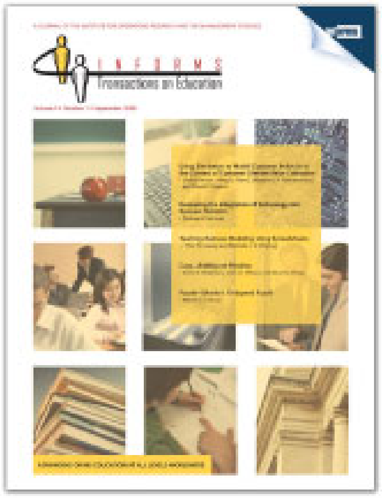

INFORMS Transactions on Education

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Exploring the Bullwhip Effect by Means of Spreadsheet Simulation

Robert N. Boute, Marc R. Lambrecht,

Vol. 10, No. 1, September 2009, pp. 1–9
ISSN 1532-0545 | 09 | 1001 | 0001

INFORMS Transactions on Education

informs®

DOI 10.1287/ited.1090.0038

Exploring the Bullwhip Effect by Means of Spreadsheet Simulation

Robert N. Boute
Operations and Technology Management Center, Vlerick Leuven Gent Management School, 3000 Leuven, Belgium, and Research Center for Operations Management, Katholieke Universiteit Leuven, 3000 Leuven, Belgium,
robert.boute@vlerick.be

Marc R. Lambrecht
Research Center for Operations Management, Katholieke Universiteit Leuven, 3000 Leuven, Belgium,
marc.lambrecht@ecom.kuleuven.be

One of the main supply chain deficiencies is the bullwhip effect: Demand fluctuations increase as one moves up the supply chain from retailer to manufacturer. The beer distribution game is widely known for illustrating these supply chain dynamics in class. In this paper we present a spreadsheet exploring the two key causes of the bullwhip effect: demand forecasting and the type of ordering policy. We restrict our attention to a single product two-echelon system and illustrate how tuning the parameters of the replenishment policy induces or reduces the bullwhip effect. We also demonstrate how bullwhip reduction (dampening the order variability) may have an adverse impact on inventory holdings or customer service. The spreadsheet can be used to help students gain insight into how inventory control policies and forecasting influence the magnitude of the bullwhip effect and the quality of customer service.

Key words: bullwhip effect; replenishment rules; forecasting techniques; spreadsheet simulation; beer distribution game
History: Received: January 2007; accepted: July 2009.

1. Introduction: Teaching Students About the Bullwhip Effect

The bullwhip effect is a well-known phenomenon in supply chains. In a simple, linear supply chain that consists of a manufacturer, a distributor, a wholesaler, and a retailer, we observe that the retailer's orders to the wholesaler display greater variability than the end-consumer sales, the wholesaler's orders to its distributor show even more oscillation, and the distributor's orders to the manufacturer are most volatile.

The bullwhip effect and its dynamics are often illustrated in class with the "beer distribution game" (beer game) developed at MIT (Sterman 1989). It is a popular simulation game and probably the most widely used game in business schools, supply chain electives, and executive seminars. Simchi-Levi et al. (2003) developed a computerized version of the beer game, and several other versions are available, ranging from manual to computerized to Web-based versions (e.g., Machuca and Barajas 1997, Chen and Samroengraja 2000, Jacobs 2000; see Wood 2007 for an overview).

Beyond the beer game, case studies are used as teaching tools to introduce the bullwhip effect, e.g., Barilla SpA (Hammond 1994), a major pasta producer in Italy, and Campbell Soup's chicken noodle soup experience (Cachon and Fisher 1997).

In this paper we explore the two key causes of the bullwhip effect: demand forecasting and the type of replenishment policy (Lee et al. 1997a). Many studies investigate the adverse effects of demand signaling and improper forecasting. Watson and Zheng (2008) use formal models to address managers' overreactions to demand changes and the misuse of forecasting. Lee et al. (1997b) prove mathematically that variance amplification occurs if the retailer adjusts his orders based on demand signals. Dejonckheere et al. (2003) and Chen et al. (2000) demonstrate that the exponential smoothing and moving average forecasting methods always lead to bullwhip, independent of the demand pattern. A recent overview of the related literature is provided by Disney and Lambrecht (2008).

When explaining these concepts in class, it is desirable to provide students with useful insights about the causes of the bullwhip effect, omitting (or at least limiting) the involved complex mathematics. To meet this challenge we developed a user-friendly and easy-to-understand spreadsheet using Microsoft Excel. Spreadsheets have been used, for example, by Munson et al. (2003) to teach the cost of uncoordinated supply chains. Our spreadsheet allows students to explore various base stock (order-up-to) replenishment policies and forecasting methods and their impact under various demand processes. We include relevant analytical results within the spreadsheets, thus making them easily available to students. We have found the spreadsheets useful in core operations management courses at both the undergraduate and MBA level and in supply chain electives.

