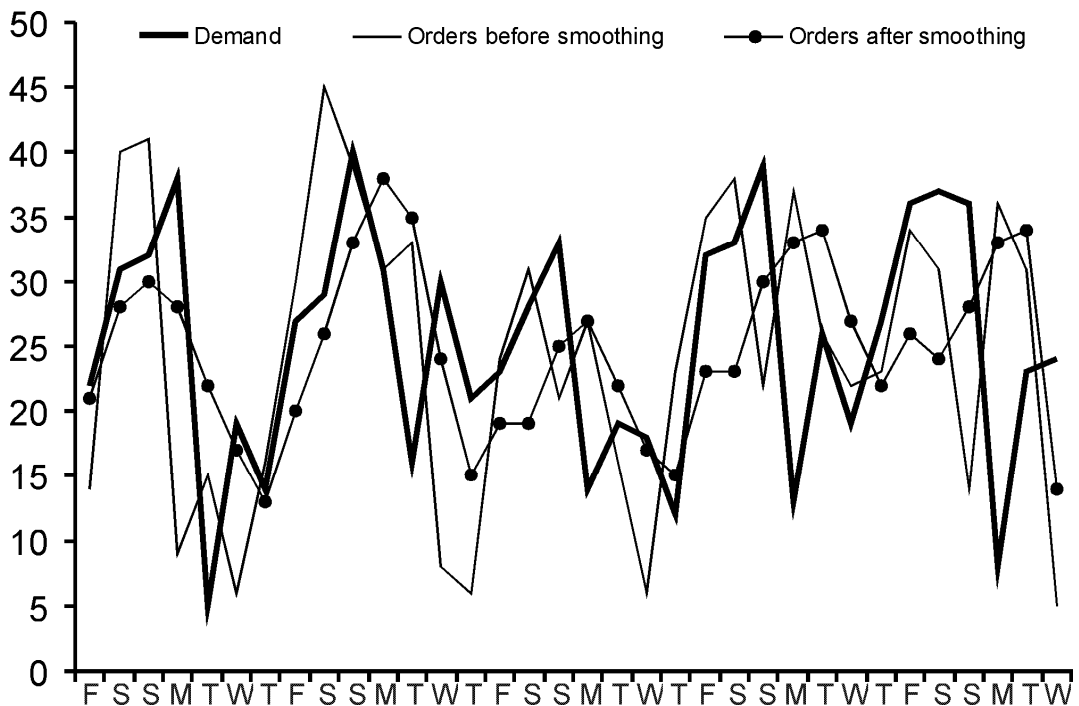

Figure 5: Smoothing in action in the UK grocery industry

In order to achieve this, the grocery retailer had accepted a slight increase in the target safety stock in their stores. That is, they assumed exogenous lead-times. But, in effect, that is all they could possibly do anyway, as they were ordering on day 1 for delivery in day 2. This meant that the suppliers had to keep a stock of finished goods. Thus if the suppliers maintained this stock with a production system that operated as a queue (that is endogenous lead-times exist), then the retailer smoothing actions will have a beneficial effect on the supplier finished goods. The retailer may still gain by this – the manufacturer may be more willing to accept on-going calls for cost reductions. This clearly illustrates the power of variance reduction techniques be it through VMI programs or smoothing replenishment policies.

7 Conclusions

The bullwhip problem has been studied by many authors in recent years. Since the bullwhip effect has a number of highly undesirable cost implications, taming the bullwhip is a dominating operations strategy. Conventional bullwhip reduction is only one side of the coin, however. In developing a replenishment rule one has to consider the impact on the inventory variance as well. More specifically, dampening the variability in orders inflates the safety stock requirements due to the increased variance of the inventory levels. As a consequence, 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.

We have shown that by treating the lead time as an 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 solution for both supply chain echelons. We also highlight alternative mechanisms to achieve the same objectives.

8   References

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.

