Measuring the bullwhip effect under a generalized demand process

Marchena, Marlene
Pontifical Catholic University of Rio de Janeiro, Department of Electrical engineering
Rua Marquês de São Vicente 225 Gávea. Rio de Janeiro cep: 22451-900, Brazil
E-mail: marchenamarlene@gmail.com

1 Introduction

In recent years, companies in various industries have been able to improve significantly their inventory management processes through integration of information technology into their forecasting and replenishment systems, and by sharing demand-related information with their supply chain partners (Aviv, 2003). However, despite the benefits resulting of the implementation of the above practices, inefficiencies still persist and are reflected in related costs.
The bullwhip effect, defined as the increase in variability along the supply chain, is a frequent and expensive phenomenon pointed out as a key driver of inefficiencies associated with SCM. It distorts the demand signals which cause instability in the supply chain and increase the cost of supplying end-customer demand.
Forrester (1958) was the first to popularize this phenomenon. Inspired by Forrester's work several researchers have studied the bullwhip effect. Sterman (1989) use the Beer Game, the most popular simulation of a simple production and distribution system developed at MIT, to demonstrated that the bullwhip effect is a significant problem with important managerial consequences. It results in unnecessary costs in supply chains such as inefficient use of production, distribution and storage capacity, recruitment and training costs, increased inventory and poor customer service levels (Metters,1997 and Lee et al. 1997b). Lee et al. (1997a,b) identified four main causes of the bullwhip effect, i.e., demand forecasting, order batching, price fluctuation and supply shortages.
Since the bullwhip effect is an expensive occurrence many attention have been devoted to it in recent years. Forecasting methods, demand processes and information sharing are the main aspects studied of this phenomenon.
Using a first-order autoregressive demand process, Chen et al. (2000a,b) investigate the impact of the simple moving average and exponential smoothing forecasts on the bullwhip effect for a simple, two-stage supply chain with one supplier and one retailer. Zhang (2004a) investigates the impact of different forecasting methods on the bullwhip effect for a simple inventory system with a first-order autoregressive demand process. By quantifying the bullwhip effect they show the impact of forecasting methods on bullwhip effect.
On regarding the demand process, a variety of time-series demand models have appeared in the literature of inventory control and SCM. By far, the first order autoregressive process,AR(1), is the most frequently adopted demand model to study the bullwhip effect (Lee et al., 1997a,b; Chen et al., 2000a,b; Zhang, 2004). Recent works use more realistic demand models such as the autoregressive moving-average demand model (Box and Jenkins, 1970). Luong et al. (2007) use an AR(2) and a general AR(p) models; Duc et al. (2008) use a $ARMA(p, q)$ model. In all these models an analytical derivation of the bullwhip effect measure is presented.
Zhang (2004b) uses an $ARMA$ model and Gilbert (2005) uses an $ARIMA$ model to study the demand evolution in supply chains. They show that the order history preserves the autoregressive structure of the demand, in both studies closed form expressions are given. In addition Zhang's work give a simple algorithm to quantifying the bullwhip effect. Inspired for this result we study the theoretical and practical applications of use this algorithm in order to measure the bullwhip effect for different demand process.

We have demonstrated that the use of a generalized form of the bullwhip effect measure makes possible to get accurate estimations of bullwhip effect. We point out that no approximation is required. Moreover, we show that for certain types of demand processes the use of MMSE considered in the model leads to significant reduction in the safety stock level. All these observations highlight the potential economic benefits resulting from the use of time series analysis.
The structure of our paper is as follows. The next section presents the inventory model. Section three presents a general $ARMA(p, q)$ case with the $ARMA(1, 1)$ case as particular case. Next, section four presents the economic implications of the bullwhip effect. The final section summarizes the main finding of the research.

2 Inventory model

In this paper we consider a simple supply chain model for a single item and an order-up-to inventory policy. We assume that there is a fixed lead time, L, between the time an order is placed and the time that it is received, shortages are back-ordered and no fixed ordering cost exists. The sequence of events during a replenishment cycle for each fixed period t is as follows: first, the retailer receives orders made L periods ago, second, the demand, d_t, is observed and satisfied, third, the retailer observes the new inventory level and finally places an order on the manufacturer.
Let O_t be the order quantity in period t and y_t be the inventory position after placing the order in period t. So following this replenishment cycle the order quantity at the end of period t can be written as:

(1) $$O_t = y_t - y_{t-1} + d_t$$

It is well know that, under the assumption of normal demand distribution and in the absence of fixed ordering cost, the order-up-to policy is optimal and the order-up-to-level, y_t, is expressed as:

$$y_t = D_t^L + z\sigma_t^L$$

where $D_t^L = \sum_{\tau=1}^{L} d_{t+\tau}$ is the total demand during lead-time and $z\sigma_t^L$ is the safety stock. The safety stock is composed by the safety factor, z, which is a fixed constant chosen to meet a required service level, and the lead-time standard deviation, σ_t^L. Note that if these values are known, the order-up-to level in any period is constant and, consequently, the order will be equal to the last observed demand, therefore, there is no bullwhip effect. However, these values are in general unknown and the retailer must estimate y_t from the observed demand as:

(2) $$y_t = \hat{D}_t^L + z\hat{\sigma}_t^L$$

Let $\mathcal{F}_t$ be the information set which represents all the information available until period t, $\mathcal{F}_t = \{d_t, d_{t-1}, ...\}$, $\hat{D}_t^L = \sum_{\tau=1}^{L} \hat{d}_t(\tau)$ is an estimate of the mean demand over L periods based on the information set, and $\hat{\sigma}_t^L = \sqrt{Var(D_t^L - \hat{D}_t^L)}$ is an estimate of the standard deviation of L periods forecast error.
The calculation of y_t based on forecasting values is one of the main causes for the variability increase along the supply chain or, in other words, the bullwhip effect. Thus, we need some measures of performance for this effect. The definition commonly used for the bullwhip effect in the literature of supply chain management is:

(3) $$M = \frac{Var(O_t)}{Var(d_t)}$$

i.e., the ratio between the variance of orders and the variance of the demand. Now, in order to get the bullwhip effect we combine equation (2) and equation (1) to rewrite the order quantity as:

$$O_t = (\hat{D}_t^L - \hat{D}_{t-1}^L) + z(\hat{\sigma}_t^L - \hat{\sigma}_{t-1}^L) + d_t \quad (4)$$

then, taken the variance in both side of the above equation we can reach to quantify this effect.
A bullwhip measure equal to one means that there is no variance amplification, larger than one means that there is a variance amplification or the bullwhip effect is present, if it is smaller than one means that the orders are less variable or smoothed if compared with the demand. The Minimum Mean Squared Error (MMSE) forecasting is applied to predict the lead time demand.

3 ARMA(p,q) case

The $MA(\infty)$ representation of a stationary $ARMA(p,q)$ demand process is written as:

$$d_t = \mu_d + \Sigma_{j=0}^{\infty} \psi_j \epsilon_{t-j} \quad (5)$$

where $\mu_d \neq 0$, $\epsilon_t \sim N(0, \sigma_\epsilon^2)$ and the ψ-weights satisfies the homogeneous difference equation that arise from match the coefficients in the follow identity and calculate them recursively,

$$(\psi_0 + \psi_1 z + \psi_2 z^2 + \cdots)(1 - \phi_1 z - \phi_2 z^2 - \cdots - \phi_p z^p) = (1 + \theta_1 z + \theta_2 z^2 + \cdots + \theta_q z^q),$$

where $\psi_j = 0$ for $j < 0$, $\psi_0 = 1$ and $\psi_j = \sum_{k=1}^{p} \phi_k \psi_{j-k} + \theta_j$ for $j \geq 1$. Note here that the homogeneous difference equation is given by

$$\psi_j - \sum_{k=1}^{p} \phi_k \psi_{j-k} = 0, j \geq max(p, q+1) \quad (6)$$

with initial conditions

$$\psi_j - \sum_{k=1}^{j} \phi_k \psi_{j-k} = \theta_j, 0 \leq j \leq max(p, q+1) \quad (7)$$

Then, considering distinct AR roots, the general solution for the ψ-weights can be read off directly as:

$$\psi_j = c_1 z_1^{-j} + \cdots + c_p z_p^{-j} \quad (8)$$

where $z_1, \ldots, z_p$ are the roots of the AR polynomial $\phi(z) = 1 - \phi_1 z - \cdots - \phi_p z^p$. The specific solution will depend on the initial conditions, as seen from equation (7). Note that these initial conditions depend on the choice of the moving average parameters as well as the autoregressive parameters.

Proposition 1 *(Zhang, 2004.) The retailer's future independent $ARMA(p,q)$ demand process can be represented by an $MA(\infty)$ process with respect its errors as in (5). Its order, O_t, to its supplier is given by*

$$O_t = \mu_d + \sum_{j=0}^{L} \psi_j \epsilon_t + \sum_{j=1}^{\infty} \psi_{L+j} \epsilon_{t-j} \quad (9)$$

Proof: See Zhang (2004b).

Proposition 2 *For a stationary ARMA(p,q) demand process, the measure for the bullwhip effect is defined by:*

$$M = 1 + \frac{2\sum_{i=0}^{L}\sum_{j=i+1}^{L}\psi_i\psi_j}{\sum_{j=0}^{\infty}\psi_j^2}, \tag{10}$$

where the $\psi_j = 0$ *for* $j < 0$, $\psi_0 = 1$, *and* $\psi_j = \sum_{k=1}^{p}\phi_k\psi_{j-k} + \theta_j$ *for* $j \geq 1$.