This paper has three objectives: (1) to help students obtain insights into bullwhip dynamics via basic spreadsheet calculations; (2) to make all relevant results from the literature available in one tool; and (3) to go beyond existing analytical results by using simulation analysis. The spreadsheet models guide students through a complicated interplay between order fluctuations, inventory fluctuations, and customer service under a variety of demand process and forecasting technique scenarios. One can easily evaluate the impact of various replenishment strategies: What often appears to be a rational policy of the decision maker may create tremendous order amplification. On the other hand, reducing the bullwhip effect may hurt customer service.

Our spreadsheet tool differs from existing models (e.g., Simchi-Levi et al. 2003) in several ways. We bring several "demand signal processing" methods together in a single spreadsheet, ranging from the early work by Lee et al. (1997a), to the traditional (moving average and exponential smoothing) forecasting methods towards the more complex mean squared error forecasting method. In addition, we extend the standard order-up-to replenishment policy to a generalized (or "smoothed") order-up-to policy, which can dampen order variability for any demand process. Finally, we consider both inventory-related costs and production-switching costs as performance measures. Both of these measures are affected by the replenishment rule, and therefore both of them should be analyzed.

In the next section we present the spreadsheet model. Section 3 analyzes the impact of the standard order-up-to policy with different forecasting techniques on the bullwhip effect. Section 4 describes a smoothed order-up-to policy and its impact on customer service. The spreadsheet, a student (user) manual, and an instructor manual including the detailed mathematics can be downloaded from the ITE website at http://ite.pubs.informs.org/.

2. Description and Use of the Spreadsheet Model

The spreadsheet is designed to illustrate the ordering dynamics between two supply chain partners. We have used it in debriefing sessions after playing the beer game and separately to illustrate the impact of the order-up-to replenishment policy in a supply chain context. The spreadsheet is particularly useful if the students have already covered basic inventory management techniques, including periodic review policies, where a variable amount of product is ordered with a fixed interval between orders (e.g., daily or weekly), as opposed to continuous-review economic order quantity (EOQ) policies, where a fixed amount of product is ordered with variable time intervals between orders.

We now outline an 80–90 minute lecture that uses the spreadsheets to demonstrate the impact of tuning the replenishment policy on supply chain performance. It is advisable to play the beer game first, although it can also be used in a class where the students never play the beer game. One may begin the lecture by briefly reviewing the periodic review order-up-to policy. If this technique has not been covered yet, the instructor may spend some time on it, as this policy is common practice in retailing and is optimal when there is no fixed ordering cost and both holding and shortage costs are proportional to the volume of on-hand inventory or shortage (Zipkin 2000). These assumptions hold in many practical cases, as well as in the standard setup of the beer game.

Once it is clear how this ordering policy works, the instructor may guide the students through the simulation table for one or two periods (see §2.1) and explain how this method simulates random demand and calculates the orders according to the chosen replenishment rule. The remainder of the lecture is then devoted to analyzing the impact of tuning the parameters of the replenishment rule (see §2.2) on the resulting ordering pattern and on total supply chain performance (§2.3). It is not necessary to go through the simulation table after each run, but the students should know they can easily check the outcome by going through the same calculations. We suggest a storyline at the end of each of the following sections, depending on exactly what the instructor wants to cover in class (§§3.6 and 4.5).

The following has worked well. The instructor asks the students to recapitulate the periodic review order-up-to technique at home and to simulate a number of scenarios before class. The same sequence of scenarios can be used as described in §§3.6 and 4.5. In class, the instructor may spend time discussing the students' findings, the impact of tuning the parameters, and the rationale behind the results.

In the remainder of this section we discuss (1) the simulation table, (2) parameter selection (input section), and (3) performance measurement (output section).

2.1. Simulation Table

Our model follows the standard beer game setup (Sterman 1989), with the following sequence of events in each period:

(1) Incoming shipments from the upstream supplier are received and placed in inventory. Assuming that the supplier has ample stock, these shipments correspond to the order placed $T_p + 1$ periods ago, where T_p is the deterministic transportation delay and there is a 1-period ordering delay;

(2) Customer demand is observed and either fulfilled (if enough inventory is available) or backlogged. A positive net stock represents inventory immediately available to meet demand, whereas a negative net stock refers to a backlog (demand that could not be fulfilled and still has to be delivered). The pipeline inventory represents the items ordered that have not yet arrived.

(3) A new order is placed to raise the inventory position to the order-up-to level. The inventory position is the sum of the net stock and the pipeline inventory.

$$\text{order quantity} = \text{order-up-to level} - \text{inventory position}. \qquad (1)$$

In the simulation table one can track how these order quantities are generated. The instructor manual provides the exact mathematics behind the calculations. In the classroom it is sufficient to provide a screenshot to illustrate the calculations for a few periods, as in Figure 1. Note that the simulation table also contains the forecast of next period's demand (discussed in §3), which is needed to calculate the order-up-to level.

The inventory costs consist of a holding cost per unit in inventory (when net stock is positive) and a shortage cost per unit backlogged (when net stock is negative). The production switching cost is incurred, for changing the level of production in a period. Assuming the production level is equal to the order quantity placed, the change in production is given by the difference in order quantity versus the previous period. Remark that this switching cost is not explicitly included in the beer game, but we use it to measure the impact of the order pattern on production costs and make trade-offs with the inventory related costs.