Bertrand, J.W.M., 1986. Balancing production level variations and inventory variations in complex production systems. *International Journal of Production Research*, **24**(5), pp 1059-1074.

Boute, R.N., 2006. *The Impact of Replenishment Rules with Endogenous Lead Times on Supply Chain Performance.* Ph.D. Thesis. K.U.Leuven.

Boute, R.N., Disney, S.M., Lambrecht, M.R. and Van Houdt, B., 2007. An integrated production and inventory model to dampen upstream demand variability in the supply chain. *European Journal of Operational Research*, **178**(1), pp 121-142.

Chen, F., Samroengraja, R., 2000. The stationary Beer Game. *Production and Operations Management*, **9** (1), 19-30.

Chaudry, M. and Templeton, J., 1983. *A first course in bulk queues.* Wiley, New York.

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., 2006. Taming the bullwhip whilst watching customer service in a single supply chain echelon. *European Journal of Operational Research* **173**, pp 151-172.

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

Holweg, M., Disney, S., Holmström, J. and Smaros, J., 2005. Supply Chain collaboration : Making Sense of the Strategy Continuum, *European Management Journal,* **23**, pp.170-181.

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

Karaesmen, F., Liberopoulos, G. and Dallery, Y., 2004. The value of advanced demand information in production/inventory systems. *Annals of Operations Research*, **126**, pp 135-157.

Jacobs, F.R., 2000. Playing the Beer Distribution Game over the Internet. *Production and Operations Management*, **9** (1), 31-39.

Kim, J.G., Chatfield, D.C., Harrison, T.P. and Hayya, J.C., 2006. Quantifying the bullwhip effect in a supply chain with stochastic lead time. *European Journal of Operational Research*, **173**, pp 617-636.

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.

Magee, J.F., 1958. *Production planning and inventory control*. McGraw-Hill. New York.

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

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

Simchi-Levi, D., Kaminsky, P. and Simchi-Levi, E., 1998. *Designing and Managing the Supply Chain.* Irwin/McGraw-Hill, New York.

Simon, H., 1952. On the application of servomechanism theory in the study of production control. *Econometrica*, 20, pp 247-268.

Sterman, J., 1989. Modeling managerial behaviour: misperceptions of feedback in a dynamic decision making experiment. *Management Science*, **35** (3), 321-339.

Tsay, A.A., Nahmias, S. and Agrawal N., 1999. Modeling supply chain contracts: A review. In *Quantitative Models for Supply Chain Management* (S. Tayur, R. Ganeshan, and M. Magazine (Ed)), pp 301-336. Kluwer Academic Publishers.

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

Warburton, R.D.H., 2004. An analytical investigation of the bullwhip effect. *Production and Operations Management*, **13**(2), pp 150-160.

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

9   Biography

ROBERT BOUTE holds an MSc degree in Commercial Engineering and a PhD degree in Applied Economics from the KULeuven (Belgium). Currently, Robert is Assistant Professor Operations Management at the Vlerick Leuven Gent Management School. His research interests are in the field of Supply Chain Management and the interaction between inventory control and production management.

STEPHEN DISNEY is a Senior Lecturer of Operations Management with the Logistics Systems Dynamics Group in the Logistics and Operations Management section of Cardiff Business School (UK). Stephen's current research interests involve the application of control theory and statistical techniques to supply chains and e-business scenarios in order to investigate their dynamic and economic performance.

MARC LAMBRECHT is a Full Professor at KULeuven (Belgium). He teaches courses in manufacturing systems analysis and inventory management with a focus on stochastic aspects of operations. He was the director of the KULeuven MBA Program for six years and chairman of the business school for four years.

BENNY VAN HOUDT holds a M.Sc. degree in Mathematics and Computer Science, and a PhD in Science from the University of Antwerp (Belgium). Currently he is Assistant Professor at the University of Antwerp. His main research interest is the performance evaluation and stochastic modeling of telecommunication networks.