Proof: From Equation (9), we have

$$Var(O_t) = \sigma_\epsilon^2(\sum_{j=0}^{L}\psi_j)^2 + \sigma_\epsilon^2\sum_{j=1}^{\infty}\psi_{L+j}^2 = \sigma_\epsilon^2(\sum_{j=0}^{\infty}\psi_j^2 + 2\sum_{i=0}^{L}\sum_{j=i+1}^{L}\psi_i\psi_j). \tag{11}$$

Since the variance of the demand process can be expressed as:

$$\sigma_d^2 = \sigma_\epsilon^2\sum_{j=0}^{\infty}\psi_j^2, \tag{12}$$

we complete the proof by substituting (11) and (12) in (3).

3.1 ARMA(1,1) case

The stationary ARMA(1,1) demand process is described as follow:

$$d_t = \mu + \phi d_{t-1} + \epsilon_t + \theta\epsilon_{t-1}. \tag{13}$$

Stationarity and invertible conditions impose $|\phi| < 1$ and $|\theta| < 1$. It can be shown that the mean and variance of the demand process are $\mu_d = \frac{\mu}{1-\phi_1}$ and $\sigma_d^2 = \frac{(1+\theta^2+2\phi\theta)\sigma_\epsilon^2}{1-\phi^2}$, respectively.

Proposition 3 *For a stationary ARMA(1,1) demand process the measure for the bullwhip effect is defined by:*

$$M(L,\phi,\theta) = 1 + \frac{2(\phi+\theta)(1-\phi^L)}{(1-\phi)(1+\theta^2+2\phi\theta)}\left[1-\phi^{L+1}+\theta\phi(1-\phi^{L-1})\right]. \tag{14}$$

Proof: Since the AR polynomial associated with (13) is $\phi(z) = 1-\phi z$, and its root, say z_1, is $z_1 = \phi^{-1}$, then the general solution for the ψ-weights can be written directly from equation (8) as $\psi_j = c\phi^j$. From (7) we find that the initial conditions are $\psi_0 = 1$ and $\psi_1 = \phi + \theta$, which combining with the general solution, results in $c = (\phi+\theta)/\phi$. Hence, $\psi_j = (\phi+\theta)\phi^{j-1}$ for $j \geq 1$. Since we know ψ_j, we can rewrite the follow relations as:

$$\begin{aligned}
\sum_{i=0}^{L}\sum_{j=i+1}^{L}\psi_i\psi_j &= \psi_0\sum_{j=1}^{L}\psi_j + \sum_{i=1}^{L}\sum_{j=i+1}^{L}\psi_i\psi_j \\
&= (\phi+\theta)\frac{1-\phi^L}{1-\phi} + \frac{\phi(\phi+\theta)^2(1-\phi^L)(1-\phi^{L-1})}{(1-\phi)(1-\phi^2)} \\
&= \frac{(\phi+\theta)(1-\phi^L)}{(1-\phi)(1-\phi^2)}\left[1-\phi^{L+1}+\theta\phi(1-\phi^{L-1})\right] \qquad (15)
\end{aligned}$$

and

$$\sum_{j=0}^{\infty}\psi_j^2 = \frac{1+\theta^2+2\phi\theta}{1-\phi^2}. \tag{16}$$

We complete the proof substituting (15) and (16) in (10). Using a generalized formula for the variance ratio, we get a similar expression to that obtained by Duc et al. (2008).

Table 1: Bullwhip, SS and SSLT generated by ARMA(0.95,0.4) demand process.*

L	Bullwhip	SS	SSLT
1	1.13711	7.299	1.645
2	1.44321	10.323	4.201
3	1.89270	12.643	7.304
4	2.46294	14.598	10.817
5	3.13393	16.322	14.652
6	3.88802	17.879	18.745
7	4.70970	19.312	23.048
8	5.58531	20.645	27.522
9	6.50289	21.898	32.137
10	7.45199	23.082	36.867

* SCperf(0.95,0.4,L,0.95)

4 Economic implications

An important economic application of the use of time series methods can be seen in the safety stock level, which is the amount of inventory that the retailer needs to keep in order to protect himself against deviations from average demand during lead time.

Let $SS = z\sigma_d\sqrt{L}$ and $SSLT = z\hat{\sigma}_t^L$ be two safety stock measures. The former is traditionally used in some operational research manuals and it is based on the standard deviation of the demand over L periods, the latter is the safety stock as defined in (2) and it is based on the standard deviation of L periods forecast error.