2.2. Parameter Selection

In the input section, users define the parameters of the customer demand process and the forecasting technique. The cells for parameters that can be changed are shaded. We protected certain cells with calculations in order to prevent accidental changes. The protection can be removed using the Unprotect Sheet command (Excel 2003: *Tools* menu, *Protection* submenu; Excel 2007: from the *Ribbon*, select the *Review* command tab). We refer to the student manual for a description of how to input the parameters and to the instructor manual for the mathematics behind the input section.

2.3. Performance Measurement

We define three types of performance measures for the simulation analysis: (1) the variance amplification ratios "bullwhip effect" and "net stock amplification"; (2) the average inventory and switching costs per period; and (3) the customer service measures "cycle service level" and "fill rate."

(1) We define the bullwhip effect as:

$$\text{Bullwhip} = \frac{\text{Variance of orders}}{\text{Variance of demand}}.$$

Bullwhip $= 1$ implies that the order variance is equal to the demand variance, or in other words, there is no variance amplification. Bullwhip > 1 indicates that the bullwhip effect is present (*amplification*), whereas bullwhip < 1 is referred to as a "smoothing" or "dampening" scenario, meaning that the orders are less variable than the demand.

We focus not only on bullwhip but also on the variance of the net stock, because it has a significant impact on customer service (the higher the variance of net stock, the more safety stock required). We measure

Figure 1 **Spreadsheet Example of a Standard Order-Up-To Policy with $T_p = 2$**

Period	Receive	Demand	Net stock	Pipeline inventory	Demand forecast	Order-up-to level	Order quantity	Inventory costs	Switching costs
10	108	110	19	225	110.36	353.60	110	9.50	6.00
11	112	113	18	223	110.89	355.18	114	9.00	8.00
12	113	122	9	224	113.11	361.85	129	4.50	30.00
13	110	120	−1	243	114.49	365.98	124	20.00	10.00
14	114	119	−6	253	115.39	368.69	122	120.00	4.00
15	129	117	6	246	115.71	369.66	118	3.00	8.00
16	124	120	10	240	116.57	372.23	122	5.00	8.00

the net stock variance amplification (NSAmp) as:

$$\text{NSAmp} = \frac{\text{Variance of net stock}}{\text{Variance of demand}}.$$

(2) The inventory and switching costs are related to the variance amplification ratios. High bullwhip implies wildly fluctuating orders, meaning that the production level has to change frequently, resulting in higher average production switching cost per period. High NSAmp results in high holding and backlog costs.

(3) We compute simulated customer service measures. The cycle service level refers to the probability that there will be no stock out within a period. The fill rate measures the fraction of total demand that is immediately fulfilled from the inventory on hand.

3. Impact of Forecasting on the Bullwhip Effect

We start our discussion with the standard order-up-to policy (in the "*standard OUT*" worksheet): we place an order equal to the difference between the order-up-to level and the inventory position (see also Equation (1)). According to standard inventory theory, the order-up-to level, which we denote S_t, covers the forecasted average demand $\widehat{D}_t^L$ during the protection interval L, as well as a safety stock SS:

$$S_t = \widehat{D}_t^L + \text{SS}. \tag{2}$$

The protection interval L equals the physical lead time plus the review period. The safety stock is sometimes expressed as a multiple z of the standard deviation of demand during the protection interval. We now review several forecasting techniques and illustrate their impact on the bullwhip effect by means of the spreadsheets.

3.1. Mean Demand Forecasting

If the decision maker knows that the demand is i.i.d., the best possible forecast of all future demands is simply the long-term average demand, $\bar{D}$. As a consequence, the forecasted lead time demand equals $\widehat{D}_t^L = L\bar{D}$, and the order-up-to level S_t given by Equation (2) remains constant over time. Hence Equation (1) becomes

$$O_t = S_t - (S_{t-1} - D_t) = D_t. \tag{3}$$

We simply place an order equal to the observed demand; we call this policy the "chase sales policy." In this setting, the variability of the replenishment orders is exactly the same as the variability of the original demand, and the bullwhip effect does not occur.

By selecting the "mean demand forecasting" technique in the spreadsheet, the user can observe how the generated orders are equal to the demand, with a bullwhip equal to 1 as a result. Although we do not discuss the net stock amplification in this section, it is worthwhile to check that number as well.

If the students have played the beer game, the instructor could ask students why they did not play the game in accordance with mean demand forecasting, or in other words, why we observe variance amplification. If the students have not played the beer game, then the instructor could ask why this policy would not work in the real world. The answer is that decision makers do not know the demand (over the lead time), and consequently they forecast demand and constantly adjust the order-up-to levels. If demand is not i.i.d. but is correlated or nonstationary, it is preferable to use the knowledge of the current demand to forecast the next period's demand. Since the true demand distribution is not directly observed (only the actual demand values are observed), many inventory theory researchers suggest the use of adaptive inventory control mechanisms. This is also how many students play the beer game. Unfortunately, these adjustments increase bullwhip. We now discuss some possible adjustments that are frequently used.

3.2. Demand Signal Processing

Lee et al. (1997a) use "demand signal processing" to refer to the use of past demand information to update demand forecasts, resulting in *adaptive* order-up-to levels. If a retailer experiences a demand surge, it will be interpreted as a signal of high future demand; the demand forecast will be adjusted and a larger order will be placed. In other words, the order-up-to level is adjusted based on the demand signal, possibly as follows:

$$S_t = S_{t-1} + \chi(D_t - D_{t-1}),$$

which results in the following order size:

$$O_t = O_{t-1} + \chi(D_t - D_{t-1}), \tag{4}$$

where $\chi \in [0, 1]$ is a constant *signaling factor*. If $\chi = 1$, the order quantities are fully adjusted by the increase (decrease) in demand from period to period.

This ordering policy can be explained to the students as follows (Cachon and Terwiesch 2006). An increase in demand could signal that demand has shifted, suggesting that the product's actual expected demand is higher than previously thought. Then, the retailer should increase his order quantity to cover additional future demand; otherwise, he will quickly stock out. In other words, it is rational for a retailer to increase his order quantity when faced with unusually high demand. These reactions by the retailer, however, contribute to the bullwhip effect. Suppose the retailer's high demand observation occurred merely due to random fluctuation. As a result, future demand will not be higher than expected even though the retailer reacted to this information by ordering more inventory. Hence, the retailer will need to reduce future orders so that the excess inventory just purchased can be drawn down. Ordering more than needed now and less than needed later implies the retailer's orders are more volatile than the retailer's demand, which is the bullwhip effect.

If we select "demand signal processing" in our spreadsheet (in the "Define a demand forecasting technique" window), we immediately observe demand amplification. If we set $\chi = 1$, bullwhip increases to approximately 5. If we react less to changes in demand (e.g., by setting $\chi = 0.2$), then the bullwhip effect remains, but it is reduced to 1.48. Observe that the switching costs also increase together with the bullwhip measure.

3.3. Moving Average Forecast

When the retailer does not know the true demand process, he can use simple methods to forecast demand, such as the moving average or exponential smoothing technique. With these methods, future demand forecasts are continuously updated using new demand realizations. Sometimes students keep track of historical demand data to forecast future demand when they play the beer game. When one adjusts the demand forecasts every period, the order-up-to level becomes *adaptive* (see Equation (2)). The computerized beer game developed by Simchi-Levi et al. (2003) offers the players different replenishment policy options. One option is an adaptive order-up-to policy based on a moving average forecast of demand.

The *moving average* forecast (MA) takes the average of the observed demand in the previous T_m periods. The forecast over the lead time demand is obtained by multiplying the next period's demand forecast by the lead time L, $\widehat{D}_t^L = L\widehat{D}_t$, which determines the order-up-to level in Equation (2).

By selecting the "moving average" forecasting technique in our spreadsheet models, we observe the impact of this forecast method on the order behavior. Assuming i.i.d. demand, 1-week periods, and a 2-week lead time, bullwhip equals 3.63 for $T_m = 4$ weeks. With $T_m = 52$ weeks (one year), we obtain a much smaller bullwhip of 1.12 and we approach the chase sales policy. Indeed, the more data we use from the past, the closer our forecast will be to the average demand and our results will coincide with mean demand forecasting.

The spreadsheets also allow us to illustrate the effect of lead time on bullwhip. When one doubles, the physical lead time to 4 weeks, for example, bullwhip increases to 6.63 with $T_m = 4$ weeks. We observe the same dynamics when demand is correlated (AR demand process). Note that the magnitude of bullwhip is impacted by the specific demand structure, but the dynamics when we start forecasting are the same, regardless of the correlative structure of the demand process. We find that bullwhip is always greater than 1 for all values of the first-order autocorrelation coefficient and the lead time. This result is worth stressing in class: no matter what the lead time is or how the demand is generated, the bullwhip is always present.

3.4. Exponential Smoothing Forecast

Exponential smoothing (ES) is another forecasting technique, in which the next period's demand forecast is adjusted with a fraction (α) of this period's forecast error. Analogously to the moving average forecasting method, we multiply the next period's demand forecast by the lead time L to obtain a lead time demand forecast.

The impact of this forecasting method can be illustrated with the spreadsheets. When demand is i.i.d. and $T_p = 2$, a forecast factor $\alpha = 0.4$ generates bullwhip $= 5.20$. A higher value of the forecast parameter α increases bullwhip, because more weight is given to the most recent observation in computing the forecast. Similar to the MA forecast, a longer lead time results in higher bullwhip.

3.5. Minimum Mean Squared Error Forecast

Finally we consider the *minimum mean squared error* (MMSE) forecasting method, which is mathematically more complex than the previous methods. With this forecasting technique, we explicitly exploit the underlying nature of the demand process to predict future demand, i.e., we explicitly take into account whether demand is an i.i.d. or an autoregressive and/or moving average process (Box and Jenkins 1976). To calculate the lead time demand forecast, we do not simply multiply the next period's forecast with the lead time, but instead explicitly forecast the demand of L periods ahead. We refer to the instructor manual for the detailed math.