Chen et al. (2000b, pp. 271) pointed out that SSLT will be greater than SS, i.e., using time series analysis, the retailer will hold more safety stock to achieve the same service level. According to the authors this is because SS captures only the uncertainty due to the random error ϵ and SSLT captures this uncertainty plus the uncertainty due to the fact that the mean demand D_t^L is estimated by $\hat{D}_t^L$, in our case using the MMSE forecasting method. We show by numerical experiments that for some special cases $SSLT$ is lower than SS regarding lead-time and service level.

It was verified that for ARMA and AR cases, high values on AR parameters and small values of lead-time result in lower $SSLT$. However, in general, there is a lead-time value for which this situation is reversed. Table 1 shows the safety stock levels SS and SSLT generated by $ARMA(0.95, 0.4)$ demand process and service level equal to 0.95 for ten different values of lead-time, $L = 1, .., 10$. For instance, for $L = 2$ we have $SS = 10.3$ and $SSLT = 4.2$, a difference of 6 units which represents a saving of 59.2% over SS. Note that this difference decreases when the lead-time increases until $L = 6$ where we have SSLT larger than SS.

It is difficult to know for which value of lead-time SSLT becomes larger than SS. In general, it depends on the AR parameters of the demand. For negative values of the AR parameters, it occurs for lower values of lead-time.

Table 1 shows that there is a benefit resulting from the use of SSLT instead of SS as a measure for the safety stock level when regarding the lead-time. This benefit was verified for special demand processes where the AR parameters are high. Moreover, if for those lead-time values where SSLT is

Table 2: SS and SSLT generated by different demand processes

Models	**Service Level**	**L=1**		**L=2**		**L=3**	
	SL	SS	SSLT	SS	SSLT	SS	SSLT
	0.90	5.687	1.282	8.043	3.273	9.850	5.691
	0.91	5.950	1.341	8.414	3.424	10.305	5.954
	0.92	6.235	1.405	8.818	3.588	10.800	6.239
	0.93	6.549	1.476	9.262	3.769	11.343	6.553
	0.94	6.899	1.555	9.757	3.971	11.950	6.904
$ARMA(0.95, 0.4)$	0.95	7.299	1.645	10.323	4.201	12.643	7.304
	0.96	7.769	1.751	10.987	4.471	13.456	7.774
	0.97	8.346	1.881	11.803	4.803	14.456	8.352
	0.98	9.114	2.054	12.889	5.245	15.785	9.120
	0.99	10.323	2.326	14.599	5.941	17.881	10.330

smaller than SS, we consider the service level, it is verified that SSLT is always smaller than SS when the service level increases.

Table 2 presents SSLT and SS generated by the same demand process for $L = 1, 2, 3$ and ten different values of service level, $SL = 0.9, 0.91, ..., 0.99$. Note that when considering the service level, the difference between SS and SSLT increases for larger values of service level differently when lead-time is regarded. For instance, for $L = 1$ and $SL = 0.97$ we have $SS = 8.35$ and $SSLT = 1.88$. There is a difference of 6.47 units which represents a saving of 77.46% over SS.

5 Summary

In this paper we quantify the bullwhip effect using Zhang's result for a stationary $ARMA(p, q)$ demand process which admits an infinite MA representation. It is well known that measuring the bullwhip effect is difficult in practice but we show that using this methodology the calculus of this ratio is simplified if compared with traditional recursive procedures. In some particular cases we obtain explicit formulas for this ratio.

With this theoretical application we present an R implementation for the bullwhip effect. We program `SCperf` function whose output gives numerical results for the bullwhip effect and other supply chain performance variables. Our function explores a variety of demand process scenarios and illustrates how tuning the parameters of the demand result in the bullwhip effect and the safety stock level.

In conclusion, when inventory cost and service level are of primary concern the MMSE forecast should be used since it leads in some cases to lowest safety stock level. Although the MMSE forecasting requires more computational effort, the `SCperf` function implements this method in an easy way. Our findings highlights the potential economic benefits resulting from the use of time series analysis but it depends on the underlying demand process.

REFERENCES

[1] Xailong Zhang, Evolution of ARMA demand in supply chains, Manufacturing and Services Operations Management, Vol 6-2, 2004b, pp:195–198.

[2] Duc, T.T.H., Luong, H.T., Kim, Y.D.,. A measure of the bullwhip effect supply chains with a mixed autoregressive moving average demand process. European Journal of Operational Research. Vol 187, 2008, pp: 243–256.

[3] F. Chen, Z. Drezner, J. Ryan and D. Simchi-Levi, The impact of exponential smoothing forecasts on the bullwhip effect. Naval Research Logistics, Vol 47 - 4, 2000b, pp: 269–286.