Because the MMSE method minimizes the variance of the forecasting error among all linear forecasting methods, it leads to the lowest average inventory-related cost among the three forecasting approaches (MA, ES, and MMSE) for a stationary demand process (Zhang 2004). It explicitly takes the demand structure into account (e.g., i.i.d. or first-order autoregressive), unlike the MA and ES techniques. It assumes, however, that the demand process parameters are known or that an infinite amount of demand data is available to estimate these parameters accurately (Hosoda 2005).

We illustrate the impact of this forecasting method with our spreadsheets, assuming, as before, that

$T_p = 2$. The results obtained are different from the previous results. In this case, when demand is negatively correlated, there is no bullwhip effect. When, for instance, the first order autoregressive parameter is $\rho = -0.5$, we obtain bullwhip $= 0.30$, meaning that order variability is dampened compared to customer demand, not amplified. Alwan et al. (2003) provide a theoretical explanation. When $\rho = 0.5$, we obtain bullwhip $= 2.64$, indicating that the bullwhip effect is present for positively correlated demand. When $\rho = 0$, the demand process is i.i.d. and the MMSE forecast equals the mean demand forecast, resulting in bullwhip $= 1$. As before, we observe that longer lead times result in higher bullwhip.

3.6. Insights for Classroom Purposes

We have contrasted the use of five different forecasting methods with the standard order-up-to policy for both i.i.d. and autoregressive demand. The findings indicate that different forecasting methods lead to different bullwhip values. Bullwhip also varies with the lead time and the demand process.

The spreadsheet helps the student to evaluate the impact of forecasting on the variability of the material flow. In class, we advise to start with forecasting demand by its long-term average, in which case there is no bullwhip effect. The instructor may then ask how realistic this policy is. If students do not immediately suggest *adaptive* forecasting, the instructor can ask what they would do if the demand doubles from one period to the next. Next, the instructor can show how demand signal processing adjusts the order-up-to level every period and why this increases bullwhip. He tells them that a way to process demand signals is to use forecasting methods, such as the simple exponential smoothing or moving average technique. The students should observe that with these methods, the order-up-to policy always results in bullwhip, independent of the demand process. The impact of lead times can also be investigated.

Finally the instructor can discuss the MMSE forecasting technique, which takes the nature of the demand process explicitly into account. This method is the winner among the forecast methods, because it chases sales when demand is i.i.d. and it dampens the order variability when demand is negatively correlated. Moreover, it minimizes the variance of the forecasting error among all linear forecasting methods, and therefore it leads to the lowest inventory costs. Nevertheless, the students should be aware that this forecast method requires an elaborate study (and more data) to estimate the parameters of the demand process, is generally more complex to calculate, and therefore (unfortunately) is used less frequently for practical purposes.

To wrap up, the instructor can point out that improper forecasting may have a devastating impact on the bullwhip effect. As a consequence, inventory, and production switching costs may increase significantly. This observation puts forecasting in a totally different perspective. A vivid discussion on proper use of forecasting and demand management techniques will arise.

4. Impact of Bullwhip Reduction on Customer Service

We have illustrated that the bullwhip effect may arise when using the standard order-up-to policy with traditional forecasting methods. In this section we introduce a smoothed order-up-to policy that avoids variance amplification.

Smoothing models have a long tradition. A smoothing policy is justified when production (ordering) and inventory costs are convex (e.g., quadratic costs) or when production switching is costly. Generally, there are one or two students who suggest smoothing the order pattern when searching for solutions to cope with the bullwhip effect. It often occurs that students who have played the beer game before, do not want to fall into the bullwhip "trap" again, and they keep their orders constant. To their own surprise, their inventory costs turn out not to be lower at all. In the debriefing of the game, it is therefore worthwhile to elaborate on smoothing strategies.

The smoothed order-up-to policy described in this section dampens order variability. Make clear to the students that this policy is a heuristic; optimality is not claimed. Finding the optimal policy is complicated (see Sobel 1969). Modigliani and Hohn (1955) offer another well-known, discrete-time smoothing policy.

4.1. Smoothed Order-Up-To Policy

We present a generalized order-up-to policy with the intention of dampening the order variability. It can be easily derived from the standard order-up-to policy. Substituting Equation (2) into Equation (1), we obtain

$$\begin{aligned} O_t &= \text{order-up-to level} - \text{inventory position} \\ &= \widehat{D}_t^L + \mathrm{SS} - \mathrm{IP}_t = L\widehat{D}_t + \mathrm{SS} - \mathrm{IP}_t \\ &= (T_p + 1)\widehat{D}_t + \mathrm{SS} - \mathrm{IP}_t \\ &= \widehat{D}_t + [T_p\widehat{D}_t + \mathrm{SS} - \mathrm{IP}_t], \end{aligned} \tag{5}$$

where $T_p\widehat{D}_t + \mathrm{SS}$ can be seen as the *desired* inventory position DIP, which is the sum of the desired pipeline stock $T_p\widehat{D}_t$ and the desired net stock or safety stock SS. IP_t denotes the inventory position at the end of period t. We refer to the difference between the desired and actual inventory position $[\mathrm{DIP} - \mathrm{IP}_t]$ as the *inventory deficit*.

Introducing a proportional controller parameter β for the inventory deficit results in the following *smoothed* order-up-to policy:

$$O_t = \widehat{D}_t + \beta[\text{DIP} - \text{IP}_t], \qquad (6)$$

with $0 < \beta < 2$. Forrester (1961) refers to $1/\beta$ as the "adjustment time." When $\beta < 1$, the user explicitly acknowledges that the deficit recovery should be spread out over time, whereas $\beta > 1$ implies an overreaction to the inventory deficit. Hence, when $\beta < 1$, the inventory deficit is only partially recovered during the next ordering period. This fractional adjustment is second nature to control engineers.

We developed a spreadsheet simulation of this smoothed inventory policy (in the "*smoothed OUT*" worksheet). This simulation is similar to the "*standard OUT*" worksheet, but with one important modification. We additionally input a value for the smoothing parameter β. Control engineers prefer to use the inverse of β, namely, $T_i = 1/\beta$, and therefore we also compute the T_i parameter in the input section.

In Figure 2 we illustrate the smoothing impact of this order policy. It shows the pattern when $\beta = 0.5$, demand is i.i.d. and forecasted with its long-term average. The fractional controller indeed has a dampened or "peak-shaving" impact on the order pattern; the resulting bullwhip is 0.33.

4.2. Trade-Off Between Bullwhip and Inventory Variance

So far we have concentrated on the order variance. Smoothing the order pattern may indeed reduce bullwhip and its corresponding production switching costs. This is, however, only one side of the coin. When students smooth the order pattern in the beer game, they do not necessarily obtain lower inventory costs. In developing a replenishment rule one has to consider the impact on the inventory variance as well, because that variance has an immediate effect on customer service: The higher the variance, the more stock will be needed to maintain customer service at the target level. We therefore measure NSAmp, which equals the ratio of the inventory variance over the demand variance. Net stock variance (let alone variance amplification) is not a common supply chain measure, but we need it to calculate the fill rate, which is a popular customer service measure (see Disney et al. 2006).

Hence, we take into consideration the two following factors: on the one hand, the bullwhip effect, which is related to order variability and switching costs; on the other hand, NSAmp, which is related to investment in inventories and customer service.

Intuitively, we expect smooth ordering patterns will result in higher inventory fluctuations because the inventory buffer absorbs the demand fluctuations, resulting in a lower fill rate. This can be illustrated with the spreadsheets. Suppose we assume i.i.d. demand, mean demand forecasting, and $T_p = 2$. A chase sales strategy with $\beta = 1$ results in bullwhip = 1 and NSAmp = 3. Smoothing with $\beta = 0.5$ reduces bullwhip to 0.33 and decreases switching costs but increases NSAmp to 3.33 and increases inventory and backlogging costs. We can smooth the order pattern, but pay the price of higher inventory fluctuations and more inventory and backlogging costs.

These observations illustrate a trade-off between the bullwhip effect and customer service (as measured by net stock variance amplification). The question we should ask is how much production rates can be smoothed to minimize production adaptation costs without increasing inventory costs too much.

Disney et al. (2004) show that when demand is i.i.d. and we forecast demand with its mean, then $\beta = 0.618$ minimizes the sum of bullwhip and NSAmp, which can be seen as a "best of both worlds" solution. This remarkable result is the "golden section," also known as the golden mean, golden ratio, or divine proportion. By adding up the bullwhip effect metric and the NSAmp metric, we assume that both factors are equally important. In the real world, companies may apply weights to bullwhip-related costs and customer service-related costs. In this case the shape of the total cost curve may be different, and the optimal smoothing parameter may no longer be "golden."

Figure 2 Generated Order Pattern When $\beta = 0.5$

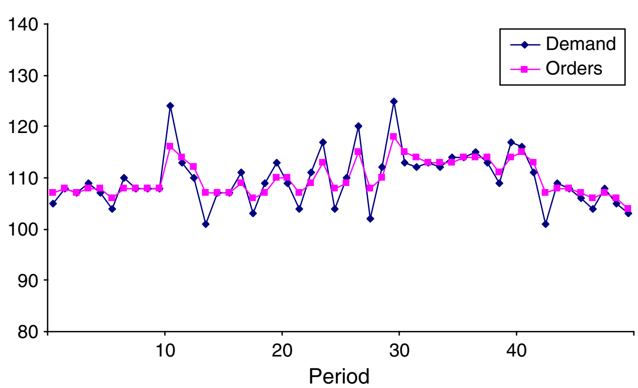

4.3. Win-Win Solutions for Some Demand Patterns

We demonstrated that bullwhip can be reduced by ordering a fraction of the inventory deficit, rather than recovering the entire deficit in one time period. When demand is i.i.d., order smoothing comes at a price: to guarantee the same fill rate, more investment in safety stock is required due to increased inventory variance. Disney et al. (2006) show that it is possible to reduce bullwhip *and* reduce inventory variance together while maintaining customer service. This is a true win-win situation resulting from the smoothing policy. However, this cannot be achieved in all cases, as it depends on the demand pattern.

Consider a stochastic demand pattern with autoregressive and moving average (ARMA) components of order 1, with ρ the correlation coefficient and δ the moving average coefficient (Box and Jenkins 1976). Depending on the specific values of ρ and δ, inventory variance can be reduced by smoothing the demand signal ($\beta < 1$). In other words, bullwhip can be reduced while reducing net stock variance. In other cases, lower inventory variability is achieved by overreacting to the ARMA signal (i.e., $\beta > 1$). In that case, variance amplification leads to lower inventory costs compared to the chase sales policy ($\beta = 1$).

Although the win-win issue is highly specialized (and can be skipped in class), these situations can be easily illustrated with the spreadsheets if an instructor desires to do so. For instance, suppose that $\rho = 0.5$, $\delta = 1.8$ and we forecast demand with its long-term average. Then, a chase sales strategy ($\beta = 1$) results in bullwhip $= 1$ and NSAmp $= 6.73$. A value of $\beta = 1.8$ increases bullwhip to 1.33 but decreases NSAmp to 5.5 (observe that smoothing with $\beta = 0.5$ decreases Bullwhip to 0.66 but increases NSAmp to 9.13). Hence, in this case lower inventory variability is achieved. When we consider another example where demand is characterized by $\rho = 0.25$ and $\delta = 0.25$, a chase sales strategy ($\beta = 1$) results in NSAmp $=$ 1.46. Smoothing with $\beta = 0.5$ decreases the inventory variability to 1.15. Inventory variance is, in this case, reduced by smoothing the demand signal, which is a win-win solution.

4.4. The Smoothed Order-Up-To Policy with Demand Forecasting

The order-up-to policy described by Equation (6) allows one to dampen order variability. Indeed, when an i.i.d. demand is forecasted with its long-term average, it is shown that for $0 < \beta < 1$ we generate a smooth replenishment pattern (dampening order variability) and for $1 < \beta < 2$ we create bullwhip (variance amplification).

However, when the smoothing rule is applied and demand is forecasted with, e.g., the moving average or exponential smoothing technique, a feedback parameter $\beta < 1$ does not always dampen order variability. For instance, when demand is i.i.d. and forecasted with exponential smoothing and a forecast parameter $\alpha = 0.5$, $\beta = 0.5$ results in bullwhip $= 2.41$. Hence the bullwhip effect is present, although the feedback parameter β is smaller than 1. We need to reduce β down to 0.2 to obtain a smooth order pattern with bullwhip less than one when using this particular forecast method. In other words, improper use of forecasting techniques may destroy the smoothing effect of the smoothed order-up-to policy.

These results are generally very complex and not always available in the literature. Using spreadsheets, one can go beyond the existing analytical results and conduct experiments to obtain insights into this complicated issue. An overview of the available analytical results in the literature is provided in the appendix of the instructor manual.

4.5. Insights for Classroom Purposes

When production is inflexible and changes in production levels are costly, standard order-up-to policies with forecasting mechanisms may not perform well. Because of the huge expenses, it may be important to avoid variance amplification or even to reduce variability of customer demand. Starting from the standard order-up-to policy, we have illustrated how to derive the smoothed order-up-to decision rule. The crucial difference compared to standard order-up-to policies is that one adjusts for only a fraction of the inventory deficit.

In using the smoothed order-up-to policy, the instructor should emphasize two aspects: the ordering behavior (as measured by the bullwhip effect), and the impact on the net stock (as measured by the net stock amplification). The insights are clearest when demand is forecasted with its long-term average and demand is i.i.d. In that case, bullwhip reduction comes at a price. In order to guarantee the same fill rate, a larger safety stock is required. The instructor can ask the students to evaluate the impact of different values of β on inventory and switching costs.

The instructor can then point out that the demand parameters impact the ordering behavior. For ARMA(1, 1) demand patterns, there are four possible scenarios, which we describe by comparing to the standard order-up-to policy: (1) *win-win*, where we remove bullwhip and reduce inventory; (2) *win-lose*, where bullwhip is removed at the expense of holding extra inventory; (3) *lose-win*, where bullwhip is endured because it results in less inventory; and (4) *lose-lose*, where both bullwhip and inventory are excessive. These scenarios depend on the statistical properties of the demand process. The exact conditions that result in the different scenarios go far beyond the scope of introductory courses, but we advise the students to experiment with the parameters to search for such scenarios. Generally, students are able to find two or three of the four scenarios.

When demand is forecasted using the exponential smoothing or moving average method, the results are much more complex. The instructor can mention that in this case, a feedback parameter $\beta < 1$ does not necessarily dampen order variability. Using the spreadsheet, students can experiment with order smoothing and forecasting and evaluate the impact of various replenishment strategies on order and inventory variability.

5. Download Information

The following files are available from the ITE website. BullwhipExplorer.xls contains two simulation models in the "*standard OUT*" and "*smoothed OUT*" worksheets, referring to the standard order-up-to policy and the smoothed order-up-to policy.

InstructorManual.doc elaborates on the mathematics behind the input section, where the user selects the parameters of the model, and the simulation table, where the user can track the calculations of how orders are generated. In addition, a summary includes the analytical results available in the literature.

StudentManual.doc provides step-by-step instructions for how to simulate using the spreadsheets.

Electronic Companion

An electronic companion to this paper is available at http://ite.pubs.informs.org/.

Acknowledgments

The spreadsheets have evolved over time, with help from Jeroen Dejonckheere, Freek Janssen, Griet Van Reeth, and Niki Goyvaerts. This research contribution is supported by contract Grants G.0051.03 and G.0547.09 from the Research Programme of the Fund for Scientific Research–Flanders (Belgium) (FWO-Vlaanderen). We are grateful to the associate editor and the two anonymous referees for their numerous suggestions to improve the readability of this paper.

References

Alwan, L. C., J. J. Liu, D. Yao. 2003. Stochastic characterization of upstream demand processes in a supply chain. *IIE Trans.* **35**(3) 207–219.

Box, G. E., G. M. Jenkins. 1976. *Time Series Analysis Forecasting and Control*. Holden-Day, San Francisco.

Cachon, G., M. Fisher. 1997. Campbell Soup's continuous product replenishment program: Evaluation and enhanced decision rules. *Production Oper. Management* **6** 266–276.

Cachon, G., C. Terwiesch. 2006. *Matching Supply with Demand: An Introduction to Operations Management*, international ed. McGraw-Hill, New York.

Chen, F., R. Samroengraja. 2000. The stationary beer game. *Production Oper. Management* **9**(1) 19–30.

Chen, F., J. Ryan, D. Simchi-Levi. 2000. The impact of exponential smoothing forecast on the bullwhip effect. *Naval Res. Logist.* **47**(4) 271–286.

Dejonckheere, J., S. M. Disney, M. R. Lambrecht, D. R. Towill. 2003. Measuring and avoiding the bullwhip effect: A control theoretic approach. *Eur. J. Oper. Res.* **147**(3) 567–590.

Disney, S. M., M. Lambrecht. 2008. On replenishment rules, forecasting and the bullwhip effect in supply chains. *Found. Trends Tech. Inform. Oper. Management* **2**(1) 1–80.

Disney, S. M., D. R. Towill, W. Van De Velde. 2004. Variance amplification and the golden ratio in production and inventory control. *Internat. J. Production Econom.* **90** 295–309.

Disney, S. M., I. Farasyn, M. Lambrecht, D. R. Towill, W. Van De Velde. 2006. Taming the bullwhip effect whilst watching customer service in a single supply chain echelon. *Eur. J. Oper. Res.* **173**(1) 151–172.

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

Hammond, J. H. 1994. Barilla SpA (A). HBS Case 9-694-046, Harvard Business School, Boston.

Hosoda, T. 2005. The principles governing the dynamics of supply chains. Ph.D. thesis, Cardiff University, Cardiff, UK.

Jacobs, F. R. 2000. Playing the beer distribution game over the Internet. *Production Oper. Management* **9**(1) 31–39.

Lee, H. L., V. Padmanabhan, S. Whang. 1997a. The bullwhip effect in supply chains. *Sloan Management Rev.* **38**(3) 93–102.

Lee, H. L., V. Padmanabhan, S. Whang. 1997b. Information distortion in a supply chain: The bullwhip effect. *Management Sci.* **43**(4) 543–558.

Machuca, J. A. D., R. P. Barajas. 1997. A computerized network version of the beer game via the Internet. *Systems Dynamics Rev.* **13**(4) 323–340.

Modigliani, F., F. Hohn. 1955. Production planning over time. *Econometrica* **23**(1) 46–66.

Munson, C. L., J. Hu, M. J. Rosenblatt. 2003. Teaching the costs of uncoordinated supply chains. *Interfaces* **33**(3) 24–39.

Simchi-Levi, D., P. Kaminsky, E. Simchi-Levi. 2003. *Designing and Managing the Supply Chain*, 2nd ed. Irwin/McGraw-Hill, New York.

Sobel, M. 1969. Production smoothing with stochastic demand infinite horizon case. *Management Sci.* **16**(3) 195–207.

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

Watson, N. H., Y.-S. Zheng. 2008. Over-reaction to demand changes due to subjective and quantitative forecasting. Working Paper 09-021, Harvard Business School, Boston.

Wood, S. C. 2007. Online games to teach operations. *INFORMS Trans. Ed.* **8**(1). http://ite.pubs.informs.org/.

Zhang, X. 2004. The impact of forecasting methods on the bullwhip effect. *Internat. J. Production Econom.* **88** 15–27.

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